Università degli Studi di Udine
Dipartimento Politecnico di Ingegneria e Architettura
Dottorato di Ricerca in Ingegneria Industriale e dell'Informazione

Ph.D. Thesis

Exploiting Sparsity for Efficient Compression and Analysis of ECG and Fetal-ECG Signals

Candidate:
Giulia Da Poian

Supervisor:
Prof. Roberto Rinaldo

External Reviewers:
Prof. Patrizio Campisi
Prof. Gari Clifford

Academic Year 2016

Author's e-mail: dapoian.giulia@spes.uniud.it

Author's address:

Dipartimento Dipartimento Politecnico di Ingegneria e Architettura
Università degli Studi di Udine
Via delle Scienze, 206
33100 Udine
Italia

Abstract

Over the last decade there has been an increasing interest in solutions for the continuous monitoring of health status with wireless, and in particular, wearable devices that provide remote analysis of physiological data. The use of wireless technologies have introduced new problems such as the transmission of a huge amount of data within the constraint of limited battery life devices. The design of an accurate and energy efficient telemonitoring system can be achieved by reducing the amount of data that should be transmitted, which is still a challenging task on devices with both computational and energy constraints. Furthermore, it is not sufficient merely to collect and transmit data, and algorithms that provide real-time analysis are needed.

In this thesis, we address the problems of compression and analysis of physiological data using the emerging frameworks of Compressive Sensing (CS) and sparse signal processing. In particular, we develop new methods and propose specific applications for compression and real-time analysis with a special focus on electrocardiogram (ECG) and fetal electrocardiogram (fECG) signals. Moreover, the proposed frameworks and results could potentially be extended to a much wider class of physiological signals.

To improve the performance of current CS frameworks, we introduce a novel sparsifying dictionary, which, when used in combination with the existing reconstruction algorithms, allows for accurate recovery of ECG signals at high compression ratios. While the CS compression is a low-complex procedure, signal recovery can be computationally expensive, and very often we are only interested in extracting certain information without necessarily needing the full reconstructed signal. This is the case of clinical evaluation based on analyzing beat-to-beat timing variation, which is calculated from the time occurring between two consecutive beats, identified by the R-peak in the ECG signal. Thus, we consider the possibility of avoiding signal recovery and directly performing beat detection in the compressed domain. To this end we propose a new method capable to provide a real-time detection of R-peaks with a limited complexity with respect to typical reconstruction procedures.

Increasing the compression ratio is the main objective, but, due to distortion, the signal quality is an issue that should always be kept under control. In particular, due to distortion induced by the sampling process it is essential to guarantee that all clinically relevant information for a given task is preserved, in order to prevent significant degradation in the performance of any standard or novel clinically relevant algorithm. Thus, in order to assess the effectiveness of the proposed dictionary and beat detector, we verify the impact of CS at different compression ratios on Atrial Fibrillation (AF) detection. We demonstrate the possibility of accurately detecting episodes of atrial fibrillation (AF) directly on the compressed measurements which has enormous potential for extending long term monitoring of transient AF and other episodic phenomena, which requires long term

monitoring and processing of data on energy-constrained devices. Moreover, the proposed dictionary allows to increase the accuracy of AF detection for a given compression ratio with respect to standard method.

We also design a framework for the compression of abdominal fECG and to obtain real time information of the fetal heart rate, providing a suitable solution for real-time, very low power fECG monitoring. Taking advantage of the sparse representation with the proposed dictionary, it is possible to increase the quality of the compressed signals, and, at the same time, perform fetal and maternal beat detection/classification. The detection scheme uses Independent Component Analysis (ICA), which we propose to compute directly in the compressed domain before signal reconstruction. The need for fast and robust reconstruction algorithms inspired us to modify an existing reconstruction algorithm and make it error-tolerant. The proposed method guarantees better immunity against inaccuracy caused by noisy original signals and possibly ill-conditioned reconstruction procedures.

Finally, we compare fECG compression using CS to a standard compression scheme using wavelets in terms of energy consumption, reconstruction quality and, more importantly, performance of fetal heart beat detection on the reconstructed signals. An actual implementation on a commercial device prove the suitability of CS as an ultra-low power compression technique for fECG signals. Indeed, CS allows for significant reductions in energy consumption in the sensor node and the detection performance is comparable to that obtained on original signals for compression ratios of up to 75%.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Motivations

Cardiac dysfunction can be associated with devastating physiological consequences such as cardiovascular disease, and, despite during the past few decades the mortality due to cardiovascular disease has decreased, it still remains the number one cause of death worldwide. Accordingly to the World Health Organization, in 2012 about 17.5 million deaths (31% of about 54.9 million deaths occurring worldwide) were due to cardiovascular disease [98].

By following a healthy lifestyle it is possible to prevent and control many of the risk factors associated with heart disease. However, in order to reduce or eliminate the complications, cardiac disease should be identified in its early or asymptomatic phase. In particular, it is important to detect heart rhythm issues (arrhythmias), which are some of the most common heart disorders. This can be done by analyzing the heart's activity through the electrocardiogram (ECG) signal. To this end, Holter systems are commonly employed to record a patient's ECG trace and monitor arrhythmias. Patients are connected to the device through a series of wires and must carry the device at all times, commonly for 24 or 48 hours. Clearly, the impact on the patient's mobility and everyday activities makes Holter systems not always easily portable for such a long time.

The use of telemonitoring systems, in particular the new concept of Wireless body sensor networks (WBSNs), promises to allow continuous and remote monitoring of physiological signals [92, 78, 10, 96]. Wireless functionality allows the devices to continuously transmit a patient's heart data to a server, and unlike Holter monitors, they may be used for several weeks at a time, overcoming the limits of wired technology that restricts patient's movements and everyday life. Furthermore, telemedicine is particularly useful since it may reduce costs of hospitalization and provide remote analysis in sparsely populated or quarantined regions.

The increasing volume of data generated by the continuous monitoring brings with it new problems. One of the major challenges in this field is the transmission of a large amount of data within the constraint of limited battery life devices. This is particularly important for implantable devices that provide accurate daily transmission of cardiac electrical data. Thus, in the design of telemonitoring systems based on WBSNs one has to consider to allow long-term clinically relevant monitoring while reducing the size of the device, allowing free-living conditions to the patients, without reducing the embedded intelligence and wireless capabilities.

To achieve low-power telemonitoring using WBSNs technologies, the use of Compres-

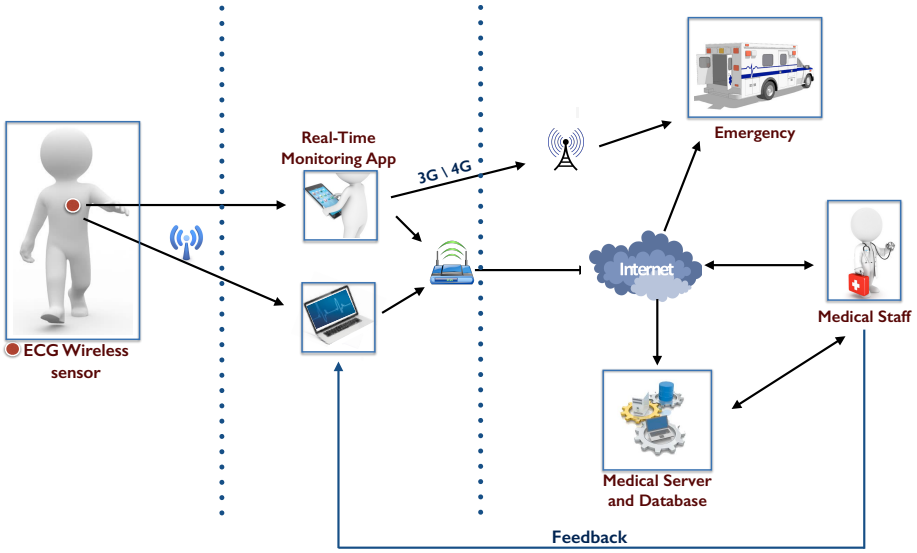

Figure 1.1: WBSN for telemonitoring applications.

sive Sensing (CS) techniques seems to be a promising framework, which may increase the lifetime of sensors reducing the amount of data to transmit and the computational cost of the compression algorithm. Furthermore, the analogue implementation of CS allows to directly acquire a compressed version of the signal, to further reduce the power consumption using sub-Nyquist sampling frequencies. Hardware designs of sub-Nyquist converters that rely on ideas of CS have been also investigated in the literature. However the digital implementation remains the most common and used due to its low computational cost.

The continuous monitoring of patient's health introduces a second issue concerning the huge amount of data generated that need to be processed and analyzed. Some remote monitoring systems are based on event-triggered transmissions, and after event detection on the wearable or implantable device data are transmitted to a mobile patient unit, and then forwarded to a physician. It is then clear that new automatic solutions for data analysis are required. Algorithms should be able to elaborate a huge amount of data giving reliable results, but in some situations it also necessary to develop solutions that work in real-time and are suitable for low-resource devices. Advances in the signal processing allow designing very sophisticated and reliable systems for event detection, both in real-time and off-line situations.

Beyond presenting some solutions that combine tools used to enable low-power compression with real-time analysis of adult the ECG signal, in this work we go further and extend the analysis to non-invasive fetal electrocardiogram (fECG) recordings. Since its introduction, fetal heart rate monitoring was expected to increase the diagnoses of fetal heart diseases, however the outputs of fetal heart rate monitors are often unreliable and difficult to interpret. In the last years, the use of sophisticated signal processing techniques and recording devices have improved the accuracy in fetal heart rate estimation. Non-invasive

fECG can theoretically be performed at every stages of the pregnancy enabling long term and continuous monitoring, but it is currently used for short and intermittent analysis. Continuous long-term fetal monitoring might provide a way to remove the limitations of subjective and intermittent evaluation of fetal wellbeing [21]. However, issues with signal quality, poor battery life and interference have prevented the use of devices for continuous monitoring of fECG, which have not yet become widely used in clinical practice.

1.2 Thesis outline

The goal of this thesis is to present solutions for ECG and fECG signal compression in low-power wireless sensors as well as to provide and analyze frameworks for their analysis, using solutions that arise in the emerging field of compressive sensing and sparse representations.

Initially, we review in Chapter 2 the current state-of-the-art methods for sparse signal representations and compressive sensing. First, we give the basic definitions and notation used in this thesis for sparse representation of signals, and we introduce the theory of CS. Then, we give an overview of some algorithms use for solving ill-posed inverse problems related both to sparse representation and CS, which will be used in the rest of this work.

Chapter 3 studies a novel dictionary used for the sparse representation of the ECG signal, which is the key to address the challenging problem of increasing the compression of ECG signal by using CS. First, we review some basic principles of ECG monitoring and interpretation, describing the model that can be used to describe the electrocardiogram signal. Based on the ECG mathematical model a new over-complete dictionary is introduced. Allowing a sparsest representation of the ECG signal it enables to achieve higher compression ratios within the CS framework. Unlike other dictionaries proposed in literature, the one introduced in this chapter is "universal", since it does not depend on a specific patient and does not require a learning stage.

In Chapter 4 a new algorithm for beat detection that works directly on the compressive sensing measurements of ECG signals is presented. Avoiding the reconstruction procedure, the adoption of this method considerably reduces the detection complexity, making the procedure suitable for low-resource and low-power devices. Without actual reconstruction, it is possible to achieve results comparable with state-of-the-art detection on the reconstructed signals, when conventional sparsifying bases are used in the reconstruction process. Using more sophisticated dictionaries, such as the one proposed in Chapter 3, allows to achieve better detection results at high compression ratios, where the proposed method performance slightly degrades. However, reconstruction with the over-complete dictionary requires considerable resources, so the proposed method can still be preferable. A possibility is to devise schemes where detection is performed in the sensor, while data are further compressed and transmitted if necessary. Note that radio transmission is in general costly in terms of energy consumption.

Motivated by the necessity to keep clinical relevant information after signal compression, Chapter 5 presents a study on the effect of compressive sensing on the performance of atrial fibrillation detection. Indeed, classical metrics based on signal quality might be not necessarily a good indicator of the actual clinical usefulness of a signal after reconstruction. Apart from assessing the performance after signal reconstruction and beat detection with a state-of-the art algorithm, we assess the performance of AF detection after applying the beat detector proposed in Chapter 4 avoiding the signal reconstruction

stage. The results of this study have widespread implications for further application of CS in a clinical context.

We move to the problem of compression and analysis to the fetal electrocardiogram in Chapter 6. The dictionary introduced for the compression of ECG signal in Chapter 3 is extended to the fECG signal in order to overcome the limitations of the classical CS framework for the fECG signals. Fetal electrocardiogram signals acquired from the mother's abdomen are typically characterized by a very low SNR. In fact, signals recorded by this method are always a mixture of noises generated, for instance, by fetal brain activity, myographic signals (both from the mother and the fetus), movement artifacts and maternal ECG. Thus, the CS frameworks that use conventional sparsifying bases, such as the Wavelet basis, generally fail to recover the signal with adequate accuracy, making the reconstructed electrocardiogram non suitable for diagnostic purposes. Moreover, in the development of the dictionary we consider the possibility to separate the maternal and fetal components during the reconstruction procedure, allowing fetal beat detection in real-time. Although the use of the dictionary allows to separate the fetal and maternal QRS complexes in the majority of the ECG traces, some further processing might be required to achieve more reliable results. To this aim, we show that it is possible to apply Independent Component Analysis (ICA) directly in the compressed domain. Recovering the independent components can increase the performance of the method. Even if ICA is applied in the compressed domain, it is possible to virtually reconstruct the same independent components that one should obtain from the original mixtures. The novel proposed dictionary is used during reconstruction as a sparsifying domain. This method can be used for in home real-time monitoring of fetal heart rate, using small low-power wearable sensors. While a first analysis can be provided by the proposed method, the signals can be further evaluated by a physician, since reconstruction preserves the clinical quality of the signals.

In Chapter 7 a robust version of the reconstruction method used in Chapter 6 is introduced to overcome some limitations due to the nature of the dictionary and sensing matrix. The improved reconstruction algorithm, based on the Smoothed-l0 (SL0) algorithm, might be suitable in all the situations in which the matrix used in the CS reconstruction process is ill-conditioned and the signals or measurements are corrupted by noise. We also show that the use of a sparse sensing matrix instead of a classical random sensing matrix, with i.i.d. components, leads to the same detection/reconstruction results. This chapter is concluded by illustrating how it is possible to take advantage of the sparse representation obtained from the dictionary used in Chapter 6 to increase the performance of classical fetal beat detection algorithms, based on ICA working on the raw fECG signals. Indeed, combining ICA and the sparse representation, it is possible to perform a fast real-time analysis of fECG.

In Chapter 8 a sensor architecture using the CS paradigm is compared to a standard compression scheme using wavelets. Unlike other works in the literature, we consider, as a figure of merit, the accuracy of fetal beat detection after reconstruction, and compare the results of different compression/transmission/reconstruction procedures as a function of the sensor energy consumption. Our results show that a properly designed CS paradigm, using the proposed over-complete dictionary at the receiver, can preserve relevant signal information and provide a detection performance comparable to that obtained on original signals for compression ratios up to about 75%. Moreover, CS scheme has similar reconstruction quality than one based on wavelets, with the advantage of a low energy

implementation in the sensor.

To conclude the thesis, in Chapter 9 we summarize our findings and we analyze possible future directions and open questions.

1.3 Summary of Contributions

The main contributions of this thesis are summarized below.

- We propose a novel dictionary that allows a better rate-distortion performance than conventional schemes based on state of the art wavelet bases. The dictionary is based on a Gaussian model of the ECG signal. The proposed sparsifying dictionary is tested and validated for the reconstruction of real ECG signals from compressed measurements, and it is shown that it provides significant gain with respect to other conventional sparsiyfing transforms.

- We introduce a accurate novel method that allows the detection of beat directly on compressive sensed ECG signals, avoiding the reconstruction procedure.

- We assess the effect of compressive sensing at different compression ratios on atrial fibrillation detection. This study evaluates the CS framework in view of an actual implementation in a clinical context and shows the proposed method is significantly superior.

- Based on the proposed dictionary, we extend the work to non-invasive fetal electrocardiogram analysis. We show that our dictionary is able to recover the signal from the compressive sensing measurements with sufficient quality. Besides, we show that is possible to implement the Independent Component Analysis (ICA) directly in the compressed domain, and we propose a scheme that is able to effectively detect fetal beats during reconstruction.

- We propose a new regularized version of the Smoothed-$l0$ reconstruction algorithm, which allows to overcome the problem of noisy signal recovery when the sensing matrix-dictionary system is ill-conditioned.

- We study and compare the energy consumption and beat detection performance of CS and wavelet based compression for fECG signals with an actual implementation on a commercial device and demonstrate significantly superior performance of the techniques proposed in this thesis.

2

Sparse Representation and Compressive Sensing

2.1 Introduction

Over the last few years there has been an increasing interest on alternative signal representations such as *sparse representation* and *sparse approximation* [22, 58, 87]. These make it possible to represent a signal in a compact way, using only a small number of coefficients from a general basis or dictionary. The key concept of *sparsity* is used by several new algorithms developed to achieve state-of-the-art results in a wide range of signal processing applications, including compression, denoising and analysis. Despite many natural signals are not sparse in the acquisition domain, it is in general possible to describe the signal using a sparse representation in a specific dictionary or basis. In particular, for physiological signals, such as the electrocardiogram (ECG) or the fetal electrocardiogram (fECG), it is necessary to develop new dictionaries to allow an efficient sparse representation.

Sparsity is also the basic concept behind Compressive Sensing (CS) [25, 53, 31, 32, 30, 9], a novel paradigm in signal processing that allows to successfully recover certain signals sampled far below the Nyquist frequency. As we will see in this work, the design of a suitable sparsifying dictionary is essential in order to achieve better results in CS of physiological signals.

Besides compression, sparse representations can be successfully applied as an alternative tool for signal analysis and classification, since they allow to separate the components of a signal when they have different shape/morphology [128, 18].

In the current chapter, we start by formulating the sparse representation problem that is fundamental to understand the compressive sensing theory introduced later. To conclude the chapter, we introduce some of the state-of-the-art algorithms that aim to solve the inverse problems arising from sparse representation and CS.

2.2 Sparse Signal Processing

2.2.1 Notation

In this work we use normal letters like s to designate scalar quantities and boldface lowercase letters, $\mathbf{s} \in \mathbb{R}^N$, to indicate N-dimensional real column vectors, each i-th entry indicated using s_i. Matrices are indicated with boldface capitals, like $\mathbf{A} \in \mathbb{R}^{M \times N}$, which represents an $M \times N$ matrix. We indicate the j-th column of the matrix $\mathbf{A}$ with $\mathbf{a}_j \in \mathbb{R}^M$.

We make frequent use of the l_1, l_2 norms, and of the l_0-norm, which is actually a pseudo norm. For $p \in [1, \infty)$ a norm of a vector $\mathbf{x}$ in the N-dimensional Euclidean vector space $\mathbb{R}^N$ is defined as

$$\| \mathbf{x} \|_p = \left(\sum_{i=1}^{N} |x_i|^p \right)^{\frac{1}{p}}. \tag{2.2.1}$$

The standard inner product in $\mathbb{R}^N$ is denoted as

$$\langle x, y \rangle = \mathbf{y}^T \mathbf{x} = \sum_{i=1}^{N} x_i y_i, \tag{2.2.2}$$

thus the l_2-norm is $\| \mathbf{x} \|_2 = \sqrt{\langle x, x \rangle}$.

The l_0 pseudo-norm, which counts the total number of non-zero elements in a vector, is defined as

$$\| \mathbf{x} \|_0 = \text{Card}(\text{Supp}(\mathbf{x})), \tag{2.2.3}$$

where the support of a vector $\text{Supp}(\mathbf{x})$ is defined as the set containing the indices of the non-zero elements, and the cardinality $\text{Card}(\cdot)$ counts the number of elements of a set.

2.2.2 Bases and Frames

The goal of sparse signal processing is to approximate a signal using a linear model, called dictionary [88, 112], which is a set of elementary signals, called atoms, in order to make its processing or analysis faster and simpler. *Sparsity* expresses the idea that a signal has a compact representation, i.e., that one can approximate it with just a few elements from the dictionary.

When the dictionary forms a basis, it is said to be *complete*, and every signal is uniquely represented as the linear combination of the dictionary atoms. A basis $\mathbf{\Psi}$ of an N-dimensional space is defined as a *complete* set of atoms $\{\psi_i\}_{i \in N}$, such that, for any signal $\mathbf{x}$, there exists a *unique* sequence of coefficients s_i satisfying

$$\mathbf{x} = \sum_{i=1}^{N} s_i \psi_i. \tag{2.2.4}$$

When the set of atoms are orthogonal, $\langle \psi_i, \psi_k \rangle = \delta_{i-k}$, then $\mathbf{\Psi}$ is called an *orthogonal* basis. When the atoms are not orthogonal, $\langle \psi_i, \psi_k \rangle \neq \delta_{i-k}$, $\mathbf{\Psi}$ is called a *biorthogonal* basis. In a finite-dimensional space, the number of representative basis atoms is the same as the dimension of the space (Fig. 2.1 (a) and (b)).

The most commonly used bases in signal processing typically originate form traditional transforms, such as the Discrete Fourier Transform (DFT), the Discrete Cosine Transform (DCT) and the Discrete Wavelet Transform (DWT). As well known, these transforms are effective for representing natural signals and images, and have fast implementations. The application of these bases to physiological signals, however, can be problematic, and one can design alternative dictionaries that allow a more compact representation. In particular, overcomplete dictionaries have more atoms than the dimension of the signal, and can represent specific signal patterns more efficiently.

When the number of atoms P is larger than the dimension N of the space, the atoms are no longer linearly independent. If $N < P$ atoms are linearly independent, the set is

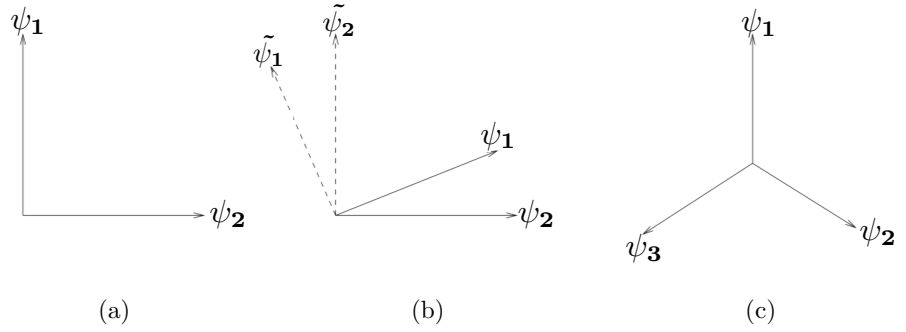

Figure 2.1: (a) An orthogonal basis, (b) a biorthogonal basis, (c) a frame in $\mathbb{R}^2$.

called a *frame* [76] or *overcomplete dictionary* $\mathbf{D} \in \mathbb{R}^{N \times P}$ (Fig. 2.1 (c)). Given a *frame* and a set of atoms $\{\mathbf{d}_i\}_{i \in P}$, any signal $\mathbf{x}$ can be represented by a sequence of coefficients s_i such that

$$\mathbf{x} = \sum_{i=1}^{P} s_i \mathbf{d}_i. \tag{2.2.5}$$

Because the atoms in $\mathbf{D}$ are linearly dependent, this signal expansion is *not unique*. Non-uniqueness gives us the possibility to choose among the infinite number of possible representations the one that most fits our purposes.

An important property of an overcomplete dictionary is its coherence, first introduced by Davis *et al.* [48], which quantifies how the atoms are correlated with each other.

Definition 2.2.1 *The coherence of a matrix* $\mathbf{D}$ *is defined as*

$$\mu(\mathbf{D}) = \max_{i \neq j} \frac{| < \mathbf{d}_i, \mathbf{d}_j > |}{||\mathbf{d}_i||_2 ||\mathbf{d}_j||_2}.$$

where $\mathbf{d}_i$ *and* $\mathbf{d}_j$ *denote columns in* $\mathbf{D}$.

It is not difficult to show that the coherence of a matrix $\mathbf{D} \in \mathbf{R}^{N \times P}$, with normalized columns, $||\mathbf{d}_j||_2 = 1$, is always in the range $\mu(\mathbf{D}) \in \left[\sqrt{\frac{P-N}{N(P-1)}}; 1\right]$. Note that when $\mathbf{D}$ is an orthogonal matrix, the coherence is exactly 0. It can be shown that incoherence is a desirable property for overcomplete dictionaries [27].

2.2.3 Sparse and Compressible Signals

A signal $\mathbf{x}$ is said to have a sparse representation, in itself or in some transform domain, if it or its transform coefficient vector $\mathbf{s}$ has a large number of coefficients equal to zero.

Formally, we say that a vector $\mathbf{s}$ is k-sparse when $\| \mathbf{s} \|_0 \leq k$. The set of all k-sparse vectors is defined as

$$\Sigma_k = \{\mathbf{s} : \| \mathbf{s} \|_0 \leq k\}. \tag{2.2.6}$$

Even if the representation is not exact, and requires more than k atoms, it is possible in general to obtain a good k-sparse approximation of the signal by setting the small

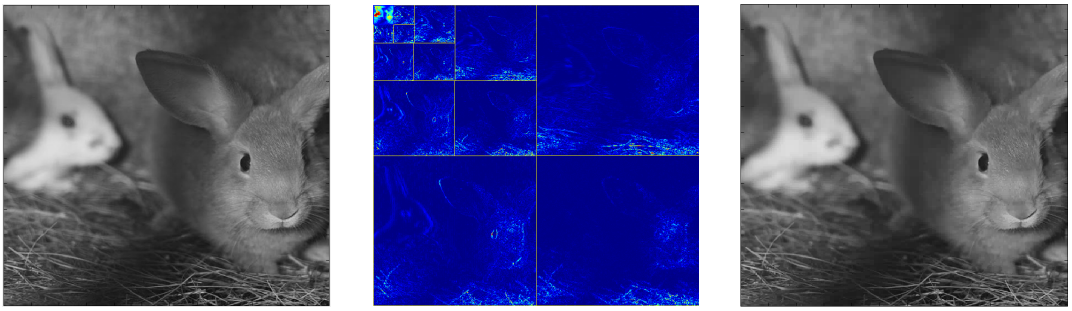

(a) (b) (c)

Figure 2.2: (a) Original image. (b) Sparse representation using wavelet transform, observe that most of the wavelet coefficients (dark pixels) are close to zero. (c) Approximation of the image obtained by keeping only the largest 10% wavelet coefficients.

coefficients to zero (e.g., by thresholding). An example is reported in Fig. 2.2 where the wavelet transform has been applied to an image. As visible in Fig. 2.2(b), most coefficients are very small (dark pixels). A good approximation can be obtained by setting the small coefficients to zero. Figure Fig. 2.2(c) shows its k-term approximation obtained by keeping only the largest 10% coefficients.

The class of approximately sparse signals, the so called *compressible signals*, is the most important in practice, since only few real-world signals are truly sparse even in the transform domain.

It is possible to quantify the compressibility of a signal by calculating the error resulting from the approximation of signal $\mathbf{s}$ by some $\hat{\mathbf{s}} \in \Sigma_k$

$$\sigma_k(\mathbf{s})_p = \min_{\hat{\mathbf{s}} \in \Sigma_k} \| \mathbf{s} - \hat{\mathbf{s}} \|_p \, . \tag{2.2.7}$$

Most natural signals and images can be represented with bases or dictionaries with coefficients that obey a power law decay. Sorting the coefficients s_i such that $|s_1| \geq |s_2| \geq \cdots \geq |s_K|$, we say that the coefficients obey a power law decay if there exist constants $C, q > 0$ such that $|s_i| \leq Ci^{-q}$ [65]. Signals which obey such property are in general compressible, and the success of compression schemes based on transform coding are based on this principle.

2.2.4 The Sparse Representation Problem

In the simplest case of an orthogonal basis $\mathbf{\Psi}$, the expansion coefficients $\mathbf{s}$ of a signal $\mathbf{x}$ can be easily computed as inner product of the signal and the atoms

$$\mathbf{s} = \mathbf{\Psi}^T \mathbf{x}. \tag{2.2.8}$$

Equation 2.2.8 is usually called the *analysis* formula, while the equation $\mathbf{x} = \mathbf{\Psi s}$ is referred as the *synthesis* formula.

Let now consider the case of an overcomplete dictionary, where the matrix $\mathbf{D} \in \mathbb{R}^{N \times P}$ is a "fat" matrix, meaning that it has more columns then rows. We assume that $rank(D) =$

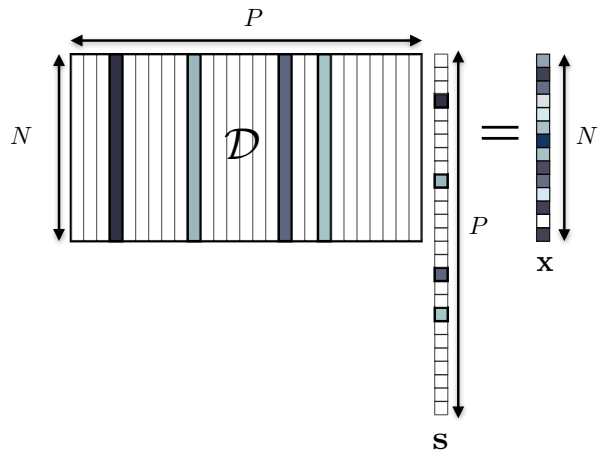

Figure 2.3: Sparse representation model. In this model colored columns, and the relative coefficients in $\mathbf{s}$, represent the selected sparse set of atoms.

N (i.e., that it is a full rank matrix). A simple approach in the analysis path is to calculate $\mathbf{s} = \tilde{\mathbf{D}}\mathbf{x}$, where $\tilde{\mathbf{D}} = \mathbf{D}^{\dagger}$ is the pseudo inverse of $\mathbf{D}$. It is well known that the resulting vector $\mathbf{s}$ is the minimum l_2-norm vector satisfying $\mathbf{x} = \mathbf{D}\mathbf{s}$. The pseudo-inverse has the closed-form expression $\mathbf{D}^{\dagger} = \mathbf{D}\left(\mathbf{D}^T\mathbf{D}\right)^{-1}$. Note however that $\mathbf{s}$ is not sparse in general.

Searching for compact representations, a more interesting alternative is to find the solution as the sparsest possible. This is more efficient from a processing and storage requirement point of view. This can be accomplished by finding the representation having as few non-zero coefficients as possible, i.e., by solving the problem

$$\min_{s} \| \mathbf{s} \|_0 \text{ s.t. } \mathbf{x} = \mathbf{D}\mathbf{s}. \tag{2.2.9}$$

Unfortunately, finding the solution to Eq. 2.2.9 is NP hard due to its combinatorial optimization nature. As we will see later for the compressive sensing problem, suboptimal solutions can be found by iterative methods like matching pursuit (MP) [88] and orthogonal matching pursuit (OMP) [104, 121], or relation methods such as basis pursuit (BP). Sec. 2.4 summarizes some of the most used algorithms used to compute sparse approximations, as well as to solve the CS-recovery problem that will be introduced in the next section.

2.3 Compressive Sensing

2.3.1 Form Sparse Signals to Sparse Sampling

The majority of traditional signal acquisition schemes are based on the classical sampling formulation, which requires to sample a signal $x(t)$ on a set of uniformly spaced time instants, at a rate at least twice the bandwidth of the signal [116] (Fig. 2.4(a)). Signal reconstruction from samples x_n is computed as $\hat{x}(t) = \sum_n x_n \phi(t/T - n)$, where $\hat{x}(t)$ is the reconstructed signal, $\phi(t)$ is the *sinc* function, and T is the sampling period. Perfect reconstruction is guaranteed for any bandlimited signal $x(t)$. This result is the well know Shannon sampling theorem, which provides a sufficient but not necessary condition for

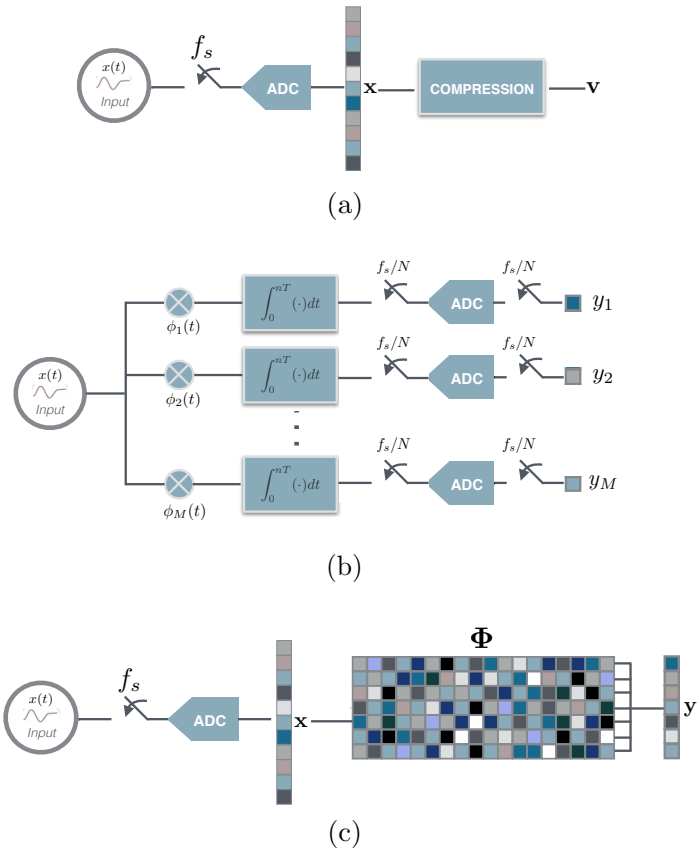

Figure 2.4: (a) Traditional 'acquire then compress' scheme, (b) possible analog implementation of compressive sensing and (c) digital compressive sensing.

perfect reconstruction. The necessity to represent the samples efficiently for transmission or recording purposes, typically requires to compress the data, possibly using costly and signal dependent schemes.

The use of compressive sensing (CS) appears to be an advantageous alternative to the standard sensing approach, since it enables to sense and compress data simultaneously while capturing all the relevant information at a rate lower than the Nyquist one. Instead of considering the signal bandwidth, compressive sensing is based on the observation that signals may have a sparse representation in some domain or can be sparse-approximated. Thus, it can be seen as a novel sampling approach in which the key element in the sparsity of the acquired signal.

CS is based on an efficient and signal-independent sampling strategy, which consists of taking M non-adaptive linear measurements of the k-sparse vector $\mathbf{x}$, with $M \sim k, M \ll N$ (Fig. 2.4(b)(c)). In order to make the discussion more concrete, even if analog implementations are possible (see Fig. 2.4(b)) for the remainder of this chapter we will restrict our attention to the discrete CS model, known as "digital" CS (Fig. 2.4(c)), which will be used in this thesis for low power compression of physiological signals.

Digital CS consists in the application of a sensing (or measurements) matrix $\mathbf{\Phi} \in$

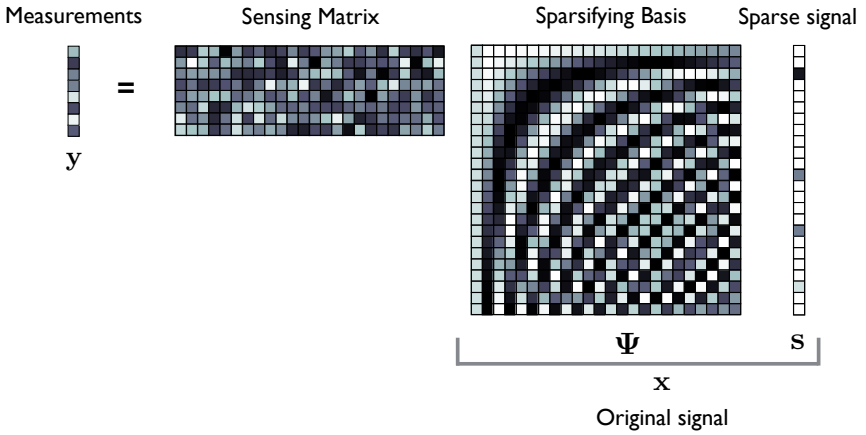

Figure 2.5: Compressive sensing framework.

$\mathbf{R}^{M\times N}$ to the signal $\mathbf{x} \in \mathbb{R}^N$ to obtain the measurement vector $\mathbf{y} \in \mathbb{R}^M$

$$\mathbf{y} = \mathbf{\Phi x}. \tag{2.3.1}$$

The sensing matrix $\mathbf{\Phi}$ is typically constructed randomly and does not depend in general on the characteristics of signal $\mathbf{x}$, making compressive sensing "universal", since it is possible to use the same acquisition/compression system for different classes of sparse/compressible signals.

Compressive sensing can be considered a "democratic" scheme since each measurement carries roughly the some amount of information and has equal priority. Thus, the reconstruction depends only on the number of measurements and not on the particular subset received. This might be particularly useful in case of unreliable communication, since receiving a subset of the measurements still allows to recover the signal, even if with a lower quality. Moreover, the measurement process is progressive and to obtain a higher signal quality it is just necessary to take more measurements. In the same way, one can further increase compression by discarding some of the measurements.

The main idea of CS goes against the common wisdom in data acquisition, substituting the linear reconstruction process of the classical sampling formulation with a non-linear one. As we will see later, under specific hypotheses on the matrix $\mathbf{\Phi}$, it is indeed possible to recover a signal as long as it is sparse or compressible.

2.3.2 Compressive Sensing Signal Recovery

The goal of the compressive sensing approach, given the compressed measurements vector $\mathbf{y}$, is to recover the original vector $\mathbf{x}$ satisfying $\mathbf{y} = \mathbf{\Phi x}$. The matrix $\mathbf{\Phi}$ used in CS is fat and has more columns than rows, thus the constraint equation defines an ill-posed problem with infinitely many solutions.
However, if we exploit sparsity as an a-priori information about the signal, it makes sense to look for the solution that has the minimum number of non-zero components, exactly as in problem (2.2.9) of Sec. 2.2.4.

When the signal $\mathbf{x}$ is not sparse in itself, but has a sparse representation in some dictionary $\mathbf{D}$, i.e., $\mathbf{x} = \mathbf{D}\mathbf{s}$, where $\mathbf{s}$ is sparse, Problem (2.2.9) becomes

$$\min_{\mathbf{s}} \| \mathbf{s} \|_0 \text{ s.t. } \mathbf{y} = \mathbf{A}\mathbf{s}, \quad \mathbf{A} = \mathbf{\Phi}\mathbf{D}, \text{ (P0)}. \tag{2.3.2}$$

The original signal $\mathbf{x}$ can be finally recovered as $\mathbf{x} = \mathbf{D}\mathbf{s}$.

A generalization of problem (P0) allows to deal with nearly sparse vectors $\mathbf{s}$ and noise. In particular, one can consider the sensing model $\mathbf{y} = \mathbf{A}\mathbf{s} + \mathbf{n}$, where $\mathbf{n}$ is an additive vector taking error into account. In this case, the reconstruction problem is formulated as

$$\min_{\mathbf{s}} \| \mathbf{s} \|_0 \text{ s.t. } \| \mathbf{A}\mathbf{s} - \mathbf{y} \|_2 \leq \epsilon, \text{ (P0,}\epsilon\text{)}, \tag{2.3.3}$$

where ϵ is a bound on the noise energy, i.e., $||\mathbf{n}||_2 \leq \epsilon$. In this case, the reconstruction gives a vector $\tilde{\mathbf{s}}$ whose distance from the original $\mathbf{s}$ is controlled by the measurement error bound ϵ (see Fig. 2.6). This leads to a robust reconstruction with respect to measurement errors.

Unfortunately, as for the sparse representation problem (2.2.9), the minimization problems 2.3.2 and 2.3.3 are NP-hard in general. Therefore, alternative approaches have been proposed in the literature.

One common approach is based on convex relaxation, replacing the l_0 with the l_1 norm, to make the problem easier to work with. Thus, Problem 2.3.2 becomes

$$\min_{\mathbf{s}} \| \mathbf{s} \|_1 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{A}\mathbf{s} \quad \text{(P1)}, \tag{2.3.4}$$

which is known as the *Basis Pursuit* (BP) problem. Similarly, Problem 2.3.3 can be rewritten as the *Basis Pursuit Denoising* (BPDN) problem

$$\min_{\mathbf{s}} \| \mathbf{s} \|_1 \text{ s.t. } \| \mathbf{A}\mathbf{s} - \mathbf{y} \|_2 \leq \epsilon \quad \text{(P1,}\epsilon\text{)} \ . \tag{2.3.5}$$

In Sec. 2.4 we will review some of the algorithms which can be used to solve the recovery problems introduced in this section.

2.3.3 Sensing Matrices and Recovery Guarantees

In the following, we assume that the sparsifying basis or dictionary is the identity matrix $\mathbf{I}$, i.e., that $\mathbf{y} = \mathbf{\Phi}\mathbf{I}\mathbf{s}$ and $\mathbf{A} = \mathbf{\Phi}$. The sensing matrix $\mathbf{\Phi}$ used in CS causes a dimensionality reduction, mapping $\mathbb{R}^N$ into $\mathbb{R}^M$, with a loss of information in general. If $\mathbf{s}$ is sparse, however, some properties of the sensing matrix, which we will describe below, such as the Null-Space Property (NSP), the Restricted Isometry Property (RIP), and incoherence, assure optimality of the signal reconstruction procedures.

Null Space Property

In the design of the sensing matrix $\mathbf{\Phi}$, we should ensure that any pair of distinct sparse vectors $\mathbf{s}_1, \mathbf{s}_2 \in \Sigma_k$, are mapped to different measurement vectors, i.e., $\mathbf{\Phi}\mathbf{s}_1 \neq \mathbf{\Phi}\mathbf{s}_2$, since otherwise it would be impossible to distinguish the two vectors based on the measurements $\mathbf{y}$. Let us denote the null space of $\mathbf{\Phi}$ as

$$\mathcal{N}(\mathbf{\Phi}) = \{\mathbf{s} | \mathbf{\Phi}\mathbf{s} = 0\}. \tag{2.3.6}$$

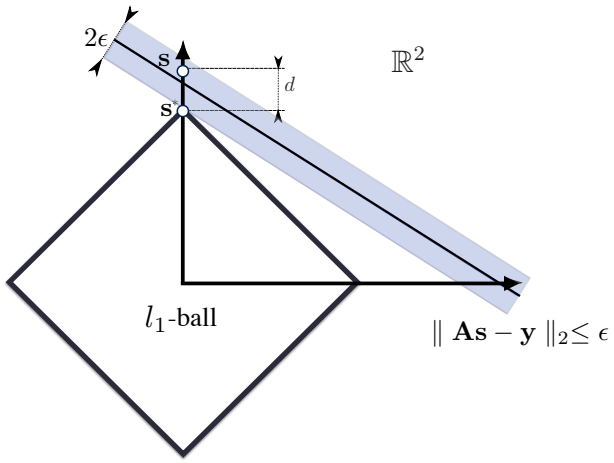

Figure 2.6: Example of k-sparse signal recovery, $k = 1$, in $\mathbb{R}^2$ in the presence of noise.

It should be noted that if $\mathbf{\Phi s}_1 = \mathbf{\Phi s}_2$, then $\mathbf{\Phi}(\mathbf{s}_1 - \mathbf{s}_2) = 0$, with $\mathbf{s}_1 - \mathbf{s}_2 \in \Sigma_{2k}$. Thus, in order to ensure that the matrix $\mathbf{\Phi}$ uniquely represents and allows the recovery of all $\mathbf{s} \in \Sigma_k$ the null space $\mathcal{N}(\mathbf{\Phi})$ must not contain vectors in Σ_{2k}.

Formally, one way to characterize this property is by using the *spark*, defined by Donoho and Elad in [54].

Definition 2.3.1 *The spark of a given matrix $\mathbf{\Phi}$ is the smallest number of columns of $\mathbf{\Phi}$ that are linearly dependent.*

Note that $spark(\mathbf{\Phi}) \in [2, M + 1]$.

Theorem 2.3.1 ([54]) *For any vector $\mathbf{y} \in \mathbb{R}^M$, there exists at most one signal $\mathbf{s} \in \Sigma_k$ such that $\mathbf{y} = \mathbf{\Phi s}$ if and only if $spark(\mathbf{\Phi}) > 2k$.*

Proof. The assumption that $spark(\Phi) > 2k$ means that any signal $\mathbf{z} = \mathbf{s}_1 - \mathbf{s}_2$, with $\mathbf{s}_1, \mathbf{s}_2 \in \Sigma_k$, can at most have $2k$ nonzero entries and hence cannot pick any linearly dependent set of columns from $\mathbf{\Phi}$. Therefore $\mathbf{s}$ cannot be in the null space of $\mathbf{\Phi}$ unless $\mathbf{s}_1 = \mathbf{s}_2$. ■

The *spark* property guarantees uniqueness of the solution for sparse signals. The *spark* can also be used to bound the minimum number of measurements required in order to allow the reconstruction of k-sparse signals from the measurements. It follows from Theorem 2.3.1 that the minimum number of measurements is $M \geq 2k$. Note that the theorem above is based on the fact that $\mathcal{N}(\Phi)$ should not contain vectors which are too sparse. A powerful property that guarantees that the solution to the tractable Problem 2.3.4 coincides with the solution to Problem 2.3.2 is the Null Space Property (NSP).

Definition 2.3.2 *A matrix $\mathbf{\Phi}$ satisfies the NSP of order k if*

$$||\mathbf{h}_{\mathcal{S}}||_1 < ||\mathbf{h}_{\bar{\mathcal{S}}}||_1,$$

holds for all non-zero $\mathbf{h}$ *in the null space of* $\mathbf{\Phi}$ *and all coordinate set* $\mathcal{S} \subset \{1, 2, \ldots, N\}$ *of cardinality* $|\mathcal{S}| \leq k$. $\mathbf{h}_{\mathcal{S}}$ *denotes a vector whose elements are equal to those of* $\mathbf{h}$ *in the coordinate set* $\mathcal{S}$ *and zero elsewhere, while* $\bar{\mathcal{S}}$ *denotes the complement of set* $\mathcal{S}$.

If a matrix $\mathbf{\Phi}$ satisfies the NSP, then the vectors in its null space are not too concentrated on a small subset of indices. This guarantees the exact recovery of all possible k-sparse signals by solving Problem 2.3.4, as stated by Theorem 2.3.2.

Theorem 2.3.2 ([26]) *Problem Eq. 2.3.4 uniquely recovers all* k*-sparse vectors* $\mathbf{s}$ *from measurements* $\mathbf{y} = \mathbf{\Phi s}$ *if and only if* $\mathbf{\Phi}$ *satisfies the NSP of order* k.

One may obtain also a stable version of the above theorem by considering compressible vectors such that $\sigma_k(\mathbf{x})_1$ is small (see Eq. 2.2.7).

Restricted Isometry Property

The null-space property can be used for exactly k-sparse and compressible signals in the noise-free setting. In the presence of noise, it is necessary to impose a stronger condition on the sensing matrix in order to guarantee signal recovery from the measurements.

To this end, Candès and Tao introduced the Restricted Isometry Property (RIP) [26].

Definition 2.3.3 ([26]) *A matrix* $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ *satisfies the Restricted Isometry Property (RIP) of order* k *with constant* $0 \leq \delta_k < 1$ *if*

$$(1 - \delta_k) \parallel \mathbf{s} \parallel_2 \leq \parallel \mathbf{\Phi s} \parallel_2 \leq (1 + \delta_k) \parallel \mathbf{s} \parallel_2, \tag{2.3.7}$$

for all k*-sparse vectors* $\mathbf{s} \in \mathbb{R}^N$.

This property guarantees that k-sparse vectors cannot be in the null space of $\mathbf{\Phi}$, and that the matrix preserves as much information as possible despite of the dimensionality reduction.

In particular, considering two k-sparse signals $\mathbf{s}_1$ and $\mathbf{s}_2$ with different support, if the matrix $\mathbf{\Phi}$ satisfies the RIP of order $2k$, then the euclidean distance $\parallel \mathbf{s}_1 - \mathbf{s}_2 \parallel_2$ is almost preserved after projection. We can write the RIP property for the difference vector $\mathbf{s} = \mathbf{s}_1 - \mathbf{s}_2$, where $\mathbf{s}$ is a $2k$-sparse vector or less, as

$$(1 - \delta_{2k}) \parallel \mathbf{s}_1 - \mathbf{s}_2 \parallel_2 \leq \parallel \mathbf{\Phi}(\mathbf{s}_1 - \mathbf{s}_2) \parallel_2 \leq (1 + \delta_{2k}) \parallel \mathbf{s}_1 - \mathbf{s}_2 \parallel_2 . \tag{2.3.8}$$

The RIP of order $2k$ requires that any subsets of at most $2k$ columns of the sensing matrix $\mathbf{\Phi}$ behaves approximately like an orthonormal system for sparse signals, preserving the Euclidean distance.

