# Compressed Sensing ECG using Restricted Boltzmann Machines

Luisa F. Polanía[a,*], Rafael I. Plaza[b,*]

[a]*American Family Mutual Insurance Company, Madison, WI 53783, USA*
[b]*Scientific Games Corporation, Cedar Falls, IA 50613, USA*

## Abstract

Recently, it has been shown that compressed sensing (CS) has the potential to lower energy consumption in wireless electrocardiogram (ECG) systems. By reducing the number of acquired measurements, the communication burden is decreased and energy is saved. In this paper, we aim at further reducing the number of necessary measurements to achieve faithful reconstruction by exploiting the representational power of restricted Boltzmann machines (RBMs) to model the probability distribution of the sparsity pattern of ECG signals. The motivation for using this approach is to capture the higher-order statistical dependencies between the coefficients of the ECG sparse representation, which in turn, leads to superior reconstruction accuracy and reduction in the number of measurements, as it is shown via experiments.

*Keywords:*
Electrocardiogram (ECG), wireless body area networks (WBAN), compressed sensing (CS), overcomplete dictionaries, restricted Boltzmann machine (RBM).

## 1. Introduction

Even though wireless body area networks (WBANs) have the potential to revolutionize health monitoring by allowing the transition from centralized health care services to ubiquitous and ambulatory health monitoring, they still need to face some challenges, such as reducing energy consumption [1]. Current Zigbee and ANT-based technologies claim that they have battery life of three years, but this is the case for low operating data rates. For continuous operation at 250 kb/s, a regular Lithium ion battery is consumed in a couple of hours [2]. Therefore, at the heart of making WBAN ready for deployment is the need for data reduction at the sensor nodes. It has been shown that the application of CS to WBAN-enabled ECG monitors results in power-efficient sensor nodes that attain high compression rates at a low computational cost when using a sparse binary sensing matrix [3]. Compressed sensing is a framework that exploits the structure of signals to acquire data at a rate proportional to the information content rather than the frequency content, therefore allowing perfect reconstruction from sub-Nyquist random linear measurements [4, 5].

The application of CS to wireless ECG systems has attracted attention in recent years. Works in this area include studies of practical design considerations for CS encoding and decoding [3, 6], as well as the development of CS reconstruction algorithms that exploit structural ECG information [7, 8, 9, 10, 11, 12]. This paper proposes to use our recently developed CS scheme [13], henceforth referred to as RBM-CS scheme, for ECG reconstruction. This scheme exploits the representational power of RBMs to model the probability distribution of the sparsity pattern of a signal class, with the ultimate goal of reducing the number of measurements.

In this paper, overcomplete dictionaries and wavelets

---

*Corresponding authors:

*Email addresses:* `lpolania@amfam.com` (Luisa F. Polanía ), `plaza@udel.edu` (Rafael I. Plaza )

are employed to sparsely represent ECG signals. Training of the dictionaries also results in a set of sparse codes associated with the training data, which are often discarded after training since the main interest lies in the dictionary. Instead, the RBM-CS scheme utilizes the sparse codes support to train a RBM, which is later used by a CS reconstruction algorithm that fully exploits the model. The reason for incorporating RBMs into the reconstruction is to exploit higher-order dependencies between sparse coefficients. The RBM-CS scheme falls within the structured compressed sensing framework [14], which aims at exploiting signal structure by considering more elaborate priors that go beyond the simplistic sparsity prior. Even though ECG signals have a rich structure, most of the previous CS ECG works only exploit signal sparsity. One exception is our recent work that incorporates prior information about the dependencies of the ECG wavelet coefficients across scales into the CS reconstruction algorithm [12]. Another work that exploits ECG structural information but using a statistical model was proposed by Zhang *et al.* [7]. Their work proposes to use the block sparse Bayesian learning framework to reconstruct fetal ECG recordings.

The performance of the RBM-CS scheme is evaluated on the MIT-BIH Arrhythmia Database [15] and the European ST-T Database [16]. Experimental results include comparisons with the Basis Pursuit denoising algorithm (BPDN), the most widely used reconstruction algorithm in the CS ECG literature [3, 6]. Experiments are also performed with the bound-optimization-based BSBL algorithm [7], referred hereafter as BO-BSBL, and with the model-based CoSaMP and model-based Iterative Hard Thresholding algorithms [12] in order to compare with other CS reconstruction algorithms that also exploit signal structure. Simulation results indicate that the RBM-CS scheme offers superior reconstruction accuracy for the low-measurement regime.

