Matrix Completion Based ECG Compression

Luisa F. Polania, Rafael E. Carrillo, Manuel Blanco-Velasco and Kenneth E. Barner

Abstract— An innovative electrocardiogram compression algorithm is presented in this paper. The proposed method is based on matrix completion, a new paradigm in signal processing that seeks to recover a low-rank matrix based on a small number of observations. The low-rank matrix is obtained via normalization of electrocardiogram records. Using matrix completion, the ECG data matrix is recovered from a few number of entries, thereby yielding high compression ratios comparable to those obtained by existing compression techniques. The proposed scheme offers a low-complexity encoder, good tolerance to quantization noise, and good quality reconstruction.

I. INTRODUCTION

An ECG is an important physiological signal for cardiac disease diagnostics. With the increasing use of modern electrocardiogram monitoring devices that generate vast amounts of data requiring huge storage capacity, ECG compression becomes mandatory to efficiently store and retrieve this data from medical database. ECG compression techniques can be classified into three categories [1]: Direct methods, transform methods and other compression methods. In the first category, the ECG samples are processed directly paying attention to the redundancy among them. In the second category, the wavelet transform-based methods play an interesting role due to their easy implementation and efficiency. The latter works in this area are characterized by hierarchical tree structures, such as embedded zero-tree wavelet (EZW) [2] and set partitioning in hierarchical tree (SPIHT) [3] protocols, which make use of the self-similarity of the wavelet transform across scales within a hierarchically decomposed wavelet tree.

An interesting line of research focuses on transforming the original one-dimensional ECG waveforms into two-dimensional information, followed by a processing stage using image processing tools. The idea of these methods is to exploit both intra-beat and inter-beat correlations. For example, Lee and Buckley [4] applied the DCT transform to an ECG data matrix composed of regular heartbeats and Bilgin [5] applied JPEG2000 compression to a similarly constructed matrix. Recently, we proposed in [6] a compressed sensing (CS) based ECG compression framework that utilizes distributed compressed sensing (see [7] and references therein) to exploit the inter-beat correlation structure. Following this line of thought, we propose a bidimensional

L.F. Polania, R.E. Carrillo and K.E. Barner are with the Dept. of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716 USA. e-mail: polania@eecis.udel.edu, carrillo@eecis.udel.edu , barner@eecis.udel.edu.

M. Blanco-Velasco is with the Dept. Teoría de la Señal y Comunicaciones, Universidad de Alcalá, Madrid, Spain. e-mail: manuel.blanco@uah.es.

ECG compression scheme based on matrix completion, that exploits both intra and inter-beat correlations.

Matrix completion refers to the problem of reconstructing a low-rank matrix from a small set of observed entries possibly corrupted by noise. Directly solving the rank minimization problem is NP hard. However, Candès and Recht in [8] extend the CS ideas (recovery of sparse signals) to the completion of low-rank matrices proposing a convex relaxation to the NP-hard problem by replacing the rank function with the nuclear norm of a matrix. Recently, many authors have proposed efficient algorithms for solving the low-rank matrix completion problem, such as Singular Value Threshholding [9] and Fixed Point Continuation algorithms (FPC) [10].

The main assumption of the matrix completion problem is that, to be recovered, the matrix has to be low-rank. To meet this requirement, the proposed method starts by forming a low-rank matrix in which each column corresponds to a length normalized ECG cycle. In the subsequent step, the sampling process forms an observed set by drawing uniformly at random from the formed matrix. The decoding process is performed using matrix completion followed by period de-normalization. The proposed scheme offers a low-complexity encoder and good tolerance to quantization noise. The performance of the proposed algorithm, in terms of reconstructed signal quality and compression ratio, is evaluated using the MIT-BIH Arrhythmia Database.

II. REVIEW OF MATRIX COMPLETION

Matrix completion refers to the problem of reconstructing a low-rank matrix M from a small set of observed entries possibly corrupted by noise.