A lower bound on the number of measurements that a matrix must have in order to achieve the RIP based only on the dimensions of the problem (N, M and k) is given in the following theorem.

Theorem 2.3.3 ([45]) *Let* $\mathbf{\Phi}$ *be an* $M \times N$ *matrix that satisfies the RIP of order* $k \leq N/2$ *with constant* $\delta \in (0, 1)$. *Then*

$$M \geq C_\delta k \log\left(\frac{N}{k}\right),$$

where $C_\delta < 1$ *is a constant depending on* δ.

From Theorem 2.3.3 it fallows that the number of required measurements only scales linearly in k up to the logarithmic factor $\log(N/k)$. This is an important result on the minimum number of measurements necessary to recover a signal even in presence of noise. We have the following theorem, which again allows recovery of signals based on the l_1-norm minimization.

Theorem 2.3.4 ([45]) *Let matrix $\mathbf{\Phi}$ satisfy the RIP of order 2k with constant $\delta_{2k} < 1/3$. Then, any k-sparse vector $\mathbf{s}$, $\mathbf{y} = \mathbf{\Phi s}$, is the unique solution of Problem 2.3.4.*

Theorem 2.3.4 provides guarantees for perfect recovery of k-sparse signals in the noise free setting. However, in a real scenario we must consider the noisy model and rather solve the problem in Eq. 2.3.5. In the following theorem, the RIP condition guarantees stable recovery of signals based on the l_1-norm minimization in presence of noise.

Theorem 2.3.5 ([28]) *Consider Problem 2.3.5, $\mathbf{y} = \mathbf{\Phi s} + \mathbf{n}$, $||\mathbf{n}||_2 \leq \epsilon$. If $\mathbf{\Phi}$ satisfies the RIP with $\delta_{2k} \leq \sqrt{2} - 1$, then the solution $\tilde{\mathbf{s}}$ of Problem 2.3.5 satisfies*

$$||\tilde{\mathbf{s}} - \mathbf{s}||_2 \leq C_0 k^{-1/2} ||\mathbf{s} - \mathbf{s}_k||_1 + C_1 \epsilon,$$

with some constants C_0 and C_1. $\mathbf{s}_k$ represents the k-sparse approximation of $\mathbf{s}$, i.e., the vector $\mathbf{s}$ with all but the k-largest entries set to zero.

These results claim that as long as the sensing matrix $\mathbf{\Phi}$ satisfies the RIP, the l_1 recovery is stable against both the measurement noise (quantified by the energy bound ϵ) and the signal noise (i.e., the error related to the best k-term approximation).

Incoherence

Another principle of compressive sensing is the mutual coherence between the sensing and sparsifying matrices or dictionaries. As a matter of fact, the NSP and RIP properties are in general impractical or even impossible to verify, due to the combinatorial nature of the definitions. In many cases, it is preferable to use the notion of mutual coherence, as defined in Def. 2.2.1, which is easier to verify with respect to NSP or RIP.

We have the following theorem.

Theorem 2.3.6 ([58]) *(Equivalence - Basis Pursuit) For the system of linear equations $\mathbf{\Phi s} = \mathbf{y}$, ($M \times N$ matrix $\mathbf{\Phi}$, full-rank, $M < N$,), if a solution $\mathbf{s}$ exists obeying*

$$\| \mathbf{s} \|_0 < \frac{1}{2} \left(1 + \frac{1}{\mu(\mathbf{\Phi})} \right),$$

that solution is both the unique solution of P1, (2.3.4) and the unique solution of P0 (2.3.2).

2.3.4 Random Sensing Matrices

So far, we summarized the relevant properties that matrices in the context of CS should satisfy. In this section we will deal with the problem of constructing such matrices. The RIP condition provides strong guarantees on the perfect recovery of sparse signals, but it can be difficult to design deterministic matrices satisfying the RIP.

It is however possible to construct random matrices with entries chosen according to a Gaussian $\mathcal{N}(0, 1/M)$, Bernoulli, $\mathrm{P}(\phi_{i,j} = \pm 1/\sqrt{M}) = 1/2$, or more generally to a sub-Gaussian distribution, according to the following definition.

Definition 2.3.4 *An $M \times N$ matrix $\mathbf{\Phi}$ is referred to as a sub-Gaussian random matrix if all the entries of $\mathbf{\Phi}$ are drawn from independent, zero mean, sub-Gaussian random variables with variance $1/M$, namely, if the entries of $\mathbf{\Phi}$, ϕ_{ij}, $1 \leq i \leq M$ $1 \leq j \leq N$, satisfy*

$$Prob(|\phi_{ij}| \geq t) \leq \beta e^{-\kappa t^2}, \ \ \forall t > 0,$$

where β and κ are positive constants.

Indeed, it has been shown that sub-Gaussian matrices satisfy the RIP with high probability [8]. In applications where the signal is sparse with respect to some basis, we require that the product matrix $\mathbf{A} = \mathbf{\Phi\Psi}$ satisfies the RIP. If the sensing matrix is chosen according to a sub-Gaussian distribution and the sparsifying domain is an orthonormal basis, then it is possible to show that $\mathbf{A}$ will also satisfy the RIP with high probability [8]. Moreover, the product of a random matrix with any fixed basis $\mathbf{\Psi}$ has low coherence in general.

2.3.5 Sparse Sensing Matrices

Random matrices with entries chosen according to a sub-Gaussian distribution are a good choice in order to satisfy the properties required for signal recovery. However, truly random matrices might be a potential obstacle in CS implementation, due to the fact that the matrix-vector multiplication to compute the measurements may require considerable computational resources. It may be preferable to consider alternative matrix constructions for low consumption implementations, such as sparse sensing matrices, i.e., matrices with only few non-zero entries. These matrices are particularly useful thanks to the resulting limited complexity and small storage requirements.

An interesting subset of sparse sensing matrices is represented by sparse binary matrices that have only few entries equal to 1 in each column, in particular d non-zeros elements in each column for a d-sparse sensing matrix. The support of the d non-zero elements is drawn uniformly at random. If $\mathbf{\Phi}$ is designed such that it is dominated by zero entries and only few entries are non-zero, then computing the product $\mathbf{\Phi x}$ takes $O(dN)$ operations, which is a significant saving when $d \ll N$.

These matrices are related to the adjacency matrices of an unbalanced expander graph [67], which consists of two classes of nodes called variable nodes and measurement nodes, corresponding to the columns and rows of the matrix, respectively. When a pair of variable and measurement nodes are connected by an edge, the binary matrix has a nonzero entry. Unfortunately, sparse matrices may no longer satisfy the RIP property, unless the number of rows is large, i.e., it is at least $\mathcal{O}(k^2)$ [34], for k-sparse signals.

To overcome this limitation, Berinde *et al.* in [16] introduced a variant of the RIP, with the l_2-norm replaced by the l_1-norm, namely, the RIP-1, defined as follows.

Definition 2.3.5 *An $M \times N$ matrix $\mathbf{\Phi}$ is said to satisfy the RIP-1 with isometry constant $\delta_{k,1}$, if for any k-sparse vector $\mathbf{s}$, we have*

$$(1 - \delta_{k,1})||\mathbf{s}||_1 \leq ||\mathbf{\Phi s}||_1 \leq ||\mathbf{s}||_1$$

As the RIP property is used to show that every matrix $\mathbf{\Phi}$ satisfying this property with isometry constant δ_k satisfies the $l_2 - l_1$ guarantee in Theorem 2.3.5, then there exists a connection between the RIP-1 and guarantee for l_1 minimization recovery [16]. Indeed, let $\tilde{\mathbf{s}}$ the solution of the problem

$$\min_{\mathbf{s}} \| \mathbf{s} \|_1 \text{ s.t. } \| \mathbf{\Phi s} - \mathbf{y} \|_1 \leq \epsilon, \tag{2.3.9}$$

then, if a matrix $\mathbf{\Phi}$ satisfy the RIP-1 it can be shown that

$$||\mathbf{s} - \tilde{\mathbf{s}}||_1 \leq C_0 k^{-1} ||\mathbf{s} - \mathbf{s}_k||_1 + C_1 \epsilon, \tag{2.3.10}$$

for some constants C_0 and C_1. Hence, the RIP-1 proves that after projection it is still possible to use l_1-minimization to recover k-sparse signals [16].

2.3.6 CS with Overcompete Dictionaries

Some of the results presented above assume that signals are sparse with respect to an orthonormal basis. Despite that, most natural signals may not be sufficiently sparse in an orthonormal basis but could be much sparser in an overcomplete dictionary $\mathbf{D}$.

Candes et al. [29] generalized CS theory for signals sparse in redundant and overcomplete dictionaries. There are two main differences with the classical CS theory presented in Sec. 2.3.3 and CS using overcomplete dictionaries. As explained, CS requires that the sensing matrix $\mathbf{\Phi}$ and in general the product matrix $\mathbf{\Phi\Psi}$ satisfies the incoherence property and the RIP. However, when the dictionary $\mathbf{D}$ is highly coherent, then the product matrix $\mathbf{A} = \mathbf{\Phi D}$ will also be coherent in general. Second, if the sparsifying transform is not an orthogonal basis, matrix $\mathbf{\Phi D}$ is not at all likely to satisfy the RIP.

It is shown in [29] that, even if the low coherence requirements are not satisfied, it is possible to successfully recover signal $\mathbf{x} = \mathbf{Ds}$. The error in the reconstruction of the signal $\| \tilde{\mathbf{x}} - \mathbf{x} \|_2$ may indeed be smaller then the error in the recovery of the sparse coefficient vector $\| \tilde{\mathbf{s}} - \mathbf{s} \|_2$. In order for this to happen, the representation $\mathbf{s}$ should have rapidly decreasing coefficients, which may be the case if a proper dictionary is chosen.

There are some basic conditions that bridge the gap between the classical compressive sensing with orthonormal bases and the CS with overcomplete dictionaries. Candes et al. [29] consider a generalization of the RIP of the sensing matrix, the so-called D-RIP. If the D-RIP is satisfied, they provide results about theoretical stability guarantees for reconstruction methods based on l_1 minimization.

Definition 2.3.6 *Let $\mathbf{D} \in \mathbb{R}^{N \times P}$ (the dictionary) with $P > N$ and let $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ (the sensing matrix). The restricted isometry constant δ_k adapted to $\mathbf{D}$ is defined as the smallest constant such that*

$$(1 - \delta_k) \| \mathbf{x} \|_2^2 \leq \| \mathbf{\Phi x} \|_2^2 \leq (1 + \delta_k) \| \mathbf{x} \|_2^2,$$

for all $\mathbf{x} \in \mathbb{R}^N$ of the form $\mathbf{x} = \mathbf{Ds}$ for some k-sparse $\mathbf{s} \in \mathbb{R}^P$.

Note that the D-RIP requirement holds for vectors $\mathbf{x}$ with a sparse representation over $\mathbf{D}$, differently from the original definition of RIP.

Basically, all random matrices that obey the standard RIP also obey the D-RIP, In particular, for any overcomplete dictionary $\mathbf{D}$, if $\mathbf{\Phi}$ has i.i.d. random entries from a Gaussian or sub-Gaussian distribution, then it satisfies the D-RIP of order k with high probability if $M = O(k \log(k/P))$.

2.4 Algorithms for Sparse Signal Recovery

This section summarizes a selection of popular sate-of-the-art algorithms for the reconstruction of sparse signals or to compute the expansion coefficients in sparse representations with over-complete dictionaries.

2.4.1 l_1-norm Algorithms

In the noiseless case, the sparse representation/recovery problems in Eq. 2.2.9 and in Eq. 2.3.2, are in fact non-convex optimization problems, and as said before they are NP-hard. They can be approximated using convex relaxation, obtianing the Basis Pursuit problem of Eq. 2.3.4. This can be recast as a linear program (LP) with equality constraints, as described by Boyd and Vandenberghe [19].

As a matter of fact, let us write $\mathbf{s}$ as $\mathbf{s} = \mathbf{u} - \mathbf{v}$, with $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$ non negative vectors, $\mathbf{u}$ with elements equal to the positive entries of $\mathbf{s}$ and zero elsewhere, and $\mathbf{v}$ with elements equal to the absolute value of the negative entries of $\mathbf{s}$ and zero elsewhere. By denoting $\mathbf{z} = [\mathbf{u}^T, \mathbf{v}^T]^T \in \mathbb{R}^{2N}$, we have $||\mathbf{s}||_1 = \mathbf{1}^T(\mathbf{u} + \mathbf{v}) = \mathbf{1}^T\mathbf{z}$ and $\mathbf{As} = \mathbf{A}(\mathbf{u} - \mathbf{v}) = [\mathbf{A}, -\mathbf{A}]\mathbf{z}$. Hence, the BP problem can be rewritten as an LP as follows

$$\min_{z} \mathbf{1^T z} \text{ s.t. } \mathbf{y} = [\mathbf{A}, -\mathbf{A}]\mathbf{z}, \ \ \mathbf{z} \geq 0.$$

The BP problem recast as an LP problem can be solved using well established numerical algorithms, such as interior-point methods [74].

When the measurements are contaminated by some form of noise or the signal is not exactly sparse, we should guarantee the stability by using an extension of the previous problem and solve Problem 2.3.5.

This problem is known as *quadratically constrained basis pursuit*, in which the constraint is related to the noise accepted in the solution. The basis pursuit denoising algorithm (BPDN) can be used to solve the the second-order cone programming (SOCP) problem in Eq. 2.3.5, by converting it into the unconstrained convex problem

$$\min_{\mathbf{s}} \| \mathbf{s} \|_1 + \lambda \| \mathbf{As} - \mathbf{y} \|_2^2 \qquad \text{(BPDN)}, \tag{2.4.1}$$

for an appropriate $\lambda \geq 0$ related to the signal power. When it is difficult to find a good estimation of the noise power to set a value for λ, and it is easier to quantify a bound on signal sparsity, the LASSO formulation can be used

$$\min_{\mathbf{s}} \| \mathbf{As} - \mathbf{y} \|_2 + \eta \| \mathbf{s} \|_1 \qquad \text{(LASSO)}. \tag{2.4.2}$$

It is not difficult to see that for an appropriate Lagrange multiplier η, the solution to LASSO is precisely the solution to the unconstrained optimization problem BPDN. Existing results about the robustness and stability of l_1-based methods for signal reconstruction

are mostly based on the properties of the matrix $\mathbf{A}$. As mentioned in Sec. 2.3.3, it is possible to show that, if the matrix $\mathbf{A}$ obeys the RIP, the reconstruction from noiseless data is exact, and similar results also hold for stable recovery in noisy settings or for compressible signals [28].

2.4.2 Greedy Algorithms

Algorithms based on l_1-norm minimization guarantee the recovery of approximately sparse solutions, and robustness against measurement noise. They rely on optimization procedures with relatively high complexity. For example, linear programming has a complexity which grows with the cubic power of the problem dimension. To overcome the complexity problem, one can use simpler approaches based on greedy algorithms, which iteratively construct a sparse approximation of the signal. This type of algorithms include Matching Pursuit (MP), proposed by Mallat in [88], and Orthogonal Matching Pursuit (OMP) [121]. Matching Pursuit is a greedy algorithm that iteratively decomposes a signal into a linear combination of a family of functions. At each iteration the algorithm chooses an atom from the dictionary $\mathbf{D}$ and then updates the residual error. The algorithm chooses the atom with the highest correlation against the current residual error $\mathbf{r}$, i.e.,

$$\lambda_j = \arg \max_{\{\mathbf{d}_i\}_{1 \leq i \leq N}} |\langle \mathbf{r}_{j-1}, \mathbf{d}_i \rangle|. \tag{2.4.3}$$

In (2.4.3), λ_j is the atom selected at the j-th iteration. The new residual error is obtained by subtracting the correlated component

$$\mathbf{r}_j = \mathbf{r}_{j-1} - \langle \mathbf{r}_{j-1}, \lambda_j \rangle \lambda_j. \tag{2.4.4}$$

The approximation at each step is $\mathbf{s}_j = \sum_{k=1}^{j} c_k \lambda_k$, where $c_k = \langle \mathbf{r}_{k-1}, \lambda_k \rangle$. The algorithm ends when the norm of the residual is lower then a desired error bound, or when a sparsity level is reached.

In the OMP the coefficients are optimized after each iteration by orthogonal projection of the signal into the subspace spanned by the selected atoms, leading to a better approximation with respect to the MP. The algorithm is summarized in Algorithm 1.

The stopping criterion should take into account measurement and computation errors $\| \mathbf{y} - \mathbf{A}\tilde{\mathbf{s}} \|_2 \leq \epsilon$ for some $\epsilon \geq 0$. If one has an estimate of the sparsity value k, the stopping criterion can be invoked when the number of iterations reaches k. It is possible to show that if the sensing matrix $\mathbf{\Phi}$ is sub-Gaussian and $M \sim k \log N$, the OMP correctly recovers a k-sparse signal with high probability. If the sparsity k is small, the OMP algorithm provides fast reconstruction of the signal, since the number of iterations is equal to k.

2.4.3 Approximation of l_0-norm

Even if greedy algorithms have an advantage over l_1-based methods in terms of computation time [43], the main drawback is a lower performance over the l_1-norm approaches [86].

In search of a fast algorithm with a performance similar to that obtained with l_1-norm minimization, it has been proposed to solve Problem 2.3.2 iteratively by approximating the l_0 pseudo-norm with a sequence of continuous cost functions $F_\sigma(\mathbf{s})$, which converge to the l_0-norm as $\sigma \to 0$ [94] [56].

Algorithm 1 OMP Algorithm

Recovery of signal $\mathbf{s}$ from measurements $\mathbf{y}$
INPUT: signal $\mathbf{y} \in \mathbb{R}^m$, matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\hat{x} = \emptyset$, sparsity level k
OUTPUT: Reconstructed signal $\tilde{\mathbf{s}}$ after k iteration, Residual $\mathbf{r}^{(k)}$
Initialize: set initial residual $\mathbf{r}^{(0)} = \mathbf{y}$

1: **while** $j \leq k$ **do**
2: $\lambda^j = \arg \max_{\{\mathbf{a}_i\}_{1 \leq i \leq n}} |\langle \mathbf{r}^{(j-1)}, \mathbf{a}_i \rangle|$. {finding the atom in $\mathbf{A}$ with maximum correlation with residual.}
3: $\mathbf{\Lambda}^{(j)} = \mathbf{\Lambda}^{(j-1)} \cup \{\lambda^j\}$
4: $\mathbf{s}^{(j)} = \arg \max_{\hat{\mathbf{s}}} \| \mathbf{y} - \mathbf{A}_{\mathbf{\Lambda}^{(j)}} \hat{\mathbf{s}} \|_2$ { where $\mathbf{A}_{\mathbf{\Lambda}^{(j)}}$ denotes the column of $\mathbf{A}$ indexed by $\mathbf{\Lambda}^{(j)}$}
5: $\mathbf{r}^{(j)} = \mathbf{y} - \mathbf{A}_{\mathbf{\Lambda}^{(j)}} \mathbf{s}^{(j)}$
6: $j = j + 1$
7: **end while**
8: $\tilde{\mathbf{s}} = \mathbf{s}^{(j)}$

The SL0 algorithm proposed in [94] solves the problem in Eq. 2.3.2 by approximating the l_0-norm with a continuous function, and optimizing the resulting cost function to provide a smooth measure of sparsity. Indeed, the l_0-norm can be approximated using a combination of Gaussian functions, for small σ values [94], as in

$$||\mathbf{s}||_0 \triangleq N - \sum_{i=1}^{N} \exp(-s_i^2/2\sigma^2). \tag{2.4.5}$$

Thus, minimizing the l_0-norm is approximately equivalent to maximize

$$F_\sigma(\mathbf{s}) = \sum_i \exp(-s_i^2/2\sigma^2),$$

where the parameter σ controls the accuracy with which F_σ approximates the l_0 norm. We have in fact

$$\lim_{\sigma \to 0} \exp(-s^2/2\sigma^2) = \begin{cases} 1 & \text{if } s = 0 \\ 0 & \text{if } s \neq 0. \end{cases} \tag{2.4.6}$$

This enables to replace the l_0-norm minimization with a convex problem, and use a classic steepest ascent algorithm to maximize $F_\sigma(\mathbf{s})$. The σ value allows a trade-off between the approximation of the l_0-norm and the smoothness of the function.

The algorithm proposed in [94] consists of two nested iterations, and the external loop is responsible to gradually decrease the σ value. Note that, when σ is sufficiently large, $\exp(-s_i^2/2\sigma^2) \approx 1 - s_i^2/2\sigma^2$, and the maximization of $F_\sigma(\mathbf{s})$ s.t. $\mathbf{y} = \mathbf{A}\mathbf{s}$ resembles the minimum l_2-norm solution of $F_\sigma(\mathbf{s})$ s.t. $\mathbf{y} = \mathbf{A}\mathbf{s}$ [94]. Therefore, the starting solution of the optimization process is usually calculated using the pseudo-inverse $\mathbf{A}^\dagger$ of $\mathbf{A}$ and set to $\mathbf{s_0} = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{y}$.

The internal loop maximizes $F_\sigma(\mathbf{s})$ on the feasible set $\{\mathbf{s}|\mathbf{y} = \mathbf{A}\mathbf{s}\}$, using a steepest ascent algorithm, and updating $\mathbf{s} \leftarrow \mathbf{s} - \mu\delta_k$ in the direction of the gradient, given by

$$\delta_k = \mathbf{s} \cdot \left[e^{-\frac{s_1^2}{2\sigma_k^2}}, \ldots, e^{-\frac{s_N^2}{2\sigma_k^2}} \right]^T. \tag{2.4.7}$$

In the previous equation, the dot indicates component by component multiplication. The next step consists in projecting $\mathbf{s}$ into the convex set to avoid trapping the algorithm in local maxima

$$\mathbf{s} = \mathbf{s} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{s} - \mathbf{y}). \tag{2.4.8}$$

Algorithm 2 SL0

Input: μ step size, $\mathbf{y}$, $\mathbf{A}$, σ_{dec}, σ_{min}, λ, K_{iter}
Initialization: $\mathbf{s}_0 \leftarrow \mathbf{A^T}(\mathbf{A}\mathbf{A^T})^{-1}\mathbf{y}$,
  $\sigma_1 = 2|\max(\mathbf{s}_0)|$
  while $\sigma_k < \sigma_{min}$ **do**
    for k=1:K$_{iter}$ **do**
      $\delta_k \leftarrow \mathbf{s} \cdot [\exp(-s_1^2/2\sigma_k^2), \ldots, \exp(-s_N^2/2\sigma_k^2)]^T$
      $\mathbf{s} \leftarrow \mathbf{s} - \mu\delta_k$
      Project $\mathbf{s}$ onto the feasible set: $\mathcal{S} = \{\mathbf{s}|\mathbf{A}\mathbf{s} = \mathbf{y}\}$
      $\mathbf{s} \leftarrow \mathbf{s} - \mathbf{A^T}(\mathbf{A}\mathbf{A^T})^{-1}(\mathbf{A}\mathbf{s} - \mathbf{y})$
    end for
    $\sigma_k \leftarrow \sigma_k\sigma_{dec}$
    $\tilde{\mathbf{s}}_\mathbf{k} \leftarrow \mathbf{s}$
  end while
Output: $\mathbf{s_{OUT}} \leftarrow \tilde{\mathbf{s}}_\mathbf{k}$

The SL0 algorithm is typically 2 to 3 times faster than the Basis Pursuit denoising algorithm (SPGL1 implementation, [122]), while resulting in many cases in the same or better accuracy [94].

3

Compressive Sensing of Electrocardiogram Signals

3.1 Introduction

Since the early 20th century the ECG signal has become an indispensable clinical tool and it is widely used in health monitoring as a non-invasive way to establish clinical diagnosis of heart diseases [62]. The importance of the ECG is still growing thanks to modern and advanced signal processing techniques that allow the evaluation of variability patterns present in rhythm or wave morphology on a beat-to-beat basis.

As mentioned in the introduction of this thesis, WBANs promise to be a key element in wireless ECG systems for long-term recording [92, 78, 10, 96]. WBANs consist of biomedical wireless sensors attached on or implanted to collect vital biomedical data from the human body and provide continuous health monitoring. Typically, wearable and implantable sensors have small batteries and, thus, a tight power budget. Consequently, to extend the battery life, it is essential to minimize the amount of transmitted data as well as the complexity of the algorithms used for data compression. In this chapter we focus on possible solutions to improve the compression of ECG signals collected by means of wearable sensors. To this purpose, compressive sensing is employed in order to compress the acquired data while keeping the acquisition device as simple as possible.

As discussed in Chapter 2, thanks to the knowledge of the signal structure, compressive sensing enables to reduce the number of measurements required by the Shannon theorem, while still being able to perfectly recover the signal. In order to increase the quality of the reconstructed signals, the choice of a good sparsifying transform is essential. For this reason we introduce a newly designed universal overcomplete dictionary composed by Gaussian functions.

In this chapter we briefly introduce the heart's electrical activity as well as the generation of the ECG signal, useful to better understand the ECG model and sparsification dictionary. In particular, we evaluate the performance of ECG compression using compressive sensing and exploiting the signal sparsity over the proposed Gaussian dictionary. The performance is evaluated in terms of quality of the reconstructed signal.

3.2 ECG Waveform Generation and Recording

The human heart is an efficient and reliable pump, that propels over 6000 liters of oxygen-rich blood throughout the body every day and beats over 4 million times a year [77]. The

pumping function is accomplished by the cardiac muscle: the myocardimum, whose coordinate contraction depends on propagation of electrical impulse generated in the sinoatrial node (SA node). The SA node has the ability to undergo spontaneous depolarization sixty to ninety times each minute when it is behaving normally and, the electricity it produces passes into the right atrium and then into the left atrium. These bioelectrical events are regulated within very tight limits to allow the coordinated propagation of excitation and contraction of the heart that is necessary for an efficient cardiac output. Abnormalities in the regulatory mechanisms often accompany cardiac disease [118].

The electrical activity of the heart can be characterized by measurements acquired from the cellular level as well as from the body surface, since a potential difference produced by the current flowing within the body are established on the surface. Within this work we consider electrocardiographic measurements collected from the body surface, which describe the different electrical phases of a cardiac cycle (see Fig. 3.1). The waveforms present in each ECG cycle and produced during depolarization and repolarization deviate from a baseline level which corresponds to the resting state of the cells. The characteristic waves are usually labeled with the letters P, QRS, T and U [55].

The P-wave corresponds to the depolarization of the right atrium and subsequently of the left atrium. In a normal ECG its amplitude is less than 300 μV and its duration is less than 120 ms.

The Q, R, and S waves are usually treated as a single composite wave known as the QRS-complex. It reflects the depolarization of ventricles, and indicates the start of ventricular contraction that pumps blood to the lungs and the rest of the body. In a normal heart beat its duration is about 70-110 ms, and it may reach an amplitude of 2-3 mV. The Q wave is the initial downward deflection of the QRS complex and the S wave is its terminal one, while the R wave is the upward deflection.

The T-wave corresponds to the repolarization of the ventricles, which is a necessary recovery process for the myocardium to depolarise and contract again. The end of the T-wave coincides with the end of ventricular contraction. The wave corresponding to the depolarization of the atrium (Ta) is usually not visible since it coincides with the QRS complex and is buried in the larger waveform. Its duration is about 300 ms. The U-wave, which appears after ventricular repolarization, may not be seen in a normal ECG.

An ECG recording is typically affected by many sources of noise. The two dominant artifacts in ECG recordings are baseline wander and electromyogram noise. The baseline wander can be caused by respiration, body movements, and poor electrode contact, and its spectral content is usually in the range between 0.05-1 Hz. It may cause problems in the detection of R peaks, indeed due to the wander, the T peak could be higher than R peak and leading to a wrong detection. The electrical activity of muscles causes the electromyographic noise (EMG noise), commonly seen in ECGs recorded during exercise as a high-frequency noise. Interference from nearby equipment is the cause of another common source of noise: the power-line interference, typically at 50/60 Hz.

3.3 Related Works

Sparse modelling of ECG signals has recently received much attention as it has shown promising results in different applications, such as compression and denoising. Since the sparse model/dictionary is the key for an efficient compression and reconstruction of the

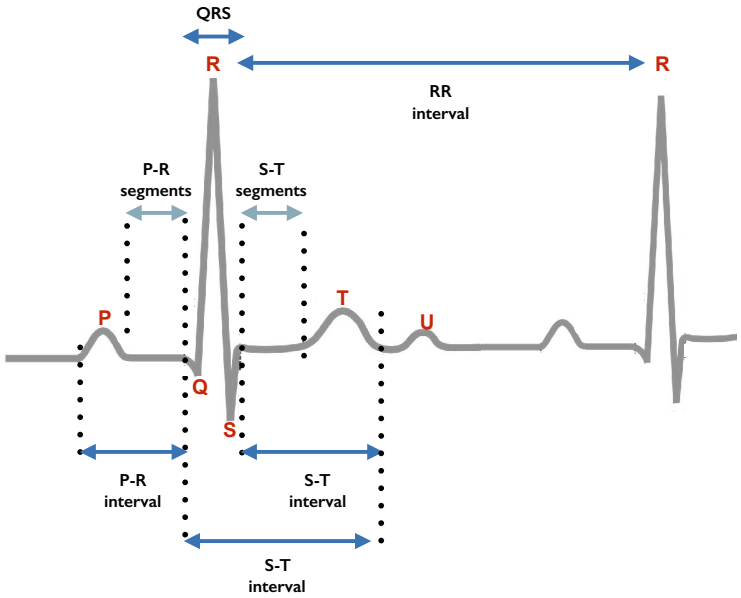

Figure 3.1: Normal futures of an electrocardiogram signal.

signal within the CS framework, over the past few years extensive work has been done to find a good sparsifying transform for ECG signals. In this context, most prior works employ orthogonal sparsifying transforms such as the wavelet transform because of their straightforward implementation. In Mamaghanian *et al.* [90], the orthogonal Daubechies wavelets are used to sparsify the ECG signal. The authors also proposed in [89] a reconstruction strategy that exploits the correlation between leads. This scenario requires to acquire multiple channels, hypothesis that dose not always apply in the context of wearable devices.

Even if the use of a wavelet basis to create sparse representations of ECG signals has been commonly employed with CS recent works seek to find alternative method to increase the generally unsatisfactory results.

An algorithm taking advantage from the sparsity on the second-order difference of the signal is proposed by Pant *et al.* [101]. This method minimize the l_p pseudo norm of the second-order difference of ECG signals and utilizes patient specific dictionaries to optimize the reconstruction quality.

Block-Sparse Bayesian Learning (BSBL-BO) has been introduced by Zhang *et al.* [127], it exploits the intra-block correlation that exists in time domain ECG signals. The reconstruction problem is solved without using any sparsifying dictionary matrix or basis, instead it exploited the block-sparsity structure of the ECG in the time domain and used BSBL-BO for the reconstruction.

Polania *et al.* [108] proposed an adaptive dictionary learning method that exploits the multi-scale sparse representation of ECG. The dictionary is built from training signals in the wavelet domain. It requires a preprocessing stage to detect the QRS complex and a period normalization. Therefore, it is necessary to store/transmit the beat information, adding computational complexity to the encoder. Moreover, the technique proposed in

[108] requires a learning procedure based on the actual patients ECG trace.

3.4 Gaussian Dictionary for ECG Sparsification

3.4.1 ECG Mathematical Model

Representing the ECG trace by its wavelet transform allows to compactly represent the signal, exploiting sparsity in the transform domain. Indeed, only a few wavelet coefficients carry most of the signal energy. However, despite the fact that various researchers have reported that ECG signals are sparse in the wavelet basis [3] [84] and that this basis is usually adopted in compressive sensing based telecardiology solutions [107] [90], the application of a more sophisticated model for ECG sparsification is desirable in order to increase the achievable compression performance.

In the work presented in this chapter, we propose to use a model for the ECG signal based on a simple class of mathematical functions, and then to design an efficient dictionary representation around this model.

We consider the Gaussian wave-based dynamical model introduced by McSharry *et al.* [91] used for generating realistic synthetic ECG signals.

This model represents the feature waves (i.e., P, Q, R, S and T waves, see Fig 3.1) of the ECG signal by using Gaussian functions, each with three parameters: amplitude, width and phase.

A one dimensional Gaussian function is defined as below

$$g_{p,b}(n) = a \exp \frac{-(n-p)^2}{2b^2}, \tag{3.4.1}$$

where a is the amplitude, p is the position/phase and b the width or shape parameter.

Using this model the morphology of an ECG cycle can be described by an ordinary differential equation

$$\dot{z}(a_i, b_i, p_i) = -\sum_{i}^{K} a_i \Delta p_i e^{\frac{-\Delta p_i^2}{2b_i^2}}, \tag{3.4.2}$$

where $\Delta p_i = n - p_i$ is the relative position, a_i is the amplitude and b_i is the width/shape of the i-th wave.

This model has been used to accomplish different tasks, such as filtering, compression and classification of the ECG [39] [115]. To fit an observed ECG cycle $x(n)$ with the model in Eq. 3.4.2 one can solve the following non linear least-square optimization problem

$$\min_{a_i, p_i, b_i} ||x(n) - z(n)||_2^2, \tag{3.4.3}$$

where $z(n) = \sum_i 2a_i \Delta p_i \exp(-\Delta p_i / 2b_i^2)$.

As a measure of sparsity that one can achieve by using this model, we may consider the number of Gaussian functions that are usually required to represent a noiseless ECG cycle. Accordingly to [37], all the symmetric waves, such as the Q, R, and S waves, can be approximated by one Gaussian function, while the asymmetric P and T waves are typically modelled by two Gaussian functions. Thus, an ECG cycle may be approximated using about K=7 functions.

Fig. 3.2 shows an actual ECG cycle approximated as a linear combination of just 5 Gaussian functions (K=5). Each wave has a different shape parameter defined by b_i and is located at p_i.

As mentioned before, the ECG signal is usually contaminated by many sources of noise and therefore, by using the model discussed in this section, one can achieve only an approximation of the original signal. Fig. 3.2 also shows the residual error between the signal and its approximation with Gaussian functions.

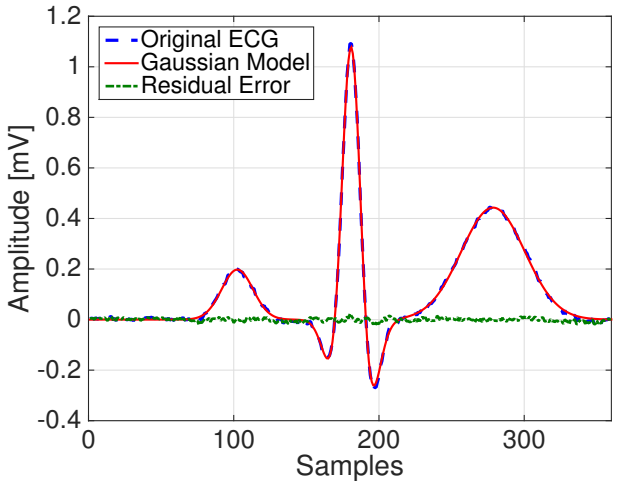

	p_i	a_i	b_i
P	102.48	0.1978	10.80
Q	169.27	-0.0564	7.35
R	181.28	3.1056	7.36
S	182.17	-2.0261	9.36
T	278.51	0.4430	21.95

Figure 3.2: Approximation of and ECG cycle sampled at 360 Hz using 5 Gaussian functions, with the relative parameters.

3.4.2 Dictionary Design

The Gaussian dictionary proposed in this work aims to be universal and does not need any training procedure. To this purpose, we need to make it independent from the specific heart beats in order to avoid any kind of synchronization or normalization on segments to be compressed. Moreover, it should be independent from the patient own signal and able to represents abnormal patterns.

So far as the main components of an ECG signal are generated by following the same model, independently from the patient, we propose to build the dictionary based on the generative model presented in the previous section (3.4.1).

Others dictionaries proposed in literature [61] [108] and based on dictionary learning require to decompose the ECG trace on cycles of the same length and centered on the R-wave. Thus, before applying compressive sensing the incoming samples need to be processed in order to find the R-wave position, increasing the computational cost. This solution is preferable in applications that require to store a large amount of data but without restricting limits on the energy consumption of the algorithm.

Based on the model previously discussed, we build a dictionary matrix $\mathcal{D}$ whose columns correspond to atoms computed as samples of the Gaussian function for suitable values of p_i and b_i.

The proposed dictionary is composed of P vectors (i.e., atoms) $\mathbf{g}_i$, $\mathcal{D} = \{\mathbf{g}_i\}_{i\in\Gamma}$, such that each ECG signal can be sparsely approximated by a subset of vectors $\{\mathbf{g}_i\}_{i\in\Lambda}$ in $\mathcal{D}$.

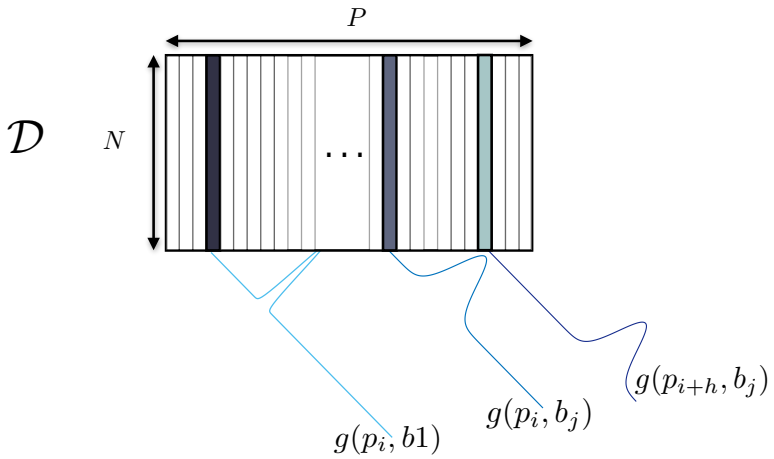

Figure 3.3: Dictionary of Gaussian like functions for ECG signal sparsification.

Several strategies were tested with different sets of values of the scale parameters and different number of atoms to give an efficient approximation of typical ECG waves. In particular, assuming $\mathbf{x}$ is an $N = 256$ ECG sample vector, the atoms are computed for scale parameters $b_i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 20, 50\}$. These parameters have been selected accordingly to the typical amplitude of the ECG characteristic waves. R-waves, as well as S and Q, are usually represented by Gaussian functions with a small width, thus using the first scale parameters among those considered for the dictionary. Whereas, the P and T waves exhibit larger width and are typically computed using the last scale parameters among those considered.

All the possible integer shift parameters values $p_i = (p_1, p_2, \ldots, p_N)$ within the vector are considered for each of the scale parameter (see Fig. 3.3). This means that an ECG cycle may be completely characterized by a subset of columns in the dictionary and the relative amplitude coefficients.

As a measure of how the proposed dictionary is suitable for ECG signal sparsification we cannot evaluate the maximum coherence of the dictionary because it is highly pessimistic, and actually it is not required in order to reconstruct the ECG trace $\mathbf{x} = \mathcal{D}\mathbf{s}$, as discussed in Chapter 2. Therefore, to show the compressibility of the ECG signal over the proposed dictionary we plot the power decay for ECG cycles represented in the sparsifying dictionary. We compute the approximation coefficients using the OMP algorithm and sort them. In Fig. 3.4(a) the absolute value of coefficients are plotted versus index sample, using logarithmic scale. The dashed line indicates the decay rate of $q = -1$. The power decay curve is computed for a fixed signal window $N = 256$ and averaged over all the ECG records from the MIT-BIH Arrhythmia Database [95, 64]. As discussed in Sec. 2.2.3, the faster the coefficients decay, the more compressible they are. As one can see in Fig. 3.4(a) the sparse representation coefficients corresponding to the use of the Gaussian dictionary exhibit a power law decay.

We compare the approximation SNR using a traditional Wavelet basis approach and the Gaussian dictionary, varying the number of retained coefficients. Fig. 3.4(b) shows the approximation SNR versus the sparseness factor, which is defined as the ratio between

the number of non zero coefficients in the sparse approximation and the length of the original signal $Sf = k/N$. The SNR using the proposed dictionary is higher than using the Wavelet basis for very low values of Sf. Thus, it is possible to use less coefficients to sparsely approximate the signal. Thus, one can expect to use less measurements in the CS compression and achieve better reconstruction quality with respect to Wavelet basis.

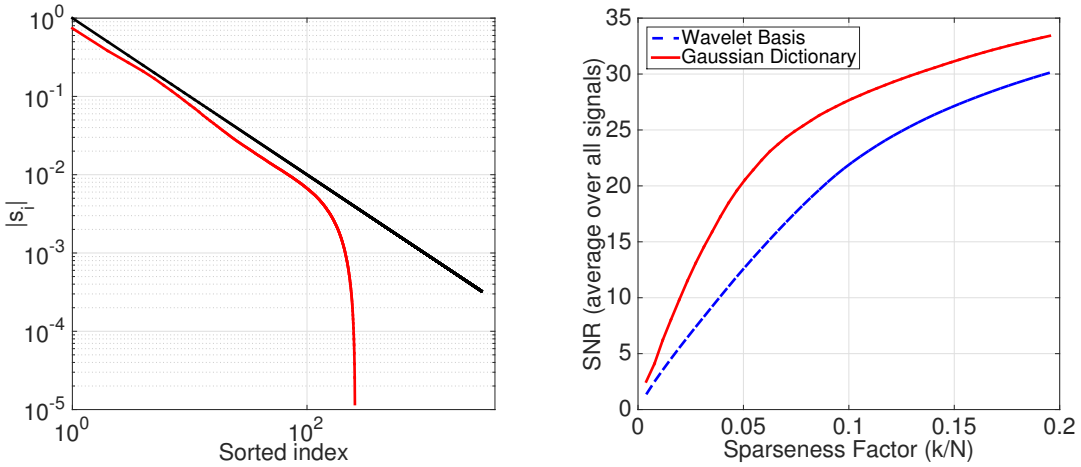

(a) (b)

Figure 3.4: (a) Power decay curves for sparse representation of ECG over the Gaussian dictionary; (b) average SNR for different sparseness factors.

3.5 ECG Signal Recovery Using the Gaussian Dictionary

3.5.1 Experimental Set-Up

Accordingly to the CS-based compression scheme introduced in Chapter 2, we collect M samples of an ECG segment $\mathbf{x}$ of length N, using sensing matrix $\mathbf{\Phi}$, and the compressed data vector is $\mathbf{y} = \mathbf{\Phi x} \in \mathbb{R}^M$. The sensing matrix used in this experiment is a Gaussian matrix $\mathbf{\Phi}$, having i.i.d. entries drawn from a standard normal distribution. Measurements $\mathbf{y}$ are quantized with an 11-bit uniform scalar quantizer.

The compression scheme should work in real-time, thus the measured vectors should be sufficiently short and we consider ECG segments $N = 256$ non overlapping samples, which is less then 1 s at $fs = 360Hz$. The proposed procedure does not require any preprocessing of the signal before compression (e.g. synchronization, normalization).

At the receiver said, we aim to reconstruct the ECG $\tilde{\mathbf{x}}$ segment from the received measurements that may be corrupted by noise $\tilde{\mathbf{y}} = \mathbf{y} + \mathbf{n}$. Here we use two algorithms to recover $\tilde{\mathbf{x}}$, i.e., the Basis Pursuit Denoising (BPDN) using the implementation provided by the popular SPGL1 [123, 122] and the Orthogonal Matching Pursuit (OMP). To validate the performance of the Gaussian dictionary (GD) we also repeat the experiments using the Wavelet basis (WT) as a sparsifying transform. In particular we adopt the orthogonal Daubechies wavelet (DB4) with 4 levels of decomposition, which provides a sparse

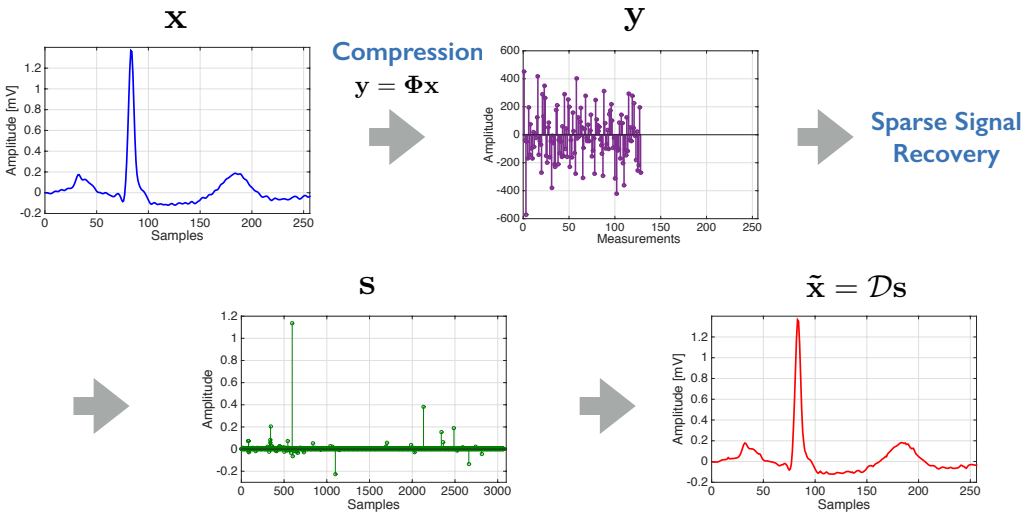

Figure 3.5: Compressive sensing applied to an ECG segment sampled at 360 Hz of length $N = 256$ and its reconstructed version.

representation for piecewise-linear signals and thus is suitable for ECG signals, leading to a relatively sparse representation with most of the coefficients close to zero.

We also compare the reconstruction results obtained by using the proposed Gaussian dictionary with respect to the BSBL-BO reconstruction procedure [127].

All simulations and time measurements are performed on a laptop with an Intel(R) core(TM) i7 2.40 GHz processor using the commercial software package MATLAB, The MathWorks Inc., Natick, MA, version 2013a.

3.5.2 Evaluation

Assessment of CS-based schemes performance are usually based on two parameters, which are also applied by Mamaghanian et al. [90], by Polania et al. [109] and by Dixon et al. [51] to asses the quality of the reconstruction. Thus, we make use of the following parameters

- Compression Ratio (CR) that takes into account the number of samples representing the original ECG signal and the number of CS ECG signal measurements, expressed by

$$CR(\%) = 100 \times \frac{N - M}{N}, \tag{3.5.1}$$

where M represents the number of samples in the CS domain and N is the number of samples in the original signal. The CR parameter can take into account quantization, by considering the number of bits necessary to represent each sample both in the CS and original signal domains.

- Percentage root-mean-square difference (PRD) that is a measure of the distortion error in the reconstructed signal $\tilde{\mathbf{x}}$ as compared to the original ECG signal $\mathbf{x}$. It is expressed by

$$PRD(\%) = \sqrt{\frac{\sum_{n=1}^{N}(x(n) - \tilde{x}(n))^2}{\sum_{n=1}^{N} x(n)^2}} \times 100, \tag{3.5.2}$$

  where it is assumed that the mean value of the original ECG signal has been subtracted. Zigel *et al.* [129] established the relation between the diagnostic distortion of the reconstructed ECG and the measured PRD. The proposed weighted diagnostic measure for ECG signal compression classify the PRD values on the basis of the signal quality perceived by a specialist. Accordingly, reconstructions with PRD values between 0% and 2% are qualified to have "very good" quality, while values between 2% and 9% are categorized as "good". For higher PRD values, quality classification is inconsistent.

3.5.3 Dataset

To validate the performance of the considered dictionary for the reconstruction of compressed sensed ECG signals we use the MIT-BIH Arrhythmia Database [64]. This public dataset contains 48 half-hour signals of two-channel ambulatory ECG recordings, obtained from 47 patients. The recordings were digitized at 360 samples per second per channel with 11-bit resolution over a 10 mV range.

We evaluate the proposed method over all database MLII lead signals. Experiments are carried out by extracting the first 5 minute long signal trace from each database record. In particular, we evaluate the reconstruction quality over a set of signals with different rhythms, wave morphologies and normal/abnormal heartbeats.

3.5.4 Experimental Results

Each five minute long ECG signal with a length of 108000 samples used in the experiment needs to be equally divided into segments of length 256 in order to be processed. For each segment, after signal recovery, we evaluate the PRD. Results are averages of the performance for all segments in each trace, and for all traces extracted from the records. Table 3.1 summarizes the performance of the CS reconstruction exploiting sparsity in the Gaussian dictionary or Wavelet basis and the two different reconstruction methods.

Fig. 3.6 compares the output PRD, averaged over all database records, for CS compression using the orthogonal wavelet transform and the proposed Gaussian dictionary at different CR values and using various recovery algorithms. It shows also the reconstruction quality achieved by using the BSBL-BO algorithm at different CR.