The structure of the paper is as follows. Section II presents a brief review of RBMs and the RBM-CS scheme. Section III validates the potential of the RBM-CS scheme to improve reconstruction quality and to reduce the number of measurements via experiments. The paper concludes in Section IV with final remarks.

## 2. Background

### 2.1. Restricted Boltzmann Machines

Restricted Boltzmann machines are a type of undirected graphical models formed by a layer of binary stochastic hidden units and a layer of stochastic visible units that, for the purpose of this work, will be Bernoulli distributed conditional on the hidden units. The visible units $v = [v_1 v_2 \ldots v_J]^T$ represent the input variables of the data that needs to be modeled. The hidden units $h = [h_1 h_2 \ldots h_P]^T$ are trained to capture the higher-order correlations that are observed at the visible units. The layers are connected via a weight matrix $W$. The structure of a RBM forms a bipartite graph with no visible-visible or hidden-hidden connections. In this work, RBMs are used to learn a joint probability distribution of training data. In a RBM, the probability distribution over visible units is defined as

$$p(v) = \sum_h p(v, h) = -\frac{1}{Z} \exp\left(-\mathrm{E}(v)\right), \qquad (1)$$

where

$$\mathrm{E}(v) = -\sum_j \log\left(1 + e^{W_{\cdot j}^T v + b_{hj}}\right) - b_v^T v, \qquad (2)$$

and $Z = \sum_{\mathbf{v}} \exp(-\mathrm{E}(\mathbf{v}))$ is the normalization term.

### 2.2. Compressed Sensing Scheme using restricted Boltzmann machines

We recently proposed a CS scheme which employs a RBM to model the probability distribution of the sparsity pattern of a signal class, the RBM-CS scheme [13]. The main advantage of this scheme is that it captures higher-order dependencies between sparse coefficients, which ultimately translates into a reduction in the number of necessary measurements to attain accurate reconstruction. Fig. 1 illustrates the block diagram of the RBM-CS scheme, which consists of two stages, namely, the training and compressed sensing stages. The training stage varies depending on the employed sparsifying transform. Details of each stage are provided in this section.

#### 2.2.1. Compressed Sensing stage

Let $x \in \mathbb{R}^N$ denote the signal to be recovered and $D \in \mathbb{R}^{N \times Q}$ the dictionary employed to represent $x$, *i.e.*
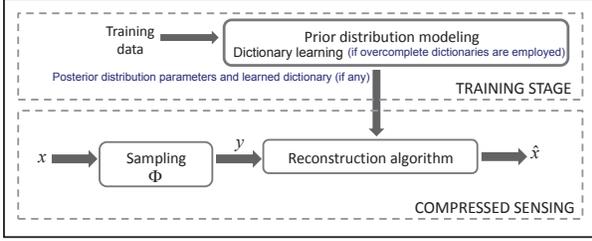
Figure 1: Block diagram of the RBM-CS scheme

$x = Ds + r$, where $s$ and $r$ are the sparse representation and the representation error. A Gaussian distribution with zero mean and covariance $\Sigma_r$ is assumed for $r$. The support of $s$, of cardinality $K$, is denoted as $\theta$. The nonzero coefficients of $s$ are denoted as $s_\theta$. A Gaussian distribution with zero mean and variance $\sigma_{s_i}^2$ is assumed for each nonzero coefficient $s_i$, $i \in \theta$. Therefore, the conditional distribution of $s_\theta$ given $\theta$ takes the form $s_\theta|\theta \sim \mathcal{N}(\mathbf{0}, \Sigma_\theta)$, where $\Sigma_\theta \in \mathbb{R}^{K \times K}$ is a diagonal matrix, whose diagonal is formed by the variances of the nonzero coefficients $\sigma_{s_i}^2$, $i \in \theta$. The sparsity pattern $S^{\{\theta\}}$ associated with the support $\theta$ is defined as $S_i^{\{\theta\}} = \mathbf{1}[i \in \theta]$ for $i = 1, \dots, N$, where $\mathbf{1}[\cdot]$ denotes the indicator function.