A full rank matrix of dimensions $(n \times n)$ has n^2 degrees of freedom and, therefore, it is not possible to estimate all its values from a small subset of entries in the matrix. However, when the matrix has a low rank (r), the degrees of freedom are $r(2n - r)$, which is much smaller than n^2. Then, in the absence of noise, it is possible to recover a low-rank matrix by solving the following optimization problem

$$\begin{aligned} \min \quad & \text{rank}(X) \\ \text{s.t.} \quad & X_{ij} = M_{ij}, \ \text{for } (i,j) \in \Omega, \end{aligned} \quad (1)$$

where $X \in \mathbb{R}^{m \times n}$ is the variable matrix, $\text{rank}(X)$ is the rank of the matrix X and Ω is the set of indices of the sample entries. However, this optimization problem is NP-hard and cannot be (expected to be) solved in polynomial time [8].

Motivated by compressed sensing [11], where minimizing the ℓ_1 norm is the tightest convex relaxation to minimizing
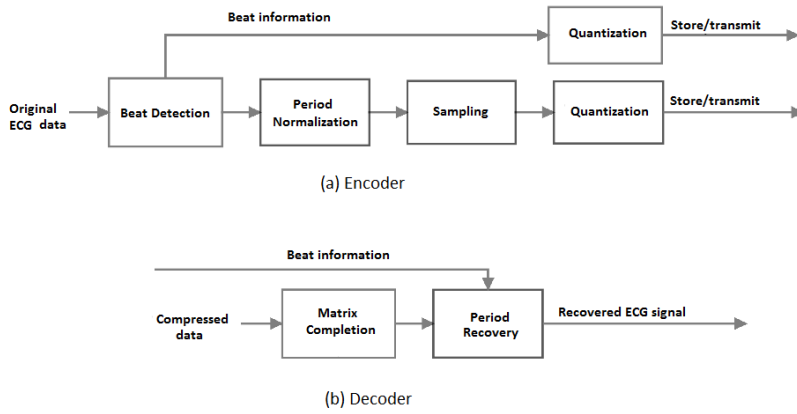
Fig. 1. Block Diagram of the proposed method. (a) Encoder. (b) Decoder.

the ℓ_0 norm, the following optimization problem is proposed

$$\begin{aligned} \min \quad & \|X\|_* \\ \text{s.t.} \quad & X_{ij} = M_{ij}, \quad \text{for } (i,j) \in \Omega, \end{aligned} \tag{2}$$

where $\|X\|_*$ stands for the nuclear norm, the sum of the singular values of X. Adopting compressed sensing ideas, rank($\cdot$) of a matrix corresponds to ℓ_0 norm of a vector, and nuclear norm to ℓ_1 norm [8]. The nuclear norm is the best convex approximation of the rank function over the unit ball of matrices.

This problem is a special case of the nuclear norm minimization problem

$$\begin{aligned} \min \quad & \|X\|_* \\ \text{s.t.} \quad & \mathcal{A}X = b, \end{aligned} \tag{3}$$

where $\mathcal{A} : \mathbb{R}^{p \times q} \to \mathbb{R}^m$ is a linear map and $b \in \mathbb{R}^m$. In the presence of noise, the constraint $\mathcal{A}(X) = b$ must be relaxed, resulting in either the problem

$$\begin{aligned} \min \quad & \|X\|_* \\ \text{s.t.} \quad & \|\mathcal{A}X - b\|_2^2 \leq \theta, \end{aligned} \tag{4}$$

or its Lagrangian form

$$\min \quad \lambda\|X\|_* + \frac{1}{2}\|\mathcal{A}X - b\|_2^2 \tag{5}$$

where θ and λ are algorithmic parameters [10].

III. COMPRESSION SCHEME

In this section, we present the proposed compression scheme. The block diagram of the proposed approach is presented in Fig. 1, where Fig. 1(a) depicts the encoder and Fig. 1(b) the decoder. Implementation details of the encoder and decoder are described below.

A. Encoder design

1) Period Normalization: The main assumption of the matrix completion problem is that the matrix to be recovered has to be low-rank. It is natural to expect that a matrix whose columns or rows are highly correlated has low rank. Motivated by this, we introduce a period normalization step in order to exploit the quasi-periodic nature of the ECG

signal, shown in the high correlation between samples of adjacent beats.