As we can see, using the Gaussian dictionary combined with the BPDN algorithm allows to improve the reconstruction quality with respect to Wavelet basis and also with respect to the BSBL-BO. Moreover, one can achieve "good" reconstruction quality for compression ratios up to CR=75% by applying the proposed dictionary and BPDN or OMP recovery algorithms. It is possible to obtain similar PRD values using the wavelet basis as sparsifiying transform in the CS reconstruction, however up to a smaller CR$\simeq$ 50%.

Fig. 3.8 shows box plots for the reconstruction quality obtained by applying the BPDN algorithm in combination with the proposed dictionary (top) and with the wavelet transform basis (bottom). The proposed method shows a smaller variation of the PRD param-

Table 3.1: Average PRD value for ECG reconstruction from CS measurements exploiting sparsity in the Gaussian dictionary (GD) or Wavelet basis (WT) and two different reconstruction methods, i.e., Basis Pursuit Denoising (BPDN) and the Orthogonal Matching Pursuit (OMP). Results for the BSBL-BO approach [127] are also reported.

Reconstruction Method	**Sparsifying Dictionary**	**Average PRD** [%]							
		CR [%]							
		40	50	60	70	75	80	85	90
BPDN (Sec. 2.4.1)	Gaussian Dictionary	2.77	3.54	4.74	5.57	6.86	8.82	12.32	18.57
	Wavelet Basis	5.55	7.71	13.61	28.91	41.28	56.86	74.38	91.09
OMP (Sec. 2.4.2)	Gaussian Dictionary	6.80	6.88	6.94	7.57	8.84	11.02	17.40	32.58
	Wavelet Basis	6.61	8.18	14.16	35.76	54.91	88.69	161.36	167.55
BSBL-BO [127]	- -	4.15	4.27	5.13	9.55	15.83	27.53	42.43	68.04

eter for all the CR values. The boxplot function from the commercial software package MATLAB Statistical Toolbox with Tukey whiskers has been used. The bottom of each box represents the 25th percentile and the top is the 75th percentile, each middle line representing the median.

From Tab. 3.1 and Fig. 3.6, we can see that using the proposed Gaussian dictionary the performance of OMP is worse than that of BPDN, especially at low compression rate. Note, however, that the reconstruction time is higher for BPDN recovery (see Fig. 3.7). When using the BSBL-BO one can compress the original signal up to CR=69% with a "good" reconstruction quality.

Finally, we provide a visual comparison of the reconstructed signals, both for the wavelet-based scheme and the proposed framework. Fig. 3.9 shows record 221 of the database [64], which comprises normal heartbeats as well as premature ventricular contractions (PVC).

It can be seen that, even for a compression ratio $CR = 80\%$, a good reconstruction is achieved by the proposed scheme, preserving detailed information for clinical diagnosis, such as amplitude and shape. When using WT the QRS complex amplitudes are reduced respect to the original one, and the shape of the T-waves are contaminated by noise.

3.6 Comparison of Sparse and Gaussian Sensing Matrices

So far we analyzed the performance of the proposed Gaussian dictionary against traditional Wavelet basis (DB4), showing the superior reconstruction quality achievable by the

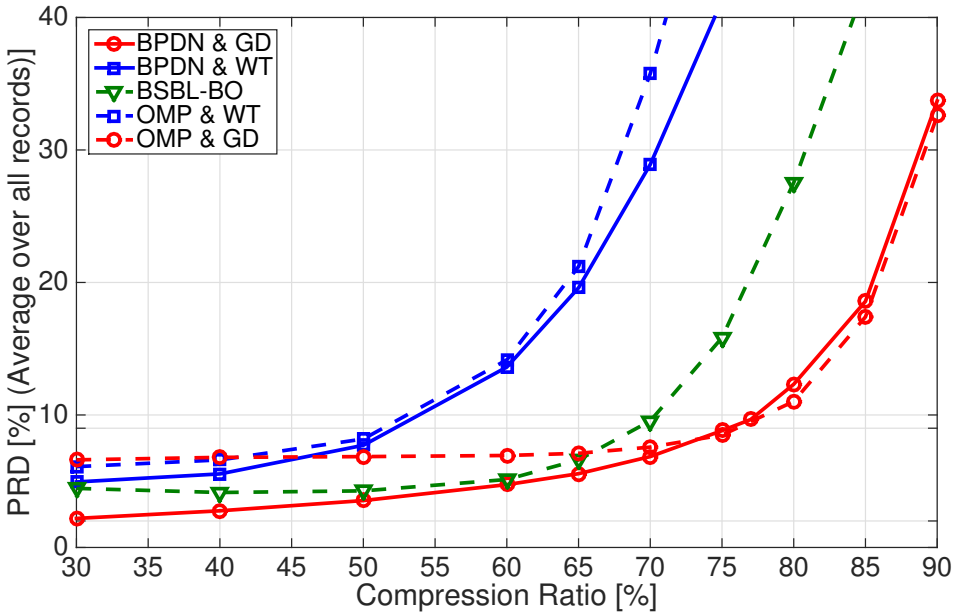

Figure 3.6: Performance comparison between different CS implementations. Average PRD for BPDN and OMP algorithm using the Gaussian dictionary (GD) and Wavelet basis (WT), and BSBL-BO algorithm. The proposed dictionary combined with BPDN reconstruction reports the best performance for CR<75%, then both OMP and BPDN combined with the GD lead to same reconstruction quality.

Table 3.2: Sensing matrix computational complexity and storage space allocation of the CS encoder

Matrix	Add	Mul	Storage Space
Gaussian	$(n-1)\times m$	$n \times m$	$n \times m$
d-sparse	$d \times n - m$	0	$n \times d$

Gaussian dictionary. Different algorithms have been employed to solve the signal recovery problem from compressed measurements obtained by using random Gaussian matrices. However, random Gaussian sensing matrices require a high computational cost, necessary to perform a large matrix multiplication on the sensor, and also high storage requirements. For a real-time and very low power implementation, we may need to replace random Gaussian matrices with more computationally tractable sensing matrices. Thus, in this section we analyze the performance of the proposed dictionary when the compression is performed by using sparse sensing matrices. In Table 3.2 the computational complexity and storage space allocation needed by the different CS sensing matrices is reported.

Sparse sensing matrices have been used for the CS of ECG signal in [90] and [127], where the authors proposed to use matrices with $d = 12$ non-zero elements per column equal to $1/\sqrt{d}$ and $d = 2$ non-zero elements equal to 1, respectively. In [90] the orthogonal Wavelet basis has been considered as the sparsifying transform, while in [127] they use the

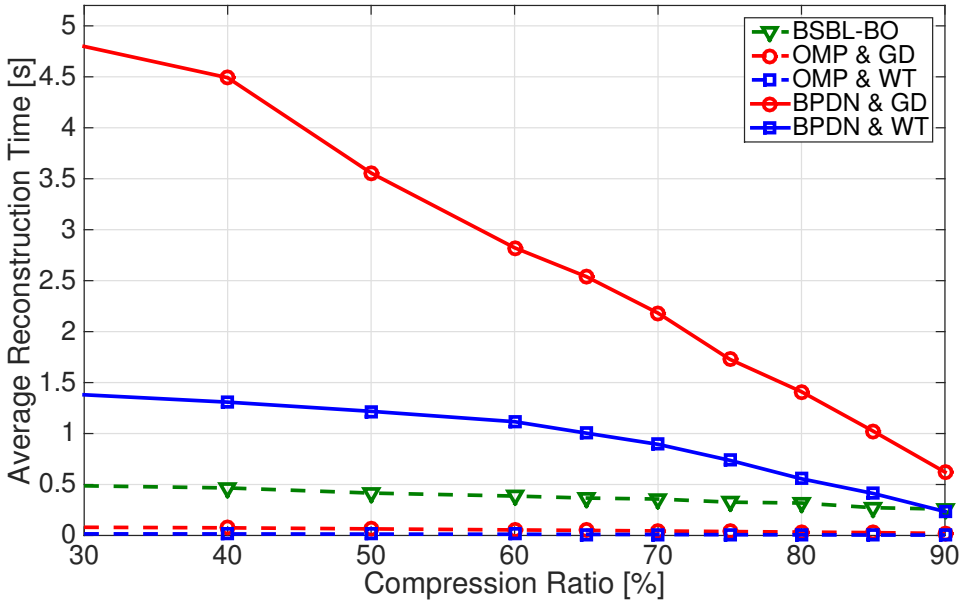

Figure 3.7: Performance comparison between different CS implementations. Average recovery in seconds required to reconstruct one block (N=256 samples) of signals. The OMP algorithm independently on the sparse representation adopted results the faster, while BPDN combined with GD shows the higher recovery time. Note that BPDN&GD at CR<80% is not suitable for real-time reconstruction.

BSBL-BO as the recovery technique.

In this section, experiments are designed to assess if sparse random matrices have comparable performance with random Gaussian sensing matrices when using the proposed Gaussian dictionary. When sparse sensing matrices are employed, there exists the problem of selecting the number of non-zero elements d. Thus, we aim to identify the minimum value of d that allows to keep the number of operations during the compression as low as possible and achieve an acceptable reconstruction quality. To do so, sparse sensing matrices with different $d \in [2, 3, \ldots, 12]$ at different compression ratios are used to compress the ECG signals and the output PRD is measured. We use the same dataset and the block size as in Section. 3.5. Fig. 3.10 reports the resulting average output PRD versus the number of non-zero elements d in the sparse binary sensing matrix at different compression ratios. The output PRD values are color coded. The obtained results show that the reconstruction quality is not affected by d, differently from the results reported in [90] in which the performance was very sensitive to the number of non-zeros elements. This can be another advantage of the proposed Gaussian dictionary with respect to the traditional Wavelet basis adopted in [90], since it allows to use less non-zero elements and further reduce the energy consumption while maintaining the same PRD. Thus, in the next experiment we use sparse binary sensing matrices with $d = 2$.

We now compare the performance of Gaussian random matrices and sparse matrices ($d = 2$). The average reconstruction quality is evaluated on the first five minutes of all the

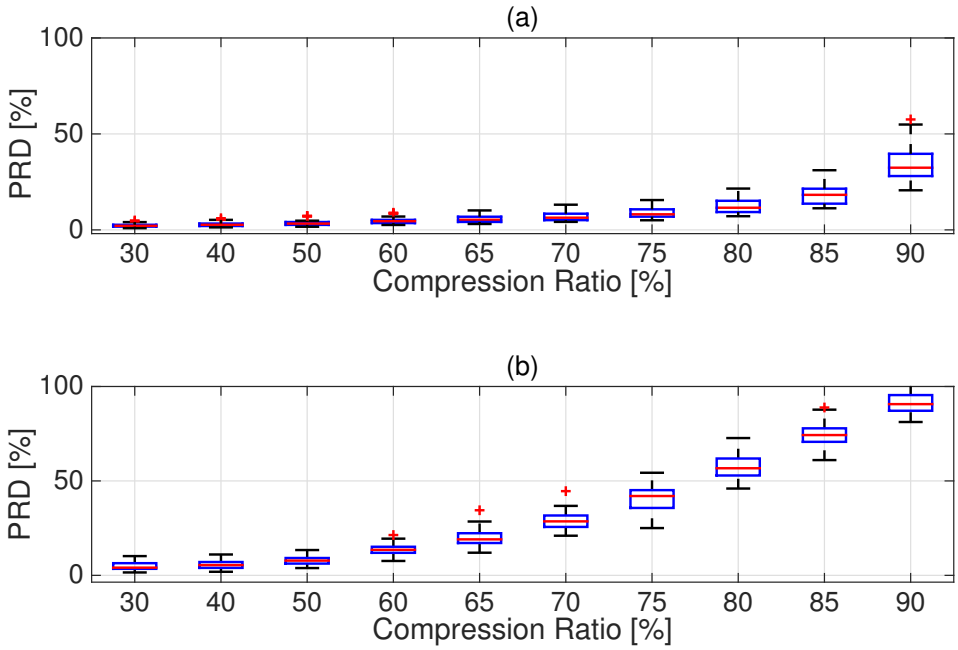

Figure 3.8: Box plots for Gaussian dictionary (GD) and BPDN reconstruction (top) and Wavelets (WT) and BPDN reconstruction (bottom).

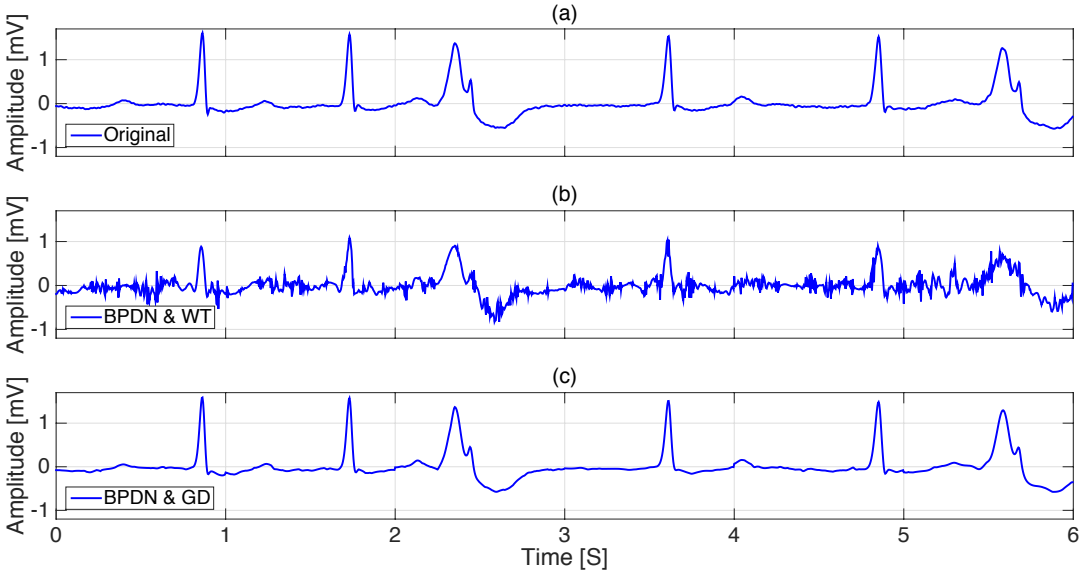

Figure 3.9: Visual evaluation of the reconstruction of record 221 compressed using the CS framework at $CR = 80\%$. (a) The original uncompressed signal, (b) recovered signal using BPDN and the Wavelet basis (WT), $PRD = 52.11\%$, and (c) recovered signal using BPDN and the proposed Gaussian dictionary (GD), $PRD = 6.51\%$.

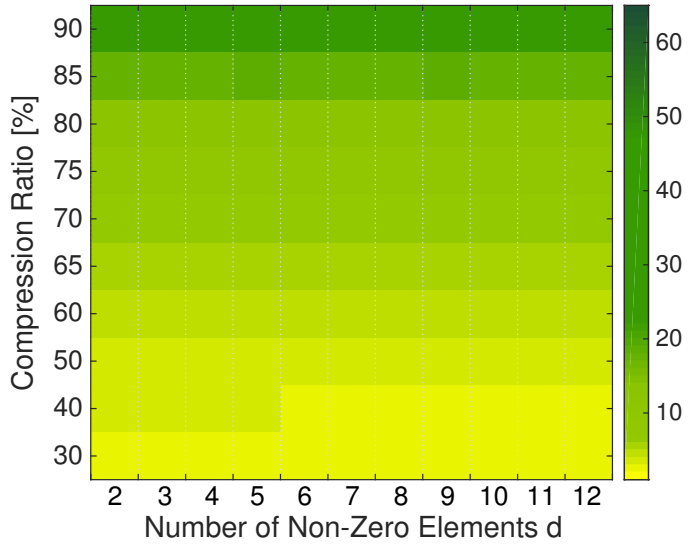

Figure 3.10: Relation between recovery quality and the number d of 1 entries at different compression ratios. The output percentage PRD values are color coded.

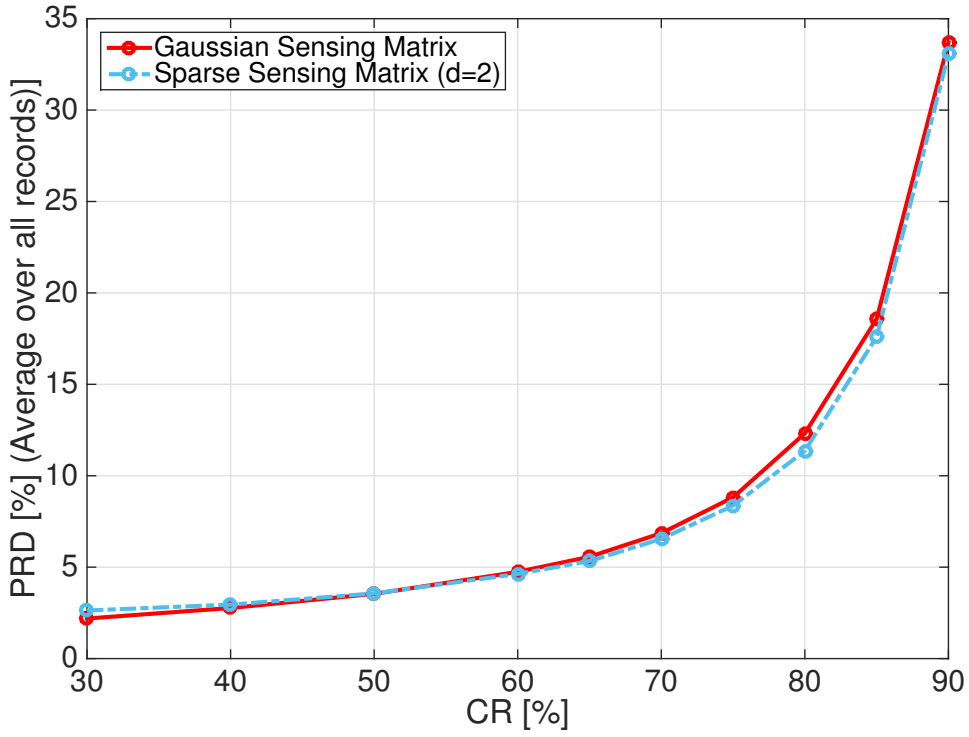

Figure 3.11: Performance comparison for different sensing matrices. Average PRD for i.i.d. Gaussian sensing matrix and sparse sensing matrix with $d = 2$ non-zero elements equal to 1 in each column. The average PRD after reconstruction using BPDN and the Gaussian dictionary are comparable for the two classes of sensing matrices.

MLII records of the MIT-BIH Arrhythmia ECG database, and we measure the average output PRD of the reconstructed segments. In each segment, a different sparse sensing matrix, with $d = 2$, is applied. Fig. 3.11 reports the average PRD at different compression ratios, and it shows very little performance difference between the sparse sensing and the Gaussian matrix.

3.7 Discussion

Results suggest that, independently of the recovery algorithm, the design of a more efficient dictionary, such as the one analyzed in this work, permits to achieve good performance at higher compression ratios. Wavelets allow a shorter reconstruction time (see Fig. 3.7) but they are not suitable for an acceptable reconstruction quality at high CR values. On the overall, the proposed dictionary combined with BPDN reconstruction gives the best performance with respect to OMP and BSBL-BO, at the expense of an increased reconstruction time, which might not be a serious issue at the receiving processing side. Moreover, we verify that the use of sparse sensing matrices with only 2 non-zero elements in each column, which is useful for low-complex implementation, leads to PRD values comparable with the ones obtained with random Gaussian sensing matrices. It should be noted that this very low number of non-zero elements is not achiviable when the Wavelet basis is adopteed, as reported in [90].

Polina *et al.* [108] proposed a method which builds an adaptive multi-scale dictionary via an offline learning procedure on the actual ECG traces. The multi-scale dictionary $\mathbf{D}$ is composed by sub-dictionaries $\mathbf{D_b}$, corresponding to a specific wavelet subband $b = 1, \ldots, B$. Each sub-dictionary is learned separately. A $B-$level wavelet transform is applied on normalized ECG cycles. Therefore, this method requires a preprocessing stage to detect R-peaks and identify heartbeats. Since heartbeats may have different durations, interpolation/decimation is used to make all the ECG cycles the same length, in order to be consistent with dictionary entries. Note that the method we propose here does not need any pre-processing stage, making computation at the encoder side very simple.

Note that Polina *et al.* [108] only used 4 patients from the MIT-BIH Arrhythmia Database without explanation for excluding the other 43 subjects. They also trained and tested their algorithm on the same records/patients, with no cross validated results. Therefore, despite their claims of a high CR (80%) with a PRD of 6%, these results cannot be considered to be at all representative of generalizable, and are likely to be highly overfit.

3.8 Conclusions

In this chapter, we proposed a dictionary for the ECG sparse modelling, which allows improvement in recovery of ECG signals from compressed sensed measurements. Experiments show that significant performance gains, in term of achieved reconstruction quality at high CR, could be obtained by using the proposed Gaussian dictionary, beyond the common Wavelet basis. In particular, our results show that for real ECG signals, the compression ratio can be increased by 25% using Gaussian based sparsity, with respect to respect to the state of the art using a wavelet basis approach.

The proposed dictionary does not require any training procedure, and is independent of a patient's specific ECG traces. It can be easily adapted to work with different length

of compression window, without requiring to be retrained.

We also showed that the choice of the sensing matrix and the reconstruction dictionary play a crucial role. The use of Gaussian matrices can make the CS procedure problematic for low-power implementation, while the use of sparse sensing matrices and the proposed Gaussian dictionary makes the procedure competitive.

Our results validate the suitability of compressed sensing for real-time continuous monitoring of ECG signals for low-power and resource-constrained sensors.

4

Beat Detection on Compressed Measurements

4.1 Introduction

This chapter addresses the problem of heart rate estimation from compressive sensing ECG recordings, avoiding the reconstruction of the entire signals. We consider a framework where the ECG signals are represented under the form of CS linear measurements and the QRS locations are estimated in the compressed domain.

As introduced in the previous chapters, compressive sensing theory states that signals, which are sparse in some domain, can be fully reconstructed using a small number of random projections, much smaller than the number required by Nyquist-rate sampling. Unfortunately, the signal recovery process involves algorithms with a relatively high computational load. Recovering long term recordings might require a long time and high computational resources. Moreover, the reconstruction process should be able to preserve clinical relevant information, necessary to assess the patient's health status.

While various methods have been reported to solve the problem of signal recovery from compressed sensed measurements, the full signal recovery might be not necessary for some clinical evaluation based on heart rate variability (HRV). The heart rate (HR) and its variability are calculated from the RR interval, which is the time occurring between two consecutive R-peaks, the QRS complex being the most prominent segment in every ECG cycle.
A large variety of methods for R-peak detection have been proposed, and the most commonly used is the one proposed by Pan and Tompkins [100] and its following modification [66]. An extensive review of most of the methods, such as derivative-based [100], wavelet transforms based [85] [35], matched filter [113], artificial neural networks [52], and many others can be found in [75].

In this chapter we present a novel R-wave detector working *directly* on the compressed sensed measurements. The proposed technique processes the signal in the compressed domain and estimates the cross correlation between the ECG and a known template. It can be used to analyze the signal and, only in the presence of detected abnormalities in the compressed domain, to recover it for further analysis. When the detection/analysis is performed on a battery operated device like a tablet or smartphone, or directly in the sensor without sending or storing the whole ECG trace, the proposed method will reduce the computational cost and the processing time, since it avoids the complexity of the recovery procedure.

The proposed correlation-based method is applied on CS data at several compression ratios to assess the feasibility and the impact of working in the compressed domain. The correlation based detection is also applied on uncompressed original data and the results are compared with those obtained after reconstructing the signals from the CS data. Finally, the benchmark Pan-Tompkins algorithm [100] is applied to the recovered data for further comparison. Experimental results show that the proposed approach provides little performance loss, with the advantage of limited complexity.

4.2 Background: Estimation with Compressive Measurements

Signal recovery from compressed measurements has been widely investigated in the past few years, leading to several accurate reconstruction algorithms with moderate complexity. However, some signal processing problems do not require full signal reconstruction, as shown in [47], where manifold-based image classification is performed on random measurements. To this end, new ways to directly process compressive measurements and extract information without having to reconstruct the original signal have been introduced in [46]. Davenport *et al.* [46] proposed a new approach to solve some signal processing problems directly in the compressive measurements domain, in particular the tasks of detecting, classifying, and estimating deterministic signals within the compressed domain. One of the problems considered in [46] is the estimation of a linear function of the signal vector $\mathbf{x}$ from compressed measurements $\mathbf{y}$. In particular, given a general and fixed test vector $\mathbf{h}$, the problem is to obtain from $\mathbf{y}$ a good estimate of $p = \langle \mathbf{x}, \mathbf{h} \rangle$. Two possible estimators are considered in [46], namely the *orthogonalized estimator* given by

$$\hat{p}_o = \frac{N}{M} \mathbf{y}^T (\mathbf{\Phi}\mathbf{\Phi}^T)^{-1} \mathbf{\Phi}\mathbf{h}, \tag{4.2.1}$$

and the so called *direct estimator*, which avoids the orthogonalization process, given by

$$\hat{p}_d = \langle \mathbf{y}, \mathbf{\Phi}\mathbf{h} \rangle. \tag{4.2.2}$$

When using a sensing matrix $\mathbf{\Phi}$ that is a *stable embedding*, which is a more general form of the RIP property, performance bounds for the direct estimator can be demonstrated.

Definition 4.2.1 *Let $\delta \in (0,1)$ and $\mathcal{U}$, $\mathcal{V} \subset \mathbf{R}^N$, $\mathbf{\Phi}$ is a* stable embedding *of $(\mathcal{U}, \mathcal{V})$ if*

$$(1-\delta)||\mathbf{u}-\mathbf{v}||_2^2 \leq ||\mathbf{\Phi}\mathbf{u}-\mathbf{\Phi}\mathbf{v}||_2^2 \leq (1+\delta)||\mathbf{u}-\mathbf{v}||_2^2,$$

for all $\mathbf{u} \in \mathcal{U}$ and $\mathbf{v} \in \mathcal{V}$.

If $\mathbf{\Phi}$ is a *stable embedding*, the following theorem provides a performance bound for the absolute error in estimating $\langle \mathbf{u}, \mathbf{v} \rangle$ from compressive measurements.

Theorem 4.2.2 (see [46, Theorem 4]) *Suppose that $\mathbf{u} \in \mathcal{U}$ and $\mathbf{v} \in \mathcal{V}$ and that $\mathbf{\Phi}$ is a* stable embedding *of $(\mathcal{U}, \mathcal{V} \cup -\mathcal{V})$, then:*

$$|\langle \mathbf{\Phi}\mathbf{u}, \mathbf{\Phi}\mathbf{v} \rangle - \langle \mathbf{u}, \mathbf{v} \rangle| \leq \delta ||\mathbf{u}||_2 ||\mathbf{v}||_2.$$

Thus, for an appropriate matrix $\mathbf{\Phi}$, using Eq. 4.2.2 one can obtain a good estimate of p by calculating a linear function of the compressed measurements. Moreover, it can be shown that the use of random $M \times N$ sensing matrices with sufficiently large M allows to satisfy the stable embedding property with high probability.

The bounds in Theorem 4.2.2 are useful also for the general case of the estimation of a linear operator represented by a matrix multiplication $\mathbf{Hx}$. We have in particular

$$\mathbf{Hx} = \begin{bmatrix} \mathbf{h}_1^T \\ \mathbf{h}_2^T \\ \mathbf{h}_3^T \\ \cdots \\ \mathbf{h}_Q^T \end{bmatrix} \mathbf{x} = \begin{bmatrix} \langle \mathbf{h}_1, \mathbf{x} \rangle \\ \langle \mathbf{h}_2, \mathbf{x} \rangle \\ \langle \mathbf{h}_3, \mathbf{x} \rangle \\ \cdots \\ \langle \mathbf{h}_Q, \mathbf{x} \rangle \end{bmatrix}, \tag{4.2.3}$$

so we can estimate the individual components in the compressed domain. In the following, we analyze an R-peak detector which calculates the correlation of the incoming signal and a known template, a problem which can be cast into the framework outlined above and which can be solved in the compressed domain.

4.3 ECG Model and Problem Formulation

We now give a description of the ECG model considered in this chapter. As already introduced in Sec. 3.2, an ECG cycle consists of several waves, labeled P, Q, R, S, and T. We can express the electrocardiogram signal $x_r(n)$, related to the recording taken during a time interval T, long enough to include at least one QRS complex, as

$$x_r(n) = x(n) + r(n) = \sum_i \alpha_i \psi(n - \theta_{Ri}) + r(n), \tag{4.3.1}$$

where α_i and θ_{Ri} are the amplitude and center of the given kernels, which we identify with the translated versions of the ECG beat QRS template. Moreover, $r(n)$ represents noise due to mismatch between the template and the actual QRS beat shape, and to other different sources, including P and T waves.

In the compressive sensing scenario, we cast $x_r(n)$ values into an N-dimensional vector $\mathbf{x}_r = \mathbf{x} + \mathbf{n}$ and observe few random measurements $\mathbf{y} \in \mathcal{R}^M$

$$\mathbf{y} = \mathbf{\Phi x} + \mathbf{n}, \tag{4.3.2}$$

where $\mathbf{n}$ represents the compressed noise and possibly additional noise due to the measurement process. Thus, given $\mathbf{y}$, we are interested in estimating the position of the R wave $\hat{\theta}_{Ri}$ in the original signal.

It should be noted that, since we have only access to the compressive measurements, no pre-processing can be performed in order to attenuate other signal components and artifacts, including the P-wave, the T-wave or baseline drift. Moreover, it is not suitable to perform some pre-processing tasks in the CS sensor, in order to reduce its complexity and avoid full rate processing.

A QRS detection method that can be successfully applied without preprocessing is based on template matching, or matched filtering. Indeed, a filter matched to the QRS

template can improve the output Signal-to-Noise ratio (SNR), attenuating the noise and emphasizing the QRS complexes. For this reason, R-peak detection methods based on matched filtering have been used extensively in the uncompressed domain (see, for example, Hamilton *et al.* [66], Kaplan [71] and Ruha *et al.* [113].) The Matched filter method seems suitable also for compressive sensing applications, as shown by Eftekhari *et al.* [57], in which a matched filter has been employed on a small number of frequency-domain observations for estimating the unknown delay and amplitude of a signal.

As we will see in the next section, the idea beyond the proposed method is to cast the matched filter approach into the compressive sensing signal processing theory introduced in Section 4.2. In particular, we propose to estimate the correlation between the compressed ECG signal and a compressed template (i.e., the compressed average QRS complex).

4.4 Compressive Beat Detection Based on Matched Filtering

Let us consider the model described in Eq. 4.3.1. Assuming white noise, the optimal strategy maximazing the signal-to-noise ratio of samples taken at time instants θ_{Ri} is to process $x_r(n)$ with a matched filter, with impulse response equal to the time reversed version of $\psi(n)$. In particular, one has to compute the correlation output

$$R_{x\psi}(n) = \sum_m x_r(m)\psi(m-n), \tag{4.4.1}$$

and time instants θ_{Ri} can be detected by searching for local maxima of the correlation $R_{x\psi}(n)$. As a matter of fact, the amplitude of $R_{x\psi}(n)$ at each time instant is a measure of how well the template matches the corresponding section of the input signal, and it achieves its maximum at each position where the observed signal best matches the template. As mentioned, R-peak detection methods based on matched filtering have been used extensively in the literature.

For finite length-N signal blocks $\mathbf{x}_r$, each correlation output can be computed as the inner product between $\mathbf{x}_r$ and an appropriately defined vector ψ_n whose non-zero elements correspond to the QRS template. Thus, the matched filter technique consists in calculating

$$R_{x\psi}(n) = \langle \mathbf{x}, \psi_n \rangle, \tag{4.4.2}$$

and selecting the indexes corresponding to maxima, in order to detect θ_{Ri}.

Using the results introduced in Sec. 4.2, in the proposed approach we estimate (4.4.2) in the compressed domain. In particular, given the measurements $\mathbf{y} = \mathbf{\Phi}\mathbf{x}_r$, we estimate the correlation $R_{x\psi}(n)$ by using the direct estimator, as follows

$$\hat{\mathbf{R}}_{\mathbf{x}\psi} = \begin{bmatrix} \langle \mathbf{y}, \mathbf{\Phi}\psi_1 \rangle \\ \langle \mathbf{y}, \mathbf{\Phi}\psi_2 \rangle \\ \langle \mathbf{y}, \mathbf{\Phi}\psi_3 \rangle \\ \cdots \\ \langle \mathbf{y}, \mathbf{\Phi}\psi_N \rangle \end{bmatrix}. \tag{4.4.3}$$

Note that each of the inner products in (4.4.3) is between length-M vectors, with no need to reconstruct the signal. As an alternative, we can use the orthogonalized estimator

$$\hat{\mathbf{R}}_{\mathbf{x}\psi,n} = \frac{N}{M}\langle \mathbf{y}, (\mathbf{\Phi}\mathbf{\Phi}^T)^{-1}\mathbf{\Phi}\psi_n \rangle. \tag{4.4.4}$$

In order to minimize the effect of baseline wander, it is convenient to subtract the signal mean from each block before computing (4.4.3). The signal mean can be estimated using the same approach suggested in [46]. As a matter of fact, we can write $\mu_{x_r} = \langle \mathbf{x}_r, \mathbf{l_m} \rangle$, where $\mathbf{l_m} = [\frac{1}{N}, \ldots, \frac{1}{N}]^t$. Using for instance the orthogonalized estimator, we have

$$\hat{\mu}_{x_r} = \langle \mathbf{y}, (\mathbf{\Phi}\mathbf{\Phi}^T)^{-1}\mathbf{\Phi}\mathbf{l_m} \rangle. \tag{4.4.5}$$

Then the signal mean can be removed from the measurements as $\mathbf{y} - \mathbf{\Phi}\hat{\mu}_{x_r}\,[1, \ldots, 1]^t$.

4.5 Summary of the Proposed Scheme

The flowchart of the proposed procedure is shown in Fig. 4.1. It consists of three main steps, namely, QRS template construction, correlation estimation using matched filtering on compressive measurements and peak-finding based on adaptive thresholding. Fig. 4.1 also includes the sensor structure: assuming a digital implementation, the sensor acquires and transmits the uncompressed ECG signal for a brief time, then switches to CS compression and transmits the compressed data. Alternatively, the first signal portion could be reconstructed at the receiver. This short uncompressed signal portion is used to identify the QRS complexes to generate the template, but it can be also useful to check that the electrodes are correctly placed and that the signals are being correctly recorded. The uncompressed signal should be sufficiently long to ensure that it contains enough beats in order to allow the template creation (i.e., 10 seconds).

Once the preliminary stage including tuning, electrode placement and patient-specific template generation is done, the long-term monitoring starts and the signal is compressed using the CS framework. The compressed measurements can be transmitted to a smartphone/computer for real time processing, or stored to be processed later.

Another possible option is to implement the detection on compressed measurements directly on the sensor, and then send only the ECG segments showing abnormal beat rate patterns. This increases power saving in all the situations in which the patient's heart rate pattern is normal. The results reported in this chapter also apply for this alternative scenario.

4.5.1 Template Generation

The selection of a good template is crucial in order to provide a reliable beat detector. In this work, the QRS template is created adaptively for each patient, using the detected QRS complexes within the first 10 seconds of the uncompressed ECG. In long-term recording, the template might be updated after a predefined time, by transmitting or reconstructing the uncompressed signal. A list of the R-peak fiducial points should be extracted from the available ECG signal by using an automatic detector, e.g., the Pan-Tompkins [100] detector. Then, the QRS complexes are extracted using a fixed time window centered at the R-peak position and with a fixed length of 100 ms, in order to contain most of the QRS complex. Finally, the QRS template is computed by taking the median of the extracted QRS complexes. An example of template is depicted in Fig. 4.3(c).

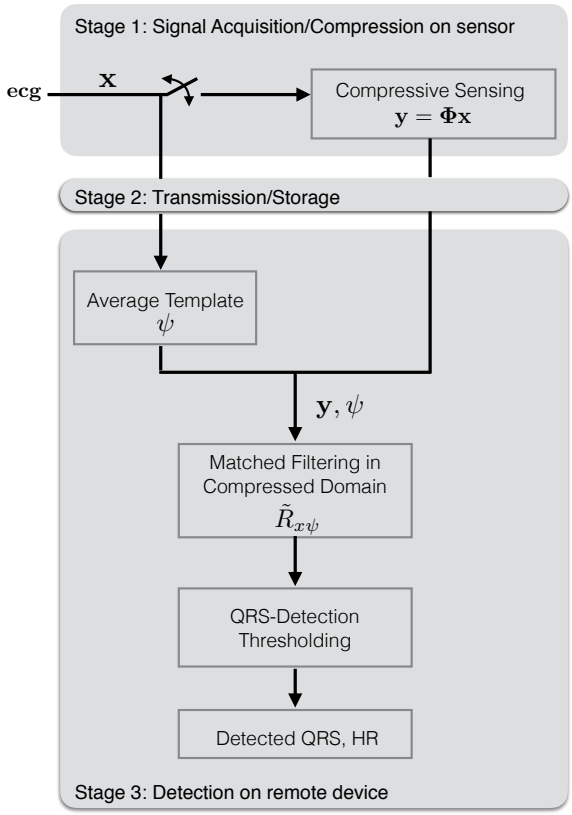

Figure 4.1: Block diagram of the proposed R-peak detection methodology using compressive sensing measurements.

4.5.2 Cross-Correlation Estimation

The main step in the proposed method consists in estimating the cross-correlation using the compressed ECG signal $\mathbf{y}$, given the reference QRS template. We assume that the sensing matrix is known at the receiver, so there is no need to transmit it. While in the experiments described later we use a different sensing matrix for each signal in order to take into account its influence on the average detection capability, in an actual implementation it is not necessary to use a different sensing matrix for each recording.

Thus, given the template ψ, the sensing matrix $\mathbf{\Phi}$ and the measurements $\mathbf{y}$, the direct or orthogonalized estimator is used to estimate the cross-correlation coefficients using Eq. 4.4.3 or Eq. 4.4.4.

An application of the proposed method on a real ECG signal is shown in Fig. 4.2, which is a section of the record number 107 of the MIT-BIH Arrhythmia Database [64]. In particular, Fig. 4.2(b) depicts the cross-correlation computed on the uncompressed data while the one obtained by using the orthogonalized estimator for a compression ratio of 50% is shown in Fig. 4.2(c).

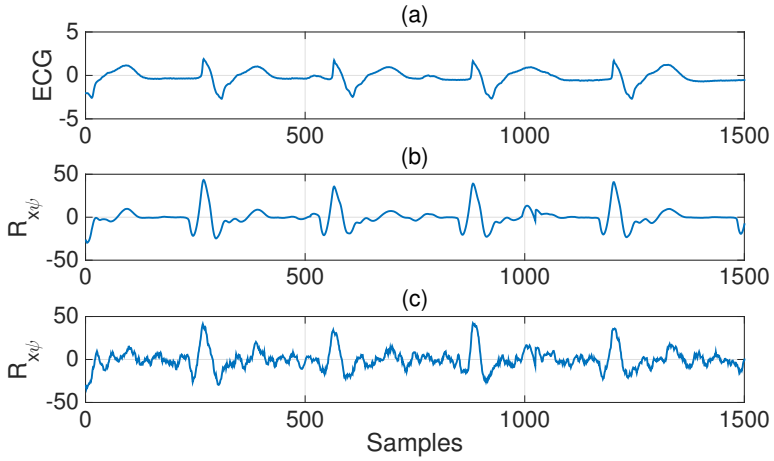

Figure 4.2: An example of correlation estimation: (a) original ECG signal sampled at 360 Hz, (b) correlation computed on the original signal, (c) correlation estimated on compressive sensing measurements (CR=50%).

4.5.3 Peak Detection

The last stage of the proposed method consists in the detection of the R-wave peaks by comparing the absolute value of the correlation against an adaptive amplitude-dependent threshold. The detection threshold is computed for each correlation window, i.e., for each measurement block, and it depends on the correlation amplitude in the current window.

As a matter of fact, as the compression ratio $1 - M/N$ increases, the estimated cross-correlation function typically becomes more noisy. Thus, applying a low threshold level, which might be appropriate for detection at low compression or on uncompressed data, is a poor choice, and it experimentally increases the number of false QRS complexes detection. Since the method should be able to provide good performance at every compression level, we use a variable threshold, which depends on the Root Mean Square (RMS) value of the cross-correlation in the current window.

After the RMS value is calculated, if it is larger than 25% of the maximum cross-correlation absolute value, the threshold is set to be 75% of the maximum value of the segment. If the RMS of the segment is less than 25% of its maximum value, the threshold is set to be 50% of the maximum value.

To avoid false detection in general, a refractory period of 200 ms is employed prior to repeating the process for the next cardiac cycle. This constraint is a physiological one due to the refractory period during which ventricular depolarization cannot occur.

Additionally decision rules for the reduction of false positive detection are applied. In particular, to avoid that a QRS peak that is situated in between two consecutive blocks is detected twice, when the distance between the last detected peak in the previous block and the first in the current one is less then 200 ms, we take the middle point as the R-peak location.
Similarly, missed beats can occur between two consecutive blocks. When the RR interval measured for R-peaks across two blocks is higher than 1.5 times the average RR computed on the previous 10 signal blocks, the threshold is adjusted to the half of its value and a

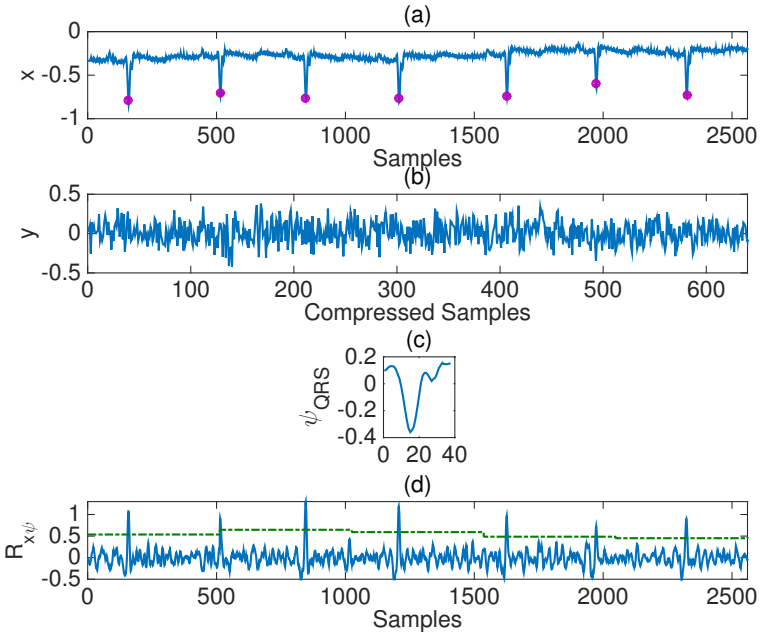

Figure 4.3: An example of beat detection: (a) original ECG signal sampled at 360 Hz and the detected beats, (b) the received compressive sensing measurements (CR=75%), (c) template used for the detection (d) correlation estimated on compressive sensing measurements using the orthogonalized estimator, and the adaptive threshold (dash line).

new search is performed on a window of 100 ms centered between the two blocks. This strategy is justified by the fact that the time interval between adjacent heart beats usually does not change so quickly [66].
The example illustrated in Fig. 4.3 helps to understand the basic operations of the proposed method. The compressed measurements Fig. 4.3(b) and the generated template Fig. 4.3(c) are used to estimate the cross-correlation $\hat{R}_{x\psi}(n)$ (Fig. 4.3(d)) with the orthogonalized estimator. Whenever $\hat{R}_{x\psi}(n)$ is higher than the adaptive threshold a peak is detected, thus providing information about the peak location in the original signal. The corresponding marker is shown in Fig. 4.3(a).

4.6 Simulation and Results

To verify the effectiveness of the proposed CS QRS detection technique, we evaluated its performance through a series of experiments.

We first evaluate and compare the detection performance of the *direct* and *orthogonalized* estimators. We then compare the performance of our technique with respect to detection on original and reconstructed signals at various compression ratios, using the same detection method based on cross-correlation and matched-filtering (MF). We also compare the results with the use of the Pan-Tompkins (PT) detection algorithm on reconstructed signals.

Experimental data for evaluating the proposed method are obtained from the MIT

Arrhythmia Database, introduced in Sec. 3.5.3.

We further define an observation window of length 512 samples (approximately 1.5 s at 360 Hz), enough to contain at least one beat in the majority of cases, and keeping the acquisition time sufficiently short for real-time monitoring. Moreover, the use of a longer window may lead to incorrect detection due to the adaptive threshold calculation.

4.6.1 Performance Evaluation

An R-peak was classified as correctly identified if the time difference between the R-peak in the reference and the R-peak as identified by the algorithm is smaller or equal to 50 ms.

For the evaluation of the proposed scheme, we use classical performance figures usually applied for the assessment of QRS detection algorithms, i.e., sensitivity (Se) and positive predictivity (P+). According to the American National Standard [2] Se and P+ are computed as

$$\mathrm{Se} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}}100, \quad \mathrm{P+} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}}100. \tag{4.6.1}$$

In the above equations, TP (true positives) is the total number of QRS correctly located by the detector, a false negative (FN) occurs when the algorithm fails to detect a true beat (actual QRS) and a false positive (FP) represents a false beat detection.

The algorithm accuracy can be also evaluated using the F measure, proposed in [12],

$$\mathrm{F} = 2\frac{\mathrm{Se\ P+}}{\mathrm{Se}+\mathrm{P+}}100 = 2\frac{\mathrm{TP}}{2\mathrm{TP}+\mathrm{FN}+\mathrm{FP}}100. \tag{4.6.2}$$

4.6.2 Direct and Orthogonalized Estimators

Table 4.1: Detection results for the *orthogonalized* and *direct* estimators and for some compression ratios.

	Orthogonalized Estimator			Direct Estimator		
CR [%]	Se [%]	P+ [%]	F [%]	Se [%]	P+ [%]	F [%]
Uncompressed	98.6±2.9	99.3±1.4	98.9±1.0	98.6±2.9	99.3±1.4	98.9±1.9
30	98.2±3.2	99.1±1.5	98.6±2.1	96.8±4.5	98.9±1.7	97.7±2.8
50	97.4±3.7	98.6±1.9	98.0±2.4	96.0±4.9	97.9±2.6	96.9±3.4
75	94.7±5.0	95.9±3.5	95.3±3.7	93.4 ±6.4	94.2±6.4	93.6±5.4
85	90.7±5.7	88.2±10.6	89.1±7.4	89.0±7.6	88.9±8.5	88.7±6.9

To assess the performance difference of the direct and orthogonalized estimators, we evaluated the performance figures described above as the number of compressive measurements varied from $M = 51$ to $M = 410$ (corresponding to a compression ratio varying from 90% to 20%). Furthermore, we also tested the proposed method on uncompressed signals ($M = N = 512$ in the figures). For each signal a different Gaussian sensing matrix $\Phi_{i,j} \in \mathcal{N}(0, 1/M)$ has been used.

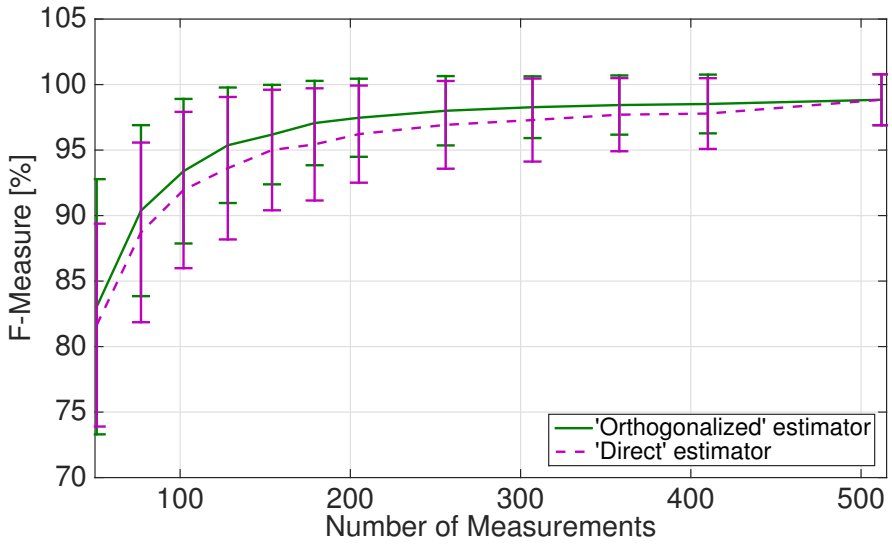

Figure 4.4: Average F measure for QRS detection using the *direct* (dashed line) and *orthogonalized* (solid line) estimator for different number of measurements. Error bars indicate standard deviation.

Fig. 4.4 shows the average F measure over all the records. As one would expect, it is possible to see that for both the estimators, the detection performance increases as M increases. However, the orthogonalized estimator has slightly better performance with respect to the direct one. The results for the two estimators and for some compression ratios are summarized in Table 4.1.