Compressed sensing addresses the recovery of signal $x$ from undersampled and noisy measurements of the form $y = \Phi x + n$, where $\Phi \in \mathbb{R}^{M \times N}$ is the sampling matrix and $n$ accounts for the additive Gaussian sampling noise of zero mean and variance $\sigma_n^2$. Defining $\Xi = \Phi D$ and $\eta = \Phi r + n$, vector $y$ takes the form $y = \Xi s + \eta$. The conditional distribution of $y$ given $\theta$ is given by

$$\begin{aligned} \mathrm{p}(y|\theta) = & C \times \det\left(\Xi_\theta^T \Sigma_\eta^{-1} \Xi_\theta \Sigma_\theta + I\right)^{-1/2} \\ & \times \exp\left\{\frac{1}{2} y^T \Sigma_\eta^{-1} \Xi_\theta P^{-1} \Xi_\theta^T \Sigma_\eta^{-1} y\right\}, \end{aligned} \tag{3}$$

where $C = \det(2\pi\Sigma_\eta)^{-1/2}\exp\left\{-\frac{1}{2}y^T\Sigma_\eta^{-1}y\right\}$ and $P = \Xi_\theta^T \Sigma_\eta^{-1} \Xi_\theta + \Sigma_\theta^{-1}$.

In the RBM-CS scheme, the MAP estimator is employed to recover $s$. The MAP estimate of $s$ requires

knowledge of the support, which is calculated as

$$\begin{aligned} \hat{\theta} & = \arg\max_\theta p(\theta|y) && (4) \\ & = \arg\max_\theta p(y|\theta)p(\theta) && (5) \\ & = \arg\max_\theta \Big(\frac{1}{2}y^T \Sigma_\eta^{-1} \Xi_\theta P^{-1} \Xi_\theta^T \Sigma_\eta^{-1} y && (6) \\ & \quad -\frac{1}{2}\log\left(\det(P\Sigma_\theta)\right) + \sum_j \log\left(1 + e^{W_{\cdot j}^T S^{\{\theta\}} + b_{hj}}\right) \\ & \quad + b_v^T S^{\{\theta\}}\Big). \end{aligned}$$

where $p(\theta)$ is calculated by using (1), the probability distribution over the visible units of the RBM.

The posterior distribution $p(s_{\hat{\theta}}|y, \hat{\theta})$ is Gaussian distributed, and therefore, the MAP estimate of $s$ is directly obtained from the mean of the posterior, $i.e.$,

$$\begin{aligned} \hat{s}_{\hat{\theta}} & = \arg\max_{s_{\hat{\theta}}} p(s_{\hat{\theta}}|y, \hat{\theta}), && (7) \\ & = \Sigma_{\hat{\theta}} \Xi_{\hat{\theta}}^T (\Xi_{\hat{\theta}} \Sigma_{\hat{\theta}} \Xi_{\hat{\theta}}^T + \Sigma_\eta)^{-1} y. \end{aligned}$$

To solve for (4), the RBM-CS scheme uses a greedy pursuit algorithm based on the Orthogonal Matching Pursuit (OMP) algorithm, referred to as RBM-OMP-like algorithm [13]. The algorithm starts by initializing the support to the empty set. Then, it searches for the element $i$ that can be added to the support in order to maximize $p(\theta|y)$ at each iteration. The algorithm stops when the number of iterations exceeds the pre-defined sparsity threshold. Once the signal support is calculated, the sparse representation $s$ is estimated via the MAP estimator (7).

### 2.2.2. Training stage

Let us first consider the case when overcomplete dictionaries are used as the sparsifying transform. In this case, the training stage is employed with dual purpose, it learns both the the dictionary and the parameters of the posterior distribution $p(\theta|y)$.

Let $G = [g^1 \dots g^B] \in \mathbb{R}^{N \times B}$ denote the set of $N$-dimensional training samples, which is referred to as training data set. The overcomplete dictionary $D = [d^1 \dots d^J] \in \mathbb{R}^{N \times J}$ $(J > N)$ is learned by solving the fol-

lowing optimization problem

$$\{\hat{D}, \hat{A}\} = \underset{D,A}{\arg\min} \|G - DA\|_F^2 \ \text{ s. t. } \ \|a^i\|_0 < K, \ \forall i, \quad (8)$$

where $A = [a^1 \ldots a^B] \in \mathbb{R}^{J \times B}$ and $K$ denote the sparse codes of $G$ and the pre-specified sparsity threshold, respectively. The representation error is defined as $E = G - \hat{D}\hat{A}$. The RBM-CS scheme uses the K-SVD algorithm proposed by Aharon *et al.* [17] to solve for (8).