The peaks of QRS waves should be detected first to identify each heartbeat. Since each ECG period can have a different duration, we normalize them to the same length, using cubic spline interpolation. Let $\mathbf{x} = [x(0)\ x(1)\ \ldots\ x(N_0-1)]$ denote an ECG cycle. The normalized cycle, denoted as $\mathbf{x}_n = [x_n(0)\ x_n(1)\ \ldots\ x_n(N_n-1)]$, is computed as follows:

$$x_n(m) = \hat{x}(t^*), \tag{6}$$

where $\hat{x}(t^*)$ is an interpolated version of the samples, and

$$t^* = (m * N_0)/(N_n), \tag{7}$$

where N_0 and N_n are the original and normalized period lengths, respectively. Once the heartbeats have been normalized, they are organized column-wise in a matrix.

Apart from period normalization, there are alternative length equalization methods, such as zero-padding, which pads the short segments with zeros, or periodic-padding, which pads the short segments with a periodic extension. The 70% largest singular values of the matrices constructed by using these three length equalization techniques are shown in Fig. 2 on a semi-log scale. Although the three methods offer approximately low-rank representations, the set of singular values falls off more rapidly in the case of period normalization. These experiments were carried out over a single-lead ECG extracted from the record 117 from the MIT-BIH Arrhythmia Database.

When reconstructing the original signal, the decoder requires the original periods as side information. Therefore, these periods are added to the header of the compressed file.

2) Sampling and Optimal Quantization: The ECG data are organized column-wise in a $N_n \times Q$ matrix, denoted as X, where Q is the number of adjacent heartbeats encoded and N_n is the normalized length of each heartbeat. To reduce the dimensionality of the ECG data, we take p samples uniformly at random of the formed ECG matrix. We denote the sample vector as
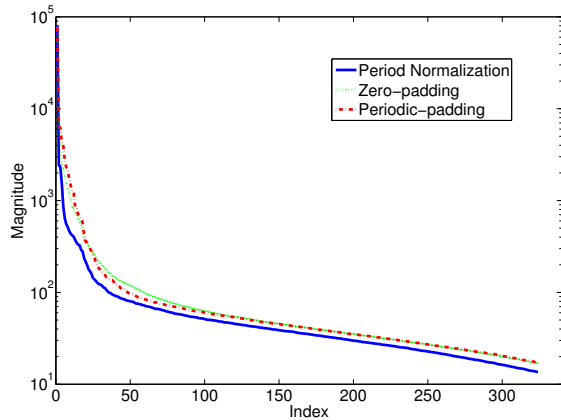
$$b = A(X) \tag{8}$$

Fig. 2. Singular values of the matrix with period normalization (70% of the largest singular values). (a) Period normalization. (b) Zero-padding. (c) Periodic extension.

where A represents the random sampling operator from $\mathbb{R}^{N_n \times Q}$ to $\mathbb{R}^p$. Although the length of the QRS complex is small compared to the length of the entire heart beat, its recovery in the decoder is important for diagnosis purposes. Therefore, we force the algorithm to use 30% of the total number of samples for the QRS sampling. To encode each sample, we utilize an optimal scalar quantizer designed with the Lloyd-Max algorithm with B_s bits per sample. Thus, the total number of bits used by the proposed method is $pB_s + QB_p$, where B_p is the number of bits used to encode the period of the heartbeats.

B. Decoder design

Although the nuclear norm minimization problem can be cast as a semidefinite programming problem, it can be computationally expensive to solve when the matrices are large. In this paper, we decide to use the algorithm proposed by Ma *et. al* in [10], which corresponds to a continuation (homotopy) technique to accelerate the convergence of the following fixed point iteration

$$Y^k = X^k - \tau g(X^k)$$
$$X^{k+1} = S_{\tau\mu}(Y^k)$$

where $g(X^k) = A^*(A(X^k) - b)$, A^* being the adjoint operator of A, and $S_v(Y)$ is a shrinkage operator defined as

$$S_v(\cdot) = U\mathrm{Diag}(\overline{\sigma})V^T,$$

where $\overline{\sigma} = \max\{\sigma - v, 0\}$ and $U\mathrm{Diag}(\sigma)V^T$ is the singular value decomposition of Y. The fixed point continuation iterative scheme offers robustness and good recoverability. The entire algorithm is detailed in [10]. The original heartbeat periods can be recovered by using the same transformation as in (6).

IV. Experimental results

This section describes experiments to evaluate the performance of the proposed scheme. The ECG data used in

the experiments are from the MIT-BIH arrhythmia database, sampled at 360 Hz with a resolution of 11 bits/sample.