Fig. 4.5 shows, for the orthogonalized estimator, the average sensitivity and positive predictivity values averaged over all the signals at different compression ratios. The ability to detect the true QRS complexes seems to be slightly affected by compression, at least for compressions as high as about 75%. Both sensitivity and positive predictivity slightly decrease for compression ratios higher than 50%, but they are still higher than 95% when the compression ratio is lower than 75%.

4.6.3 Detection on Reconstructed ECG

As part of our experiments we compare the proposed method, working in the compressed domain, with schemes where the signal is first reconstructed from the CS measurements, and then the P&T QRS or the MF is used in order to detect the QRS complexes.

As reference, the two detection method on the raw signals lead to a F measure of 98.7% and 98.9% for P&T and MF, respectively.

ECG signals are recovered with the Basis Pursuit Denoising BPDN algorithm provided in the SPGL1 solver [123], which is the reconstruction algorithm selected by many authors, such as Mamaghanian *et al.* [90] and Dixon *et al.* [51]. To enable a quicker reconstruction and also a possible power consumption reduction at the decoding stage, we also use the OMP and the Smoothed-l0 (SL0) [94] algorithms, which have a lower complexity.

As the sparsifying transform $\mathbf{\Psi}$, we select the orthogonal Daubechies−4 wavelets (WT), which can can effectively provide a sparse representations of the ECG, as suggested in [107].

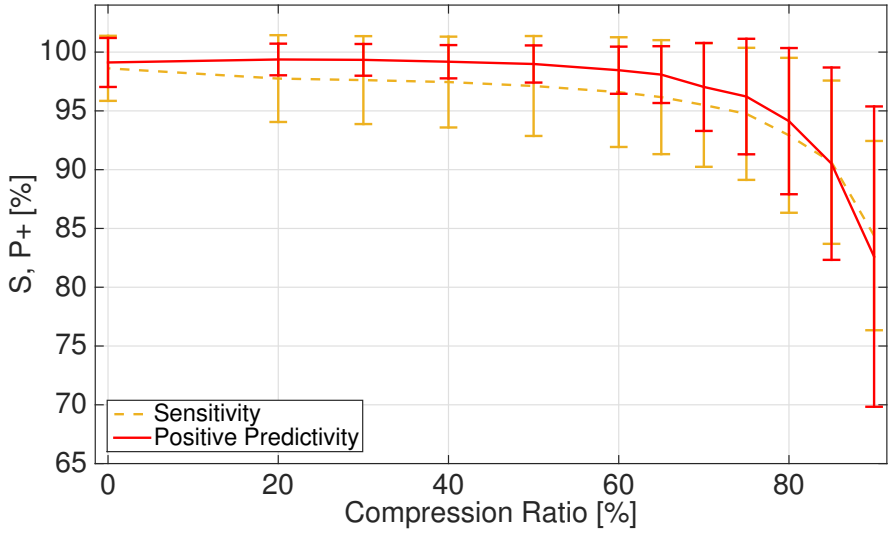

Figure 4.5: Average Sensitivity Se (solid line) and Positive Predictivity P+ (dash line) for QRS detection using the orthogonalized estimator at different compression ratios. Error bars indicate standard deviation.

The results of this experiment are displayed in Fig. 4.6(a) and (b), which show a series of plots comparing the average F measure at different compression ratios. In particular, in Fig. 4.6(a) we compare the performance of the proposed method with respect to MF detection that works in quasi real-time on reconstructed signals. Whereas, Fig. 4.6(b) shows the performance of P&T detection applied off-line on the whole reconstructed signal. We can see from this figures that the proposed QRS detection procedure, using the orthogonalized estimator, has a performance similar to the methods that require reconstruction. At compression ratios larger than about 65%, the proposed method allows a slower performance degradation. In particular, for compression ratios higher than 65% the proposed detection method in the compressed domain gives better results, also with respect to detection using the P&T algorithm on the reconstructed signals. As a matter of fact, we observed that the reconstruction at higher compression ratios, using the wavelet basis as the sparsifying domain, may introduce artifacts that are attenuated by the matched filtering detection procedure, while these artifacts might be incorrectly classified as R peaks by the P&T algorithm.

In the CS framework, the reconstruction quality depends not only on the signal itself, but also on the sparsifying basis, or possibly overcomplete dictionary, used in solving the reconstruction problems.Indeed, one can design dictionaries that are more likely to lead to sparse solutions with a consequent better reconstruction for a given compression ratio. One can argue that the choice of the wavelet basis as the sparsifying domain is not the best choice, even if it is widely used, since overcomplete dictionaries increase the reconstruction complexity [90] [93]. Applying overcomplete dictionaries that are patient based or more general overcomplete dictionaries based on signal models, can improve the reconstruction quality at a given compression, and consequently provide better detection performance [108] [40] [42]. However, note from Fig. 4.6(a) and (b) that the detection performance is not far from that obtained with uncompressed signals (compression ratio equal to 0% in

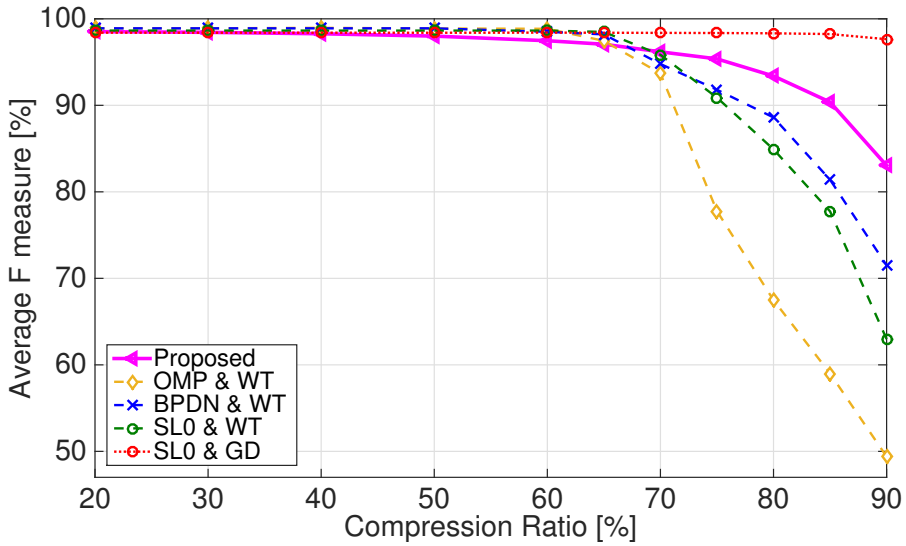

(a)

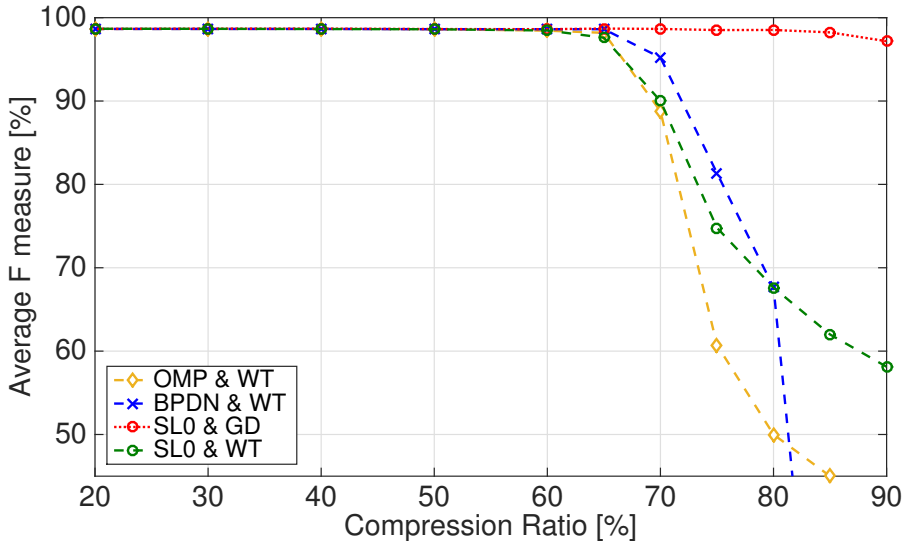

(b)

Figure 4.6: Comparison of the average F measure for the proposed method working directly on compressive measurements with respect to detection after signal reconstruction with (a) online detectors, i.e., the proposed method working directly on compressive measurements and matched filtering (MF), and (b) offline detector, i.e., the Pan-Tompkins (PT). Signal reconstruction in performed using a variety of algorithms (BPDN, OMP, SL0) combined with Daubechies−4 wavelets (WT) and Gaussian dictionary (GD).

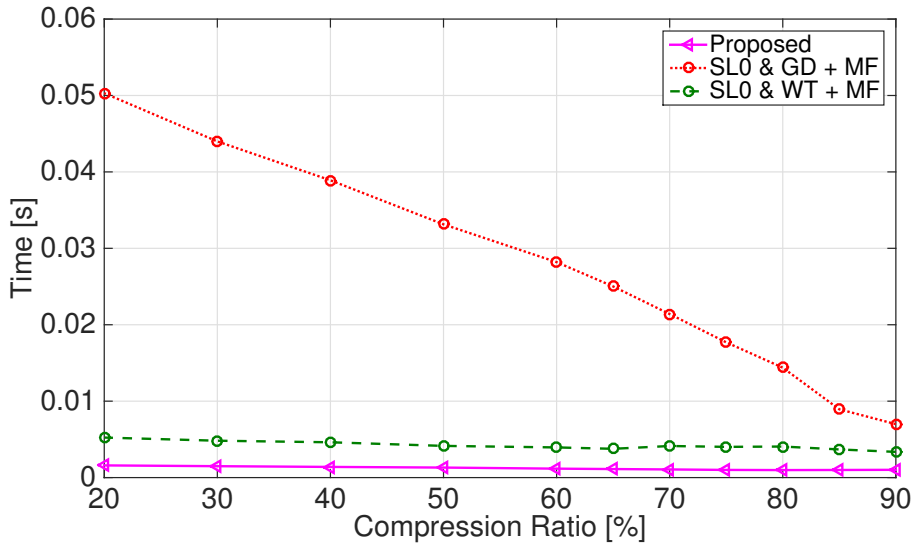

Figure 4.7: Average time required to process 1 s of signal with the proposed method and using SL0 recovery with WT or GD followed by MF.

the figure), at least for compression values up to 50-60%.

To verify the detection performance while using a sparsifying domain that increases the sparsity of the signal, we carried out experiments using an overcomplete Gaussian dictionary (GD) based on the morphology of the ECG signal, designed such that it preserves the shape of QRS complexes as well as of the P and T waves [40]. Since the use of an overcomplete dictionary has a great impact on the reconstruction complexity, we report in Fig. 4.6 (a) and (b) the results obtained with the SL0 reconstruction algorithm, which has a much lower complexity than BPDN and similar reconstruction quality performance[1]. As mentioned before, the GD is specifically designed to preserve the shape of the main waves, such as the QRS complexes, thus its application increases the detection performance in term of F measure as shown in Fig. 4.6 (a) and (b), with an average F=97.6% and F=97.2% for a compression ratio 90% for the MF and P&T, respectively.

As expected, as the use of the GD allows better reconstruction quality than the WT, the detection performance increases at high compression ratios. However, this comes at the price of increased computational complexity.

Fig. 4.7 reports the average time required to process 1 s of the ECG signal using the proposed method and applying the MF detector after signal reconstruction using the SL0 algorithm. We show results only for the SL0 algorithm since it is the one that has the lowest execution time, without significant reconstruction quality loss. As we can see, the proposed method is always the fastest, requiring a processing time always less then 2.9 seconds to process 30 minute of signal, corresponding to less than 0.0023 s to process a 1.5 s signal block. The reconstruction with the GD requires at least 12.6 s for a 30 minutes record,corresponding to 0.007 s for each block. Decreasing the compression ratio, the time increases up to 90.6 s, corresponding to a CR=20% (i.e., 0.05 s per block). All

[1]Using the BPDN algorithm followed by the P&T or MF algorithms provides very similar detection results, which are not shown in Fig. 4.6 (a) and (b).

the simulations are written in Matlab, running on an Intel Core i7 processor, equipped with 16 GB memory. Fig. 4.6 and Fig. 4.7 allow to conclude that the proposed method, operating in the compressed domain, allows significant detection complexity reduction with similar or better performance than schemes using reconstructed signals, when the wavelet sparsifying domain is used. The performance is also similar to the one obtained with original or reconstructed signals using an overcomplete dictionary, followed by P&T detection, at least for compressions as high as about 60-70%, with a relevant complexity reduction.

4.7 Conclusions

In this chapter we presented a method that allows R-peak detection *directly* on the compressive sensing measurements, in order to estimate the peak positions and heart rate avoiding the signal recovery stage. The proposed method is based on the estimated cross correlation with a template, combined with further processing procedures to detect the QRS complexes. Experimental results on real ECG signal demonstrated that the proposed method achieves an accuracy comparable with those obtained on the original signal up to compression ratios of about 60-70%. Moreover, the average Sensitivity and Positive Predictivity values are about 95% even for a compression ratio as high as CR=75%. In Chapter 5 we will consider other metrics, such as the ability to detect atrial fibrillation still operating in the compressed domain.

5

Accuracy of Atrial Fibrillation Detection on Compressed Sensed ECG Signals

5.1 Introduction

So far, we have seen that CS approaches to electrocardiogram compression can provide efficient real-time and low-complexity encoding. In doing so, it is important to assess the downstream effect of the compression on any signal processing and classification algorithm. The reconstruction process within the CS framework enables an accurate approximation of the original signal. However, in the reconstruction of physiological signals, it is essential to guarantee that all clinically relevant information for a given task is preserved, in order to prevent significant degradation in the performance of any standard (or novel) algorithm.
In this chapter we investigate the impact of CS on Atrial Fibrillation (AF) analysis, which is the most common super-ventricular arrhythmia and consists in an abnormal electrical activity arising in the atrium [24]. Although it is not a lethal disease, it may lead to very disabling complications such as cardiac failure and atrial thrombosis, with the subsequent risk of a stroke [80]. In order to diagnose arrhythmia, it is necessary to document the heart rhythm at the time of symptoms (e.g., palpitations, syncope, chest pain) with electrocardiography. In presence of AF, the ECG trace is characterized by absent P-waves and irregularity of the ventricular response. AF patients exhibit irregular rhythms at rates between 100 and 175 beats per minute, while a normal sinus rhythm has a resting heart rate between 60 and 100 beats per minute.
If a suspected arrhythmia cannot be detected and documented on a resting ECG during initial evaluation, the cardiac rhythm may be recorded for 24 or 48 hours using portable Holter monitoring devices. Despite the undeniable benefits of such medium term ECG monitoring, some arrhythmias might be not detected (because they are too infrequent or asymptomatic/'silent'). The extension of ECG recording to 7-day or even longer using wearable or implantable devices can assist in detecting such episodes.

The aim of the work described in this chapter is to quantify the performance of an AF detection algorithm on CS reconstructed signals at different compression ratios. Tipically, after CS signal reconstruction, the Pan-Tompkins ('P&T') R-peak detection algorithm [100] is applied and the resulting RR interval series are employed for the identification of AF episodes. The reconstruction process might introduce distortions leading to inaccurate R-peak detection and consequently to a degradation in the ability of identifying AF. In

addition, since the recovery process involves algorithms with a high computational load, recovering a whole long term recording might require long time and high resources. Thus, we also evaluate the reliability of the QRS detector proposed in Chapter 4, which operates directly on the compressed measurements. In such a way, it is possible to perform AF detection without recovering the original ECG signal. For this scenario, we also assess the performance of the AF detector at different compression ratios.

5.2 Method

5.2.1 Method Description

The work-flow adopted for the AF evaluation process is reported in Fig. 5.1. It clarifies the AF evaluation procedure to assess the effect of different CS compression ratios on AF detection.

First, we consider the uncompressed scenario, i.e, AF detection based on the QRS annotations directly available in the MIT Atrial Fibrillation Database (MIT AF DB) [120], as well as AF detection based on the detected QRS locations from the uncompressed ECG signals using the Pan-Tompkins (P&T) detection algorithm [100]. An AF detection method based on multi-feature extraction and a Support Vector Machine (SVM), described in Sec. 5.2.4, is applied on segments of 30 consecutive beats to perform AF detection.

For the assessment of AF accuracy on compressed ECG signals, we consider three different scenarios, in addition to different values of compression ratios. The first two scenarios require to reconstruct the ECG signals from compressed measurements as explained in Sec. 5.2.3, where two different sparsifying bases are adopted. Then, R-peak detection is performed using the P&T algorithm on the reconstructed signals. The third scenario evaluated in this study is motivated by the desire to simplify the detection process after compression. In particular, this scenario does not require signal reconstruction and the R-peaks are directly detected using an algorithm that operates on the compressed ECG signals. The detector, based on matched filtering, is described in Chapter 4, and is herein referred to as Compressed Sensing Matched Filtering (CSMF). After extraction of the RR interval time series, the multi-feature SVM detector is also applied for AF detection. Additionally, we test the performance of an AF detector based on the normalized fuzzy entropy measure, NFEn, [82], the results of which are reported and discussed in Appendix A.1.

Finally, in order to verify that clinically relevant information is preserved, AF detection accuracy is assessed in a range of compression ratios, i.e., 10, 20, 30, 40, 50, 60, 65, 70, 75, 80, 85 and 90%. In addition, we also evaluated the accuracy of QRS detection in the three different scenarios. In this way, we can verify the relation between a good R-peak detection and the ability of correctly classify an AF episode.

5.2.2 Data

For the analysis in this chapter, ECG signals from the MIT Atrial Fibrillation Dataset [120], freely accessible on PhysioNet [64], are used. This database contains 25 ECG recordings with a duration of approximately 10 hours each, sampled at 250 Hz, 12 bit resolution, with accompanying expert beat annotations. Among the records, 23 records include raw two-channel ECG signals. Records 00735 and 03665 are represented only by the rhythm and beat annotation files and were therefore excluded in this study. QRS annotations

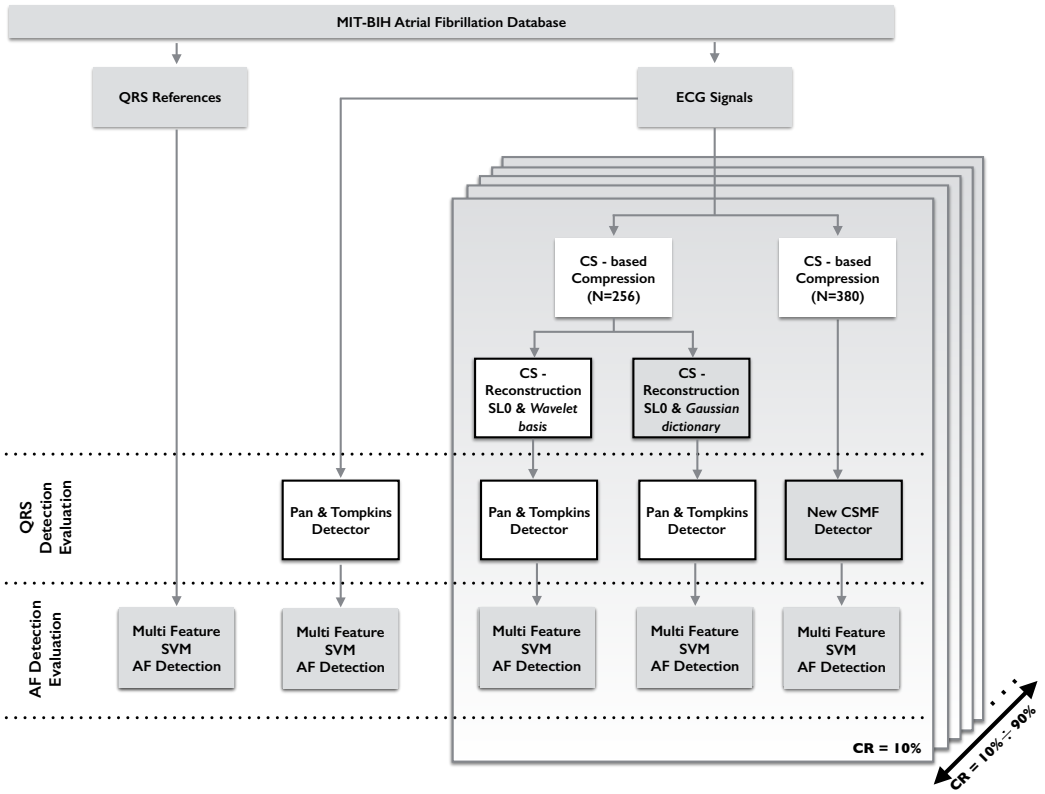

Figure 5.1: General flowchart of the AF accuracy evaluation method employed in this work.

were derived using an automated detector. The rhythm annotations were manually annotated and contain 4 types of rhythms: AF (atrial fibrillation), AFL (atrial flutter), J (AV junctional rhythm), and N (used to indicate all other rhythms). The RR interval series corresponding to the latter three rhythm types (AFL, J and N) were merged as non-AF rhythms in this study, to create AF and non-AF rhythm types.

5.2.3 Setting the CS Parameters

In this work, when signal compression is followed by reconstruction, CS is applied to non-overlapping windows (blocks) of length $N = 256$ samples, which corresponds to almost one second in the MIT AF DB ECG data. We chose N equal to a power of two to allow the use of a dyadic wavelet matrix $\mathbf{\Psi}$ as the sparsifying basis, as described below. Moreover, using windows corresponding to approximately one second, makes the compression process suitable for low-delay real-time applications.

Each signal is compressed using a different random sensing matrix with i.i.d. entries drawn from the normal distribution, $\mathbf{\Phi}_{i,j}$, $j \in \mathcal{N}(0, 1/M)$. The resulting measurements $\mathbf{y}$, $y_j = \sum_{i=1}^{N} \phi_{ji} x_i$ are quantized with a 12-bit uniform scalar quantizer. Signal reconstruction from CS measurements is performed using the SL0 algorithm [94]. Since the reconstruction process is based on the signal sparsity assumption, we need to use a sparsifying transform $\mathbf{\Psi}$. To this end, we employ the orthogonal Daubechies-4 wavelets (Db4),

and a 5-level decomposition, which can effectively provide a sparse representations of the ECG, as suggested in [107].

We also compare the wavelet basis to the use of an over-complete Gaussian dictionary introduced in Sec. 3.4.2, which provides a sparse representation of ECG signals. Since it is based on the ECG morphology, it preserves the shape of QRS complexes as well as of P and T waves, increasing the quality of reconstructed signals.

In this study we also verify the performance of a beat detector that operates directly on the compressed sensed measurements and does not require signal reconstruction. In particular we use CSMF, introduced in Chapter. 4. For this method, it is required that the signal block contains at least one heart beat in the majority of cases, so we set $N = 380$, corresponding to about 1.5 s block duration. The sensing matrices are still drawn from an i.i.d. standard normal distribution, $\mathbf{\Phi}_{i,j}, j \in \mathcal{N}(0, 1/M)$, and the measurements quantized with 12-bit resolution.

5.2.4 AF Detection Using an SVM Method

Generally, AF detectors are based on two approaches. One is based on atrial activity analysis and it focuses on the absence of P waves in the ECG signal. However, the P-wave has relatively low amplitude, and the ambulatory ECG often exhibits movement-related noise resembling the P-wave, which can lead to many false positives. The second approach is based on ventricular response analysis, and it is based on the predictability of the beat-to-beat intervals of the ventricular contractions. These RR intervals are derived from the most obvious large amplitude feature in the ECG, the R-peak. This approach is robust to artifacts, and is suitable for analysis of ECG recorded by wearable devices [33].

In this study, we used a state-of-the-art method developed by Li *et al.* [79] for the ventricular response-based AF detection, which was developed on the MIT AF DB described in Sec. 5.2.2. The AF classification step is based on a Support Vector Machine (SVM) applied to 8 features that quantify irregularity in the RR interval time series. The SVM is trained by considering 30 s long signal windows, manually marked as AF and non-AF rhythms [79].

5.2.5 Evaluation Metrics

Evaluation on QRS Detection Accuracy

A QRS is correctly identified if the time difference between the annotated QRS in the reference and the detected R-peak by the algorithm is smaller or equal to 50 ms, according to the recommendation of the American National Standard for ambulatory ECG analyzers (ANSI/AAMI EC38-1994) [2]. We compute the sensitivity ($\mathrm{Se_{QRS}}$) and positive predictivity ($\mathrm{P+_{QRS}}$) for QRS detection as reported in Sec. 4.6.1. We indicate with $\mathrm{TP_{QRS}}$ the true positives, i.e., the QRS number correctly located by the detector, $\mathrm{FN_{QRS}}$ the false negative, and $\mathrm{FP_{QRS}}$ the false positive.

Accurate R-peak detection is crucial for a reliable analysis of AF episodes. In order to test QRS detection accuracy, record 07126 was excluded since its reference QRS annotations are not consistent with the ECG signal. As a reference for the performance of detection in compressed/reconstructed signals, we applied the P&T QRS detector [100] on the remaining 22 raw original ECG signals, obtaining $\mathrm{Se_{QRS}}$= 96.38%, $\mathrm{P+_{QRS}}$= 90.38%.

Evaluation on AF Detection Accuracy

For this purpose, the RR series are classified into: AF episodes and non-AF episodes. The accuracy of AF episodes classification adopted in this work use the following metrics:

$$\text{Sensitivity: } Se = \frac{TP}{TP + FN}, \tag{5.2.1}$$

$$\text{Specificity: } Sp = \frac{TN}{TN + FP}, \tag{5.2.2}$$

$$\text{Accuracy: } Acc = \frac{TP + TN}{TP + FP + TN + FN}, \tag{5.2.3}$$

$$\text{Positive predicitivity value: } PPV = \frac{TP}{TP + FP}, \tag{5.2.4}$$

$$\text{Negative predicitivty value: } NPV = \frac{TN}{TN + FN}, \tag{5.2.5}$$

$$\text{Youden index: } J = Se + Sp - 1. \tag{5.2.6}$$

where TP, FN, FP and TN denote the true positive, false negative, false positive and true negative detections, respectively. All the measures were computed on all the RR interval series within the dataset, including noisy segments to represent a real world scenario.

As a reference for the performance of AF detection in compressed/reconstructed signals, we perform classification based on the manually annotated QRS positions or detected by P&T on the raw original ECG signals. In these cases we obtain an accuracy Acc =96.53% on hand annotated signals and 95.28% using peak detection with the P&T detector.

5.3 Results

5.3.1 QRS Detection Performance

Fig. 5.2(a-c) show TP_{QRS}, FN_{QRS} and FP_{QRS} for QRS detection using the considered three approaches, namely the CSMF R-peak detector in the compressed domain, the P&T detector after signal reconstruction using the wavelet basis (WT) and the Gaussian dictionary (GD). Fig. 5.2(d) shows the total number of detected QRS, i.e., the sum of TP_{QRS} and FP_{QRS} as a function of the compression ratio.

Fig. 5.3 (a-b) illustrate the results of QRS detection sensitivity (Se_{QRS}) and positive predictivity (P+_{QRS}) as a function of the compression ratio. At low CR levels (CR<60%), QRS detection using the three CS approaches gives similar results.

For CR=10%, the CSMF method results are Se_{QRS}=96.61% and P+_{QRS}=97.06%. The P&T method run on the reconstructed ECG signals using the wavelet basis results in Se_{QRS}=97.01% and P+_{QRS}=97.54%. The P&T method run on the reconstructed ECG signals using the Gaussian dictionary results in Se_{QRS}=96.98% and P+_{QRS}=97.48%. These results are slightly higher than those obtained with P&T-based QRS detection on the raw

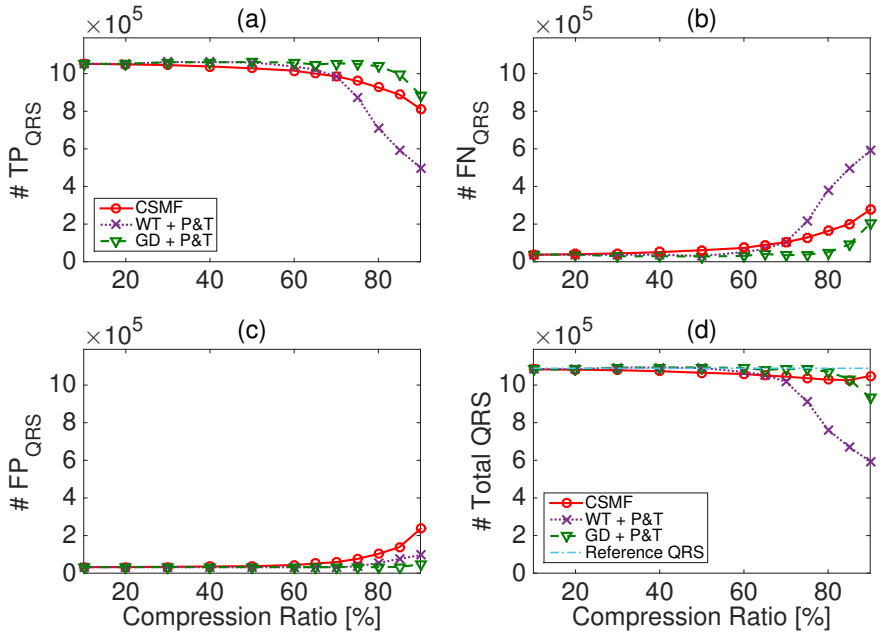

Figure 5.2: Numbers of (a) TP_{QRS}, (b) FN_{QRS} and (c) FP_{QRS} QRS (d) Total number of the detected QRS (TP_{QRS}+FP_{QRS}) varying the compression ratio for the CSMF detection or Pan-Tompkins (P&T) detection after reconstruction using Wavelet Transform (WT) and Gaussian dictionary (GD).

ECG signals (see Section 5.2.5), and can be explained by the filtering properties of the CS approach at low CR rates.

At CR levels higher than about 60-70%, the QRS detection accuracy of all three CS approaches declines rapidly. It is worth noting the difference between the wavelet basis recovery and Gaussian dictionary recovery. Fig. 5.2(d) clearly shows that recovery performed using the wavelet basis leads to many missed QRS detections at high CR levels. Thus, as we can see in Fig. 5.3(a-b), detection on signals reconstructed using the wavelet basis has lower sensitivity and positive predictivity values than on signals recovered using the Gaussian dictionary. Indeed, the reconstruction process using the wavelet basis typically introduces artifacts that lead to incorrect QRS detections. At the same CR level, the reconstructed ECG signals using the Gaussian dictionary, based on a model of QRS waveforms, exhibit less artifacts, leading to better accuracy for QRS detection. This can be seen in Fig. 5.4, which shows an example from record 05121. In particular, Fig. 5.4(a) shows the raw ECG signal and the corresponding annotated QRS complexes marked with triangles. Fig. 5.4(b) and (c) depict the reconstructed ECG signals using the Gaussian dictionary and the wavelet basis, respectively, at CR=75%. Triangles in (b) and (c) represent the detected QRS using the P&T method on the reconstrucetd signals. It can be seen that the artifacts present in the reconstructed ECG signal using the wavelet basis cause wrong beat detection. Obviously, if one or more QRS complexes are missed or wrongly detected, the resulting RR interval series and consequently AF classification performance, are compromised. In addition, for high CR levels, the proposed CSMF method gives higher sensitivity, but lower positive predictivity for QRS detection, than

the method based on the wavelet basis. We note that recostruction using the Gaussian dictionary gives the best results.

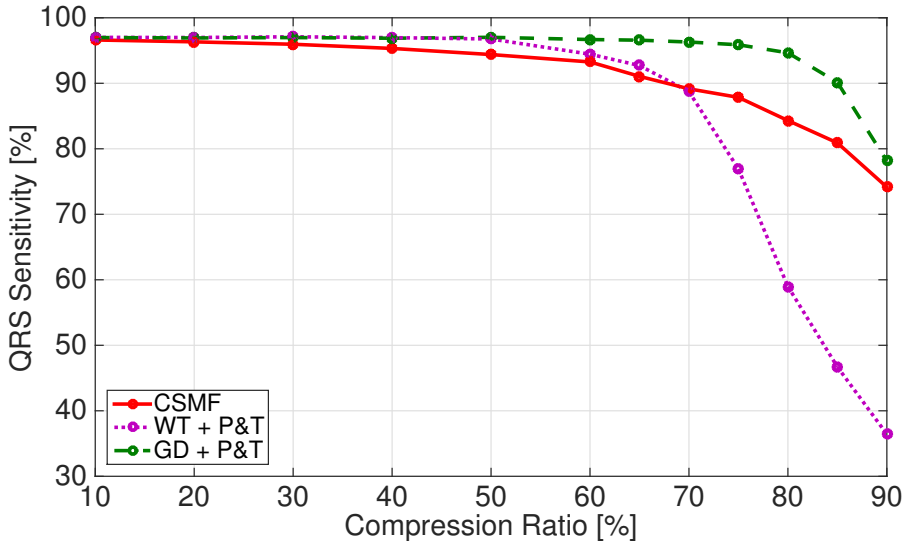

(a)

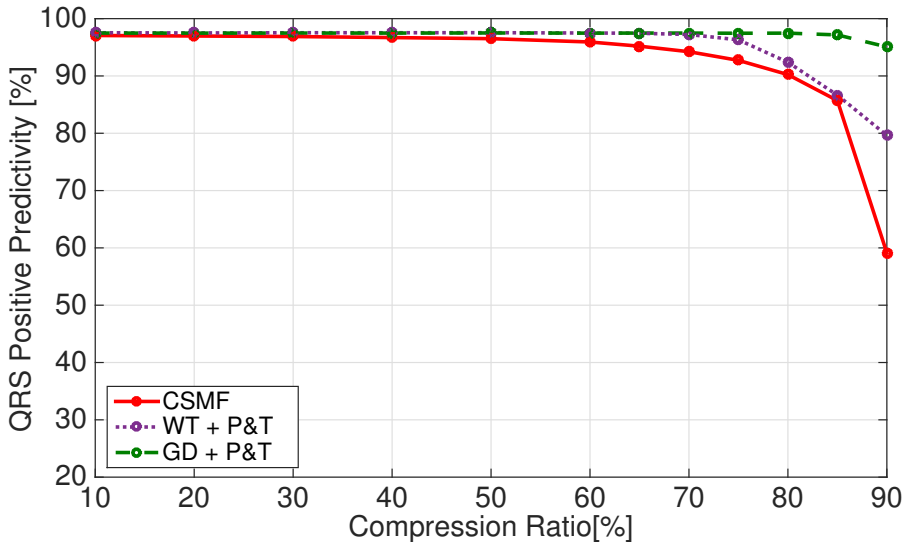

(b)

Figure 5.3: QRS detection (a) Sensitivity and (b) Positive Predictivity versus CR.

5.3.2 AF Detection Performance

As mentioned, the detector operates on segments of 30 consecutive beats within the manually annotated AF and non-AF time intervals. The reference total number of AF segments is given by the number of segments belonging to AF time intervals and obtained from the annotated QRS complexes. The reference total number of non-AF segments is computed

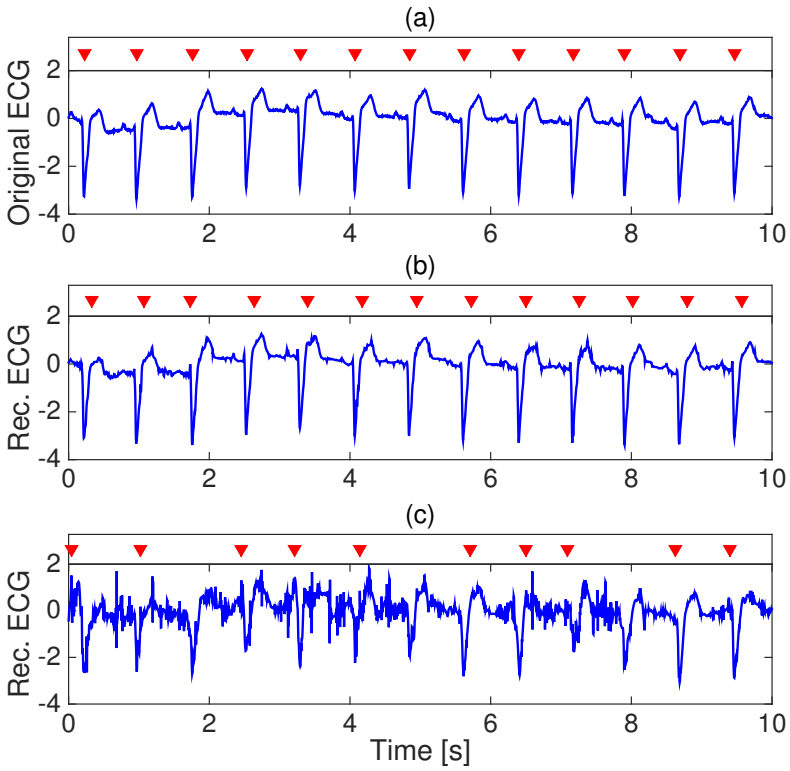

Figure 5.4: (a) Original ECG signal sampled at 250 Hz (first 10 s of record 05121) and reference QRS annotations. (b) Signal recovered at 75% compression using Gaussian Dictionary and QRS positions detected using the P&T algorithm. (c) Signal recovered at 75% compression using Wavelet Transform (WT) and QRS positions detected using the P&T algorithm.

similarly. It is important to note that, due to compression and errors in QRS detection, the total number of segments in AF time intervals for a given technique, computed as the sum of TP and FN classification decisions, is in general different from the reference value. The same happens for non-AF segments, defined as the sum of TN and FP after classification. Of course, if many QRS complexes are missed, we expect a large difference with respect to the reference values. Fig. 5.5(a) and (b) compare the total number of AF and non-AF segments for a given technique with the reference values. It can be seen that reconstruction with the wavelet basis exhibits a significant drop at compression ratios higher than 60%. This is consistent with Fig. 5.2, where it can be seen that at high CR, the method based on signal recovery using WT missed many QRS complexes, whereas the GD and the CSMF methods could detect almost as many beats as given by the reference annotated QRS complexes.

Fig. 5.5(c-f) show TP, FN, FP and TN values for AF detection as a function of CR, for the three CS scenarios considered in this study.

The number of correctly classified AF segments, represented by the TP value, is re-

ported in Fig. 5.5(c). It can be seen that TP starts to rapidly drop at CR>60% when AF detection is performed after reconstruction with the wavelet basis. Instead, when reconstruction is performed using the gaussian dictionary, TP starts to significantly decrease at CR> 75%. The CSMF technique, applied directly in the compressed domain, results in a relatively small performance loss up to CR=50% with a rapid decline at higher compression.

Figures 5.5(d-f) similarly show that for CR up to 60% the techniques have similar performance, while the WT and CSMF techniques degrade at higher compression. The method based on reconstruction with the Gaussian dictionary exhibits very good performance for CR up to about 75%. Note that the SVM detector tends to classify a segment with an RR pattern not consistent with AF as a non-AF segment, which is a safe harbour approach. Indeed, if a patient needs treatment, many AF segments would be present, and it is likely that eventually a positive trigger would be seen. This explains the larger FN values at high compression ratios.

An example of detected QRS complexes by the three CS techniques for non-AF and AF episodes is reported in Fig. 5.6. The figure shows a sample of record 04746, and, in particular, the AF episode occurs at time 1:13:10. The detected RR series after compression (CR=80%) and reconstruction shows that WT leads to inaccurate QRS locations. In this example, the use of the Gaussian dictionary enables a better QRS detection that allows to correctly classify the normal rhythm and the AF episode. This also applies for CSMF detection on the CS measurements.

Table 5.1 summarizes the results for the SVM-based AF detector on the MIT-BIH Atrial Fibrillation Database in a variety of scenarios, as described in Sec. 5.2. In particular, the second columns specifies the technique used for the evaluation. 'Reference QRS' refers to the application of the SVM-based AF detector on annotated QRS complexes, while 'Raw signals' refers to classification after QRS detection using the P&T procedure on uncompressed signals, as specified in the fourth column. It is clear that the corresponding performance values help to quantify the performance of the CS-based techniques. Note that the CSMF method does not require signal reconstruction, as specified in the third column, where the reconstruction technique is indicated.

It can be seen from Table 5.1 that the application of the AF detector on annotated QRS complexes ('Reference QRS') results in a very high Sp value equal to 99.14% and a relative low Se value equal to 93.22%. As mentioned before, this can be justified by the fact that the classifier tends to mark as non-AF those segments which do not exhibit a typical AF pattern.

As expected, classification after QRS detection by employing the P&T algorithm on the uncompressed ECG signals also results in a high Sp of 97.68% and a relative low Se of 92.97%.

An overall picture of the accuracy of AF detection performance as a function of CR is given in Fig. 5.7. It can be seen from the figure and from Table 5.1 that, for CR values up to about 50%, the AF detector applied to QRS complexes derived from the compressed signals, using the CS techniques described in this work, gives results comparable to those achieved when employing a standard QRS detector on the raw uncompressed signals. At a compression ratio equal to 60%, we have less than 1% loss for the WT and GD based techniques, and about a 5% loss for CSMF. At a compression ratio equal to 75%. the GD method guarantees a small performance loss of about 1.2%.

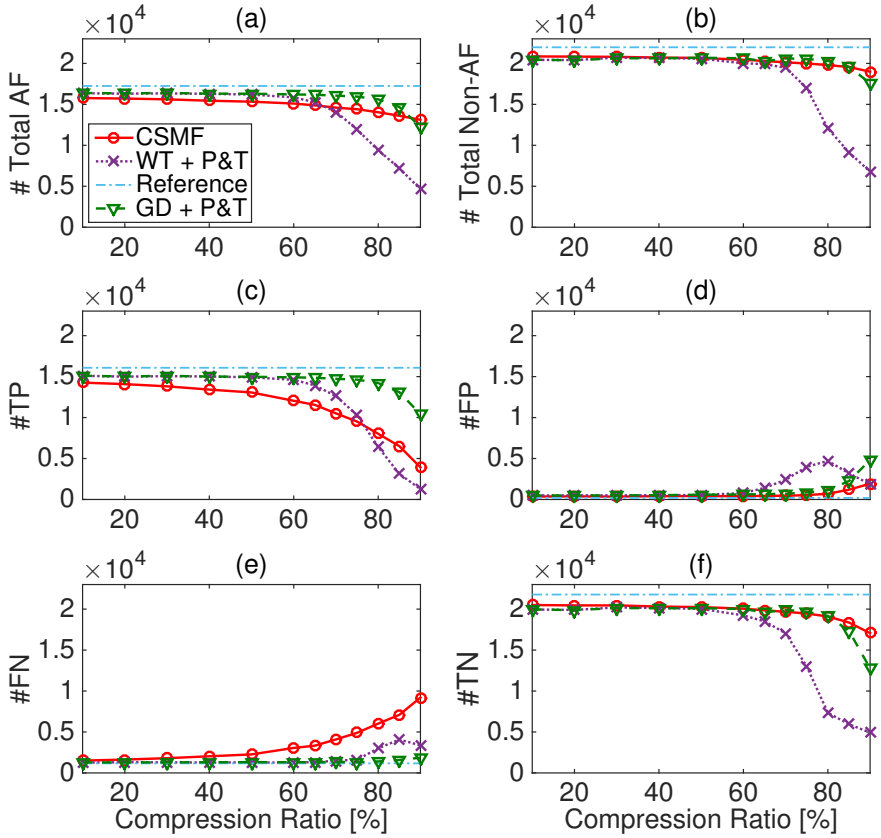

Figure 5.5: Total number of (a) AF episodes (TP+FN), (b) non-AF episodes (TN+FP). Number of (c) True Positive (TP), (d) False Positive (FP), (e) False Negative (FN) and (f) True Negative (TN) detections for the SVM AF classifier operating on QRS detected using CSMF or (P&T) detection after reconstruction using Wavelet Transform (WT) and Gaussian Dictionary (GD). The reference numbers (dash-dot line) refer to the AF episodes and non-AF episodes found on the reference QRS.

5.4 Discussion

To detect AF automatically and reliably is a challenging task even on raw uncompressed ECG data. In our study, we investigated the effect of CS-based ECG compression on the accuracy of an AF detector applied to the processed data, for a wide range of compression ratios between 10% and 90%. To this end, two different sparsifying representations, in combination with the SL0 algorithm, were used to reconstruct the ECG signals from the CS measurements. Afterwards, the P&T algorithm was employed for QRS detection. Furthermore, we also describe a newly introduced beat detector that allows direct processing of the compressed measurements, without any signal reconstruction. Finally, the RR interval series obtained from the three different CS scenarios at different CR levels was used to perform AF detection using a previously reported state-of-the-art SVM-based model.

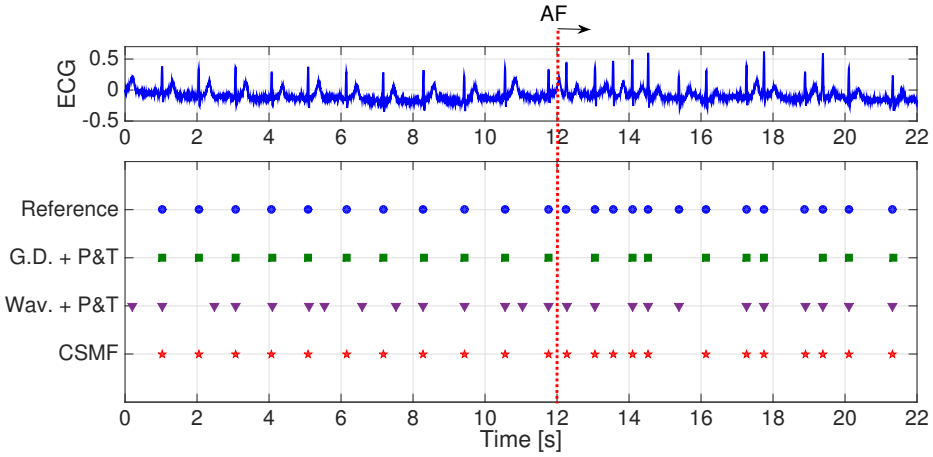

Figure 5.6: From top to bottom: ECG signal (04746) at time 1:13:10 and corresponding RR series of a non AF and AF episodes from the annotations file. The RR series detected after signal reconstruction using wavelet transform (WT) and Gaussian dictionary (GD). RR series from the compressed beat detector (CSMF).

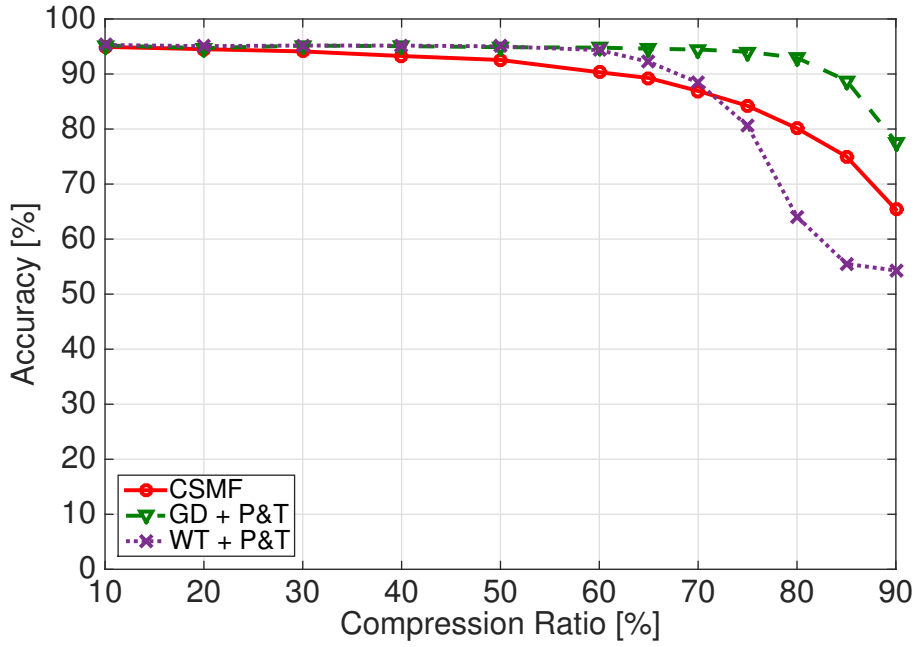

Figure 5.7: Output AF detection Accuracy versus CR when analysis is performed in the compressed domain using the proposed CSMF beat detector and after signal reconstruction using the Pan & Tompkins (P&T) beat detector.

Table 5.1: Results of the performance metrics on uncompressed and compressed ECG signals using different reconstruction/detection methods for the SVM-based AF detector.