Let $u^i$ denote the sparsity pattern of the sparse code $a^i$, $i = 1, \ldots, B$. The $j$th element of $u^i$ is defined as $u_j^i = \mathbf{1}[j \in \text{supp}(a^i)]$. The set of vectors $U = [u^1 \ldots u^B]$ are employed to train the RBM model. In the RBM-CS scheme, contrastive divergence [18] is used for learning the parameters of the RBM model. Furthermore, the set of sparse codes $A$ are employed to learn the variances of the sparse coefficients:

$$\hat{\sigma}_{s_i}^2 = \frac{\sum_{j=1}^{B} \left(a_i^j\right)^2}{\sum_{j=1}^{B} \mathbf{1}[i \in \text{supp}(a^j)]}, i = 1 \ldots N \quad (9)$$

The RBM-CS scheme assumes that the sampling noise variance is known *a priori*. It also assumes independence between the representation error coefficients $r_i$ and $r_j$ for $i \neq j$, which implies that the the covariance matrix $\Sigma_r$ is a diagonal matrix, whose diagonal is formed by the variances of the representation error coefficients $\sigma_{r_i}^2$, $\forall i$. The representation error of the learned dictionary $E = [e^1 \ldots e^B]$ is employed to estimate each diagonal element of $\Sigma_r$. That is,

$$\hat{\sigma}_{r_i}^2 = \frac{1}{B} \sum_{j=1}^{B} \left(e_i^j\right)^2. \quad (10)$$

The estimate of $\Sigma_\eta$ is directly calculated as $\hat{\Sigma}_\eta = \Phi \hat{\Sigma}_r \Phi^T + \sigma_n^2 I$. In the case when the sparsifying transform is an orthonormal basis, the training stage is only used to learn the parameters of the posterior distribution $p(\theta|y)$, following the same procedure as in the case of overcomplete dictionaries.

## 3. Experimental Results

To validate the potential of the RBM-CS scheme to improve the performance of CS ECG systems, experiments are performed on records from the MIT-BIH Arrhythmia Database [15] and the European ST-T database [16].

The entries of the sampling matrix $\Phi$ are independently sampled from a symmetric Bernoulli distribution $(\text{P}(\Phi_{i,j} = \pm 1/\sqrt{M} = 1/2))$ in order to build an efficient hardware implementation. The use of Bernoulli matrices, as compared to other sub-Gaussian matrices, results in simpler circuit complexity, data storage, and computation requirements [19]. Presented results correspond to averages of 50 repetitions of each experiment, with a different realization of the random measurement matrix at each time. Both the wavelet transform and learned overcomplete dictionaries are employed as sparsifying transforms in the experiments.

The reconstruction SNR (R-SNR), precision, recall, and the percentage similarity are used as performance measures for our experiments. The R-SNR is defined as

$$\text{R-SNR} = 10\log_{10}\frac{\|x\|_2^2}{\|x - \hat{x}\|_2^2}, \quad (11)$$

where $x$ and $\hat{x}$ denote the $N$-dimensional original and reconstructed signals, respectively. Precision is defined as *Precision* = TP/(TP + FP) and recall is defined as *Recall* = TP/(TP + FN), where TP, FP, and FN denote true positives, false positives and false negatives, respectively. The percentage similarity, also used in [20], is defined as PSim = $100 - \left(\frac{|y - \bar{y}|}{y} \times 100\right)$, where $y$ is the value of a calculated metric using the original signal and $\bar{y}$ is the value of the same metric on the resulting signal after reconstruction. PSim is used for the precision and recall metrics in this paper. Precision, recall, and the percentage similarity are used for the performance evaluation of QRS detection in Section 3.1.3.

### 3.1. Experiments on the MIT-BIH Arrhythmia Database

Every file in the MIT-BIH Arrhythmia Database contains two 30-min long lead recordings sampled at 360 Hz with 11 bits per sample of resolution. Single leads from records 100, 101, 102, 107, 109, 111, 115, 117, 118 and 119 are employed for the experiments in this section. This data set consists of a variety of signals with different morphologies, rhythms and abnormal heartbeats.
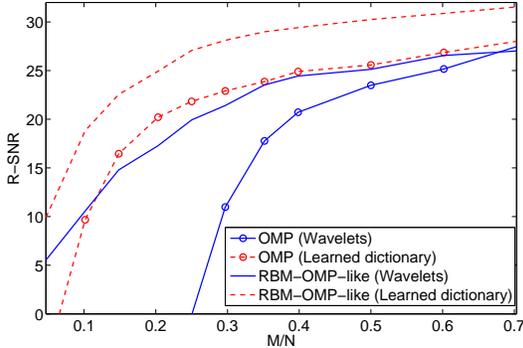
4