Let x and $\hat{x}$ be the N-dimensional vectors representing the original and reconstructed signals, respectively. The quality of the reconstructed signals is evaluated via the percent root mean square difference (PRD), defined as $PRD = (\|x - \hat{x}\|/\|x\|) \cdot 100$, where the operator $\|\cdot\|$ denotes the Euclidean norm. The compression performance is evaluated by using the compression ratio (CR), defined as

$$CR = \frac{11 \times N}{(p \times B_s) + (Q \times B_p)},$$

where N is the length of the input signal, 11 is the number of bits to encode each sample, B_s are the bits representing the measurements, p is the number of measurements, Q is the number of ECG cycles and B_p are the bits representing each period.

In the first experiment, we compare the performance of the proposed algorithm with conventional SPIHT [3] for the records 100, 117 and 119. For each experiment, 30 repetitions are averaged. We use eight bits for the quantization of the measurements, $B_s = 8$, and eight bits for the heartbeat periods, $B_p = 8$. The dimensions of the resulting ECG images depend on the number of detected periods and their average length.

Figure 3 shows the results of the first experiment. An overall evaluation of the records reveals that if the target CR is high, the proposed method is preferred since it significantly outperforms SPIHT, in this case. The heartbeats in the records 117 and 100 are very regular and, therefore, it is possible to get a good approximation of the low-rank matrix after the period normalization step. The algorithm also shows a good performance in the case of signals with extremely varying periods, such as the record 119.

The proposed algorithm is also compared to other ECG coders through their reported performance in the literature. In Table I, we show PRD comparisons of different coding algorithms. As the results show, the proposed scheme exhibits better performance than well-known methods, such as those based on wavelet transforms [12], wavelet packets [2], [13] or the DCT transform [4]. For high compression ratios, the introduced method also outperforms algorithms, such as SPIHT, according to our first experiment. When compared to JPEG2000 [5], the ECG compression method based on matrix completion exhibits lower performance in terms of compression. Conversely, our method offers a less complex and computationally demanding encoder. The complexity of the encoder is dominated by the cost of the period normalization stage, assuming that fast sampling operators are available. Suppose L represents the cost of the period normalization for a single heartbeat, then the computational cost of the encoder is $\mathcal{O}(QL)$.

In this work, visual study of error signal is also considered. The waveforms of record 117 given by the proposed compression scheme are visually evaluated in Fig. 4. The reconstructed signal remains close to the original signal and the error is equally distributed along the horizontal axis. This
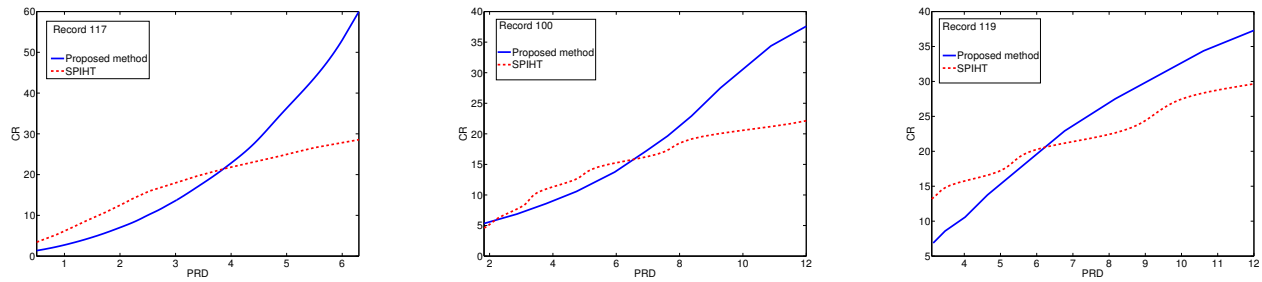
Fig. 3. Comparison of the proposed method with SPIHT. L: Record 117, M: Record 100, R: Record 119.

TABLE I
COMPARISON OF DIFFERENT ECG COMPRESSION ALGORITHMS

Algorithm	Record	CR	PRD
Lu et. al [3]	117	8:1	1.18
Hilton [2]	117	8:1	2.6
Djohan et. al [12]	117	8:1	3.9
Bilgin et. al [5]	117	8:1	0.86
Proposed method	117	8:1	2.18
Bilgin et. al [5]	117	10:1	1.03
Proposed method	117	10:1	2.5
Bradie et. al [13]	100	23.61:1	9.92
Proposed method	100	23.61:1	8.4
Bradie et. al [13]	119	19.65:1	6.15
Proposed method	119	19.65:1	6.01

implies that the proposed method performs well locally and does not incorporate outliers in the reconstruction.