Metric		CS Reconstruction	QRS Detector	Compression Ratio													
				0%	10%	20%	30%	40%	50%	60%	65%	70%	75%	80%	85%	90%
	Reference QRS	-	-	93.22	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	92.97	-	-	-	-	-	-	-	-	-	-	-	-
Se [%]	CS ECG signals	SL0 & WT	P&T	-	92.35	91.97	92.11	92.21	92.06	92.16	91.29	89.52	86.39	68.09	43.38	27.41
	CS ECG signals	SL0 & GD	P&T	-	92.15	91.35	92.10	92.02	91.91	91.93	91.90	91.67	91.75	91.02	89.25	84.83
	CS ECG signals	No	CSMF	-	90.47	89.63	88.47	86.85	85.24	79.79	77.63	72.00	66.01	57.10	47.87	29.88
	Reference QRS	-	-	99.14	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	97.68	-	-	-	-	-	-	-	-	-	-	-	-
Sp [%]	CS ECG signals	SL0 & WT	P&T	-	97.61	97.54	97.59	97.49	97.36	96.08	92.91	87.65	76.50	60.94	65.04	72.65
	CS ECG signals	SL0 & GD	P&T	-	97.64	97.38	97.59	97.47	97.21	97.04	96.76	96.61	95.83	94.41	88.50	72.66
	CS ECG signals	No	CSMF	-	98.30	98.19	98.33	98.06	97.97	98.10	97.80	97.74	97.34	96.46	93.89	89.88
	Reference QRS	-	-	96.53	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	95.28	-	-	-	-	-	-	-	-	-	-	-	-
Acc [%]	CS ECG signals	SL0 & WT	P&T	-	95.27	95.07	95.17	95.17	95.02	94.35	92.20	88.44	80.58	64.07	55.49	54.29
	CS ECG signals	SL0 & GD	P&T	-	95.20	94.72	95.16	95.06	94.88	94.79	94.60	94.45	94.05	92.93	88.82	77.66
	CS ECG signals	No	CSMF	-	94.93	94.51	94.10	93.27	92.55	90.32	89.27	86.92	84.22	80.13	75.01	65.37
	Reference QRS	-	-	98.83	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	96.96	-	-	-	-	-	-	-	-	-	-	-	-
PPV [%]	CS ECG signals	SL0 & WT	P&T	-	96.87	96.77	96.79	96.67	96.48	94.91	90.79	84.00	72.09	57.63	49.44	40.63
	CS ECG signals	SL0 & GD	P&T	-	96.90	96.49	96.81	96.64	96.29	96.07	95.77	95.47	94.47	92.60	85.32	68.43
	CS ECG signals	No	CSMF	-	97.57	97.39	97.55	97.09	96.89	96.87	96.27	95.86	94.71	91.97	84.50	67.10
	Reference QRS	-	-	94.91	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	94.04	-	-	-	-	-	-	-	-	-	-	-	-
NPV [%]	CS ECG signals	SL0 & WT	P&T	-	94.10	93.83	94.00	94.08	93.97	93.93	93.30	92.03	88.89	70.99	59.31	59.44
	CS ECG signals	SL0 & GD	P&T	-	93.95	93.44	93.95	93.92	93.85	93.87	93.73	93.70	93.73	93.19	91.66	87.27
	CS ECG signals	No	CSMF	-	93.17	92.62	91.91	90.91	89.95	86.79	85.64	82.80	79.90	76.03	72.13	64.98
	Reference QRS	-	-	92.36	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	P&T	89.95	-	-	-	-	-	-	-	-	-	-	-	-
J [%]	CS ECG signals	SL0 & WT	P&T	-	89.96	89.51	89.70	89.71	89.42	88.24	84.20	77.17	62.89	29.03	8.42	0.06
	CS ECG signals	SL0 & GD	P&T	-	89.79	88.73	89.68	89.49	89.13	88.98	88.66	88.28	87.58	85.43	77.74	57.49
	CS ECG signals	No	CSMF	-	88.77	87.82	86.80	84.91	83.21	77.89	75.43	69.74	63.36	53.56	41.76	19.76

All three CS scenarios, i.e., reconstruction with the wavelet basis or the Gaussian dictionary followed by a standard Pan & Tompkins detector, and the direct detection on compressed measurements (CSMF), exhibit similar characteristics for what concerns the AF classification quality metrics. In particular, at low CR levels, the AF detection results are comparable with those obtained on raw uncompressed ECG signals. However, for high CR values, the AF detection accuracy for the three methods decreases, as can be seen in Fig. 5.7. The results reveal that AF detection based on the new CSMF method has an acceptable performance loss, with respect to the techniques that require signal reconstruction, up to compression ratios of about 60%. This technique performs better than the WT-based method for CR higher than 70%. Indeed, CSMF reaches an accuracy equal to 92.55% at a 50% compression ratio, while at CR=90%, its accuracy decreases to 65.37%. AF detection after reconstruction using WT allows slightly better results up to CR= 70%. However, its performance rapidly decreases at higher CR levels and reaches an accuracy equal to 54.29% at CR=90%. The best performance is achieved by the method with signal reconstruction using the Gaussian dictionary, which allows to reach accuracies 94.05% and 77.66% at CR=75% and CR=90%, respectively.

This study also highlights some drawbacks related to each of the three CS scenarios. One major drawback of the CSMF method is the quickly increasing number of false negative AF detections FN when CR> 50%. This is due to the SVM-detector tendency to favour non-AF classifications. The WT method has a reconstruction quality that decreases rapidly at high compression ratios, thus compromising classification quality. Overall, it

appears that using the Gaussian dictionary for signal reconstruction enables a good AF detection up to a CR level of 75-80%, at the expense of increased reconstruction complexity.

The *democracy* property of compressive sensing consists in the fact that each measurement carries the same amount of information. Thus, the reconstruction quality depends only on how many measurements are received and not on the particular received subset. This allows to modulate the compression ratio by simply discarding or retaining some measurements. Considering the trade-off between AF classification accuracy, execution time, and compression, one could envision a two-stage processing system where the CSMF method is employed in the sensor for mild compression ratios up to CR=60%. The system then switches to a higher compression, when transmission or recording is needed after AF episodes are detected, in view of reconstruction with GD followed by P&T detection. The switch simply consists in transmitting or recording fewer measurements. In a concrete scheme, one could acquire a compressed version of the signal using the analog CS implementation [63, 14] with a low CR, e.g., 50%, and use CSMF for AF detection. Then, the CR can be increased (up to 80%) by keeping a subset of the measurements in order to save/transmit a lower amount of data, still allowing accurate AF classification when the reconstruction is performed using the GD method.

As one could expect, there is a relationship between the AF detection accuracy and QRS detection accuracy. In particular Se_{QRS} starts to rapidly decline for CR>60%, similarly to what happens for the AF detection accuracy, when reconstruction using the WT and CSMF methods are used. Furthermore, reconstruction using the Gaussian dictionary allows to obtain similar results at a higher CR=75%.

Many studies related to CS-based ECG compression limited their assessment to the reconstruction quality, without evaluating the actual impact that signal reconstruction has on preserving relevant clinical information. In this study, we show that CS can be successfully employed as a compression technique for ECG signals when the final goal is to perform AF detection. As for reconstruction quality, we also show that the reliability of detected QRS complexes significantly depends on the sparsifying basis adopted for reconstruction.

5.5 Conclusions

The results of this study show that AF classification performed after CS-based compression allows to correctly detect AF episodes when the compression ratio is lower than 60% or 75%, depending on the reconstruction/detection method adopted. In particular, we found that acceptable results are obtained for compression ratios up to 60% when AF classification is performed on signal reconstructed using wavelets as the sparsifying basis, or when the CSMF method is used. However, when a specifically designed sparsifying dictionary, as the one proposed in Chapter 3, is used during signal reconstruction, good results are obtained for CR values as high as 75%. These findings have positive implications concerning the acquisition and compression of ECG signals for clinical purposes using low-power wearable devices. Moreover, the possibility to correctly identify an AF episode directly on the compressed measurements represents a good opportunity for future long-term monitoring applications that need to process the data on energy-constrained devices.

6

Joint Reconstruction and Detection of Fetal Beats from CS Fetal ECG Measurements

6.1 Introduction

In this chapter, we focus on non-invasive abdominal fetal ECG (fECG) signal processing and compression. Processing of such physiological signals has attracted much research interest due to the possibility to assess fetal health conditions during pregnancy. fECG signal analysis can provide early detection of fetal heart disorders, making it possible to treat them with drug administration or to pre-schedule the delivery [72]. According to the World Health Organization, in 2015 about 303,000 newborns died within 4 weeks after birth due to congenital anomalies. Heart defects are the most common severe congenital anomaly affecting approximately $6-11$ per 1000 newborns. About $20-30\%$ of these heart defects are severe and potentially life threatening, requiring surgery within the first year of life [124].
Fetal heart rate (fHR) is commonly used to indirectly assess fetus conditions during pregnancy and labor [102], and among the different approaches (e.g., Doppler ultrasound, fetal magnetocardiogram, phonocardiography) the examination of the fECG from non-invasive measurements on the mother's abdomen plays an important role. Indeed, like for adults, the fECG allows to visualize the electrical activity of fetal heart, which conveys important information about the health and conditions of the fetus.

Nowadays, fetal heart monitoring is mainly carried out for short time periods in a hospital environment; however, long-term continuous monitoring can provide a rapid objective assessment of fetal well-being. Its implementation is showing an increasing interest as demonstrated by the different wearable devices that have been developed in recent years [59] [21]. The advances in the field of wireless wearable sensors for human daily activity monitoring, have made it possible to use this technology also for the continuous monitoring of fetal electrocardiogram. As for adult ECG, also for non-invasive abdominal fECG the use of compressive sensing is vital in order to increase the battery life of the wearable sensors. Moreover, in the specific case of abdominal fECG, we are usually required to record more than one channel, in order to enable further signal processing, thus increasing the amount of data to be transmitted/stored. The CS framework introduced in Chapter 2 and applied to low-power monitoring of adult ECG in Chapter 3, may be used as a viable solution for the compression of fECG.

A first study on the usability of CS for the telemonitoring of fECG signals has been investigated in [127], by Zhang *et al.*. As reported by the authors, due to the particular complexity of the abdominal fECG signals, current CS frameworks using linear sparsifying bases, generally fail to adequately recover the signal at high compression ratios, making the reconstructed electrocardiogram non suitable for diagnostic purposes. Thus, they proposed to use a new reconstruction algorithm, namely block sparse Bayesian learning (BSBL), to overcome the problem. Using the scheme proposed in [127], the authors report that compression ratios up to 60% allow the extraction of fECG signals without significative performance loss with respect to using the original non-compressed signals.

In this chapter, we design a novel joint reconstruction/detection signal processing technique for the compression and analysis of non-invasive fECG through abdominal electrodes. The proposed technique allows real-time analysis of the collected signals [42]. It can be implemented into wearable sensors allowing at-home continuous monitoring of fetal health.

Motivated by the observation that the use of a properly designed dictionary allows to have better performance then BSBL algorithm for adult ECG signals (see Sec. 3.5.4) we introduce a novel dictionary for fECG sparsification. The design of such dictionary is based on the morphology of fetal beats, which leads to effective performance in realistic settings and can also be used for the separation of maternal and fetal beats.

6.2 Fetal ECG Signal Processing

Fetal ECG signals can be collected invasively by using an intrauterine electrode during labor, with the recording electrodes in direct contact with the fetal skin or scalp [99]. Alternatively, non-invasive methods may be used by collecting the signals with non-invasive electrodes placed on the mother's abdomen surface. These methods can be used at any stage of pregnancy [105] and they promise greater prospect for long-term monitoring of fHR and fetal well-being. Abdominal fECG can be recorded from the maternal abdomen as early as at the twentieth week after conception [105].

The mechanical functioning of the fetal heart differs from that of the adult heart, however, its beat-to-beat electrical activity is similar. In common with the adult electrocardiogram, fetal ECG contains three easily identifiable components, namely the P wave, QRS complex and the T wave, representing the atria and ventricles contraction and relaxation activities. Even if the morphology of adult and fetal ECG patterns are similar, the fetal wave amplitude changes considerably throughout gestation. Moreover, despite the similarities in the electrical properties of fetal and adult cardiac systems, there are some differences in their RR-interval. Indeed, the fetal heart beat is almost twice as fast as the adult heartbeat, with considerable variations corresponding to different stages of fetal cardiac development. The range of normal fHR varies significantly with gestational age. For example, at the end of the first month of pregnancy, the normal average fetal heart rate is about 110 bpm and increases to 175 bpm after the ninth week. In the following weeks, the rate starts to slow down, reaching 110-160 bpm at delivery [119].

Despite the fact that non-invasive abdominal fECG signals can be conveniently recorded, their analysis represents a challenge for clinical technicians as well as for engineers, due to the complexity of the trace. Indeed, there are three main problems concerning non-invasive measurement of fECG. The first one is the low signal to noise ratio (SNR). In

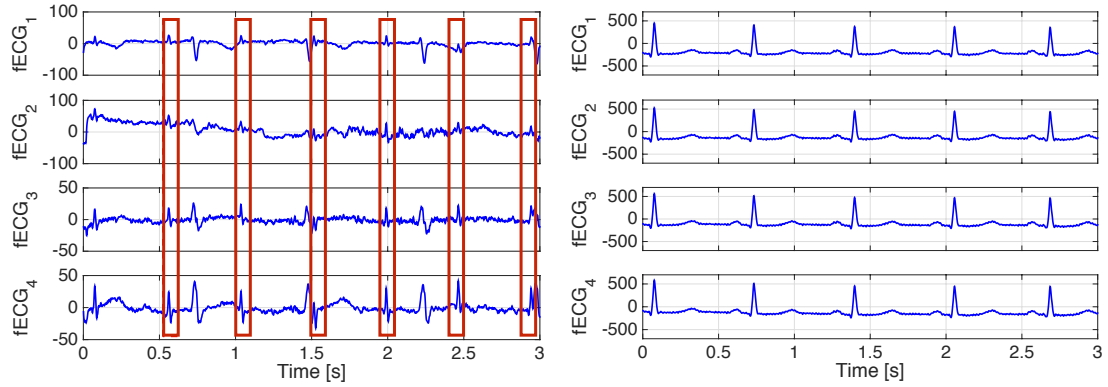

(a) (b)

Figure 6.1: Two abdominal recordings from the Physionet Challenge dataset A sampled at 1 KHz. In (a) 4 channels of signal a22, the fECG QRS complexes are clearly visible (red boxes) while in (b), 4 channels of signal a32 of the dataset, the fECG cannot be identified by visual inspection. It is also clear, from record a22, that different sources of noise may affect the signals.

fact, recorded signals are always a mixture of noises generated, for instance, by fetal brain activity, myographic signals (both from the mother and the fetus), movement artifacts and maternal ECG. Moreover, the fetal component is usually smaller compared to the maternal one, with a voltage amplitude of 1-10 μV compared to the maternal ECG (mECG), which can have an amplitude as high as $1000\mu V$ [20]. The lower amplitude is due to the fact that the fetal heart is smaller than the adult one, and that the signal is typically attenuated by tissues in the path to the measuring electrodes. Fig. 6.1 shows two abdominal recordings from the Physionet Challenge database [117] [64]. In Fig. 6.1(a), the fetal QRS complex is barely visible, while it can be seen clearly in Fig. 6.1(b). Note that the duration and morphology of the fetal QRS complex is indeed similar to the maternal one, but with a smaller amplitude and QRS width.

The fetal component of the abdominal fECG trace is often very weak. Its detection can be a difficult task and there is need for signal processing techniques to improve SNR and eliminate the maternal component from the signal.
The problem of extraction of the fetal ECG trace from non-invasive abdominal fECG has been typically addressed in the literature by using a combination of different techniques [38]. Indeed, extraction methods are typically based on a multi-step approach. Typically, the first step consists of a pre-processing stage, whose objective is to remove baseline wandering and power-line interference.
The second step estimates the maternal beats by using some form of processing such as filtering and template subtraction [50] [83]. Blind and semi-blind signal separation techniques can also be used to this end [6] [125]. The idea is that signals, in a multi-channel abdominal recording, are mixtures of uncorrelated independent sources corresponding to the mother's heart beat, the fetal one and noise. Two of the most commonly applied blind source separation techniques are the Independent Component Analysis (ICA) [114],

and the Principal Component Analysis (PCA). Note that PCA allows to decompose the input signals into statistically uncorrelated components, which do not necessarily represent independent sources. Due to the significant noise usually affecting the signals, ICA and PCA have some limitations when used alone, and other pre-/post-signal processing techniques have to be adopted. For instance, ICA has been used together with the wavelet decomposition in [7].
After identification, the maternal component is removed from the abdominal recording by using an adaptive or matched filter trained from several maternal reference channels [99] [126] [50]. The same method can be used to extract directly the fetal QRS waves [60, 103].
Methods based on linear decomposition have been also employed for fetal ECG separation and enhancement. The abdominal signals are decomposed into different components by using suitable basis functions. The basis functions can be selected from classes that are somehow in coherence with the time, frequency, or scale characteristics of the fetal components. Wavelet decomposition [73, 5], and matching pursuits [4], are among these methods.
Finally, fetal QRS are typically estimated by using a modification of existing adult QRS detection techniques. After this final step, it is possible to further process the detected beats to constrain the estimated fHR and RR time series within physiological or statistical limits based on heuristics.
Behar *et al.* [13] compared the results relative to a wide variety of state-of-the-art methods for non-invasive fetal ECG analysis and proposed to use a combination of several techniques to improve performance.

6.3 Proposed Joint Reconstruction/Detection Framework

The jointly Reconstruction/Detection framework that we analyze in this chapter is summarized in the block diagram of Fig. 6.2.

First, since fetal heart rate estimation requires the separation of the fetal and maternal beats from the acquired ECG signals, we model the recorded ECG signals as a mixture of independent components (ICs) given by the mother's and fetal heart beat signals, as well as other noise sources. In particular, as shown in Sec. 6.4 it is possible to perform ICA directly in the compressed sensed domain, with no performance loss than using the original signals. Our scheme operates on small blocks of compressed coefficients, and estimates the compressed independent components, thus reducing computational complexity and permitting low delay detection and reconstruction.

In order to recover the ICs, we extend the work proposed in Chapter 3, to overcame the limitations of classical CS framework for fECG signals, introducing a new universal dictionary that permits to successfully increase compression while maintaining a reconstruction quality which does not affect the detection performance. Similarly to the dictionary used for adult ECG sparsification, the one used for abdominal fECG reconstruction comprises two classes of Gaussian-like functions [91], modeling the fetal and mother ECG, respectively. The dictionary is described in Sec 6.5.

Finally, from the reconstructed compressed independent components, we further enhance the maternal and fetal beats by checking the activated atoms during reconstruction within the CS framework. Indeed, some of the ICs may still contain both maternal and fetal beats.

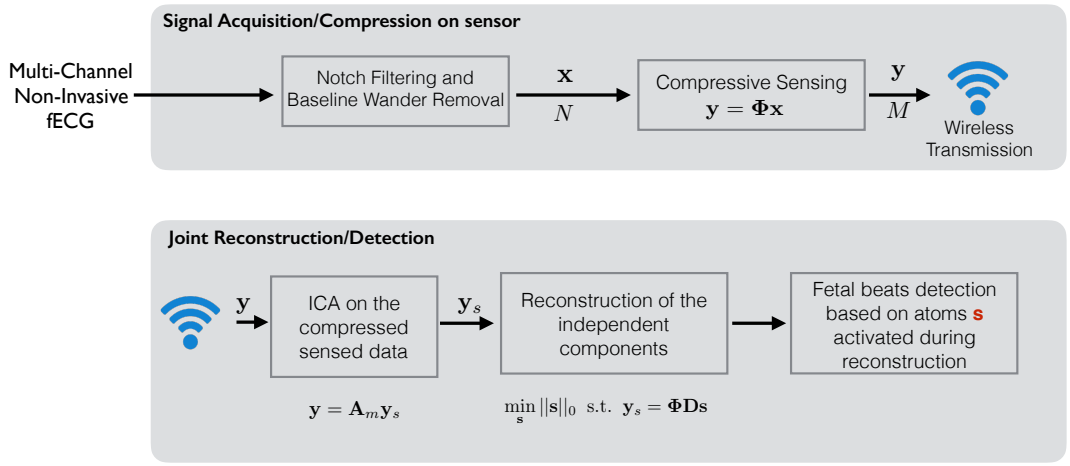

Figure 6.2: The proposed joint reconstruction/detection framework for fetal ECG signals.

Although the focus of the framework proposed in this chapter is to deign a low-power and low-complexity acquisition/detection real-time system, the performance loss with respect to non real-time procedures, which apply off-line post-processing techniques and several detection refinement stages, is acceptable, besides the fact that the proposed approach does not require training, it is completely automatic and can be used for real-time analysis with a small delay.

6.4 Source Separation in the Compressed Domain

Among the different approaches proposed for fECG extraction from multi-channel recordings, blind source separation techniques, such as the ICA, combined with other methods (e.g., maternal QRS cancellation), seem to allow reliable results [13]. In the following, we summarize the ICA technique and propose its application in the domain of compressed signals $\mathbf{y}$.

6.4.1 Independent Component Analysis

Let $x_i(k)$, $i = 1, ..., S$, be one of the multi-channel recorded signals. In the framework of Independent Component Analysis, $x_i(k)$ is modeled as a linear combination of independent processes corresponding to S unknown sources, i.e.,

$$x_i(k) = a_{i1}c_1(k) + a_{i2}s_2(k) + \cdots + a_{iS}c_S(k). \tag{6.4.1}$$

Each sample $x_i(k)$ is therefore an instance of a random variable

$$x_i = \sum_{j=1}^{S} a_{ij}c_j, \tag{6.4.2}$$

where c_j are assumed to be independent and non-Gaussian (except for one of the variables at most). By collecting variables (6.4.2) in vector $\mathbf{x}_m = [x_1, ..., x_P]^T$, we can write the

ICA model in matrix-vector notation as

$$\mathbf{x}_m = \mathbf{A}_m \mathbf{c}, \tag{6.4.3}$$

where $\mathbf{x}_m$ are the observed linear mixtures, $\mathbf{c} = [c_1, ..., c_S]^T$ are the hidden source signals and $\mathbf{A}_m$ is the unknown mixing matrix.

ICA algorithms estimate matrix $\mathbf{W} = \mathbf{A}_m^{-1}$ and obtain the independent components $\mathbf{c} = \mathbf{W}\mathbf{x}_m$. ICA is based on the fact that, according to the central limit theorem, each variable x_i, since it is a linear combination of independent variables, tends to have a Gaussian distribution, opposite to the distribution of the source components c_j. According to this observation, separation algorithms for ICA estimation are based on the maximization of the non-gaussianity of $\mathbf{w}^\mathbf{T}\mathbf{x}_m$, which allows to find one row $\mathbf{w}^\mathbf{T}$ of matrix $\mathbf{W}$. This gives one of the independent components. The other ICs can be estimated by finding all the local maxima of the optimization problem in the S-dimensional space of vectors $\mathbf{w}$.

6.4.2 ICA in the Compressed Domain

As explained before, the approach proposed in this work uses the CS framework as a compression technique, and therefore the information available at the receiver side consists of the compressed measurements $\mathbf{y}$, one vector for each channel.

Considering the CS model, each measurement $\mathbf{y}_i$ of the i-th channel corresponds to the compressed version of the i-th acquired mixture signal $\mathbf{x}_i$.

As said before, detection schemes based on ICA would require to reconstruct the signals before application of the technique. However, we propose to apply ICA directly in the compressed domain. Indeed, the projection of the data into a lower dimensional space, as provided by CS, reduces the dimension of the problem and may be very useful in order to speed-up the ICA algorithm. As we will see in this chapter, the high-dimensional problem can be solved in the compressed domain, with a limited quality loss.

Define $\mathbf{x}_i = [x_i(n_0), ..., x_i(n_0 + N - 1)]^T$ as a block of N consecutive samples of the $i-$th multi-channel recorded signal. According to (6.4.1), we can write

$$\mathbf{x}_i = a_{i1}\mathbf{c}_1 + \cdots + a_{iS}\mathbf{c}_S, \tag{6.4.4}$$

where $\mathbf{c}_j = [c_j(n_0), ..., c_j(n_0 + N - 1)]^T$. With the help of a random sensing matrix $\mathbf{\Phi}$, the CS framework computes measurements for a block as

$$\mathbf{y}_i = \mathbf{\Phi}\mathbf{x}_i = a_{i1}\mathbf{\Phi}\mathbf{c}_1 + \cdots + a_{iS}\mathbf{\Phi}\mathbf{c}_S. \tag{6.4.5}$$

Across various blocks, the statistics we observe allow to model the elements of $\mathbf{y}_i$ as samples of a process

$$y_i(h) = a_{i1}\Phi_h\mathbf{c}_1 + \cdots + a_{iS}\Phi_h\mathbf{c}_S, \tag{6.4.6}$$

where Φ_h is one row of the random matrix $\mathbf{\Phi}$, and $\mathbf{c}_j$ is an $N-$dimensional random vector of samples $c_j(k)$. Each sample $y_i(h)$ is therefore an instance of a random variable

$$y_i = \sum_{j=1}^{P} a_{ij} y_{cj}, \tag{6.4.7}$$

where variables $y_{cj} = \Phi_h\mathbf{c}_j$, obtained by sensing the unknown sources, can be supposed to be independent, since they are functions of the independent source processes. Moreover,

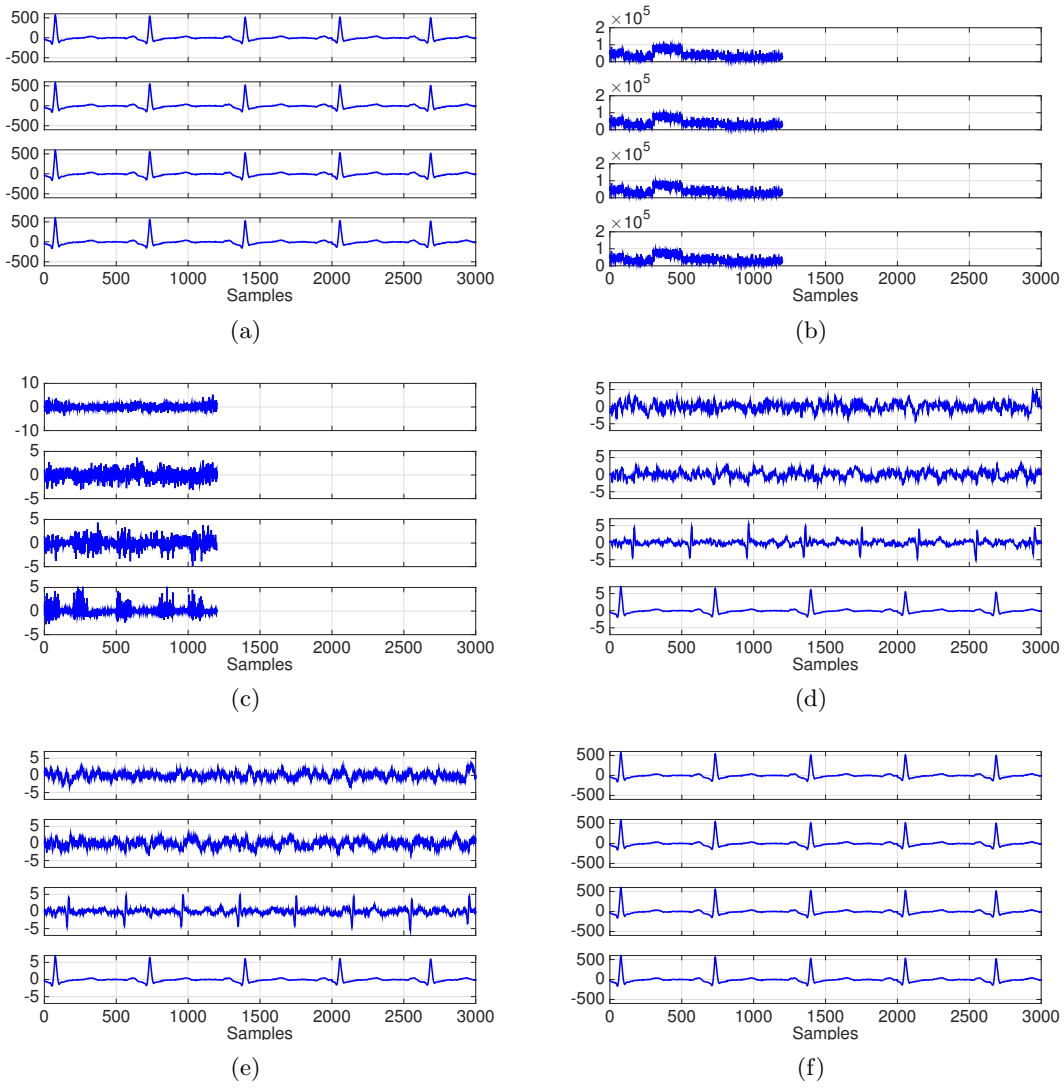

Figure 6.3: (a) First 2 s of original 4 channels abdominal fetal ECG record a32 of the Challenge dataset A sampled at 1 KHz) (b) $k = 8$ consecutive blocks of compressed sensed measurements of each channel (record a32), corresponding to the first 2 s. The compression ratio is CR=60% (c) Estimated 4 independent components, which are still a compressed version of the original ICs (d) Reconstructed ICs from the compressed ICs (e) ICs corresponding to ICA applied to the original signal. (f) Reconstructed signals obtained from the reconstructed ICs thanks to the estimated mixing matrix.

variables y_{cj} are in general non-Gaussian. Assume for instance that the elements of Φ_h, one row of the sensing matrix $\mathbf{\Phi}$, are i.i.d. normalized Gaussian random variables. Then, given $\mathbf{c}_j$, variable y_{cj} is Gaussian with variance

$$\sigma^2 = \sum_{k=k_0}^{k_0+N-1} c_j^2(k).$$

However, the distribution of y_{cj} is obtained by averaging over the distribution of $\mathbf{c}_j$, and is given by a non-Gaussian linear superposition of Gaussian components. Of course, non-gaussianity arises also when $\mathbf{\Phi}$ is generated with distributions other than Gaussian.

As an example, Fig. 6.3(a) shows the first two seconds of record a32 of the 4-channel abdominal recordings from the Physionet Challenge dataset. The signals, sampled at 1 kHz, are divided into blocks of length $N = 250$ samples. Fig. 6.3(b) shows the corresponding CS measurements, consisting of 8 consecutive blocks $\mathbf{y}_i$ of length $M = 100$ (60% compression ratio). Fig. 6.3(c) shows the ICs y_{cj} of the measurements, which are estimated by performing ICA in the compressed domain. Finally, Fig. 6.3(d) shows the signal independent components after CS reconstruction. It may happen that ICA fails to separate the mother's and fetal components exactly. The typical situation is that one of the independent components reveals the mother's beat, while one of the others is still a mixture of the mother's and fetal ECG, the other components corresponding to noise. The next section describes the proposed procedure to detect the two signals during the reconstruction process. In Fig. 6.3(e) we also show the result of ICA on the original signals. It can be seen that, apart from the different order and possible sign changes, the signals are almost identical to those in Fig. 6.3(d), thus confirming that the proposed component analysis can be safely applied in the compressed domain.

In summary, according to (6.4.7), we can write in matrix notation $\mathbf{y} = \mathbf{A}\mathbf{y_c}$ i.e., the measurements can be expressed as a combination, with the *same* mixing matrix of the original model, of non-Gaussian independent components, obtained as CS measurements of the unknown sources. We can therefore use ICA to estimate $\mathbf{W}$ and the compressed sensed unknown sources, in order to apply the separation procedure within the CS reconstruction framework.

6.5 Gaussian Dictionary for fECG Sparsification

Compressive Sensing is typically used combined with an orthogonal basis for signal sparsification, such as the DWT or DCT. However, the use of an appropriate overcomplete dictionary may lead to a sparser representation of signals, improving compression as well as the identification of signal patterns matched to the dictionary atoms.

Following the ideas of using a Gaussian-like functions dictionary for the sparsification of the adult ECG, proposed in Chapter 3 based on the ECG signal approximation introduced by McSharry *et al.* [91], we propose the use of an appropriate overcomplete Gaussian dictionary $\mathcal{D}$ for the abdominal fetal ECG signals sparsification.

In the design of the proposed dictionary, we take into account not only the aspects of efficient signal representation for compression purposes, but also the possibility to help the separation between maternal and fetal beats. As a matter of fact, if atoms in the dictionary approximate different characteristics of the two signal classes, fetal and maternal,

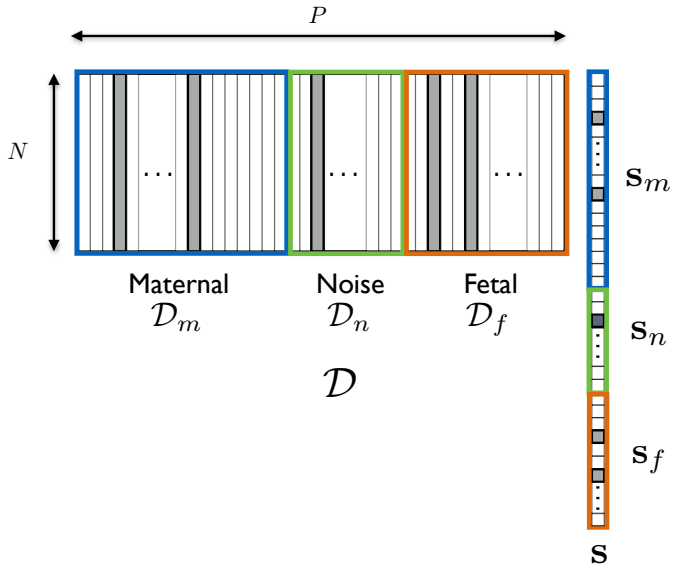

Figure 6.4: The proposed overcomplete dictionary $\mathcal{D}$ represented as an $N \times P$ matrix with columns $\mathbf{g}_i$, and split into three parts, for maternal, noise and fetal waves approximation. In this model any sparse set of atoms (columns in gray) can be selected.

it is possible to separate them by considering which class of atoms are activated during reconstruction. This turn to be very helpful when, after ICA, the resulting independent components are still a mixture of maternal and fetal beats.
The proposed dictionary is composed of P vectors $\mathcal{D} = \{\mathbf{g}_i\}_{i \in P}$, such that each fECG signal can be sparsely approximated by a subset of vectors $\{\mathbf{g}_i\}_{i \in K}$ in $\mathcal{D}$, where $|K|$ represents the support of the best k-term approximation $\mathbf{x}_K$ of $\mathbf{x}$

$$\mathbf{x}_K = \sum_{i \in K} s_i \mathbf{g}_i. \tag{6.5.1}$$

As for the dictionary proposed in for adult ECG sparsification the vectors $\mathbf{g}_i$ are Gaussian-shaped with elements

$$\mathbf{g}_i(n) = \exp\left(\frac{n - p_i}{b_i}\right)^2, \tag{6.5.2}$$

where b_i is a shape (or scale) parameter, which is chosen to give an efficient approximation of typical ECG waves, such as the P, Q, R, S and T waves. Note that the dictionary is universal, and it can be used to represent different kinds of beats (normal and abnormal) as for the Gaussian dictionary proposed in Chapter 3.

We use a total of 17 different shape parameter, 11 of them are used for maternal waves approximation $\mathcal{D}_m$, 4 for fetal ECG waveforms approximation $\mathcal{D}_f$ and 2 for noise $\mathcal{D}_n$. Hence, the over-complete dictionary is given by the concatenation of the three sub-dictionaries: $\mathcal{D} = [\mathcal{D}_m|\mathcal{D}_n|\mathcal{D}_f]$.
Different sets of values of the scale parameter and different number of atoms have been

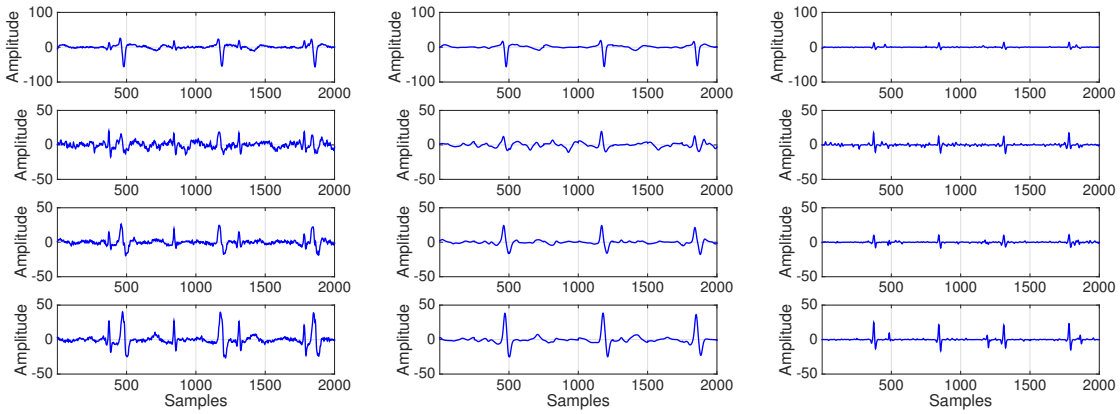

Figure 6.5: Separation using the proposed Gaussian dictionary. (a) Original 4 channel a04 record sampled at 1 KHz, with both maternal and fetal QRS complexes visible. (b) Reconstruction of the maternal ECG using coefficients $\mathbf{s}_m$ related to $\mathcal{D}_m$. (c) Reconstruction of the fetal ECG considering only atoms belonging to $\mathbf{s}_f$ related to $\mathcal{D}_f$.

tested for dictionary construction, in order to model the maternal and fetal ECGs. In particular, for the type of recordings considered in the experiments, which are ECG traces sampled at 1 kHz, we observed that scale parameters $b_i \in \{5, 6, 7, 8, 9, 10, 12, 15, 20, 30, 50\}$, ensure a good approximation of both symmetric (Q, R, S) and asymmetric (P and T) maternal waves. For fetal ECG approximation, atoms are computed using scale parameters $b_i \in \{2.5, 3, 3.5, 4\}$. For the noise components we use $b_i \in \{1.6, 2\}$. For each b_i, all the possible shift parameters values p_i within the signal segment $\mathbf{x}$ are considered. In the proposed scheme, since signal segments $\mathbf{x}$ with a duration of $N = 250$ samples are not synchronized with heart beats, the atoms corresponding to waves close to the segment boundaries are built using windowed Gaussian functions. Fig 6.4, reports the resulting partitioning of the proposed over-complete dictionary.
Thus, the atoms used for reconstruction using the proposed dictionary provide information about the location and nature of ECG waveforms, which can be used for classification. As a matter of fact, one would expect that the reconstruction of a fetal or mother's ECG waveform would use only atoms from the corresponding section of the dictionary, i.e. $\mathcal{D}_m, \mathcal{D}_f, \mathcal{D}_n$. Of course, an abdominal ECG signal of a pregnant woman requires a set of atoms (columns in gray in Fig. 6.4) belonging to all three sections, so additional processing is needed. Fig. 6.1 also shows that direct analysis of the reconstructed signals may not be sufficient to detect the fetal beats.

To make the discussion more concrete, we give an example of the separation capabilities of the proposed dictionary. In this example we do not apply any other separation technique before reconstruction. The only information we use to separate the maternal and fetal components are the atoms selected to recover the signal and their location with respect to the dictionary section. Fig. 6.5 (a) reports the reconstructed signal a04 after CS compression at CR=75%, using the SL0 algorithm and the proposed dictionary. Fig. 6.5 (b) shows the approximation of the maternal ECG signal using only the atoms activated

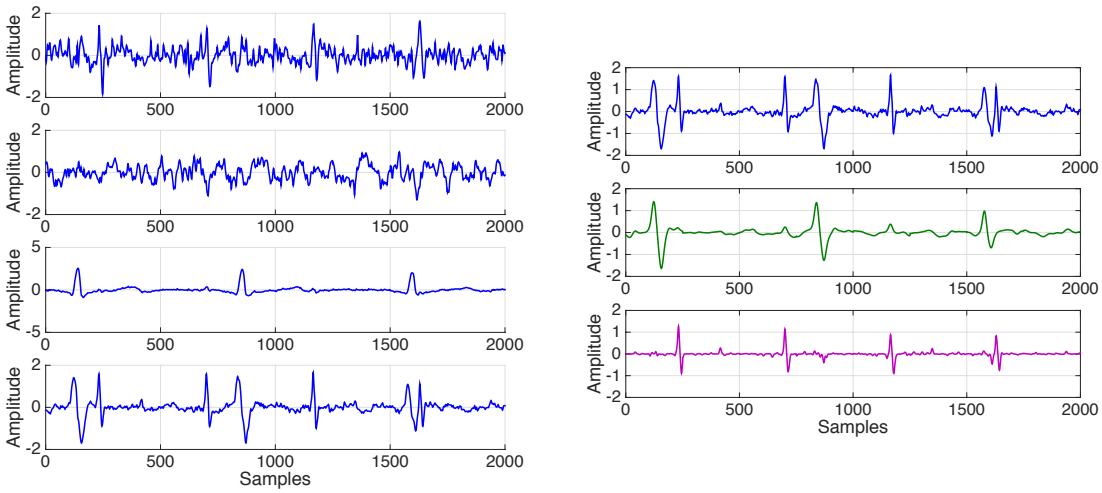

Figure 6.6: Reconstructed ICs for the first 2 s of record a08 of the Challenge dataset A, with ICA applied in the CS domain. (b) Further separation of maternal and fetal traces from the 4-th independent component by exploiting the dictionary-based classification procedure.

in $\mathcal{D}_m$. As it can be seen, the maternal beats in each channel are well separated from the other components. Fig. 6.5 (c) shows the approximation of the fetal QRS complexes using the atoms belonging activated in $\mathcal{D}_f$, and the fetal ECG components in each trace is the most prominent. It should be noted that in correspondence of the maternal peaks some residual signal may still be present, however, this can be treated as noise and will not interfere with the fetal beat detection.
From this example we can see that the maternal component and the fetal ones can be approximated by different parts of the dictionary. However, when the fetal components are hidden by the maternal ones, due to a very low fetal beat amplitude, ICA should be applied in order to enhance the fetal trace. Fig. 6.6(a) shows the ICs corresponding to the four channels of record a08 of the Challenge dataset A. While the two plots from above correspond to noise components, the mother's beats are clearly visible in the third plot. However, as we can see in the bottom plot of Fig. 6.6(a), the corresponding independent component can still contain a mixture of the fetal and mother's beats. In such a case, the proposed classification procedure can be used to further separate the maternal and fetal components. Indeed, by looking at the atoms used for the reconstruction, we can actually separate the mother's and fetal beats in the mixed signal, as shown in Fig. 6.6(b).

6.6 Detection and Classification of Fetal and Maternal Beats

After applying the ICA in the compressed domain and perform the reconstruction of the compressed ICs the final step of the proposed approach involves the detection of the maternal and fetal QRS complexes and their classification. The reconstruction process is

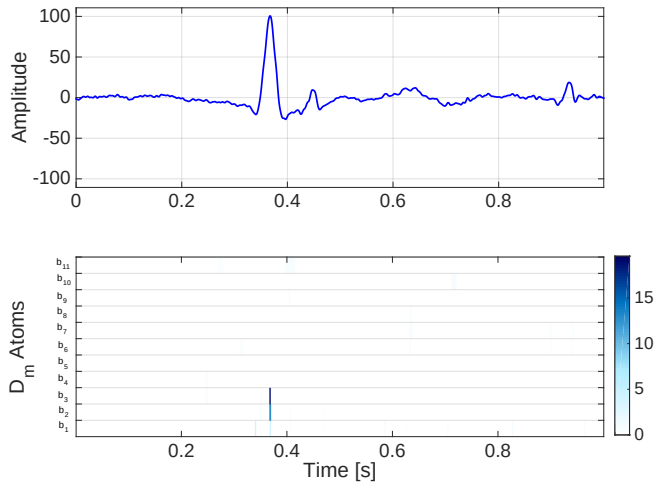

Figure 6.7: Absolute value of atoms s_i activated during reconstruction of one independent component of signal a23 of the Physionet Challenge database, in the dictionary section corresponding to maternal ECG approximation $\mathcal{D}_m$. Colors represent the magnitude of the atoms as indicated in the vertical colorbars.

based on the assumption that the independent components, resulting from the ICA analysis, share the same structure of the original signals, i.e., they are sparse in the Gaussian dictionary used for approximation.
Two assumptions have been made for the localization of maternal beats. The first one is that at least one of the ICs contains maternal beats. The second one is that these components are reconstructed using atoms mainly belonging to the maternal dictionary section. Fig. 6.7 shows the atoms in the maternal dictionary activated during reconstruction of one independent component of signal a23 of the Physionet Challenge database.
Based on the sparse representation, only atoms with a non-negligible magnitude in the maternal dictionary are considered as potential maternal QRS complexes. Note that fetal QRS complexes, clearly visible in Fig. 6.7, do not significantly activate atoms from the maternal dictionary.

Within a signal block of 2 seconds, an R peak position candidate p at time t is selected if the corresponding atom magnitude $|s_i|$ (see (6.5.1)) is greater than a threshold $T = \alpha A_{max}$, where A_{max} is the maximum atom magnitude in the current signal segment. In the experiments, we set $\alpha = 0.75$. If two detected atoms correspond to time positions less then 0.3 s apart, the one with smaller $|s_i|$ is discarded, since they are supposed to be redundant.

Detection is performed on three of the four channels simultaneously. The channel with the lowest kurtosis is supposed to be noise and is not considered. For beat position detection, we select the channel with the largest absolute s_i value.

The detection of fetal beats is carried out in a similar way as for maternal detection, but considering the dictionary section used for fetal approximation. However, in this case, some of the atoms activated during reconstruction might be related to the mother's beats. Not considering the atoms located in correspondence of the previously detected mother's

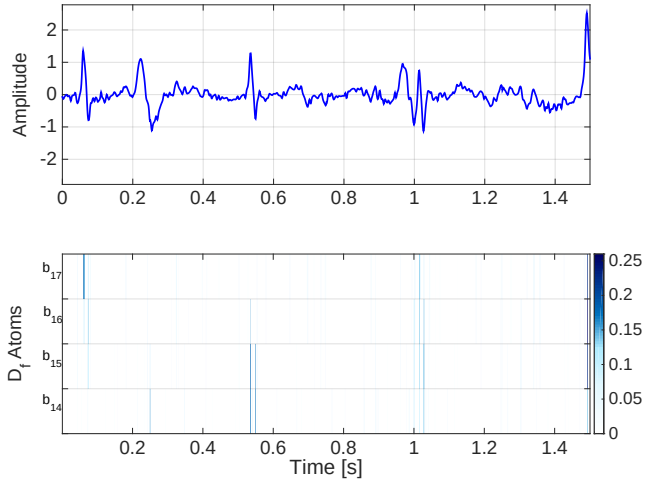

(a)

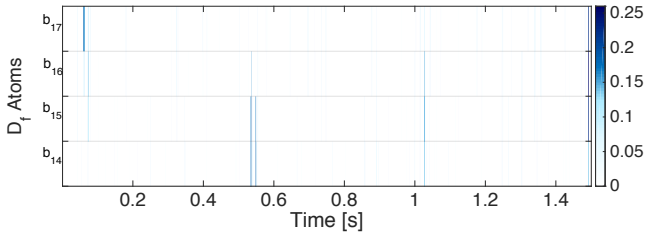

(b)

Figure 6.8: One of the IC's of signal a22. (a) Absolute value of atoms (s_i) activated in the fetal dictionary (b) absolute value of atoms resulting after attenuation. Colors represent the magnitude of the atoms as indicated in the vertical colorbars.

beats is not efficient, since maternal and fetal beats may overlap in time.
A possible solution is the attenuation of these values by multiplying them by a constant value $0 < \beta < 1$, For simulations we set this constant equal to $\beta = 0.25$. This ensures that when there is a time overlap between the fetal and maternal beats, we are still able to detect the fetal one (see Fig. 6.8).
The threshold value used for the detection of fetal QRS complexes should take into account the lower power of the fetal signal, and it is set at $T = 0.5A_{max}$ in the experiments. The time interval used to classify close peaks as redundant is reduced to 0.25 s, given that the fetus heart rate is faster than the maternal normal sinus rhythm (i.e., about 110 to 160 beats per minute (bpm) [97]). We compute the time differences between the detected atom positions in the three channels, and select the one with the lowest difference variance, corresponding to the most regular beat rate.

6.7 Evaluation of the Proposed Method

6.7.1 Experimental Set-Up

Fig. 6.2 summarizes the approach implemented for this work, which consists of five main steps, namely, (1) compression of the raw abdominal ECG signals using Compressive Sensing, (2) application of ICA on the compressed measurements, (3) reconstruction of the Independent Components via their sparse representation, (4) detection of the maternal beats, (5) detection of the fetal beats.

Before calculating the CS measures, raw signals are pre-processed using a zero phase high-pass Butterworth digital filter for baseline wander removal, with cut-off frequency equal to 2 Hz. To remove power-line interference, a second order notch filter at 50 Hz or 60 Hz is applied, depending on the fact that signals come from European or US recordings. Original signals are sampled at 1 kHz.