Figure 2: Comparison of the reconstruction of ECG signals using the OMP and the RBM-OMP-like algorithm for both wavelets and learned overcomplete dictionaries.
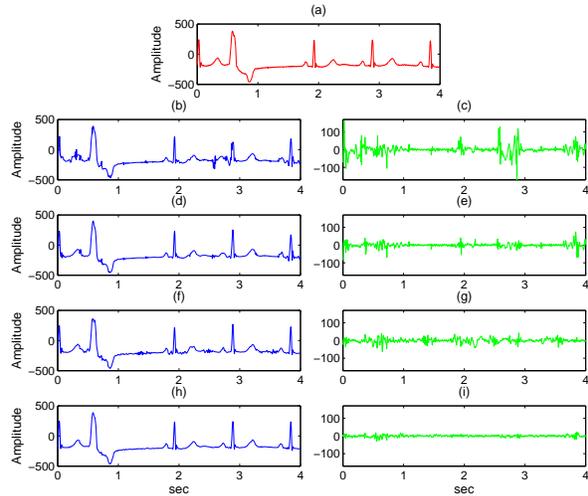


Figure 3: Visual evaluation of the RBM-OMP-like reconstruction algorithm using overcomplete dictionaries and the wavelet basis as sparsifying transforms. Record 119, $M = 0.35N$. First row: (a) Original signal. Second and third rows: (b-e) Reconstructed signals using a wavelet basis as the sparsifying transform (b-c) OMP reconstruction and error, R-SNR=18.32, (d-e) RBM-OMP-like reconstruction and error, R-SNR=25. Fourth and fifth rows: (f-i) Reconstructed signals using an overcomplete dictionary as the sparsifying transform (f-g) OMP reconstruction and error, R-SNR=24.59, (h-i) RBM-OMP-like reconstruction and error, R-SNR=31.67.

### 3.1.1. Performance Evaluation

The first experiment compares the ECG reconstruction performance of the RBM-OMP-like algorithm with the traditional OMP algorithm. The training and testing data sets for this experiment consist of 92800 and 5000 segments of size $N = 128$, extracted from the selected ECG records from the MIT-BIH Arrhythmia Database, respectively. The same number of segments is extracted from each record. A RBM with the same number of hidden units as of visible units is employed to model the probability distribution of the sparsity pattern. Compressed measurements are artificially contaminated with Gaussian noise of variance $\sigma_n^2 = 0.25$. The number of dictionary atoms is set to $3N$. The Daubechies-4 wavelet transform, using a decomposition level $L = 4$, is employed. The sparsity threshold is set to $K = 0.08N$ and $K = 0.1N$ for overcomplete dictionaries and wavelets, respectively. The results of the comparison in Fig. 2 indicate that the RBM-OMP-like algorithm has superior reconstruction performance than OMP and requires significantly less number of measurements to achieve accurate reconstruction. The performance gain is larger in the case of overcomplete dictionaries than in the case of wavelets.

Fig. 3 visually illustrates the reconstruction of an ECG signal using the RBM-OMP-like algorithm. Fig. 3(a) corresponds to a 4-second duration segment from record 119, which contains ventricular ectopic beats. The signal is divided into segments of 128 samples. Each seg-

ment is sampled and reconstructed separately. Then, the segments are concatenated to reconstruct the original signal. For this experiment, the number of measurements is set to $M = 0.35N$. Figs. 3(b-e) and Figs. 3(f-i) are obtained when wavelets and overcomplete dictionaries are employed as the sparsifying transform, respectively. Figs. 3(b-c) and Figs. 3(f-g) show the obtained reconstructed (left) and error (right) signals when using the OMP algorithm. Similarly, Figs. 3(d-e) and Figs. 3(h-i) show the obtained reconstructed (left) and error (right) signals when using the RBM-OMP-like algorithm. The recovered signals using the RBM-OMP-like algorithm are better estimates of the original signals than those obtained with the traditional OMP algorithm. It should be noted that using overcomplete dictionaries leads to preservation of detailed information for clinical diagnosis and less number of artifacts in the reconstruction.