Like most ECG compression algorithms that operate on 2-D arrays, the proposed scheme relies on the assumption that the periods are gathered and organized in a matrix. This will cause some delay, at least one 2-D array collection time. It is clear that this behavior compromises strict real-time data transmission. However, the proposed scheme is useful in cases when one is only concerned with data storage, or minor delays are tolerable.

V. CONCLUSION

We applied the matrix completion problem to ECG compression by constructing an approximately low rank matrix with ECG data. The results indicate that the presented algorithm compares favorably to other methods in the literature. The proposed scheme offers the low complexity encoder needed by holters and other ECG acquisition devices. Future directions for research will include the incorporation of an entropy coding stage and the design of an adaptive sampling method that will allow further reductions in the number of samples.

REFERENCES

[1] M. Blanco-Velasco, F. Cruz-Roldán, F. López, Á. Bravo, and D. Martínez, "A low computational complexity algorithm for ECG signal compression," *Medical Eng. Phys.*, vol. 26, no. 7, pp. 553 –568, Sept. 2004.
[2] M.L. Hilton, "Wavelet and wavelet packet compression of electrocardiograms," *IEEE T-BME*, vol. 44, no. 5, pp. 394 –402, May 1997.
[3] Z. Lu, D. Y. Kim, and W. A. Pearlman, "Wavelet compression of ECG signals by the set partitioning in hierarchical trees algorithm," *IEEE T-BME*, vol. 47, no. 7, pp. 849 –856, July 2000.
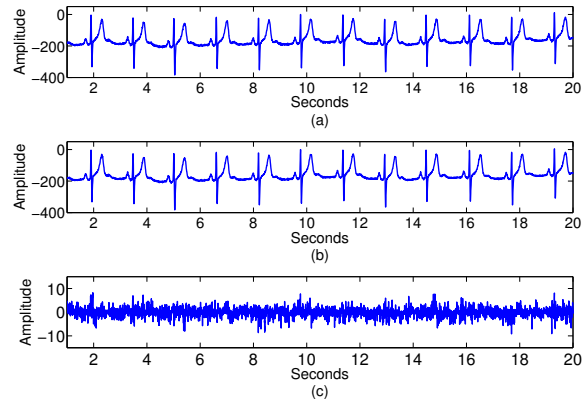
Fig. 4. Compression waveform of record 117 for CR=10, PRD=2.51. (a) Original signal. (b) Reconstructed signal. (c) Reconstruction error.

[4] H. Lee and K.M. Buckley, "ECG data compression using cut and align beats approach and 2-D transforms," *IEEE T-BME*, vol. 46, no. 5, pp. 556 –564, May 1999.
[5] A. Bilgin, M.W. Marcellin, and M.I. Altbach, "Compression of electrocardiogram signals using JPEG2000," *IEEE Transactions on Consumer Electronics*, vol. 49, no. 4, pp. 833 – 840, 2003.
[6] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco, and K. E. Barner, "Compressed sensing based method for ECG compression," in *Proc., IEEE ICASSP*, Prague, Czech Republic, May 2011.
[7] D. Baron, M. Duarte, M. Wakin, S. Sarvotham, and R. Baraniuk, "Distributed compressive sensing," *Preprint*, Jan. 2009.
[8] E.J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
[9] E.J. Candès J.-F. Cai and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. on Optimization*, vol. 20, pp. 1956–1982, March 2010.
[10] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," Tech. Rep., Department of IEOR, Columbia University, 2008.
[11] E.J. Candès and M.B. Wakin, "An introduction to compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21 –30, Mar. 2008.
[12] A. Djohan, T.Q. Nguyen, and W.J. Tompkins, "ECG compression using discrete symmetric wavelet transform," in *IEEE 17th Annual Conference Engineering in Medicine and Biology Society*, Sept. 1995, vol. 1, pp. 167 –168 vol.1.
[13] B. Bradie, "Wavelet packet-based compression of single lead ECG," *IEEE T-BME*, vol. 43, no. 5, pp. 493 –501, May 1996.