Compressive Sensing is applied on signal blocks of length $N = 250$ samples, using a sparse sensing matrix $\mathbf{\Phi}$, with only two non-zero elements in each column, whose row position is randomly chosen. This class of sensing matrices is particularly interesting for low-power applications and allows an efficient implementation in hardware. In the experiments, we consider signal blocks of $N = 250$ samples, and take $M = 62$ measurements, corresponding to a compression ratio CR=75%. With a 63×250 sensing matrix, we require 437 additions. The choice of a compression ratio of 75% is based on the results reported in [40], where it has been shown that it is the maximum CR that ensures a good reconstruction quality of the signals. However, for the sake of completeness, in Section 6.7.5, we evaluate the performance of the proposed framework at different CRs. Compressed data are represented with 16 bits.

In the proposed framework, we reconstruct signal ICs and not the original mixed signals. However, these can be recovered by simply multiplying the estimated mixing matrix and the reconstructed ICs. As mentioned, in this work we use the reconstruction algorithm proposed in [94], the Smoothed $l0$ (SL0) algorithm, which allows real-time implementation. Detection and classification are applied on the sparse representation and the fHR is computed every 2 s. This time represent a good trad-off keeping the analysis window short enough to allow a real-time analysis and still have an accurate detection.

6.7.2 Standard fECG Evaluation Datasets

The robustness of signal processing methods for non-invasive fECG extraction should be evaluated quantitatively using standard databases. The evaluation of the proposed framework is performed on three public database.

The first one is the Abdominal and Direct Fetal Electrocardiogram Database [64] [70] denoted as the *Silesia dataset*. It contains 5 multichannel fECGs, obtained from 5 different women, each one consisting of four abdominal fECG recordings as well as one fECG recording taken directly from a scalp electrode and used as reference. The position of the electrodes was constant during all recordings. The sampling rate is 1 kHz with a resolution of 16 bits. For this dataset, data are already pre-processed by digital filtering for removal of power-line interference and baseline drift. The R-wave locations were automatically determined in the direct fECG signal and these locations were verified by a group of expert cardiologists.

The second and the third datasets we consider are set A and set B of the Physionet Challenge dataset [64]. In particular, set A contains 75 records and both records and expert annotations are available. Whereas, for set B, which contains 100 records, only the records were made public and the scoring of the performance is performed online. For the non-invasive datasets, the annotations were obtained from fECG QRS estimates derived manually. Due to the inaccuracy of reference annotations, records a38, a46, a52, a54, a71, a74 of dataset A are discarded, as suggested in [13]. Both datasets consist of 1-minute long 4-lead abdominal fECG recordings, all sampled at 1 kHz with a resolution of 16 bit, but recorded with different hardware, following separate protocols and on a number of pregnant women.

6.7.3 Standard Evaluation Metrics

For the evaluation of the proposed scheme, we use classical performance figures usually applied for the assessment of QRS detection algorithms, i.e., sensitivity (Se) and positive predictivity (P+), as well as the F measure, defined in Sec. 4.6.1.

Additionally, we apply the scoring methods proposed in [38], using two metrics, i.e., fetal heart rate measurement and RR interval measurement. The first one, denoted here as *HRmeas* (bpm^2), is used to assess the ability of the algorithm to provide valid fHR estimation. It is based on the squared difference between matched reference (fHR) and detected fHR^d measurements every 5 s (12 instances for 1 min long signals)

$$HRmeas = \frac{1}{12}\sum_{i=1}^{12}(fHR_i - fHR_i^d)^2. \tag{6.7.1}$$

To asses the ability of the algorithm to extract the correct fetal QRS locations with respect to the reference markers, the RR measure metric is used, which is calculated from the differences between matched reference RR and test RR^d intervals. This metric is denoted here by *RRmeas* (ms)

$$RRmeas = \sqrt{\frac{1}{I-1}\sum_{i=1}^{I-1}(RR_i - RR_i^d)^2}, \tag{6.7.2}$$

where I is the total number of fetal QRS complexes in the reference. These measures correspond to the Physionet Challenge scores and were obtained with the same code used by the Challenge scorer. Additionally, we apply the scoring methods proposed in [38], using two metrics, i.e., fetal heart rate measurement and RR interval measurement.

Since we are applying a compression technique, i.e. compressive sensing, reconstruction quality is also evaluated using the PRD metric, as defined in Sec. 3.5.2. As for adult ECG we consider reconstructions with PRD values between 0% and 2% as "very good" quality, while values between 2% and 9% are categorized as "good".

Finally, we consider the total time required by the algorithm for beat classification, including reconstruction of all the 4 channels, in order to asses the possibility to implement the proposed framework in a real-time application. The average time required by the algorithm, for a 1 minute long signal, is approximately 3.7 s. The reconstruction program is written in Matlab, running on an Intel Core i7 processor, equipped with 16 GB memory.

6.7.4 Results for Detection Performance

Table 6.1 reports a full evaluation of the proposed framework in terms of the number of correctly found beats (TP), sensitivity Se, positive predictivity P+ and F measures, for all the 5 records in the *Silesia dataset*. The compression ratio resulting from Compressive Sensing is 75%, since we take 63 measurements every 250 signal samples at 1 kHz. For each signal the first row of the table shows results obtained within the first minute, whereas the second row shows results for the whole five minutes long signals. The same randomly selected sparse sensing matrix has been used for all traces and signal segments. For this dataset and the selected sensing matrix, the average value for sensitivity is Se=92.5% within the first minute, and Se=90% for the five minutes long signals. The positive predictivity average values are P+=92% and P+=88% within the first minute and for five minutes, respectively. Average *HRmeas* and *RRmeas* values within the first minute are 10.34 bpm^2 and 12.33 ms, respectively. For 5 minutes long signals we obtain *HRmeas* = 53.55 bpm^2 and *RRmeas* = 16.48 ms.

Table 6.1: Results, after one minute and after five minutes, for signals from the Abdominal and Direct Fetal ECG Dataset (Silesia). The same sensing matrix has been used to sense all signals.

Record	**t**	**TP**	**FP**	**FN**	**Se** %	**P+** %	**F** %	**mean PRD** %	***HRmeas*** (bpm^2)	***RRmeas*** (ms)
r01	1'	129	0	0	100	100	100	9.85	0.0048	1.49
	5'	628	21	15	98	97	97.5	10.29	22.1578	13.10
r04	1'	110	17	15	88	87	87.5	5.52	23.6373	12.30
	5'	508	142	123	80	78	79	5.93	75.9085	22.21
r07	1'	122	3	5	96	97	96.5	6.20	14.7479	14.74
	5'	5721	63	55	91	90	90.5	5.21	33.0583	16.04
r08	1'	123	10	9	93	92	92.5	10.24	0.0038	12.03
	5'	638	18	12	98	97	97.5	10.60	7.2327	10.12
r10	1'	109	20	19	85	84	84.5	10.55	13.30	20.98
	5'	519	154	117	82	77	79.5	10.24	129.4236	20.92

Evaluation of the proposed detection method on the dataset A, using a fixed randomly selected sparse sensing matrix $\mathbf{\Phi_A}$ for all the signals of the dataset, gives an average sensitivity Se=78% and an average positive predicitivity value P+=77%. Concerning the scoring method proposed by the Physionet Challenge the average *HRmeas*, for the same sensing matrix $\mathbf{\Phi_A}$, is 138.65 bpm^2 and the *RRmeas* is 20.92 ms. A full evaluation for all the records in set A is reported in Appendix A.2. Since the Challenge dataset contains signals from different databases, recorded using different instrumentation and methods, results have an inhomogeneous distribution. Therefore, we analyzed the distribution among the signals of the dataset, obtaining a minimum value for sensitivity equal to 15% (signal a18) and a maximum value 100% (signal a32), while positive predictivity values range from 21% up to 99%. The median value of the sensitivity distribution is about 87.7%, while the positive predictivity median value is 85.5%. Fig. 6.9(a) shows the boxplot of the sensitivity and positive predictivity values. The distribution of HRmeas and RRmeas for

the proposed procedure is presented in Fig. 6.9(b)(c). Median value for the *HRmeas* and *RRmeas*, are 31.58 bpm^2 and 17.96 ms, respectively.

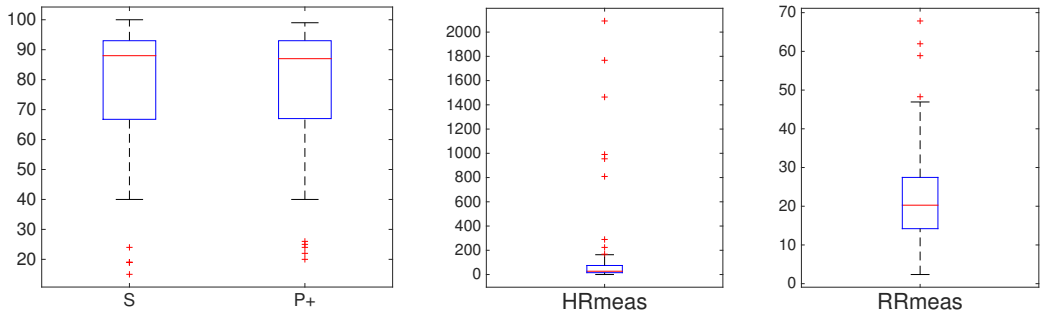

(a) (b) (c)

Figure 6.9: Distribution of (a) of the sensitivity S (%) and the positive predictivity P+ (%), (b) HRmeas (bpm^2)and (c) RRmeas (ms) for Challenge set A.

To assess the influence of different sensing matrices on the ability to correctly classify the fetal beats, we tested 50 different sensing matrices on the Challenge dataset A, for a compression ratio CR=75% and the entire dataset. Average results obtained for sensitivity are Se=78±1% and for positive predictivity P+=77±1%. Again using 50 different sensing matrices, the values for *HRmeas* are 133.16 ± 9.04 bpm^2 and 21.41 ± 0.6 ms for *RRmeas*. The influence of the sensing matrix appears to be moderate.

Finally, we tested the proposed framework on dataset B of the Challenge. Also for this dataset, the proposed method has been tested using 50 different sensing matrices. (Average and standard deviation results for each signal are reported in Table A.3). The scores are *HRmeas*=188 ± 13 bpm^2 and *RRmeas*=24.52 ± 0.26 ms.

To allow the comparison with some off-line methods considered in Section 6.8, we add a limited complexity post-processing stage after real-time detection, in order to correct the estimated fHR and RR time series (we call this variation of the proposed method *smoothed* in Section 6.8). It operates on 1 min long blocks of the detected fetal channel and consists in the removal of beats that are too short to be physiologically possible. It also checks for missed beats, using an approach similar to the one proposed in [50]. After this post-processing stage, the average scores obtained using dataset B, with 50 randomly chosen sensing matrices, are 136 ± 11 bpm^2 and 17.23 ± 0,41 ms for *HRmeas* and *RRmeas*, respectively.

6.7.5 Effect of Compression Ratio

Fig. 6.10 shows a comparison between the reconstruction obtained with the proposed Gaussian dictionary and a wavelet-based sparsifying basis, as commonly adopted for CS implementation [90]. In accordance with the results presented in [40], relative to standard ECG traces and a simpler dictionary, the use of the proposed Gaussian dictionary for mother and fetal ECG compression allows a great improvement in the reconstruction quality (on average, about 75% CR with the Gaussian dictionary against a smaller 50% CR with wavelets, for a "good" reconstruction quality PRD=9%). On average, the proposed

dictionary seems to ensure good reconstruction quality for signals in the *Silesia* dataset (see Table 6.1), with a mean PRD value of 8.5%. Even if some signals have a PRD greater than 9%, the ability of the algorithm to correctly detect and classify the fetal beats does not seem to be affected. Concerning the reconstruction quality on the Challenge dataset A, results show a mean PRD value of 7.5 %, satisfying the reconstruction requirements for diagnostic purposes.

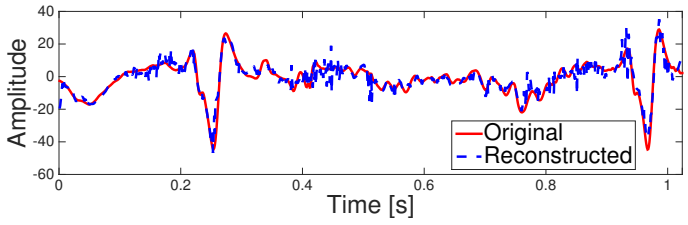

(a)

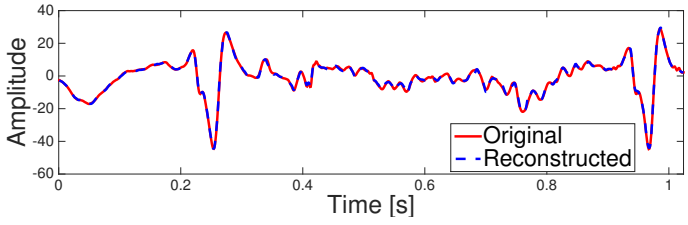

(b)

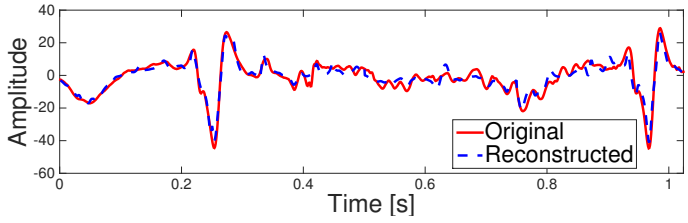

(c)

Figure 6.10: (a) and (b) - Reconstruction of 1 s of record a21, channel 1, of Challenge A dataset for the same compression ratio CR=75%. In (a) a wavelet basis is used (PRD=28.19%), while in (b) the Gaussian overcompelete dictionary is used (PRD=5.48%). (c) Reconstruction of record a21, channel 1, with compression ratio CR= 90% and the Gaussian Dictionary, PRD = 29.25%.

In Fig. 6.11, the performance of our framework at six different compression ratios, ranging from 60% to 90%, are reported. Results are averages of the performance for the entire dataset A, using 50 different sensing matrices at each compression ratio. Both mean and standard deviation values are shown for each CR. As it can be seen in Fig.6.11(a), the reconstruction quality changes with the CR, as well as the reconstruction time, from 6.7 ms for a CR=60% to 4.6 ms when the CR is 90%, see Fig.6.11(b).

The detection performance in term of sensitivity, positive predictivity, *HRmeas* and *RRmeas* shows a negligible variation when the reconstruction ensures a good quality (i.e., for CR<75%), as shown in Fig.6.11(c) and (d) for *HRmeas* and *RRmeas*. In particular, the average sensitivity value ranges from 79%, for CR=65%, to 76% for CR=80%. The

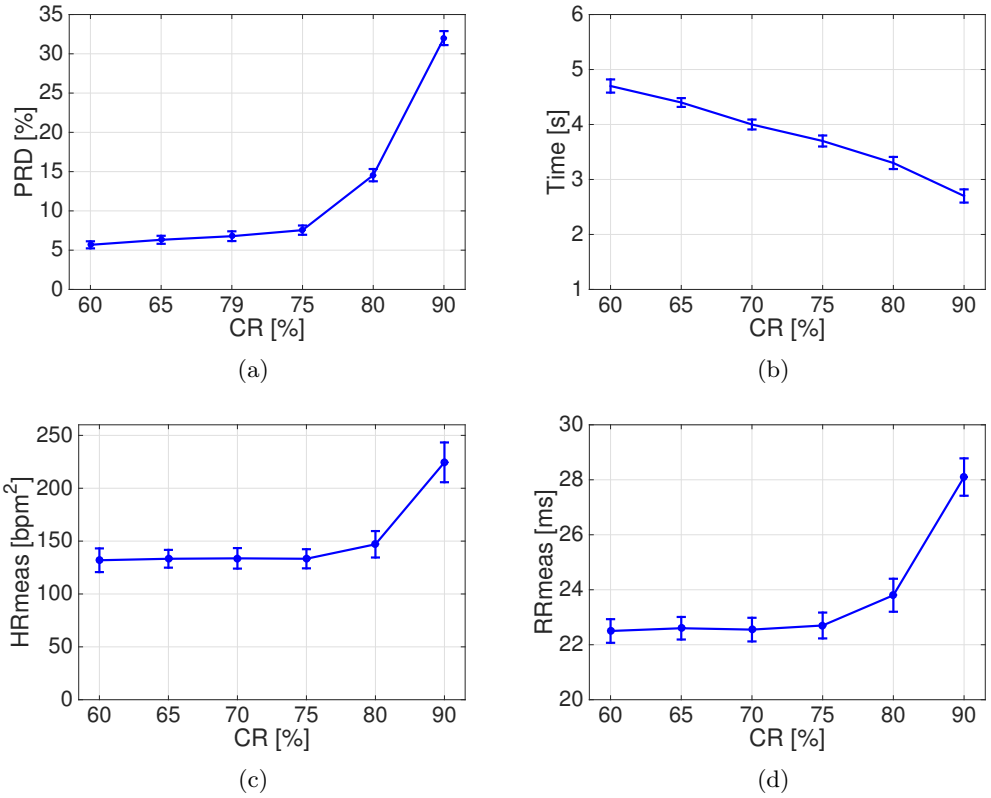

Figure 6.11: Average reconstruction performance (i.e., average PRD (a) and reconstruction time (b) values) at six different compression ratios ranging from 60% to 90%, for dataset A; (c) *HRmeas* and (d) *RRmeas* at different CRs. Error bars indicate standard deviation.

positive predictivity value shows a similar trend. Also *HRmeas* and *RRmeas* show small variations (i.e., a variation of 3 bmp^2 for *HRmeas*, and less then 1 ms for *RRmeas*). Higher compression ratios (CR>80%) cause a loss of detection performance, both in terms of S and P+ and in terms of *HRmeas*and *RRmeas*. At CR=90%, sensitivity decreases to 61% and P+ decreases to 59%, while *HRmeas* and *RRmeas* increase to 225 bpm^2 and 28 ms, respectively.

6.8 Discussion and Relation to other fECG extraction methods

Unlike others methods proposed in literature, which are designed for off-line analysis after pre-processing of the signal, our framework allows detection of fetal beats, as well as maternal beats, by processing short 2 s signal blocks during signal reconstruction within the CS framework. Due to its relatively low complexity, the proposed procedure can be implemented in real time at the receiver. From the best of our knowledge, this is the first time that a framework for the low-power CS compression of fetal abdominal ECG is proposed combined with a beat detector for fHR estimation.

Data from this *Silesia* dataset have been previously used in [5], for the validation of a wavelet-based method. Sensitivity values reported in [5] range from Se=85% to Se=95% for the five minute long signals. Instead, when it is evaluated on each minute, the sensitivity ranges from about 72% up to 98%. From the results reported in Table 6.1, it is apparent that the proposed framework allows a similar performance, besides the fact that it operates on compressed data which are acquired with very little complexity. The *HRmeas* and the *RRmeas* values on this dataset have been previously reported in [111] using two different approaches, with average values *HRmeas*=122.5 bpm^2 and *RRmeas*=31.46 ms for the global approach, and *HRmeas*=11.21 bpm^2 and *RRmeas*=26.64 ms for the time adaptive approach. Thus, on this dataset, our framework (*HRmeas*=53.55 bpm^2 and *RRmeas*=16.48 ms) outperforms the global approach proposed by Rodrigues *et al.* [111] and is comparable with the second one.

In [13], several techniques have been tested on dataset A of the Chellenge dataset. One of the methods operates on the full original signals, computes ICA and uses an adapted version of the standard Pan and Tompkins QRS detector [100]. Besides baseline wander removal and notch filtering, the method does not employ additional pre/post-processing techniques, similarly to what we do in the proposed approach. Compared to this method, our solution increases detection performance by 10% for the mean sensitivity value (reported result for the same dataset is 69.1%). Moreover, the percentage improvement for the positive predictivity is 18% (reported result, P+=60%). When PCA is used instead of ICA in [13] our improvement is 21% and 31% for S and P+, respectively. With respect to the results reported in [13], the mean *HRmeas* achieved by the proposed approach outperforms the methods based on ICA and PCA (mean *HRmeas* value 2852.1 bpm^2 for the ICA based method and mean *HRmeas* value 3892.1 bpm^2 for the PCA method). References [38] [13] also report results obtained with different techniques comprising a combination of ICA processing stages and template subtraction. These techniques operate on the full length pre-processed original signals with no compression. Our method has comparable results with respect to each individual technique considered in [13], even if the most elaborate, which use a combination of individual techniques, can achieve considerably better results than the proposed method.

Methods reviewed in [38] usually follow a four step approach, which includes pre-processing, estimation of maternal component and its removal, estimation of fHR and RR time series, and a post-processing final stage where detected beats are corrected based on prior knowledge. Table 6.2 summarizes and compares the performance of some approaches proposed in [38], with respect to the one proposed in this thesis, for dataset B. As it can be seen, the proposed method achieves results comparable to some off-line methods, even though it works on compressed signals and, using short signal blocks, can be implemented in real-time with a short delay. Regarding execution time, as shown in Fig. 6.2, the proposed approach includes reconstruction from compressed samples, according to the fact that the intended focus of this work is a low-power and low-complexity acquisition/detection real-time system, while existing algorithms operate on non-compressed data. A fair comparison would need to include compression/decompression procedures in the existing schemes. In Table 6.2 we report results by neglecting the time required by the reconstruction stage (see Fig. 2). We evaluate the CPU execution time of Matlab publicly available implementations, running on an Intel Core i7 processor, equipped with 16 GB memory. For the proposed techniques, we report average values, using 50 different sensing matrices. To compute ICA in the compressed domain and to detect the fetal QRS

Table 6.2: Comparison of some detection methods on dataset B. In case of no publicly available implementation, the execution time is set to N.A.

Method	***HRmeas*** bpm^2	***RRmeas*** ms	Realtime	Compression	Execution Time s
ICA-Template Adaption [6]	20.4	4.6	NO	NO	N.A.
FUSE - smooth [13]	29.6	4.7	NO	NO	1.8
ICA-TS-ICA upsampling [125]	34	5.1	NO	NO	1.55
Wiener filter [111]	124.8	14.4	NO	NO	4.8
TS-ICA upsampling [50]	134.5	12.4	NO	NO	200
Proposed method - smoothed	136	17.2	NO	YES	0.64
Proposed method	188	24.5	YES	YES	0.64
ICA-Extended Kalman [6]	219	7.7	NO	NO	N.A.
TS and PCA[36]	305.7	23.1	NO	NO	N.A.
Wavelet based [5]	513.1	35.3	NO	NO	N.A.
TS_{PCA} [13]	759.42	21.86	NO	NO	1.09

complexes the proposed algorithm takes about 0.64 s for a one minute long signal, while the total time, including reconstruction, is about 3.7 s, as mentioned before.

Pre-processing may increase detection performance, especially when impulsive artifact removal is applied. Impulsive artifacts may affect the performance of ICA decomposition, and this may be one of the reasons why the proposed method has for some signal a worse performance than other methods [13]. Moreover, we believe that the proposed method fails on some abdominal signals of the Challenge A dataset due to the poor quality of the signals. As mentioned in [50] and [111], for some signals, such as a02, a09, a18 or a29, due to the low signal quality, it is not possible to correctly identify the fetal beats, obtaining unreliable results. For practical purposes, the choice of the optimal number and position of electrodes is important since the efficiency of the signal processing methods strongly depends on the signal quality and on the number of channels. This observation is confirmed by the different average performance obtained with the Silesia and Challenge datasets.

6.9 Conclusions

In this chapter, a framework for the CS compression of abdominal fetal ECG recordings jointly with real time beat detection and classification has been presented.

We proposed to take advantage of the reduced dimension of the problem to apply a classical separation technique, i.e., the independent component analysis (ICA), directly in the compressed domain, with a consequent reduced computational cost. By applying the ICA in the compressed domain we obtained a good estimation of the original independent components (ICs), which are reconstructed from the compressed ICs by following the same procedure used to reconstruct compressed sensed signals. To this end, we introduced a new dictionary for the sparsification of maternal and fetal components that can be used to recover the compressed measurements as well as the compressed independent components. The fetal beast detection is based on the atoms used for the reconstruction of the fetal trace in the independent components.

Evaluation of the proposed method has been done on three open datasets, showing

promising results for both correct detection of fetal beats (Se=92.5% for Silesia dataset and Se=78% for the Challenge dataset A) and reconstruction quality (PRD=8.5% and PRD=7.5%). These results allow us to conclude that the proposed framework may be used for compression of abdominal fECG and to obtain real time information of the fetal heart rate, providing a suitable solution for low-power telemonitoring applications.

7

Improving the Smoothed $l0$ Algorithm

7.1 Introduction

In this chapter, we return to the problem of signal reconstruction from compressed measurements. As seen before, the Smoothed-l_0 algorithm (SL0) introduced in Sec. 2.4.3, relaxes the original NP-hard problem by using smooth approximations of the l_0 norm. SL0 is particularly interesting for its low complexity, which allows for real-time signal reconstruction [42]. However, as we will see in this chapter, the SL0 algorithm may fail in case of noisy signals or ill-conditioned sensing/dictionary-based reconstruction procedures. We describe a variant of the SL0 technique in this more challenging setting and we show that, while maintaining the same computational cost of the original algorithm, the proposed modification significantly improves the reconstruction quality, both for synthetic and real-world ECG signals. We also show that the proposed algorithm allows robust heart beat classification when sparse matrices, implementable with very low computational complexity, are used for compressed sensing of the ECG signal. Furthermore, similarly to the analysis performed in Chapter 3, we show that the use of sparse sensing matrices with only 2 non-zero elements in each column compares successfully with the use of random Gaussian matrices even when using the proposed variant of the SL0 on fetal ECG signals.

7.2 Regularization of Smoothed $l0$ Alogorithm

As described in Sec. 2.4.3, the SL0 algorithm approximating the l_0-norm with a continuous function, and the minimiazation of the l_0-norm is approximately equivalent to maximaize $F_\sigma = \exp(-s_i^2/2\sigma^2) \approx 1 - s_i^2/2\sigma^2$ s.t. $\mathbf{y} = \mathbf{As}$. This problem resembles the minimum $l2$-norm solution and the starting solution of the optimization process is usually calculated using the pseudo-inverse $\mathbf{A}^\dagger$ of $\mathbf{A}$ and set $\mathbf{s_0} = \mathbf{A}^T(\mathbf{AA}^T)^{-1}\mathbf{y}$.

Note that the algorithm requires that matrix $\mathbf{A}$ has full rank M . When the number of measurements M increases (i.e., the compression ratio decreases), this requirement may become critical, making the reconstruction problem ill-conditioned and sensitive to noise.

In real scenarios where the sparse signal or the measurements are affected by noise, if the compound matrix $\mathbf{A}$ is ill-conditioned, then application of $\mathbf{A}^\dagger$ amplifies the error and results in poor reconstruction, even using the Robust SL0 proposed in [56]. Ill-conditioned matrix $\mathbf{A}$ may no longer satisfy the requirements imposed by traditional compressed sens-

ing assumption and the recovery algorithm becomes sensitive to noise: in this situations small perturbations in the signal or in the measurements deviate the reconstruction from the original signal. Introducing a regularization term in the optimization problem enables a stable recovery of $\mathbf{x} = \mathbf{Ds}$.

As in the SL0 algorithm, we approximate the $l0$-norm by using (2.4.5), and the algorithm again consists in two nested iterations. The internal loop seeks the maximum of F_σ in the feasible set $\{\mathbf{s}| \parallel \mathbf{y} - \mathbf{As} \parallel_2 \leq \epsilon\}$. At each step we compute $\tilde{\mathbf{s}} = \mathbf{s} - \mu\delta_k$ and project $\tilde{\mathbf{s}}$ by solving

$$\min_{\hat{s}} \parallel \hat{\mathbf{s}} - \tilde{\mathbf{s}} \parallel_2 \qquad \text{s.t.} \qquad \parallel \mathbf{A}\hat{\mathbf{s}} - \mathbf{y} \parallel_2 \leq \epsilon. \tag{7.2.1}$$

Using the Lagrangian function of Eq. (7.2.1), the problem can be rewritten as

$$\min_{\hat{s}} \parallel \mathbf{A}\hat{\mathbf{s}} - \mathbf{y} \parallel_2^2 + \lambda \parallel \hat{\mathbf{s}} - \tilde{\mathbf{s}} \parallel_2^2, \tag{7.2.2}$$

where λ is the regularization parameter. The solution, whose full derivation is given in Appendix B, is

$$\hat{\mathbf{s}} = \tilde{\mathbf{s}} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T + \lambda\mathbf{I}_M)^{-1}(\mathbf{A}\tilde{\mathbf{s}} - \mathbf{y}). \tag{7.2.3}$$

As for the SL0 algorithm, for large σ values, the solution is equal to the $l2$ norm solution subject to $\parallel \mathbf{y} - \mathbf{As} \parallel_2 \leq \epsilon$. Solving the problem

$$\min_{s} \parallel \mathbf{As} - \mathbf{y} \parallel_2^2 + \lambda \parallel \mathbf{s} \parallel_2^2, \tag{7.2.4}$$

we set the initial solution of the algorithm to $\mathbf{s_0} = \mathbf{A^T}(\mathbf{AA^T} + \lambda\mathbf{I}_M)^{-1}\mathbf{y}$.

The proposed λSL0 algorithm is summarized in Algorithm 3. The value of the regularization parameter λ represents a compromise between the two terms of the cost function. When the noise norm ϵ is small, $\lambda \to 0$, and the algorithm reduces to the original SL0 for the noiseless case. We carried out some experiments (results are omitted due to lack of space) and we observed that the value of λ is not critical and should be $\sim$10$-$100 times the expected noise ϵ.

Algorithm 3 λSL0

Input: μ step size, $\mathbf{y}$, $\mathbf{A}$, σ_{dec}, σ_{min}, λ, K_{iter}
Initialization: $\mathbf{s}_0 \leftarrow \mathbf{A^T}((\mathbf{AA^T}) + \lambda\mathbf{I}_M)^{-1}\mathbf{y}$,
  $\sigma_1 = 2|\max(\mathbf{s}_0)|$
  while $\sigma_k < \sigma_{min}$ **do**
    for k=1:K$_{iter}$ **do**
      $\delta_k \leftarrow \mathbf{s} \cdot [e^{-\frac{s_1^2}{2\sigma_k^2}}, \ldots, e^{-\frac{s_N^2}{2\sigma_k^2}}]^T$
      $\mathbf{s} \leftarrow \mathbf{s} - \mu\delta_k$
      Project $\mathbf{s}$ onto the feasible set: $\{\mathbf{s}| \parallel \mathbf{As} - \mathbf{y} \parallel_2 \leq \epsilon\}$
      $\mathbf{s} \leftarrow \mathbf{s} - \mathbf{A^T}((\mathbf{AA^T}) + \lambda\mathbf{I}_M)^{-1}(\mathbf{As} - \mathbf{y})$
    end for
    $\sigma_k \leftarrow \sigma_k\sigma_{dec}$
    $\tilde{\mathbf{s}}_\mathbf{k} \leftarrow \mathbf{s}$
  end while
Output: $\mathbf{s_{OUT}} \leftarrow \tilde{\mathbf{s}}_\mathbf{k}$

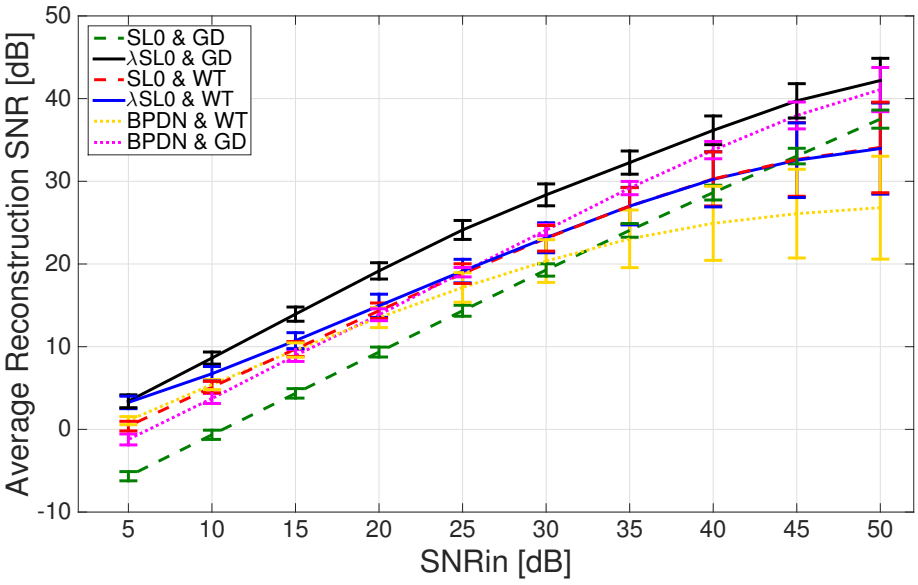

Figure 7.1: Reconstruction SNR versus input SNR obtained from 100 trials for simulated fECG signals, at CR=50%, using the SL0, λSL0 and BPBN (SPGL1) algorithms with the Wavelet (WT) or the Gaussian Dictionary (GD).

7.3 Performance of λSL0

In this section, the effect of noise on the reconstruction performance is experimentally analyzed. We compare the performance of the proposed algorithm with the original SL0 and the BPDN-SPGL1 algorithms [123]. The signals used in these experiments are simulated fECG signals [11] with length $N = 256$. As sparsifying dictionaries we use the dictionary of Gaussian like functions proposed in Sec. 6.5, and the Wavelet basis with Daubechies' length-4 filters. The sensing matrix elements are drawn as independent Gaussian random variables [27]. We repeat the experiment 100 times with different source signals at different noise levels, and using each time a different random sensing matrix. The reported SNR value is the average of these simulations. In Fig. 7.1, we report the reconstruction SNR as a function of the input SNR (SNRin) when gaussian noise is added to the simulated fECG traces, for a compression ratio CR=50% ($M = 128$). Compared to the original SL0 algorithm, λSL0 allows to achieve better reconstruction quality, especially when the Gaussian dictionary is used. An improvement can be also appreciated when the Wavelet basis is used, especially at lower SNRin values. Note that the use of the Gaussian Dictionary gives much better performance than Wavelets also when BPDN is used.

In addition to the previous experiments, we assess the reconstruction performance as the compression ratio changes. In Fig. 7.2, it is possible to see that the average SNR achieved by the λSL0 algorithm combined with the Gaussian Dictionary outperforms the SL0 method, especially at low compression ratios (M large) and is comparable with respect to the traditional BPDN algorithm, which has a higher complexity. At higher compression ratios (CR$>$ 50%) the Wavelet basis achieves a lower performance independently of the reconstruction algorithm.

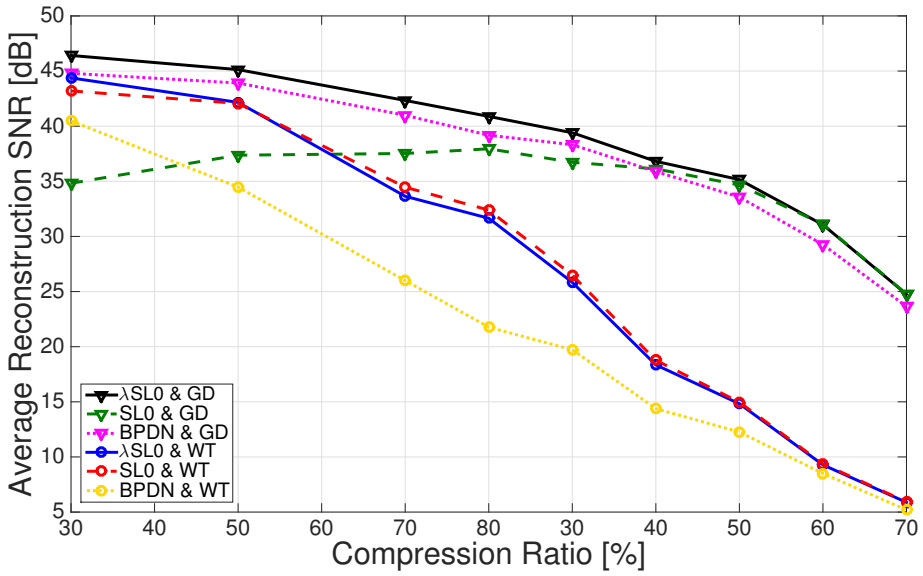

Figure 7.2: Reconstruction SNR versus CR obtained from 100 trials for simulated fECG signals, using using the SL0, λSL0 and BPBN (SPGL1) algorithms with the Wavelet (WT) or the Gaussian Dictionary (GD)

As a measure of the computational cost of the algorithms we use the CPU time, setting the same parameters ($K_{iter} = 3$ an $\sigma_{dec} = 0.5$) for SL0 and λSL0. Experiments show an average reconstruction time for the λSL0 algorithm ranging from 0.07 s, when CR=30%, to 0.01 s, when CR=80%. Thus, it maintains approximately the same computational cost of the original SL0 algorithm (ranging from 0.03 s to 0.01 s), while being much faster than the BPDN algorithm (1.6 s to 0.6 s). Programs are written in Matlab, running on an Intel Core i7 processor, equipped with 16 GB memory.

7.4 Application to Joint Compression and Beat Detection in fECG

In this section, we analyze the performance of λSL0 for fetal beat detection in CS-compressed real-world fECG signals, in particular signals belonging to set A of Challenge database (see Sec. 6.7.2) (excluding badly annotated signal as in [13] and signals containing saturated values [1]).

The metrics used for the assessment of detection quality are the classical figures, i.e. Sensitivity (S) and Positive Predictivity (P+) described in Sec. 6.7.3. Reconstruction quality is assessed by using the PRD value (see Sec. 3.5.2). We also analyze the influence of different sensing matrices, in particular sparse matrices which allow very low complexity of the CS sensor.

[1] Some signals in set A contain invalid values corresponding to invalid output of the A/D converter. In this chapter signal segments containing saturated values, i.e., a02, a09, a11, a16 and a18, were not considered differently from Chapter 6, leading to different overall performance.

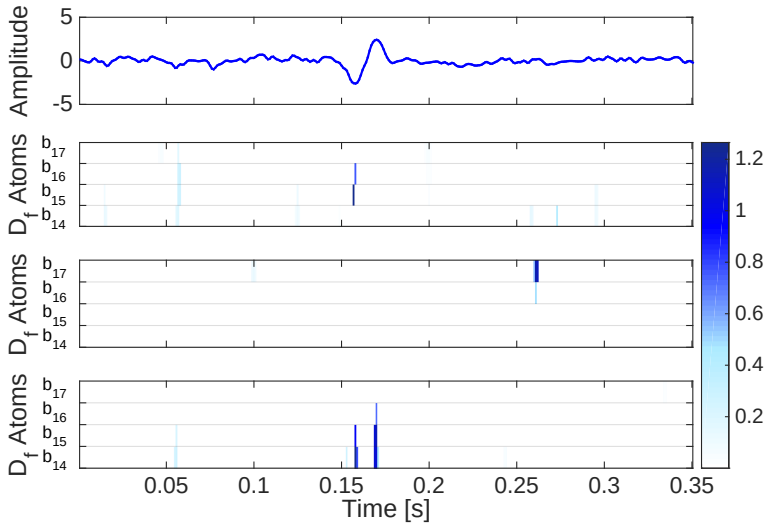

Figure 7.3: Sparse decomposition of the independent component in (a) using the SL0 algorithm, for (b) CR=75% and (c) CR=40% and (e) using the λSL0 algorithm for CR=40%. In the graphs, different intensities represent the weight of the activated atoms.

7.4.1 fECG Reconstruction and Fetal Beat Detection

In Chapter 6 a framework for the compression of multichannel abdominal fECG and joint detection of fetal beats has been proposed. The separation is based on the atoms, belonging to the fetal or mother's sub-dictionaries, activated during the reconstruction process. The compression of the signal is based on Compressive Sensing and uses a binary sparse sensing matrix, containing only $d = 2$ ones in random positions in each column, in order to reduce the sensor complexity [17].

In accordance with the analysis of the previous section 6.7.4, we found out experimentally that, for compression ratios greater than 50%, the detection performance is preserved, while at lower compression ratios, besides the increased information available, fetal beat detection may fail. This problem arises when the compound matrix $\mathbf{A}$ resulting form the multiplication of the sensing matrix by the dictionary matrix is poorly conditioned and the independent components (ICs) are noisy.

As an example, we show in Fig. 7.3 (a) a portion of the IC related to the fetal trace of signal a32 of the 4-channel Physionet Challenge dataset [1], where the fetal beat is clearly visible. When the compression ratio is CR=75%, the reconstruction quality value PRD (percentage root-mean-square difference) is about 6% (average of the 4 channels) and the detection performance in terms of sensitivity (Se) and positive predictivity (P+) are Se=100% and P+=99.34%, respectively. At lower compression ratios detection fails, and for CR=40% we have PRD=0.47%, Se=60% and P+=68%. Fig. 7.3 (b) and (c) show the positions of the fetal activated dictionary atoms for the two cases CR=75% and CR=40%, respectively. As we can see, when CR=40%, the algorithm fails to find the correct sparse representation, leading to a wrong beat detection. Instead, Fig. 7.3 (d) shows the activated atoms when the λSL0 algorithm is used instead of SL0. For the whole signal and using the same sensing matrix at CR=40%, λSL0 achieves Se=100% and P+=99%. We repeat

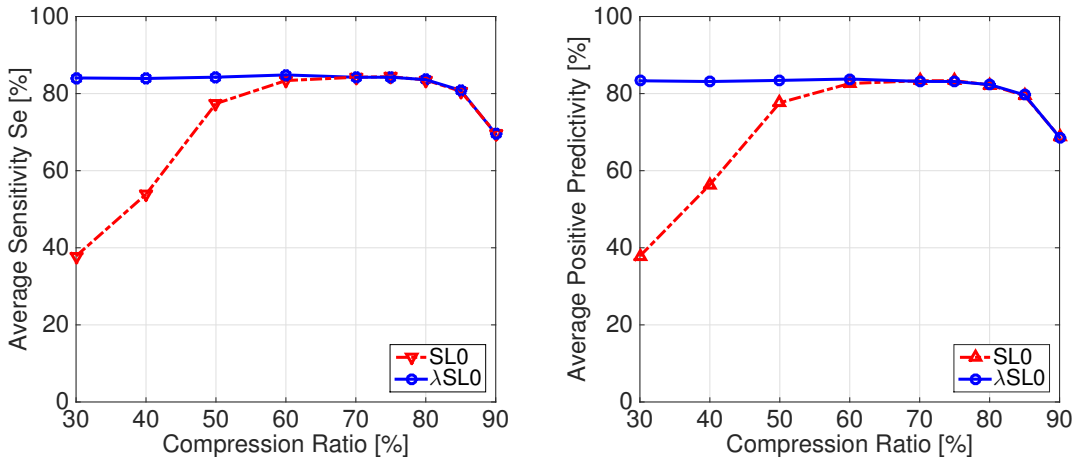

(a) (b)

Figure 7.4: Detection performance for SL0 and λSL0 algorithm. The vertical coordinate gives (a) the average Sensitivity and (b) the average Positive Predictivity at different CR values.

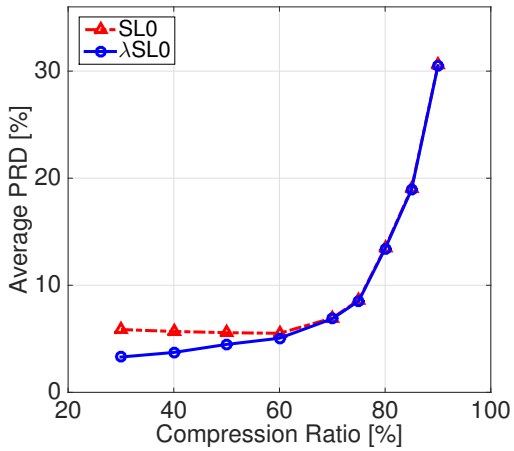

Figure 7.5: Comparison of average PRD when using the SL0 and the λSL0 algorithm at different CRs.

the experiment 20 times with different random sparse binary matrices ($d = 2$), for all the signals in Challenge data set A (see Sec. 6.7.2). The reported sensitivity value is the average of these simulations. Fig. 7.4 (a-b) shows that the detection performance of the proposed algorithm is almost independent of the CR, while the SL0 algorithm fails at lower CRs.

In Fig. 7.5 it can be seen that the proposed λSL0 algorithm outperforms the original algorithm, in terms of average reconstruction quality PRD as introduced in Sec.3.5.2, at low compression ratios.

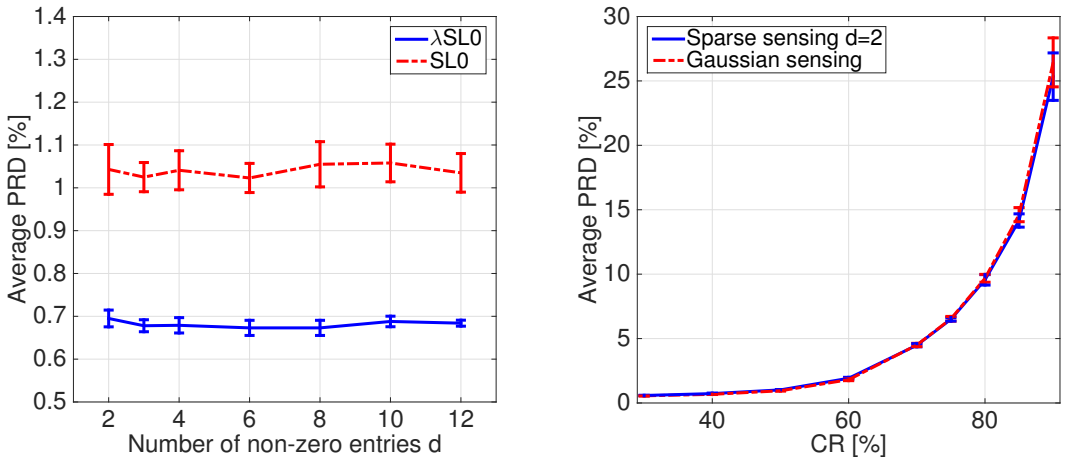

(a) (b)

Figure 7.6: Average recovery quality for signal a25 of dataset A. The vertical coordinate gives the average PRD and the error bar gives the standard deviation. (a) Effects of the number of non-zero entries in each column of the sensing matrix at CR=40% for SL0 and λSL0. (b) Comparison of average PRD using sparse sensing matrices with $d = 2$ and random Gaussian sensing matrices at different CR (λSL0 algorithm).

7.4.2 Influence of the Sensing Matrix

Table 7.1: Average performance of detection and reconstruction for SL0 and λSL0 for dataset A.

		Reconstruction Method					
		SL0			λSL0		
CR %	Sensing Matrix	Se [%]	P+ [%]	PRD [%]	Se [%]	P+ [%]	PRD [%]
40	Sparse d=2	46.6	48.7	5.27	84.9	84.2	3.77
	Gaussian	43.8	45.6	5.58	84.4	83.4	3.72
50	Sparse d=2	77.6	77.6	5.93	84.6	83.4	4.49
	Gaussian	75.6	76.1	5.71	84.7	83.2	4.27
75	Sparse d=2	84.3	83.5	8.81	84.5	83.3	8.14
	Gaussian	84 .1	83.2	8.80	84.3	83.2	8.02

In this section, we analyze the performance of the reconstruction procedure, using SL0 and λSL0 algorithms, when different sensing matrices, generated from i.i.d. Gaussian random variables, or sparse matrices with different d values, are used. There are no theoretical guidelines for choosing the optimal number of non-zero elements d, therefore it has been determined experimentally. In the following experiments, the signal is divided into $N = 250$ sample long segments, which are compressed independently. For each d value and compression ratio, the experiments were repeated using 20 randomly generated different sensing matrices, and the average performance is reported. Results for signal a25 of the

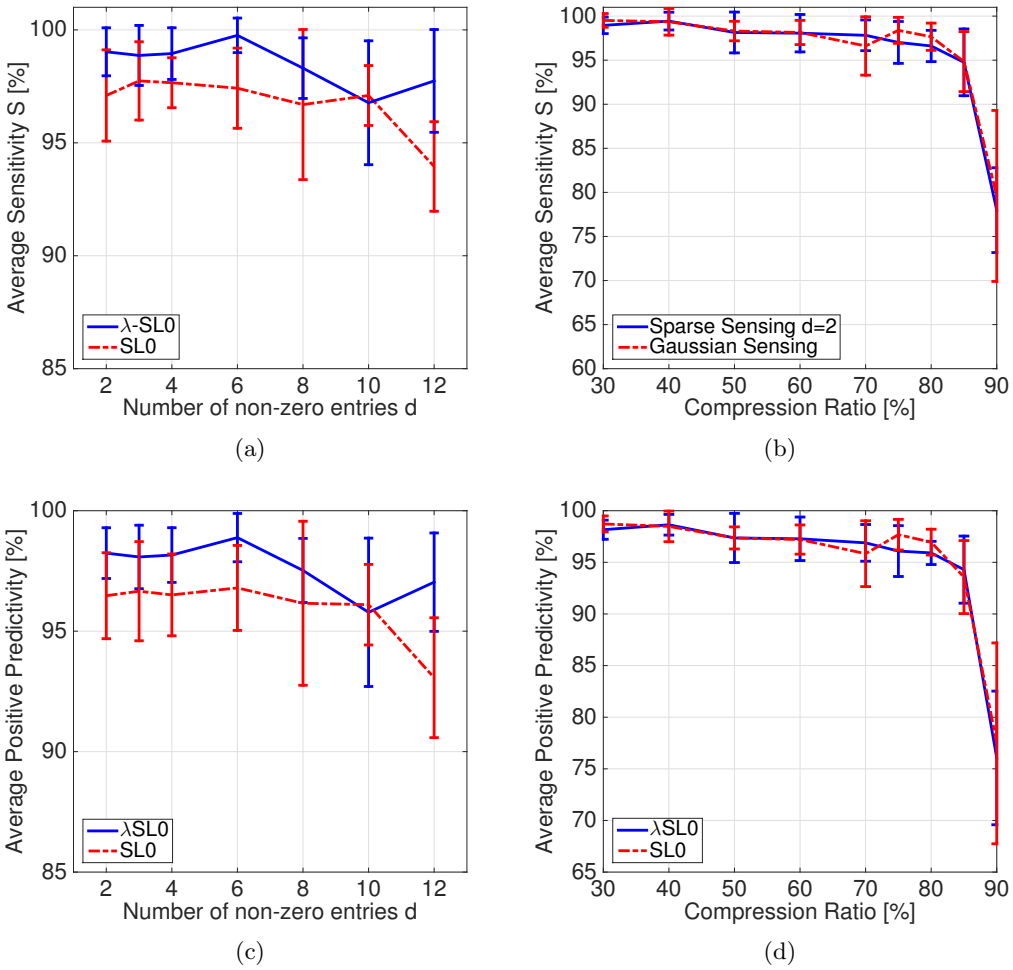

Figure 7.7: (a) Effects of the number of non-zero entries in each column of the sensing matrix on detection performance. The vertical coordinate gives (a) the average Sensitivity and (c) the average Positive Predictivity for signal a25 of dataset A for CR=40%, the error bar gives the standard deviation. (c) Comparison of average Sensitivity and (d) Positive Predictivity when using sparse sensing matrices with $d = 2$ and random Gaussian sensing matrices at different CR (λSL0 algorithm).

Challenge dataset A are shown in Fig. 7.6 and Fig. 7.7. As we can see from Fig. 7.6(a), for a compression ratio CR=40%, increasing the number of ones in each column of the sensing matrix, does not improve the reconstruction quality, both for the SL0 and λSL0 reconstruction algorithms. Note however that λSL0 outperforms SL0. Fig. 7.6(b) shows a comparison of the reconstruction quality obtained with sparse matrices, $d = 2$, with respect to the quality obtained using Gaussian random matrices. Although the theoretical reconstruction performance for i.i.d. Gaussian sensing matrices is well established, we can see experimentally that, for the class of signals we are considering, sparse matrices have similar performance. The use of a sparse sensing matrix with $d = 2$ provides almost identical reconstruction results, despite the very low complexity implementation.

Finally, Table 7.1 summarizes the average reconstruction and detection performance for Challenge dataset A, at different compression ratios, when using a sparse sensing matrix with $d = 2$ and an i.i.d. Gaussian sensing matrix. Experiments are repeated for the SL0 and λSL0 algorithms. Both detection and reconstruction are mostly independent from the sensing matrix, while it is apparent that the λSL0 algorithm allows robust detection and reconstruction.

7.5 Conclusions

In this chapter we proposed a regularized version of the SL0 algorithm. Experimental results confirm that the proposed algorithm has good performance, while preserving the low computation complexity of the original approach. The application of λSL0 to the joint compression and detection framework of fetal ECG also demonstrates that the proposed modification can efficiently reconstruct the signals and correctly detect the beats in the presence of noise and for different compression ratios. Moreover, we have shown that the use of sparse sensing matrices with only 2 non-zero elements in each column, compares successfully with random Gaussian matrices, while permitting a very low complexity implementation.

8

Evaluation of Energy and Detection Accuracy of Compressed Sensed fECG

8.1 Introduction

Within this thesis we have seen that advances in the field of signal processing have made it possible to develop new techniques for an efficient and reliable continuous monitoring of ECG and fECG signals. The issue of reducing power consumption in wearable devices has been addressed by many researchers, and some works also proposed solutions for the specific case of fetal ECG monitoring. However, results that are focused on energy saving do not analyze algorithms and solutions for the analysis of collected signals. In this chapter, we compare different compression schemes, and evaluate their overall performance and ability to preserve relevant diagnostic features as a function of power consumption.
In particular we analyze how effective compressive sensing could be for fECG monitoring in battery constrained devices, with limited computational capacity, compared with classical compression techniques, in particular based on the Discrete Wavelet Transform (DWT). We evaluate the performance via actual implementation of the CS and wavelet compression paradigms on the Shimmer platform [23] for ECG signal monitoring.
Within this chapter we prove that, for the analysis of abdominal recordings of fetal ECG signals, which can be difficult to process due to the low-amplitude of the fetal beats, CS can provide significant advantages with respect to conventional CS schemes based on wavelets, and with respect to DWT schemes in the signal domain. In particular, we show that a CS scheme based on reconstruction with the over-complete dictionary proposed in Sec 6.5, instead of the wavelet basis considered in [127], has similar reconstruction quality than one based on wavelet compression, proving that the CS paradigm is suitable for fECG acquisition, with the advantage of a low power implementation in the sensor. Unlike other works in the literature, we consider, as a figure of merit, the accuracy of fetal beat detection after reconstruction, and compare the results of different compression/transmission/reconstruction procedures as a function of the sensor power consumption.

8.2 Related Work

Energy saving by using compressive sensing for adult electrocardiogram signal monitoring with wireless sensors has been studied in [90]. A sparsifying wavelet basis is used at the receiver. Authors show that compressive sensing is a competitive low-complexity compression paradigm with respect to state-of-the-art DWT-based compression. As reported in the paper, despite an inferior compression performance in terms of signal reconstruction quality, the overall energy efficiency achieved by CS outperforms the DWT-based ECG compression scheme, thanks to the lower complexity and lower CPU execution time. Accordingly with the results reported in [90], is it possible to achieve a life extension of the battery up to 37% with an acceptable reconstruction quality. Results are based on an actual implementation on the Shimmer 2r platform [23] for adult ECG signal monitoring, considering a sampling frequency of 256 Hz and 11 bit resolution.

The application of the compressive sensing paradigm for fetal electrocardiogram signal acquisition, however, introduces some issues. Indeed, as discussed in Sec. 6.2 non-invasive records are mixture of the fetal, the mother's heart beats and noise. As stated by Zhang *et al.* [127], the use of the traditional wavelet basis is not suitable for the reconstruction of fECG signals from compressive sensing measurements. In [127], a new method, namely Block Sparse Bayesian Learning (BSBL), was introduced to overcome the limitations of the traditional compressive sensing framework. Exploiting the spatial, temporal and dynamic structure of signals, it enables to reconstruct non-sparse signals with high quality. In particular, the reconstruction process proposed in [127] does not destroy the interdependence structure of multichannel recordings, in order to allow the application of fetal beat detection algorithms, usually based on independent component analysis or Blind Source Separation (BSS) [49]. Besides considering the mean squared error (MSE) between reconstructed and original signals, performance assessment of BSBL is based on the Pearson correlation between the detected fetal beats in the original and reconstructed signals.

As evidenced by previous works, compressive sensing seems to be an effective tool for low power compression when compared to traditional schemes such as wavelet based compression. In the following, CS energy saving is analyzed in the context of fetal ECG detection performance. We show that the use of an appropriate over-complete dictionary, such as the one proposed in Chapter 6, at the receiver can make CS adequate for fetal ECG reconstruction, allowing significant power saving for a given reconstruction quality or beat detection accuracy with respect to traditional DWT based compression.

8.3 Method

8.3.1 Evaluation Framework

The work-flow adopted for the CS paradigm evaluation in fetal ECG acquisition is shown in Fig. 8.1. The aim of this work is to compare the energy and reconstruction/detection performance of two encoding procedures, namely CS acquisition and DWT compression (Section 8.3.5).

Energy consumption for the two compression schemes is evaluated at different compression ratios as described in Sec. 8.3.2. Both encoding techniques operate on non-overlapping signal blocks of N=256 samples, for signals sampled at 1 kHz with 16-bit resolution.

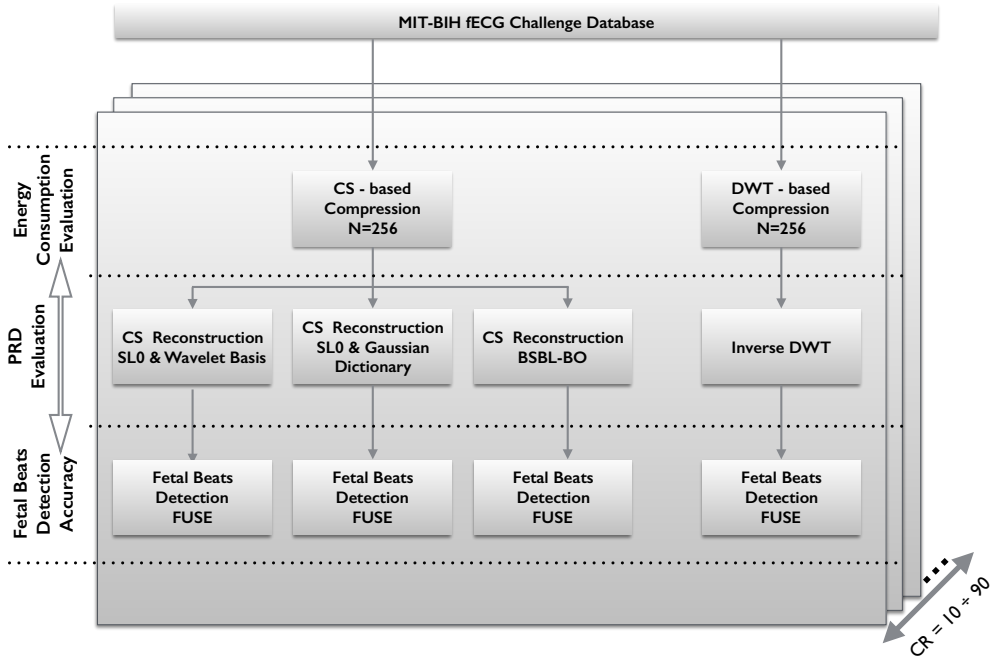

Figure 8.1: Proposed evaluation method

Since the aim of this work is to comparatively evaluate the CS-based compression schemes and DWT on a standard database, we do not use real-time ECG acquisition to asses the recovery and detection quality. Thus, experiments are carried out using non-invasive fetal ECG signals from set-A of the public database described in Sec. 6.7.2.

For the assessment of fetal beat detection accuracy on compressed ECG signals, we consider four different scenarios. The first three scenarios require to reconstruct the ECG signals from compressed sensed measurements. To this end, we adopt two different sparsifying matrices in the decoder reconstruction process. In particular, we use the traditional Daubechies (DB4) wavelet basis with a 5-level decomposition, and the over-complete Gaussian dictionary proposed in Sec. 6.5. The over-complete dictionary is specifically designed to preserve the relevant waves of both maternal and fetal electrocardiogram signals [42]. The reconstruction algorithm adopted to solve the inverse problem is the λSL0 described in Chapter 7, which allows to achieve good performance in presence of noisy signals and ill-conditioned matrices [41]. Moreover, we also consider for comparison the BSBL-BO algorithm, used for fECG reconstruction in [127].

In the fourth scenario, the ECG signal is compressed using a DWT-based method (Section 8.3.5) and, at the decoder, a standard inverse DWT is applied to reconstruct the signals on the basis of the received coefficients.

After signal reconstruction form the two compression schemes (CS or DWT) the FUSE method [13][1] for fetal ECG extraction and beats detection is applied. The fetal extraction method is a combination of template subtraction and principal/independent com-

[1]Authors have made available the entire FUSE code at https://physionet.org/challenge/2013/sources/.

ponent analysis, while the final fetal QRS complex detection is performed by using a Pan-Tompkins QRS detector on all the obtained channels [100], then the one with the smoothest fetal heart rate time series is selected. Assessment of the performance of the different scenarios is evaluated using the metrics reported in Sec. 8.3.3.

Finally, we combine the energy consumption with the detection performance in order to establish the actual energy saving that one can achieve while guaranteeing a certain detection accuracy.

8.3.2 Energy Consumption Evaluation

Energy requirements of the two different compression algorithms are evaluated on the basis of the actual number of MCU cycles and transmission bitrate required by implementation in a commercial acquisition device. Since energy consumption does not depend on the actual signals, we evaluate the cost of the algorithms in a real implementation, although experiments, for comparison purposes, are carried out off-line on signals of the public database described above.

The hardware considered in this work is the one present on Shimmer devices[2] [23] powered by a rechargeable Li-polymer battery, with an internal ECG daughter board, validated for ambulatory and research purposes. In our experiments, the sampling rate of the device is set to 1 kHz. The device includes the low-power Texas Instrument 16-bit MSP430F5438 micro-controller [69], and a low-power CC2420 IEEE 802.15.4 [68] compliant radio. The MSP430 is a 16-bit word processor, and the compression performance of both algorithms is evaluated using 16 bit arithmetic. Code Composer Studio (CCS) has been used to generate the firmware binary code. One of the functionalities of the CCS development kit allows to count MCU cycles for the running code.

We consider the energy cost E_{comp} required by the compression algorithms to process one signal block of $N = 256$ samples. E_{comp} can be expressed as

$$E_{comp} = N_{cyc} E_{cyc}, \tag{8.3.1}$$

where N_{cyc} is the number of MCU cycles to encode one signal block.

The energy consumption per clock cycle E_{cyc} can be easily calculated for the considered micro-controller, which in active mode consumes 312 μA/MHz when the MCU operates at 8 MHz and the supply voltage is +3 V, namely

$$E_{cyc}(@8\text{ MHz}) = 312 \cdot 3 \cdot 10^{-12} = 0.936\text{ nJ/cycle}. \tag{8.3.2}$$

Since the two compression schemes may require different bitrates for the same reconstruction quality or beat detection capability, we also consider in the following the transmission cost, which usually has the highest impact on the overall energy consumption. To this purpose, we take into account the Texas Instruments CC2420 radio specifications [68]. In [68], it is reported that the energy consumption per transmitted bit is $\mathrm{E}_{bit} = 230$ nJ/bit.

The transmission energy, E_{tx}, to transmit one signal block equals therefore

$$E_{tx} = N_{bit} E_{bit}, \tag{8.3.3}$$

[2] http://www.shimmersensing.com/.

where N_{bit} is the number of bits necessary to encode the block.

Finally, he total energy required to process and transmit a signal block is given by.

$$E_{TOT} = E_{comp} + E_{tx} = N_{cyc}E_{cyc} + N_{bit}E_{bit}. \tag{8.3.4}$$

In the following, we do not consider the energy required by the ADC, since the cost is the same in both scenarios.

8.3.3 Reconstruction Quality Assessment

In order to assess the quality of the reconstructed signals, we use the traditional PRD quality metric, defined as in Sec 3.5.2, usually employed for the assessment of adult ECG reconstruction quality. In the experiments described below, the PRD value is computed for each reconstructed signal block in every channel and then the average and standard deviation are reported.

Moreover, the results of fetal beat detection are used as a reconstruction quality measure. In particular, we evaluate the Sensitivity (Se) and the Positive Predictivity (P+), as introduced in Sec. 4.6.1.

8.3.4 CS Implementation

We implement CS compression by using sparse sensing matrices, with two non zero elements in each column, similarly to the one used in Chapter 6. In particular, the non zero elements are equal to 1, thus the compression stage reduces to the sum of signal samples indexed by the matrix elements, that can be implemented using a single accumulator and no floating-point multiplication. Moreover, instead of storing the whole sensing matrix, it is possible to store just the positions of the non-zero elements.

8.3.5 DWT-Based Compression Implementation

The DWT [44] allows to hierarchically decompose an input signal into a series of successively lower frequency approximations and their associated detail signals. As for Compressive Sensing, DWT-based compression is still based on the sparsity principle, since most of the wavelet coefficients of natural smooth signals have a small amplitude, so that the signal is approximately sparse in the wavelet domain. Indeed, smallest wavelet coefficients can be neglected without much signal quality loss, as shown in Fig. 8.2 where only 10% of the original coefficients are kept. Thanks to this property, the wavelet transform is widely used for the compression of signals and images.

Differently from the CS framework, where the transmitting encoder simply computes random projections, and an optimization problem is solved at the receiving decoder using knowledge of the sparsifying representation, a DWT-based compression scheme needs to compute the transform in the sensor, then exploit sparsity by transmitting a subset of the computed coefficients.

The DWT linear transform operates on a signal vector $\mathbf{x}$ of length N, which is typically a power of two, and allows to separate data into different frequency components. The DWT is computed with a cascade of filters followed by a factor 2 subsampling. The resulting low-pass coefficients represent a rough subsampled approximation of the original signal, while the high-pass coefficients represent detail information. Due to subsampling, the number

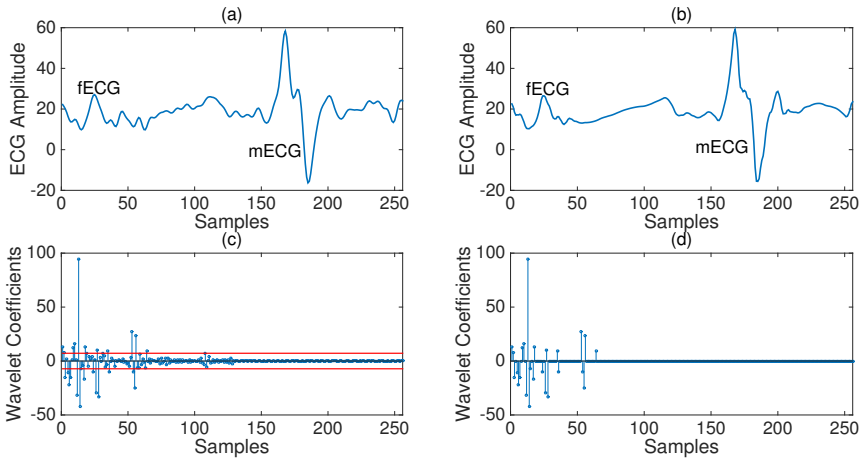

Figure 8.2: (a) Original ECG signal sampled at 1 KHz and (c) corresponding wavelet coefficients, the lines correspond to the threshold level to select the 10% largest coefficients. Plots (b) and (d) show the reconstructed signal and the 10% largest coefficients used for reconstruction.

of wavelet coefficients is N for an N-length input signal vector (when using appropriate extension of the signals at the borders.)

In this work, we use the orthogonal Daubechies wavelet (DB4) with 8-tap filters, which provides a sparse representation for piecewise-linear signals and thus is suitable for ECG signals, leading to a relatively sparse representation with most of the coefficients close to zero (see Fig. 8.2). In the implementation, the block size and the number of layers of the wavelet transform needs to be appropriately selected based on the desired frequency resolution and the coding delay. Our choice considers a 4-level wavelet decomposition and 256 sample frames. Before applying the DWT, the ECG signal block is preprocessed to remove the mean value.

In order to have accurate power consumption estimates in a concrete scenario [90], we implemented the algorithms on a general-purpose MSP430 [69] microcontroller that does not include a floating-point unit. Therefore, all computations are performed in fixed-point, which is suitable for real-time embedded applications. Filter coefficients are represented with 9 bits, and the original abdominal ECG samples in $\mathbf{x}$ and the corresponding wavelet coefficients $\alpha = [d1, d1, d3, d4, a4]$ are both represented with $B = 16$ bits.

As mentioned, the main advantage of using the DWT representation is that the signal can be compactly represented by few large coefficients. This allows us to set to zero the small coefficients, thus reducing the number of coefficients to encode. Signal compression is performed by keeping the largest DWT coefficients α, as suggested in [15]. In this work, the number of retained coefficients is selected based on the desired compression ratio CR. We only encode N_{coeff} non-zero entries and their corresponding positions in the N-length coefficient vector. In addition, we need to send the mean of the ECG vector, encoded with B_{mean}=16 bits. Since we work with a signal window of length $N = 256$, we need B_{idx}=8 bits to represent the index of each retained coefficient. The number of coefficients

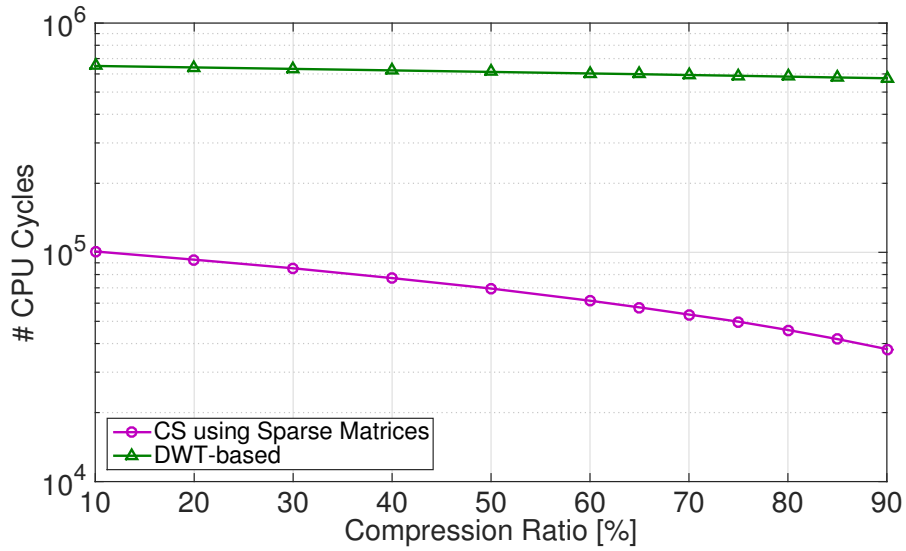

Figure 8.3: Number of CPU cycles required to compress a signal block (N=256 samples) using CS with sparse random matrices or DWT-based compression as a function of the compression factor.

is computed as follows

$$N_{coeff} = \lceil \frac{NB(\frac{100-\mathrm{CR}}{100}) - B_{mean}}{B + B_{idx}} \rceil. \tag{8.3.5}$$

In particular, given a compression ratio and the relative number of coefficients to retain N_{coeff}, only the first largest N_{coeff} coefficients in absolute value are kept. The largest coefficients are selected using the ordering algorithm *merge sort*, which has a computational complexity $\mathrm{O}(N \log_2 N)$.

In summary, two vectors are used to code the DWT coefficients, one contains only the nonzero wavelet coefficients, and the other one contains the corresponding positions. Both need to be sent to the receiver in order to recover the signal.

8.4 Experimental Results

In Fig. 8.3, the number of cycles required by the microcontroller to perform the compression of a single $N = 256$ signal block for the two coding schemes is shown. The computational workload required for the wavelet scheme, which includes filtering and multiplications, is significantly higher than that required by the CS scheme, requiring additions only. As an example, to compress at CR=50% 1 s of one channel ECG data, the wavelet based compression code executes in about 306 ms, whereas the CS code only requires about 35 ms. Note that DWT does not allow real-time processing for 4-channel recordings sampled at 1 kHz, and would require buffering at the sensor or using a smaller sampling frequency.

The required total energy, including the one for transmission, is shown in Fig. 8.4 for the two algorithms at various compression ratios. In particular, results refer to the

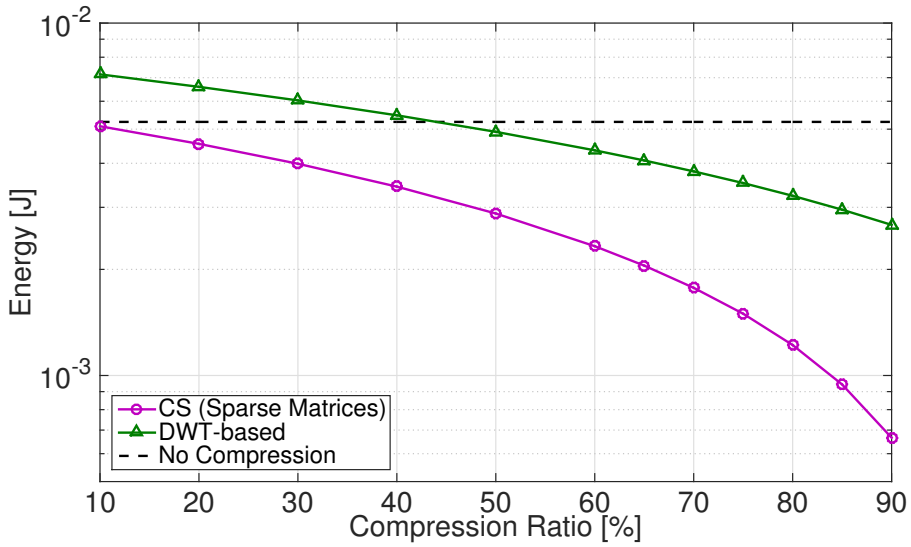

Figure 8.4: Energy required to compress and transmit one $N = 256$ signal block for each channel in a 4-channel recording (4 blocks in total), using the CS or DWT-based schemes.

energy required for the compression and transmission of four blocks of length N=256, corresponding to 250 ms in 4-channel recording.

In terms of energy, we can see that compressive sensing is more energy-efficient than transmitting the uncompressed original signal (for one signal block, $E_{comp} = 0, E_{tx} = NBE_{bit}$) for compression ratios CR> 10%, while the DWT-based compression becomes favorable for CR> 45%. Fig. 8.4 also shows that increasing the compression ratio, the gap between the two compression techniques increases, leading to higher energy-saving for the CS-based method.

It is clear that CS allows for the reduction of the encoding complexity in the sensor node, allowing a reduction of the overall energy consumption. However, we need to assess how the different compression schemes impact on the reconstruction quality of the signal at the receiver. In particular, for the CS-based scheme, several methods can be used for reconstruction. In Fig. 8.5 and Fig. 8.6 we report the average PRD value, for reconstruction quality assessment, and the average Sensitivity and Positive Predictivity, to verify the accuracy of the detection algorithms resulting from the CS and DWT-based schemes.

As it can be seen from Fig. 8.5, the DWT-based scheme allows to have good reconstruction quality up to CR=80%. In previous works, e.g., [90], the performance of DWT-based compression was compared to a CS scheme where the optimization reconstruction problem was solved at the receiver using a wavelet basis as the sparsifying matrix $\mathbf{\Psi}$, showing that the quality of the recovered signal was in favour of the DWT-based scheme. Fig. 8.5 confirms that CS reconstruction using the wavelet basis at the receiver (dotted line) has lower performance, i.e., higher PRD values, than DWT-based compression (dashed line). Using the specifically designed dictionary [42], however, the performance of the DWT-based scheme and CS scheme (continuous line in Fig. 8.5) become similar in terms of average PRD value. Indeed, both algorithms allow compression up to CR= 80% maintaining a good reconstruction quality. The CS-based approach, however, requires significantly lower

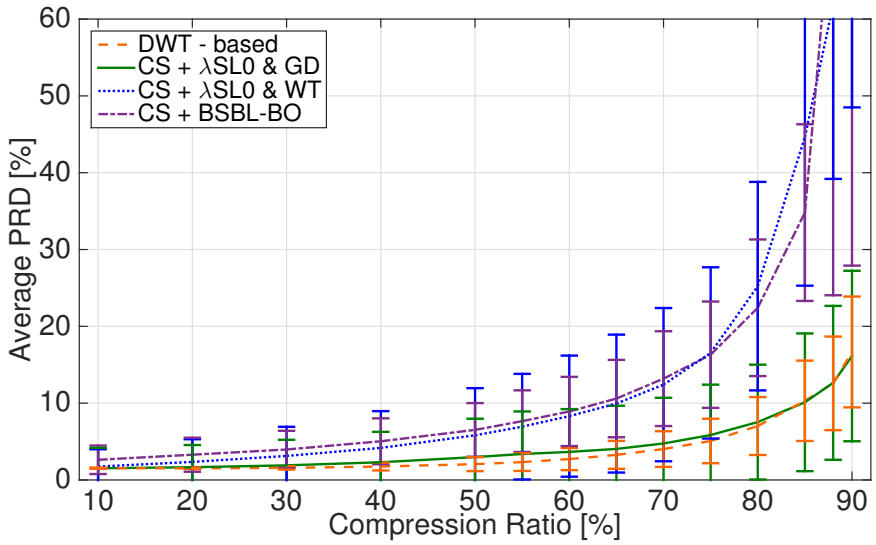

Figure 8.5: Average PRD value for different compression/ reconstruction schemes. Error bars indicate standard deviation.

energy as we will see below. Fig. 8.5 also shows PRD values for the BSBL reconstruction technique in the CS scenario (dashed-dotted line). The performance is similar to the one obtained using the wavelet sparsifying basis at the receiver. However, we will confirm below that BSBL better preserves signal characteristics and allows for improved detection performance after signal reconstruction.

As a matter of fact, PRD values can provide in some cases a rough estimate of the signal quality, especially when the interest is preserving clinically relevant aspects. As an example, we report in Fig. 8.7 one ECG trace and the corresponding PRD values for each signal block. It can be noticed that blocks with the largest PRD values do not contain relevant ECG information. Clearly, a noisy block cannot be sparse in the wavelet basis or Gaussian Dictionary representations, thus worsening the performance of CS-based schemes.

To compare how well the different techniques can preserve relevant signal characteristics, Fig. 8.6(a) and Fig. 8.6(b) show the average Sensitivity and Positive predicitivity measures obtained from the application of the detection algorithm on the reconstructed signals. Both indicators show that the different techniques have approximately the same performance for compression ratios less than 50%, with S and P+ values very similar to those obtained on the uncompressed signals, i.e., S=98.9% and P+=97.7%. Note that the BSBL reconstruction technique (dashed-dotted line) outperforms CS reconstruction with wavelet basis at the receiver (dotted line), confirming that BSBL can preserve dependency among ECG channels, which is exploited by ICA in the detection algorithm. However, performance achieved adopting the BSBL method is still lower than that obtained with the DWT-based scheme (dashed line) and CS with Gaussian dictionary reconstruction at the receiver (continuous line). For these techniques, the S and P+ values are approximately constant up to CR=75%, with values similar to those obtained with uncompressed signals. In summary, the DWT-based scheme and the CS scheme with Gaussian dictionary at the

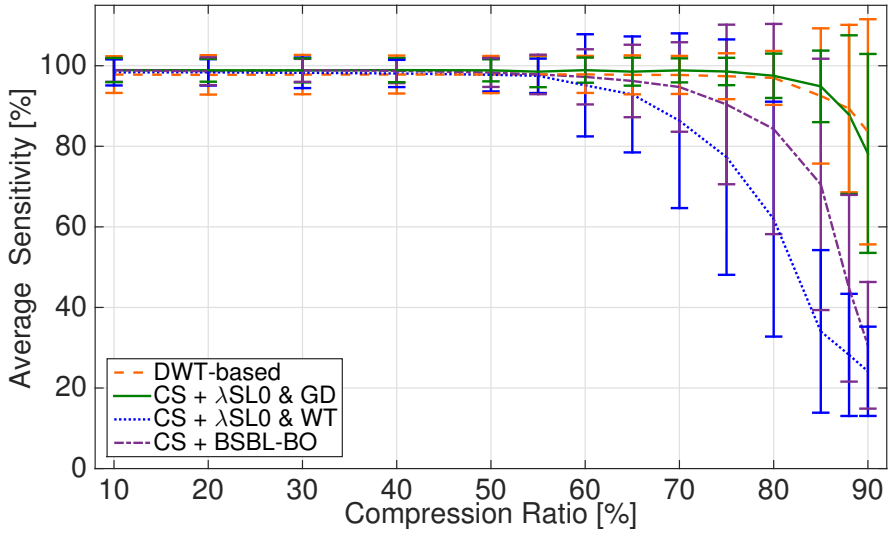

(a)

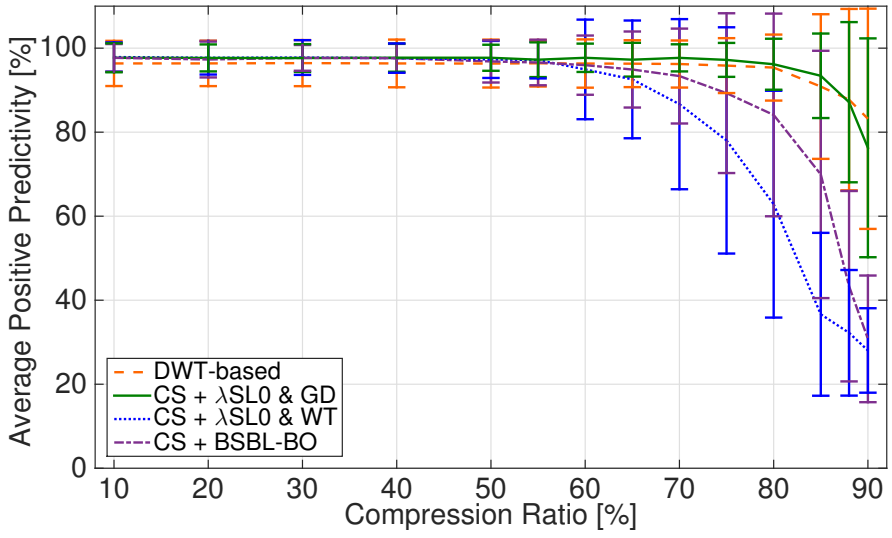

(b)

Figure 8.6: (a) Average Sensitivity value and (b) average Positive Predicitivity value for different compression/reconstruction schemes. Error bars indicate standard deviation.

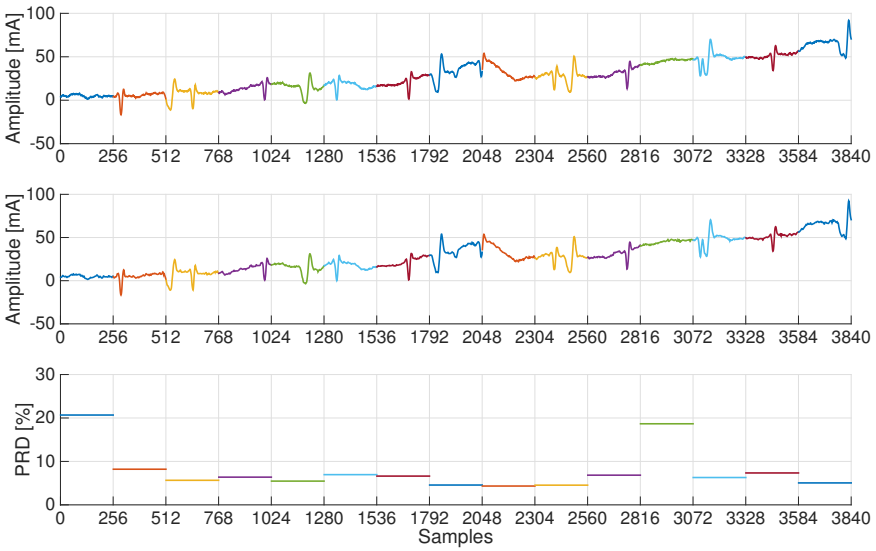

Figure 8.7: (Top) Original signal a28 of the Challenge dataset A sampled at 1 KHz and (Middle) reconstructed record after CS compression at CR=70% using the Gaussian dictionary for sparsification. (Bottom) Corresponding PRD value for each window (different colors represent different windows).

receiver appear to have comparable performance in terms of PRD, S and P+ metrics. However, the CS scheme allows significant energy-saving in the encoding sensor.

Fig. 8.8 shows the energy required by the two schemes as a function of the PRD value. In particular, we report the energy values necessary to compress and transmit the entire 4-channel, 1 minute long, signals. The energy required to achieve a desired fetal beat detection Sensitivity and Positive Predictivity is reported in Fig. 8.9 (a) and (b), respectively. It is apparent from the figures that the considered CS scheme allows significant energy saving for all the considered figures of merit. For instance, for PRD=9%, less than 0.3 J are required by the CS scheme, while the DWT-based scheme requires about 0.7 J. A sensitivity value S=95% requires about 0.2 J and 0.7 J for DWT-based and CS schemes, respectively. Similar values are required to have P+=95% for the two schemes.

8.5 Conclusions

In this chapter, we evaluated energy consumption of two acquisition schemes for multi-channel abdominal fECG signals, one based on DWT and the other based on the emerging CS paradigm. Experimental results with an actual implementation on a commercial device, show that compressive sensing allows to significantly reduce energy consumption in the sensor node. Moreover, it is advantageous with respect to transmission of the uncompressed signals for compression ratios higher than 10% (the DWT-based scheme becomes preferable only for CR>45%). We compared the quality of the recovered signal in terms of PRD values, and also, more importantly, by testing the performance of a state-of-the-art

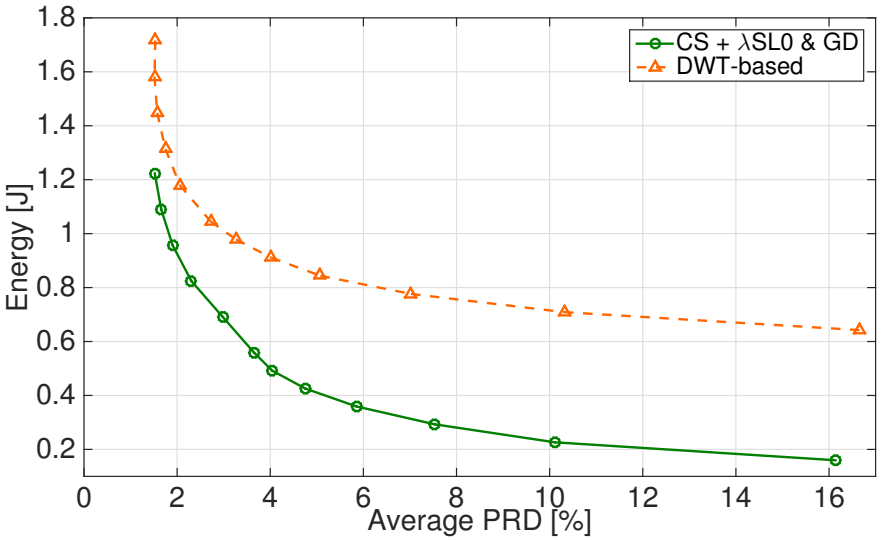

Figure 8.8: Energy required by the DWT-based and CS schemes to achieve a desired PRD value. Energy values refer to a 4-channel, 1 minute long, signal.

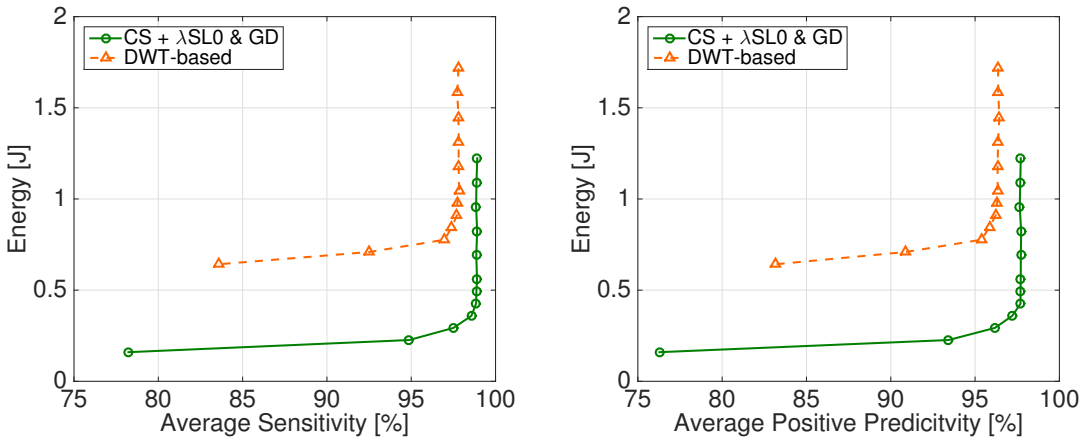

(a) (b)

Figure 8.9: Energy required by the DWT-based and CS schemes to achieve a desired (a) Average Sensitivity value and (b) average Positive Predicitivity value. Energy values refer to a 4-channel, 1 minute long, signal.

fetal beat detector on the recovered traces. We found that compressive sensing, using a suitable dictionary for signal sparsification at the receiver, can achieve the same results of the DWT-based scheme, but with significantly lower energy consumption. In particular, we showed that the detection performance obtained with the CS scheme is comparable to that obtained on original signals for compression ratios up to about 75%. This study confirms that compressive sensing, by moving complexity to the receiving end, where the reconstruction optimization algorithm, together with other processing tasks, is run, is indeed suitable for fECG monitoring in low power applications, and allows the use of sensors with limited complexity.

9

Conclusions and Future Works

The aim of this dissertation was to develop tools to improve current compression and analysis methods for long-term monitoring of cardiac signals acquired by means of wearable sensors. To this end, taking advantage of sparse signal approximation and compressive sensing theory, we have designed and analyzed new frameworks.

The first contribution provided by this work aims to improve the quality of ECG signals reconstructed from compressed sensed measurements. Under the sparse approximation model, signals are assumed to be well represented as sparse linear combinations of atoms from a predetermined dictionary and we expect that designing an appropriate dictionary will lead to better compression/reconstruction performance. Hence, we employed Gaussian-like functions, which have been successfully applied for the approximation of the characteristic waves of the ECG, to design a new dictionary for its sparsification. As a result, our proposed dictionary guarantees the preservation of the signal structure while approximating each ECG cycle using a linear combination of just few dictionary atoms. Since the proposed dictionary is based on a signal model, it turns out to be universal and, thus, it can be used independently of the particular patient ECG trace. Numerical experiments, conducted on actual ECG recordings from a public database, have demonstrated the good performance of the proposed dictionary in terms of reconstruction quality. Moreover, we have shown that, with respect to the use of sparsifying transforms such as wavelets, it is possible to take samples at higher compression ratios while maintaining the same reconstruction quality.

Very often, for the analysis of the ECG signal, we are only interested in extracting certain information, such as the heart rate, while the reconstruction of the full signal from compressive measurements is not of particular interest in an early stage. Moreover, the reconstruction process can be computationally too expensive to be implemented on some devices or to be done in real-time. Thus, we proposed a method to extract heart rate information directly from the compressed measurements. In particular, we have focused on the beat detection problem that can be recast as a compressed signal processing problem. To this end, a method that use a compressive sensing matched filtering approach combined with further processing procedures, have been proposed for the detection of the QRS complexes. Experimental results on real ECG signals have demonstrated that the proposed method achieves an accuracy comparable with that obtained on original signals up to compression ratios of about 60-70%.

Although there are generally accepted standards for the evaluation metrics of signal reconstruction quality, such as the PRD, there exists no clear interpretation on the clinical relevance of these metrics. Therefore, to understand the impact of compressive sensing on clinical analysis of ECG signals, we evaluated its effect at different compression ratios

for atrial fibrillation detection. We have shown that the performance is dependent on the compression ratio as well as on the reconstruction procedure selected. In particular, when the Gaussian dictionary proposed in this work is used as the sparsifying transform, it is possible to achieve accurate atrial fibrillation detection in signals that have been compressed up to 80%. In addition, we have shown that the proposed algorithm for beat detection on compressed measurements allows an accurate atrial fibrillation detection up to CR=50% when working in the compressed domain with a very low complexity procedure. The proposed compressive detector results in a trade-off between effective atrial fibrillation detection and low computational cost. The described work is definitely a first important step toward the analysis of compressed signals for clinical purpose.

The second part of the thesis describes a new solution for the compression and real-time processing of non-invasive fetal ECG signals. We have focused in particular on the compressive sensing approach, which is a difficult task for this class of signals due to the particular complexity of the abdominal fECG trace. The use of orthogonal bases generally fails to recover the signal at high compression ratios. Thus, we have introduced a novel dictionary for fetal ECG sparsification, similar to the one proposed for adult ECG, which can also provide a separation of the different components, such as the maternal and the fetal one. In order to provide a real-time analysis within the reconstruction process, we have proposed a scheme based on ICA performed in the compressed domain, which reduces the dimension of the problem and speeds up the ICA algorithm. Mother's and fetal beat detection is thus based on the atoms activated during the reconstruction of the independent components. Results suggest that the proposed method is suitable for real-time reconstruction and joint detection/classification of the maternal and fetal beats.

In case of noisy original signals and possibly ill-conditioned sensing/reconstruction procedures we have introduced a modification of the low-complexity compressive sensing reconstruction Smoothed$l0$ (SL0) algorithm, in order to guarantee the robustness of the reconstruction/detection framework. As show by experimental results the modification significantly improves the reconstruction quality, both for synthetic and real-world fetal ECG signals.

Finally, we evaluated the energy consumption of two acquisition schemes for multichannel abdominal fECG signals, one based on DWT and the other based on CS paradigm. We have compared the quality of the recovered signal in terms of reconstruction quality, and also, more importantly, by testing the performance of a state-of-the-art fetal beat detector on the recovered traces. We have found that compressive sensing, using the proposed dictionary for signal sparsification at the receiver, can achieve the same results of the DWT-based scheme, but with significantly lower energy consumption. Thus, we have quantified the potential of the CS paradigm for low-complexity and energy efficient fECG sensing and compression for storage or transmission, considering also detection accuracy aspects.

In summary, in this thesis we have addressed several important problems for the continuous monitoring of physiological signals. We believe that the research efforts in this thesis give some new insights to develop next generation personal health care monitoring systems able to provide concrete improvements in detecting and preventing cardiac diseases. However, there are many roads to follow for future works on the topics of this dissertation. Indeed, the use of sparse signal processing as well as compressive sensing for physiological signals is a relatively new research field and, despite this thesis brings several contributions, there is still much room to provide new reliable methods.

In particular, in Chapter 4 we proposed a method for the analysis of ECG signals in the compressed sensed domain, which is restricted to the simple case of beat detection and atrial fibrillation analysis. However, the application of this method for the detection of other characteristic waves, such as the P or T wave, as well as of general patterns that identify abnormality in the ECG trace, is definitely a very interesting problem. Investigating the improvement on the battery lifetime by implementing the proposed methods on an actual wearable device for event detection is certainly another interesting study.

Concerning the analysis of the effect of compressive sensing on atrial fibrillation detection of Chapter 5, the results suggest that algorithms for the analysis of ECG based on compressed heart rate estimation should be modified by taking in account the fact that the analysis is done in the compressed domain. Indeed, in this work, we employed an atrial fibrillation detection designed and optimized for uncompressed signals. It would be interesting to design the compressed detector and analysis algorithm together in order to achieve optimal performance.

A

Supplementary Results

A.1 Results for AF Detection Using Entropy-based Method

In Chapter 5 we evaluated the accuracy of an SVM-based Atrial Fibrillation (AF) detector applied on ECG signals that have been perviously compressed by compressive sensing. In this section we summarize the results for an AF detector based on the normalized fuzzy entropy measure, NFEn, introduced by Liu *et al.* [82]. This method is based on entropy, which refers to the degree of regularity or irregularity of a time series, and it is estimated by counting how many "template" patterns repeat. Repeated patterns imply increased regularity in the time series and lead to low entropy values [106, 110]. Typical entropy measures are approximate entropy (ApEn) [106] and sample entropy (SampEn) [110].

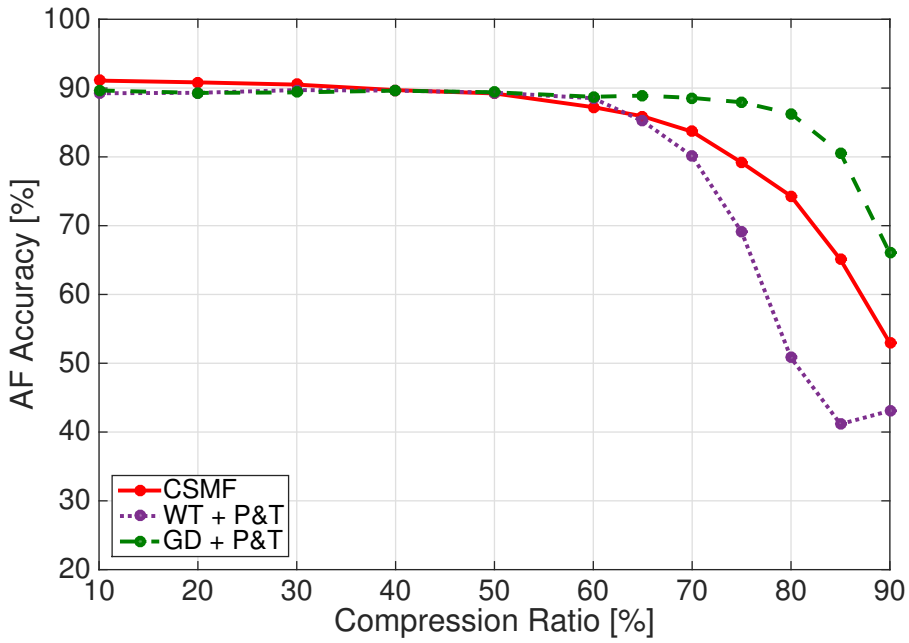

Figure A.1: Output AF Accuracy versus CR.

Both SampEn and CoSampEn have the problem of weak statistical stability due to the rigid determination rule (0-1 determination) [81]. Liu *et al.* therefore replaced the 0-1 determination with a fuzzy rule, proposing a fuzzy measure entropy (FuzzyMEn) [81].

Recently, Liu *et al.* proposed a new entropy-based AF detector, based on the normalized fuzzy entropy (NFEn), which combines the advantages of both FuzzyMEn and CoSampEn methods, and verified its improved performance [82].

To assess the accuracy of the NFEn method, we employed the same evaluation scheme described in Sec. 5.2, by using detection windows of 30-beat RR segments. Tab. A.1 provides the results in terms of the considered evaluation metrics, which have been described in Sec. 5.2.5, for the NFEn-based AF detector on the MIT-BIH Atrial Fibrillation Database (see Sec. 5.2.2). We consider a variety of setting scenarios, as described in Sec. 5.2. An overall picture of the AF detection performance, for varying CR, is given by the AF accuracy measure, reported in Fig. A.1.

All the considered scenarios, i.e., reconstruction with the Wavelet basis (WT) or with the Gaussian dictionary (GD) followed by Pan-Tompkins (P&T) detection, or direct detection on compressed measurements (CSMF), show similar trends for the AF metrics. In particular, at low CR levels the AF detection results are comparable with those obtained on raw uncompressed ECG signals. However, for high values of CR, the AF detection accuracy decreases, as can be seen in Fig. A.1. The results of this study reveal that, when using the NFEn AF classifier, the CSMF method outperforms the detection performed after signal reconstruction using the Wavelet basis. Indeed, classification in the compressed domain reaches about 89% of correctly classified signals for a compression ratio of about 50%. As the compression ratio gets closer to 90%, it decreases reaching 52% of correctly classified AF episodes. The classification after reconstruction using the Wavelet basis achieves similar results up to CR= 50%, however, its performance rapidly decreases at high CR reaching 43% accuracy at CR=90%.

Similarly to the SVM-based AF detector, the best performance is achieved by classification after signal reconstruction with the Gaussian dictionary, which allows to reach about 90% and 66% of correctly classified AF episodes at CR=65% and CR=90%, respectively.

Table A.1: Results of the performance metrics on uncompressed and compressed ECG signals using different reconstruction/detection methods for the NFEn AF detector.

Metric		CS Reconstruction	QRS Detector	Compression Ratio												
				0%	10%	20%	30%	40%	50%	60%	65%	70%	75%	80%	85%	90%
Se [%]	Reference QRS	-	-	96.69	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan&Tomp.	98.21	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	96.87	97.03	97.79	97.85	98.40	99.15	99.39	99.70	99.83	98.77	99.93	99.93
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	97.27	97.20	97.16	97.69	97.84	97.96	98.43	99.02	99.06	99.38	99.59	99.87
	CS ECG signals	No	CSMF	-	96.73	96.84	96.71	96.98	97.30	97.63	97.88	98.22	98.62	99.10	99.53	99.90
Sp [%]	Reference QRS	-	-	88.37	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan&Tomp.	84.64	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	81.06	81.02	81.11	80.91	79.87	77.66	71.58	62.96	46.31	21.74	7.07	2.50
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	81.45	80.83	81.02	81.04	80.46	79.12	78.99	77.92	76.74	73.52	63.77	42.57
	CS ECG signals	No	CSMF	-	86.80	86.28	85.83	84.26	83.21	79.51	76.56	73.86	66.22	56.32	42.96	18.28
Acc [%]	Reference QRS	-	-	92.12	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan &Tomp.	90.19	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	89.25	89.31	89.71	89.63	89.34	88.52	85.20	80.11	69.10	50.79	41.16	43.09
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	89.66	89.30	89.37	89.63	89.39	88.74	88.91	88.57	87.94	86.29	80.57	65.99
	CS ECG signals	No	CSMF	-	91.09	90.83	90.50	89.71	89.22	87.23	85.57	84.13	79.81	74.24	66.37	52.09
PPV [%]	Reference QRS	-	-	87.20	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan&Tomp.	81.55	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	84.60	84.61	84.64	84.47	83.63	81.94	77.05	70.21	57.96	43.32	38.41	42.26
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	84.96	84.46	84.58	84.58	84.11	83.02	82.98	82.05	81.10	78.53	70.82	54.59
	CS ECG signals	No	CSMF	-	84.78	84.24	83.73	82.17	81.20	77.94	75.35	73.26	67.84	62.06	55.19	46.36
NPV [%]	Reference QRS	-	-	97.02	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan&Tomp.	98.56	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	96.02	96.21	97.18	97.25	97.95	98.89	99.19	99.59	99.73	96.70	99.46	97.97
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	96.52	96.42	96.38	97.06	97.24	97.38	97.98	98.73	98.78	99.18	99.43	99.78
	CS ECG signals	No	CSMF	-	97.21	97.30	97.19	97.39	97.63	97.84	98.01	98.27	98.51	98.86	99.24	99.61
J [%]	Reference QRS	-	-	85.06	-	-	-	-	-	-	-	-	-	-	-	-
	Raw signals	-	Pan&Tomp.	82.85	-	-	-	-	-	-	-	-	-	-	-	-
	CS ECG signals	SL0 & Wavelet	Pan&Tomp.	-	77.93	78.05	78.90	78.76	78.27	76.80	70.97	62.66	46.14	20.52	7.01	2.42
	CS ECG signals	SL0 & GD	Pan&Tomp.	-	80.03	79.34	79.43	79.45	79.08	77.93	77.74	76.69	75.57	72.77	63.31	42.44
	CS ECG signals	No	CSMF	-	83.53	83.12	82.53	81.25	80.50	77.14	74.44	72.08	64.84	55.42	42.50	18.18

A.2 Results for Fetal Beat Detection

Table A.2 reports a full evaluation of the joint reconstruction/detection method for fECG proposed in Chapter 6 for Challenge set A signals. All the signals have been compressed using the same sensing matrix at CR=75%. In Table A.2 we report the number of correctly found beats (TP), false negative (FN) and false positive (FP) values, as well as the Sensitivity Se, Positive Predictivity P+ and F measures. For each signal, HRmeas and RRmeas values are also reported (see Sec.6.7.3).

Table A.3 reports the average and standard deviation for sensitivity Se, positive predictive P+, *HRmeas* and *RRmeas* for the entire Challenge dataset A. Results are average values after using 50 different sensing matrices for each signal (CR=75%).

Table A.2: Results for Challenge dataset A. The number of TP, FP, FN are calculated for a window size of 100 ms (± 50 ms around the reference R peak).

Num	TP	FP	FN	S (%)	P+ (%)	F (%)	HR (bpm)	*HRmeas* (bpm^2)	*RRmeas* [ms]	PRD %
1	130	11	14	90	92	91	143	18.10	24.32	2.45
2	32	90	127	20	26	22.6	122	1464.17	46.93	6.62
3	120	10	7	94	92	93	130	11.22	12.28	10.44
4	126	4	2	98	97	97.5	130	16.46	16.81	9.77
5	119	12	9	93	91	92	131	29.26	12.07	11.47
6	97	48	62	61	67	63.9	145	289.25	36.46	3.43
7	76	57	53	59	57	58	133	37.93	31.22	4.72
8	122	5	5	96	96	96	127	3.64	12.36	10.52
9	28	75	101	22	27	24.2	103	990.59	58.89	3.70
10	144	22	30	83	87	85	166	90.93	26.20	3.31
11	27	77	112	19	26	22	104	1767.46	61.96	4.81
12	129	8	8	94	94	94	137	7.33	10.97	13.07
13	113	16	12	90	88	89	129	14.36	19.23	5.63
14	92	38	30	75	71	73	130	72.12	25.72	12.86
15	127	6	6	95	95	95	133	27.17	18.98	11.61
16	23	81	106	18	22	20	104	809.00	48.28	6.74
17	114	17	17	87	87	87	131	8.09	20.27	12.48
18	23	85	126	15	21	17.5	108	2091.73	67.85	5.43
19	115	14	11	91	89	90	129	21.01	20.77	6.71
20	96	37	34	74	72	73	133	37.01	28.23	7.72
21	95	46	49	66	67	66.5	141	167.64	29.64	4.34
22	121	7	4	97	95	96	128	12.35	16.55	9.93
23	86	45	39	69	66	67.5	131	120.76	29.39	5.90
24	114	9	8	93	93	93	123	14.55	14.13	5.94
25	112	16	12	90	88	89	128	28.01	16.26	5.81
26	98	40	39	72	71	71.5	138	72.62	24.09	3.09
27	38	67	96	28	36	31.5	105	954.73	43.08	1.65

Num	TP	FP	FN	S (%)	P+ (%)	F (%)	HR (bpm)	*HRmeas* (bpm^2)	*RRmeas* [ms]	PRD %
28	160	3	6	96	98	97	163	16.25	27.18	5.15
29	81	52	45	64	61	62.5	133	30.13	26.85	5.04
30	118	27	21	85	81	83	145	21.97	15.30	5.36
31	130	4	6	96	97	96.5	134	26.04	13.75	12.41
32	151	1	0	100	99	99.5	152	0.01	2.36	2.70
33	126	14	13	91	90	90.5	140	94.24	90.97	4.58
34	86	51	44	66	63	64.5	137	22.98	22.62	13.91
35	152	6	10	94	96	95	158	102.08	21.81	5.89
36	164	4	3	98	98	98	168	0.0001	5.65	4.71
37	121	17	20	86	88	87	138	53.16	21.92	9.09
39	124	10	10	93	93	93	134	17.18	13.18	9.49
40	127	17	20	86	88	87	144	115.09	23.12	22.80
41	100	38	35	74	72	73	138	25.95	16.62	8.01
42	140	9	10	93	94	93.5	149	14.56	13.84	23.33
43	152	4	5	97	97	97	156	17.64	12.44	2.82
44	158	3	4	98	98	98	161	7.31	15.19	5.86
45	71	69	68	51	51	51	140	84.18	29.29	8.91
46	106	28	24	82	79	80.5	134	23.80	19.37	6.90
47	79	57	64	55	58	56.5	136	223.76	35.46	3.04
48	64	76	68	48	46	47	140	56.21	28.68	2.92
49	144	4	3	98	97	97.5	148	0.05	5.29	4.31
50	130	11	11	92	92	92	141	46.74	23.18	2.06
51	42	99	94	31	30	30.5	141	163.44	29.66	4.94
53	138	14	14	91	91	91	152	15.70	14.23	26.13
55	90	45	52	63	67	65	135	104.64	29.24	8.68
56	80	61	52	61	57	58.9	141	80.50	30.79	4.72
57	138	6	9	94	96	95	144	39.04	19.20	3.53
58	99	42	37	73	70	71.5	141	20.25	16.75	10.80
59	150	2	2	99	99	99	152	7.61	8.26	3.71
60	104	44	43	71	70	70.5	148	13.36	23.58	1.38
61	130	13	9	94	91	92.5	143	56.39	17.74	21.94
62	138	6	5	97	96	96.5	144	10.95	12.60	9.46
63	108	32	29	79	77	78	140	41.36	22.11	3.23
64	121	18	14	90	87	88.5	139	26.02	19.21	3.27
65	131	11	12	92	92	92	142	15.51	17.64	19.09
66	103	32	26	80	76	78	135	33.41	20.00	12.31
67	129	19	24	84	87	85.5	148	45.73	24.85	7.53
68	132	7	6	96	95	95.5	139	7.28	7.25	1.47
69	144	3	4	97	98	97.5	147	5.32	9.37	6.54
70	116	29	24	83	80	81.5	145	15.60	19.32	6.06
72	165	2	1	99	99	99	167	10.97	11.12	5.51
73	94	49	44	68	66	67	143	25.13	23.47	6.36
75	107	28	26	80	79	79.5	135	37.73	24.90	2.69

Table A.3: Average and standard deviation for sensitivity Se, positive predictive P+, *HRmeas* and *RRmeas* for the entire dataset A, using 50 different sensing matrices.

Num	S (%)	P+ (%)	*HRmeas* (bpm^2)	*RRmeas* $[ms]$
1	91.4 ± 1.8	91.9 ± 1.5	31.637 ± 14.025	15.15 ± 1.66
2	18.7 ± 0.7	23.7 ± 0.7	1387.799 ± 42.077	38.24 ± 0.65
3	77.4 ± 5.1	74.7 ± 4.4	53.773 ± 19.344	21.82 ± 2.27
4	98.4 ± 1.4	96.7 ± 1.7	8.914 ± 4.281	14.66 ± 0.39
5	99.0 ± 0.9	98.2 ± 0.9	0.013 ± 0.015	3.49 ± 1.74
6	62.9 ± 7.6	68.0 ± 5.4	282.952 ± 163.946	34.29 ± 3.30
7	46.8 ± 2.2	47.3 ± 2.7	78.955 ± 33.894	36.78 ± 3.16
8	98.2 ± 0.5	96.2 ± 1.3	10.938 ± 6.361	7.87 ± 2.13
9	24.3 ± 0.5	28.0 ± 0.7	419.425 ± 83.098	40.53 ± 3.84
10	64.4 ± 22.8	71.8 ± 19.1	587.196 ± 666.376	38.00 ± 13.42
11	22.8 ± 1.7	30.0 ± 2.7	1245.240 ± 134.335	55.24 ± 6.72
12	87.6 ± 4.1	85.9 ± 4.2	18.518 ± 16.538	13.15 ± 2.03
13	89.1 ± 0.5	86.3 ± 0.4	21.079 ± 8.915	15.32 ± 2.42
14	62.0 ± 8.4	55.2 ± 9.9	295.594 ± 201.337	24.97 ± 5.90
15	96.7 ± 1.9	95.1 ± 1.9	6.590 ± 3.864	12.81 ± 1.23
16	21.7 ± 3.4	26.0 ± 3.6	511.232 ± 119.087	46.22 ± 7.98
17	91.3 ± 1.2	90.2 ± 1.3	5.796 ± 1.883	12.40 ± 1.38
18	15.0 ± 1.0	19.6 ± 1.7	1392.199 ± 269.465	46.79 ± 8.60
19	92.6 ± 2.0	92.1 ± 2.1	2.381 ± 4.102	18.44 ± 1.00
20	86.4 ± 1.2	85.8 ± 2.8	59.073 ± 36.949	23.95 ± 3.85
21	76.6 ± 5.1	75.8 ± 6.4	49.677 ± 15.553	24.94 ± 2.65
22	96.8 ± 1.4	96.3 ± 1.2	0.163 ± 0.231	5.66 ± 2.22
23	88.8 ± 0.8	85.6 ± 1.3	50.513 ± 52.352	17.15 ± 1.45
24	90.4 ± 6.8	87.7 ± 8.9	43.769 ± 18.238	16.23 ± 4.67
25	93.3 ± 4.1	91.3 ± 4.4	12.143 ± 9.957	10.95 ± 3.07
26	70.3 ± 3.7	68.3 ± 3.4	40.813 ± 28.200	23.32 ± 3.09
27	67.2 ± 5.4	66.0 ± 6.8	71.491 ± 16.335	25.86 ± 1.74
28	95.8 ± 2.1	95.4 ± 1.8	17.019 ± 6.467	24.84 ± 0.50
29	64.3 ± 4.2	56.8 ± 0.7	181.981 ± 142.427	25.48 ± 2.18
30	78.4 ± 8.7	78.5 ± 5.8	53.937 ± 17.130	27.72 ± 9.77
31	91.2 ± 3.4	91.4 ± 3.8	33.423 ± 15.463	16.77 ± 3.41
32	99.6 ± 0.4	98.9 ± 0.4	0.014 ± 0.007	3.13 ± 0.67
33	94.0 ± 3.0	91.6 ± 2.5	190.177 ± 57.670	86.23 ± 6.35
34	89.2 ± 4.8	85.9 ± 5.1	21.612 ± 28.115	18.22 ± 4.35
35	96.1 ± 1.3	95.7 ± 1.1	27.375 ± 13.849	14.11 ± 3.80
36	96.0 ± 3.3	96.2 ± 2.5	10.749 ± 12.868	9.80 ± 4.71

Num	S (%)	P+ (%)	*HRmeas* (bpm^2)	*RRmeas* [ms]
37	82.7 ± 5.7	83.1 ± 5.2	44.015 ± 13.336	21.46 ± 5.23
39	80.8 ± 3.5	78.9 ± 2.2	58.627 ± 18.776	19.34 ± 5.35
40	85.0 ± 6.5	85.2 ± 5.4	55.283 ± 55.524	20.50 ± 5.35
41	82.5 ± 5.8	81.0 ± 7.2	20.799 ± 18.214	15.71 ± 1.19
42	93.8 ± 3.3	93.4 ± 3.4	12.262 ± 4.347	11.05 ± 1.79
43	98.5 ± 1.0	98.1 ± 1.1	2.647 ± 4.159	5.80 ± 3.08
44	98.8 ± 0.6	98.6 ± 0.3	4.866 ± 8.427	7.68 ± 6.99
45	47.0 ± 9.2	44.9 ± 8.9	71.105 ± 35.402	24.74 ± 2.47
47	68.1 ± 1.5	70.1 ± 1.7	216.643 ± 70.009	37.20 ± 4.70
48	50.8 ± 6.8	48.1 ± 5.1	99.191 ± 53.109	31.25 ± 7.49
49	98.4 ± 1.0	97.5 ± 1.4	2.450 ± 4.213	5.17 ± 2.73
50	86.3 ± 4.2	85.1 ± 5.0	21.329 ± 22.239	14.34 ± 0.97
51	44.6 ± 7.9	41.0 ± 8.9	306.998 ± 167.849	32.27 ± 5.48
53	82.5 ± 2.0	82.3 ± 1.6	57.305 ± 27.779	22.08 ± 7.15
55	71.6 ± 2.1	73.0 ± 2.7	71.804 ± 30.489	30.48 ± 5.79
56	50.0 ± 2.0	44.1 ± 1.9	258.552 ± 81.460	29.27 ± 3.10
57	95.7 ± 3.5	95.3 ± 3.1	0.171 ± 0.285	7.47 ± 5.09
58	85.5 ± 2.2	84.7 ± 2.6	7.610 ± 5.220	13.50 ± 0.40
59	97.1 ± 2.1	97.4 ± 2.3	9.094 ± 6.072	8.67 ± 2.29
60	73.9 ± 7.3	72.8 ± 7.1	23.841 ± 11.033	23.18 ± 4.23
61	92.8 ± 4.0	90.5 ± 5.3	50.942 ± 49.446	17.35 ± 3.14
62	88.3 ± 0.8	89.0 ± 0.4	12.662 ± 7.444	18.25 ± 1.42
63	87.1 ± 3.7	85.4 ± 3.3	19.906 ± 8.169	15.04 ± 1.93
64	90.9 ± 1.9	89.5 ± 1.7	4.637 ± 4.070	11.77 ± 2.04
65	95.1 ± 3.0	93.8 ± 2.5	11.223 ± 6.326	9.64 ± 1.81
66	79.3 ± 5.2	75.7 ± 6.5	59.095 ± 72.762	19.96 ± 0.36
67	89.1 ± 2.0	87.2 ± 2.0	67.421 ± 74.678	19.06 ± 2.05
68	92.0 ± 3.2	89.2 ± 3.4	28.692 ± 12.324	11.14 ± 1.59
69	95.5 ± 2.4	94.7 ± 3.0	3.816 ± 3.821	9.76 ± 2.06
70	83.1 ± 3.9	82.9 ± 4.3	35.129 ± 16.702	22.23 ± 3.91
72	97.6 ± 1.0	98.2 ± 0.6	35.207 ± 43.781	10.73 ± 4.91
73	73.2 ± 1.2	70.6 ± 1.4	57.280 ± 46.832	21.20 ± 3.59
75	82.0 ± 3.9	79.6 ± 4.4	33.886 ± 21.982	20.40 ± 5.82

B

Mathematical Derivations

B.1 Derivation of Problem 7.2.1

Problem 7.2.1 can be solved by considering the unconstrained convex problem

$$\min_{\hat{s}} \parallel \mathbf{A}\hat{\mathbf{s}} - \mathbf{y} \parallel_2^2 + \lambda \parallel \hat{\mathbf{s}} - \tilde{\mathbf{s}} \parallel_2^2 .$$

It can be shown that this formulation is indeed equivalent to the original Problem 7.2.1 for an appropriate choice of the parameter λ. This problem can be solved analytically by expressing the objective function as the convex quadratic function

$$L(\mathbf{s}, \lambda) = \mathbf{s^T A^T A s} - 2\mathbf{y^T A s} + \mathbf{y^T y} + \lambda \mathbf{s^T s} - 2\lambda \tilde{\mathbf{s}}^\mathbf{T} \mathbf{s} + \lambda \tilde{\mathbf{s}}^\mathbf{T} \tilde{\mathbf{s}}.$$

A point $\mathbf{s}$ minimizes L if and only if

$$\nabla_s L = 2\mathbf{A^T A s} - 2\mathbf{A^T y} + 2\lambda \mathbf{s} - 2\lambda \tilde{\mathbf{s}} = 0,$$

which yields

$$\mathbf{s} = (\mathbf{A^T A} + \lambda \mathbf{I})^{-1} (\mathbf{A^T y} + \lambda \tilde{\mathbf{s}}).$$

By using the following inverse matrix formula

$$(\mathbf{I} + \mathbf{AB})^{-1} = \mathbf{I} - \mathbf{A}(\mathbf{I} + \mathbf{BA})^{-1}\mathbf{B},$$

we can write

$$\begin{aligned}
(\lambda \mathbf{I} + \mathbf{AB})^{-1} &= [\lambda(\mathbf{I} + \lambda^{-1}\mathbf{AB})]^{-1} \\
&= (\mathbf{I} + \lambda^{-1}\mathbf{AB})^{-1}\lambda^{-1} \\
&= [\mathbf{I} - \lambda^{-1}\mathbf{A}(\mathbf{I} + \lambda^{-1}\mathbf{BA})^{-1}\mathbf{B}]\lambda^{-1} \\
&= \{\mathbf{I} - \lambda^{-1}\mathbf{A}[\lambda^{-1}(\lambda\mathbf{I} + \mathbf{BA})]^{-1}\mathbf{B}\}\lambda^{-1} \\
&= [\mathbf{I} - \lambda\lambda^{-1}\mathbf{A}(\lambda\mathbf{I} + \mathbf{BA})^{-1}\mathbf{B}]\lambda^{-1} \\
&= \lambda^{-1}[\mathbf{I} - \mathbf{A}(\lambda\mathbf{I} + \mathbf{BA})^{-1}\mathbf{B}], \qquad \text{(B.1.1)}
\end{aligned}$$

and thus we have

$$\mathbf{s} = \lambda^{-1} \left[\mathbf{I} - \mathbf{A^T}(\lambda\mathbf{I} + \mathbf{AA^T})^{-1}\mathbf{A}\right] (\mathbf{A^T y} + \lambda\tilde{\mathbf{s}})$$

$$
\begin{aligned}
&= \tilde{\mathbf{s}} - \lambda^{-1}\mathbf{A}^{\mathbf{T}}(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\mathbf{A}(\mathbf{A}^{\mathbf{T}}\mathbf{y} + \lambda\tilde{\mathbf{s}}) + \lambda^{-1}\mathbf{A}^{\mathbf{T}}\mathbf{y} \\
&= \tilde{\mathbf{s}} - \lambda^{-1}\mathbf{A}^{\mathbf{T}}\left[(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\mathbf{A}\mathbf{A}^{\mathbf{T}}\mathbf{y} + (\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\lambda\mathbf{A}\tilde{\mathbf{s}} - \mathbf{y}\right] \\
&= \tilde{\mathbf{s}} - \lambda^{-1}\mathbf{A}^{\mathbf{T}}\{\left[(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\mathbf{A}\mathbf{A}^{\mathbf{T}} - \mathbf{I}\right]\mathbf{y} + (\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\lambda\mathbf{A}\tilde{\mathbf{s}}\}.
\end{aligned}
\tag{B.1.2}
$$

Finally, using the same inverse matrix formula, we can write

$$
(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\mathbf{A}\mathbf{A}^{\mathbf{T}} - \mathbf{I} = -\lambda(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1},
$$

which gives us the following solution

$$
\begin{aligned}
\mathbf{s} &= \tilde{\mathbf{s}} - \lambda^{-1}\mathbf{A}^{\mathbf{T}}\left[-\lambda(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\mathbf{y} + (\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}\lambda\mathbf{A}\tilde{\mathbf{s}}\right] \\
&= \tilde{\mathbf{s}} - \mathbf{A}^{\mathbf{T}}(\lambda\mathbf{I} + \mathbf{A}\mathbf{A}^{\mathbf{T}})^{-1}(\mathbf{A}\tilde{\mathbf{s}} - \mathbf{y}).
\end{aligned}
\tag{B.1.3}
$$

Bibliography

[1] Physionet Challenge 2013. `http:www.physionet.org/challenge/2013/`.

[2] American National Standard: Ambulatory Electrocardiographs, ANSI/AAMI, EC38, 1994.

[3] Paul S Addison. Wavelet transforms and the ECG: a review. *Physiological Measurement*, 26(5):155–199, 2005.

[4] Metin Akay and Eduard Mulder. Examining fetal heart-rate variability using matching pursuits. *IEEE Engineering in Medicine and Biology Magazine*, 15(5):64–67, 1996.

[5] Rute Almeida, Hernâni Gonçalves, João Bernardes, and Ana Paula Rocha. Fetal QRS detection and heart rate estimation: a wavelet-based approach. *Physiological Measurement*, 35(8):1723–1735, 2014.

[6] Fernando Andreotti, Maik Riedl, Tilo Himmelsbach, Daniel Wedekind, Niels Wessel, Holger Stepan, Claudia Schmieder, Alexander Jank, Hagen Malberg, and Sebastian Zaunseder. Robust fetal ECG extraction and detection from abdominal leads. *Physiological Measurement*, 35(8):1551–1567, 2014.

[7] Bruno Azzerboni, Fabio La Foresta, Nadia Mammone, Francesco Carlo Morabito, Località Feo, and Vito Reggio. A new approach based on wavelet-ICA algorithms for fetal electrocardiogram extraction. In *Proc. 13th Eur. Symp. Artif. Neural Netw*, 2005.

[8] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

[9] Richard G Baraniuk. Compressive sensing. *IEEE Signal Processing Magazine*, 24(4):118–121, 2007.

[10] K Baskaran. A survey on futuristic health care system: WBANs. *Procedia Engineering*, 30:889–896, 2012.

[11] Joachim Behar, Fernando Andreotti, Sebastian Zaunseder, Qiao Li, Julien Oster, and Gari D Clifford. An ECG simulator for generating maternal-foetal activity mixtures on abdominal ECG recordings. *Physiological Measurement*, 35(8):1537, 2014.

[12] Joachim Behar, Alistair Johnson, Gari D Clifford, and Julien Oster. A comparison of single channel fetal ECG extraction methods. *Annals of biomedical engineering*, 42(6):1340–1353, 2014.

[13] Joachim Behar, Julien Oster, and Gari D Clifford. Combining and benchmarking methods of foetal ECG extraction without maternal or scalp electrode data. *Physiological Measurement*, 35(8):1569, 2014.

[14] David E Bellasi and Luca Benini. Energy-efficiency analysis of analog and digital compressive sensing in wireless sensors. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 62(11):2718–2729, 2015.

[15] R Benzid, F Marir, A Boussaad, M Benyoucef, and D Arar. Fixed percentage of wavelet coefficients to be zeroed for ECG compression. *Electronics Letters*, 39(11):830–831, 2003.

[16] Radu Berinde, Anna C Gilbert, Piotr Indyk, Howard Karloff, and Martin J Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 798–805. IEEE, 2008.

[17] Radu Berinde and Piotr Indyk. Sparse recovery using sparse random matrices. *MIT-CSAIL Technical Report*, 2008.

[18] Pau Bofill and Michael Zibulevsky. Underdetermined blind source separation using sparse representations. *Signal Processing*, 81(11):2353–2362, 2001.

[19] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[20] Joseph D Bronzino. *Biomedical engineering handbook*, volume 2. CRC Press, 1999.

[21] Rebecca Brown, Jayawan HB Wijekoon, Anura Fernando, Edward D Johnstone, and Alexander EP Heazell. Continuous objective recording of fetal heart rate and fetal movements could reliably identify fetal compromise, which could reduce stillbirth rates by facilitating timely management. *Medical Hypotheses*, 83(3):410–417, 2014.

[22] Alfred M Bruckstein, David L Donoho, and Michael Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM review*, 51(1):34–81, 2009.

[23] Adrian Burns, Barry R Greene, Michael J McGrath, Terrance J O'Shea, Benjamin Kuris, Steven M Ayer, Florin Stroiescu, and Victor Cionca. Shimmer–a wireless sensor platform for noninvasive biomedical research. *IEEE Sensors Journal*, 10(9):1527–1534, 2010.

[24] A J Camm, P Kirchhof, G Y Lip, U Schotten, I Savelieva, S Ernst, I C Van Gelder, N Al-Attar, G Hindricks, B Prendergast, H Heidbuchel, O Alfieri, A Angelini, D Atar, P Colonna, R De Caterina, J De Sutter, A Goette, B Gorenek, M Heldal, S H Hohloser, P Kolh, J Y Le Heuzey, P Ponikowski, and F H Rutte. Guidelines for the management of atrial fibrillation: the task force for the management of atrial fibrillation of the european society of cardiology (ESC). *European Heart Journal*, 31(19):2369–2429, 2010.

[25] Emmanuel J Candè and Michael B Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008.

[26] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.

[27] Emmanuel J Candès. Compressive sampling. In *Proceedings of the International Congress of Mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.

[28] Emmanuel J Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9):589–592, 2008.

[29] Emmanuel J Candès, Yonina C Eldar, Deanna Needell, and Paige Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2011.

[30] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.

[31] Emmanuel J Candès, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.

[32] Emmanuel J Candès and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.

[33] M Carrara, L Carozzi, T J Moss, M De Pasquale, S Cerutti, M Ferrario, D E Lake, and J R Moorman. Heart rate dynamics distinguish among atrial fibrillation, normal sinus rhythm and sinus rhythm with frequent ectopy. *Physiological Measurement*, 36(9):1873–1888, 2015.

[34] Venkat Chandar. A negative result concerning explicit matrices with the restricted isometry property. *preprint*, 2008.

[35] Szi-Wen Chen, Hsiao-Chen Chen, and Hsiao-Lung Chan. A real-time QRS detection method based on moving-averaging incorporating with wavelet denoising. *Computer Methods and Programs in Biomedicine*, 82(3):187–195, 2006.

[36] I Christov, I Simova, and R Abächerli. Extraction of the fetal ECG in noninvasive recordings by signal decompositions. *Physiological Measurement*, 35(8):1713–1721, 2014.

[37] Gari D Clifford. A novel framework for signal representation and source separation: Applications to filtering and segmentation of biosignals. *Journal of Biological Systems*, 14(02):169–183, 2006.

[38] Gari D Clifford, Ikaro Silva, Joachim Behar, and George B Moody. Non-invasive fetal ECG analysis. *Physiological Measurement*, 35(8):1521–1536, 2014.

[39] GD Clifford, A Shoeb, PE McSharry, and BA Janz. Model-based filtering, compression and classification of the ECG. *International Journal of Bioelectromagnetism*, 7(1):158–161, 2005.

[40] Giulia Da Poian, Riccardo Bernardini, and Roberto Rinaldo. Gaussian dictionary for compressive sensing of the ECG signal. In *Biometric Measurements and Systems for Security and Medical Applications (BIOMS) Proceedings, 2014 IEEE Workshop on*, pages 80–85. IEEE, 2014.

[41] Giulia Da Poian, Riccardo Bernardini, and Roberto Rinaldo. Robust reconstruction for CS-based fetal beats detection. In *Signal Processing Conference (EUSIPCO), 2016 24th European*, pages 1303–1307. IEEE, 2016.

[42] Giulia Da Poian, Riccardo Bernardini, and Roberto Rinaldo. Separation and analysis of fetal-ECG signals from compressed sensed abdominal ECG recordings. *Biomedical Engineering, IEEE Transactions on*, 63(6):1269–1279, 2016.

[43] Wei Dai and Olgica Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *Information Theory, IEEE Transactions on*, 55(5):2230–2249, 2009.

[44] Ingrid Daubechies. The wavelet transform, time-frequency localization and signal analysis. *Information Theory, IEEE Transactions on*, 36(5):961–1005, 1990.

[45] Mark A Davenport. *Random observations on random observations: Sparse signal acquisition and processing.* PhD thesis, Citeseer, 2010.

[46] Mark A Davenport, Petros T Boufounos, Michael B Wakin, and Richard G Baraniuk. Signal processing with compressive measurements. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):445–460, 2010.

[47] Mark A Davenport, Marco F Duarte, Michael B Wakin, Jason N Laska, Dharmpal Takhar, Kevin F Kelly, and Richard G Baraniuk. The smashed filter for compressive classification and target recognition. In *Electronic Imaging 2007*, pages 64980H–64980H. International Society for Optics and Photonics, 2007.

[48] Geoff Davis, Stephane Mallat, and Marco Avellaneda. Adaptive greedy approximations. *Constructive Approximation*, 13(1):57–98, 1997.

[49] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. Fetal electrocardiogram extraction by blind source subspace separation. *Biomedical Engineering, IEEE transactions on*, 47(5):567–572, 2000.

[50] Alessia Dessì, Danilo Pani, and Luigi Raffo. An advanced algorithm for fetal heart rate estimation from non-invasive low electrode density recordings. *Physiological Measurement*, 35(8):1621–1636, 2014.

[51] Anna MR Dixon, Emily G Allstot, Daibashish Gangopadhyay, and David J Allstot. Compressed sensing system considerations for ECG and EMG wireless biosensors. *Biomedical Circuits and Systems, IEEE Transactions on*, 6(2):156–166, 2012.

[52] Zümray Dokur, Tamer Ölmez, Ertugrul Yazgan, and Okan K Ersoy. Detection of ECG waveforms by neural networks. *Medical Engineering & Physics*, 19(8):738–741, 1997.

[53] David L Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

[54] David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via $l1$ minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.

[55] Anthony Dupre, Sarah Vincent, and Paul A Iaizzo. Basic ECG theory, recordings, and interpretation. In *Handbook of cardiac anatomy, physiology, and devices*, pages 191–201. Springer, 2005.

[56] Armin Eftekhari et al. Robust-SL0 for stable sparse representation in noisy settings. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, pages 3433–3436. IEEE, 2009.

[57] Armin Eftekhari, Justin Romberg, and Michael B Wakin. Matched filtering from limited frequency samples. *Information Theory, IEEE Transactions on*, 59(6):3475–3496, 2013.

[58] Michael Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer Publishing Company, Incorporated, 1st edition, 2010.

[59] A. Fanelli, M. Ferrario, L. Piccini, G. Andreoni, G. Matrone, G. Magenes, and M. G. Signorini. Prototype of a wearable system for remote fetal monitoring during pregnancy. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 5815–5818, 2010.

[60] Andrew G Favret. Computer matched filter location of fetal R-waves. *Medical and Biological Engineering*, 6(5):467–475, 1968.

[61] Monica Fira, Liviu Goraş, Nicolae Cleju, and Constantin Barabasa. Results on ECG compressed sensing using specific dictionaries and its validation. In *Information Technology Interfaces (ITI), Proceedings of the ITI 2012 34th International Conference on*, pages 423–428. IEEE, 2012.

[62] W Bruce Fye. A history of the origin, evolution, and impact of electrocardiography. *The American Journal of Cardiology*, 73(13):937–949, 1994.

[63] Daibashish Gangopadhyay, Emily G Allstot, Anna MR Dixon, Karthik Natarajan, Subhanshu Gupta, and David J Allstot. Compressed sensing analog front-end for bio-sensor applications. *IEEE Journal of Solid-State Circuits*, 49(2):426–438, 2014.

[64] Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, 2000.

[65] Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Signal Processing Magazine*, 18(5):9–21, 2001.

[66] Patrick S Hamilton and Willis J Tompkins. Quantitative investigation of qrs detection rules using the MIT/BIH arrhythmia database. *Biomedical Engineering, IEEE Transactions on*, (12):1157–1165, 1986.

[67] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[68] Texas Instruments. CC2420: 2.4 Ghz IEEE 802.15. 4/ZigBee-ready RF transceiver. Available at http://www.ti.com/lit/gpn/cc2420, 2006.

[69] Texas Instruments. MSP430F5438 Datasheet, 2008.

[70] Janusz Jezewski, Adam Matonia, Tomasz Kupka, Dawid Roj, and Robert Czabanski. Determination of fetal heart rate from abdominal signals: evaluation of beat-to-beat accuracy in relation to the direct fetal electrocardiogram. *Biomedizinische Technik/Biomedical Engineering*, 57(5):383–394, 2012.

[71] Daniel T Kaplan. Simultaneous QRS detection and feature extraction using simple matched filter basis functions. In *Computers in Cardiology 1990, Proceedings.*, pages 503–506. IEEE, 1990.

[72] Raymond G Kennedy. Electronic fetal heart rate monitoring: retrospective reflections on a twentieth-century technology. *Journal of the Royal Society of Medicine*, 91(5):244–250, 1998.

[73] Ali Khamene and Shahriar Negahdaripour. A new method for the extraction of fetal ECG from the composite abdominal signal. *Biomedical Engineering, IEEE Transactions on*, 47(4):507–516, 2000.

[74] Seung-Jean Kim, Kwangmoo Koh, Michael Lustig, Stephen Boyd, and Dimitry Gorinevsky. An interior-point method for large-scale-regularized least squares. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):606–617, 2007.

[75] B-U Kohler, Carsten Hennig, and Reinhold Orglmeister. The principles of software QRS detection. *IEEE Engineering in Medicine and Biology Magazine*, 21(1):42–57, 2002.

[76] Jelena Kovacevic and Amina Chebira. Life beyond bases: The advent of frames (part I). *IEEE Signal Processing Magazine*, 24(4):86–104, 2007.

[77] Vinay Kumar, Abul K Abbas, Nelson Fausto, and Jon C Aster. *Robbins and Cotran Pathologic Basis of Disease*. Elsevier Health Sciences, 2014.

[78] Benoît Latré, Bart Braem, Ingrid Moerman, Chris Blondia, and Piet Demeester. A survey on wireless body area networks. *Wireless Networks*, 17(1):1–18, 2011.

[79] Q Li, C Y Liu, J Oster, and G D Clifford. Signal processing and feature selection preprocessing for classification in noisy healthcare data. *In book: Machine Learning for Healthcare Technologies*, Chapter 3:33–58, 2016.

[80] G Y H Lip, L Fauchier, S B Freedman, I Van Gelder, A Natale, C Gianni, S Nattel, T Potpara, M Rienstra, H Tse, and D A Lane. Atrial fibrillation. *Nature Reviews Disease Primers*, 2:16016, 2016.

[81] C Y Liu, K Li, L N Zhao, F Liu, D C Zheng, C C Liu, and S T Liu. Analysis of heart rate variability using fuzzy measure entropy. *Computers in Biology and Medicine*, 43(2):100–108, 2013.

[82] C Y Liu, J Oster, E Reinertsen, Q Li, L N Zhao, S Nemati, and G D Clifford. A review of methods of entropy for atrial fibrillation detection. *BioMedical Engineering OnLine*, submitted, 2016.

[83] Chengyu Liu, Peng Li, Costanzo Di Maria, Lina Zhao, Henggui Zhang, and Zhiqing Chen. A multi-step method with signal quality assessment and fine-tuning procedure to locate maternal and fetal qrs complexes from abdominal ecg recordings. *Physiological Measurement*, 35(8):1665, 2014.

[84] Zhitao Lu, Dong Youn Kim, and William A Pearlman. Wavelet compression of ECG signals by the set partitioning in hierarchical trees algorithm. *Biomedical Engineering, IEEE Transactions on*, 47(7):849–856, 2000.

[85] João PV Madeiro, Paulo C Cortez, Francisco I Oliveira, and Robson S Siqueira. A new approach to qrs segmentation based on wavelet bases and adaptive threshold technique. *Medical Engineering & Physics*, 29(1):26–37, 2007.

[86] Arian Maleki and David L Donoho. Optimally tuned iterative reconstruction algorithms for compressed sensing. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):330–341, 2010.

[87] Stephane Mallat. *A wavelet tour of signal processing: the sparse way*. Academic Press, 2008.

[88] Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41(12):3397–3415, 1993.

[89] Hossein Mamaghanian, Giovanni Ansaloni, David Atienza, and Pierre Vandergheynst. Power-efficient joint compressed sensing of multi-lead ECG signals. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4409–4412. IEEE, 2014.

[90] Hossein Mamaghanian, Nadia Khaled, David Atienza, and Pierre Vandergheynst. Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes. *Biomedical Engineering, IEEE Transactions on*, 58(9):2456–2466, 2011.

[91] Patrick E McSharry, Gari D Clifford, Lionel Tarassenko, and Leonard A Smith. A dynamical model for generating synthetic electrocardiogram signals. *Biomedical Engineering, IEEE Transactions on*, 50(3):289–294, 2003.

[92] Aleksandar Milenković, Chris Otto, and Emil Jovanov. Wireless sensor networks for personal health monitoring: Issues and an implementation. *Computer Communications*, 29(13):2521–2533, 2006.

[93] Akanksha Mishra, Falgun Thakkar, Chintan Modi, and Rahul Kher. ECG signal compression using compressive sensing and wavelet transform. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3404–3407. IEEE, 2012.

[94] Hosein Mohimani, Massoud Babaie-Zadeh, and Christian Jutten. A fast approach for overcomplete sparse decomposition based on smoothed norm. *Signal Processing, IEEE Transactions on*, 57(1):289–301, 2009.

[95] George B Moody and Roger G Mark. The impact of the MIT-BIH arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50, 2001.

[96] Samaneh Movassaghi, Mehran Abolhasan, Justin Lipman, David Smith, and Abbas Jamalipour. Wireless body area networks: A survey. *IEEE Communications Surveys & Tutorials*, 16(3):1658–1686, 2014.

[97] American College of Obstetricians and Gynecologists. ACOG practice bulletin No. 106: Intrapartum fetal heart rate monitoring: nomenclature, interpretation, and general management principles. *Obstetrics and Gynecology*, 114(1):192–202, 2009.

[98] World Health Organization et al. Global health observatory (GHO) data., 2015.

[99] NJ Outram, EC Ifeachor, PWJ Van Eetvelt, and JSH Curnow. Techniques for optimal enhancement and feature extraction of fetal electrocardiogram. *IEE Proceedings-Science, Measurement and Technology*, 142(6):482–489, 1995.

[100] Jiapu Pan and Willis J. Tompkins. A real-time QRS detection algorithm. *Biomedical Engineering, IEEE Transactions on*, BME-32(3):230–236, March 1985.

[101] Jeevan K Pant and Sridhar Krishnan. Compressive sensing of electrocardiogram signals by promoting sparsity on the second-order difference and by using dictionary learning. *Biomedical Circuits and Systems, IEEE Transactions on*, 8(2):293–302, 2014.

[102] Julian T Parer. *Handbook of fetal heart rate monitoring.* WB Saunders Company, 1997.

[103] YC Park, KY Lee, DH Youn, NH Kim, WK Kim, and SH Park. On detecting the presence of fetal R-wave using the moving averaged magnitude difference algorithm. *Biomedical Engineering, IEEE Transactions on*, 39(8):868–871, 1992.

[104] Yagyensh Chandra Pati, Ramin Rezaiifar, and PS Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, pages 40–44. IEEE, 1993.

[105] Maria Peters, John Crowe, Jean-Francois Piéri, Hendrik Quartero, Barrie Hayes-Gill, David James, Jeroen Stinstra, and Simon Shakespeare. Monitoring the fetal heart non-invasively: a review of methods. *Journal of Perinatal Medicine*, 29(5):408–416, 2001.

[106] Steven M Pincus and Ary L Goldberger. Physiological time-series analysis: what does regularity quantify? *American Journal of Physiology-Heart and Circulatory Physiology*, 266(4):H1643–H1656, 1994.

[107] Eduardo Correia Pinheiro, Octavian Adrian Postolache, and Pedro Silva Girao. Implementation of compressed sensing in telecardiology sensor networks. *International Journal of Telemedicine and Applications*, 2010:7, 2010.

[108] Luisa F Polania and Kenneth E Barner. Multi-scale dictionary learning for compressive sensing ECG. In *Digital Signal Processing and Signal Processing Education Meeting (DSP/SPE), 2013 IEEE*, pages 36–41. IEEE, 2013.

[109] Luisa F Polania, Rafael E Carrillo, Manuel Blanco-Velasco, and Kenneth E Barner. Compressed sensing based method for ECG compression. In *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, pages 761–764. IEEE, 2011.

[110] Joshua S Richman and J Randall Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *American Journal of Physiology-Heart and Circulatory Physiology*, 278(6):H2039–H2049, 2000.

[111] Rui Rodrigues. Fetal beat detection in abdominal ECG recordings: global and time adaptive approaches. *Physiological Measurement*, 35(8):1699–1711, 2014.

[112] Ron Rubinstein, Alfred M Bruckstein, and Michael Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045–1057, 2010.

[113] Antti Ruha, Sami Sallinen, and Seppo Nissila. A real-time microprocessor QRS detector system with a 1-ms timing accuracy for the measurement of ambulatory HRV. *Biomedical Engineering, IEEE Transactions on*, 44(3):159–167, 1997.

[114] Reza Sameni, Christian Jutten, and Mohammad B Shamsollahi. What ICA provides for ECG processing: Application to noninvasive fetal ECG extraction. In *Signal Processing and Information Technology, 2006 IEEE International Symposium on*, pages 656–661. IEEE, 2006.

[115] Reza Sameni, Mohammad B Shamsollahi, Christian Jutten, and Gari D Clifford. A nonlinear bayesian filtering framework for ECG denoising. *Biomedical Engineering, IEEE Transactions on*, 54(12):2172–2185, 2007.

[116] C.E. Shannon. A mathematical theory of communication. *Bell System Technical Journal, The*, 27(3):379–423, July 1948.

[117] Ivanovitch Silva, Joachim Behar, Reza Sameni, Tingting Zhu, Julien Oster, Gari D Clifford, and George B Moody. Noninvasive fetal ECG: the physionet/computing in cardiology challenge 2013. In *Computing in Cardiology Conference (CinC), 2013*, pages 149–152. IEEE, 2013.

[118] Nicholas Sperelakis, Yoshihisa Kurachi, Andre Terzic, and Michael V Cohen. *Heart physiology and pathophysiology*. Academic Press, 2000.

[119] E Malcolm Symonds, Daljit Sahota, and Allan Chang. *Fetal electrocardiography*. World Scientific, 2001.

[120] K Tateno and L Glass. A method for detection of atrial fibrillation using RR intervals. In *Computers in Cardiology 2000*, pages 391–394. IEEE, 2000.

[121] Joel A Tropp and Anna C Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 53(12):4655–4666, 2007.

[122] E. van den Berg and M. P. Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.

[123] E Van Den Berg and MP Friedlander. SPGL1: A solver for large-scale sparse reconstruction, 2007.

[124] Denise van der Linde, Elisabeth EM Konings, Maarten A Slager, Maarten Witsenburg, Willem A Helbing, Johanna JM Takkenberg, and Jolien W Roos-Hesselink. Birth prevalence of congenital heart disease worldwide: a systematic review and meta-analysis. *Journal of the American College of Cardiology*, 58(21):2241–2247, 2011.

[125] M Varanini, G Tartarisco, L Billeci, A Macerata, G Pioggia, and R Balocchi. An efficient unsupervised fetal QRS complex detection from abdominal maternal ECG. *Physiological Measurement*, 35(8):1607–1619, 2014.

[126] R Vullings, CHL Peters, RJ Sluijter, M Mischi, SG Oei, and JWM Bergmans. Dynamic segmentation and linear prediction for maternal ECG removal in antenatal abdominal recordings. *Physiological Measurement*, 30(3):291–307, 2009.

[127] Zhilin Zhang, Tzyy-Ping Jung, Scott Makeig, and Bhaskar D Rao. Compressed sensing for energy-efficient wireless telemonitoring of noninvasive fetal ECG via block sparse Bayesian learning. *IEEE Transactions on Biomedical Engineering*, 60(2):300–309, 2013.

[128] Michael Zibulevsky and Barak A Pearlmutter. Blind source separation by sparse decomposition in a signal dictionary. *Neural Computation*, 13(4):863–882, 2001.

[129] Yaniv Zigel, Arnon Cohen, and Amos Katz. The weighted diagnostic distortion (WDD) measure for ECG signal compression. *Biomedical Engineering, IEEE Transactions on*, 47(11):1422–1430, 2000.