The second experiment aims at comparing the performance of the RBM-OMP-like algorithm with previously

applied CS algorithms to the problem of ECG reconstruction, such as BPDN [3, 6], model-based CoSaMP, denoted as MB-CoSaMP, model-based Iterative Hard Thresholding, denoted as MB-IHT [12], and the BO-BSBL algorithm [7]. Results are shown in Fig. 4. RBM-OMP-like algorithm is the only algorithm that uses both wavelets and overcomplete dictionaries in this experiment. The other algorithms only use wavelets as the sparsifying transform. The second experiment uses the same training and testing datasets of the first experiment.

The most commonly used reconstruction algorithm in the CS ECG literature is BPDN. MB-CoSaMP and MB-IHT exploit the connected subtree structure formed by the largest (in magnitude) wavelet coefficients. The BO-BSBL algorithm was previously employed to reconstruct non-invasive fetal ECG [7], however, it can also be successfully applied to the recovery of adult ECG due to the clustering property of the ECG wavelet coefficients that motivates the use of the block-sparsity model assumed in the BO-BSBL algorithm. For the low-measurement range ($M < 0.4N$), the RBM-OMP-like algorithm using wavelets outperforms MB-CoSaMP, MB-IHT, and the BPDN algorithms. The RBM-OMP-like algorithm using an overcomplete dictionary exhibits the best performance for $M < 0.55N$. It is only outperformed by the BO-BSBL algorithm when a large number of measurements are available, which is not the case of interest for WBAN-enabled ECG monitors. The IEEE 802.15.4 MAC protocol (Zigbee) is the standard for WBANs and is able to operate at 250 kbps (2.4 GHz), 40 kbps (915 MHz) and 20 kbps (868 MHz) [2]. The data rate of 20 kbps is of special interest since it leads to low energy consumption and extended battery life. In [21], it was reported that WBAN-enabled ECG typically have a raw data rate of 10-100 kbps. To satisfy the upper limit of 100 kbps with Zigbee operating at 20 kbps, $M/N = 0.2$ would be needed.

### 3.1.2. Qualitative Observations

This section presents a qualitative assessment of the trained RBM models. Fig. 5(a) illustrates the visible layer bias terms of the RBM using overcomplete dictionaries. Note that the bias terms are negative since the elements of the sparse representation are zero most of the time. Fig. 5(b) illustrates the dictionary atom sharing the same index as the largest bias term. Note the resemblance of the dictionary atom with the QRS complex of an ECG cycle.
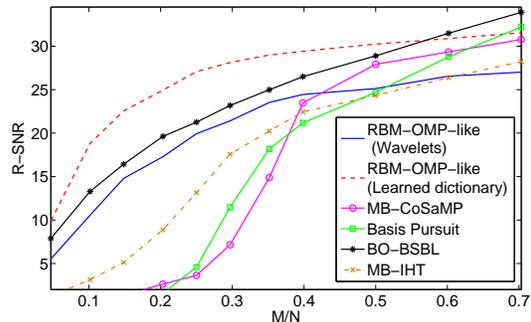


Figure 4: Comparison of the reconstruction of ECG signals using the RBM-OMP-like algorithm, Basis Pursuit Denoising, MB-CoSaMP, MB-IHT, and the BO-BSBL algorithm.

Fig. 6(a) illustrates the visible layer bias terms of the RBM using wavelets. Most of the bias terms are negative. Fig. 6(b) illustrates the magnitude of the wavelet representation of a segment from record 117, where the dashed lines separate each wavelet subband. Sparse representation coefficients sharing the same indexes as the largest bias terms of the visible layer are likely to belong to the best $K$-term sparse approximation. This is consistent with the result in Fig. 6(a) showing that the largest bias term indexes correspond to the indexes of the scaling wavelet coefficients (8 first indexes of the wavelet representation), which accumulate most of the ECG signal energy [9], and therefore, are likely to belong to the best $K$-term sparse approximation. Note that the magnitude of the wavelet coefficients tends to decay across scales. A similar behavior is observed for the bias terms of the visible layer (Fig. 6(a)). The detail coefficients of the ECG wavelet representation at the highest level are often disregarded as they are expected to have low magnitude [22, 12]. However, Fig. 6(a) shows that the 127th bias term, corresponding (in location) to the highest-level wavelet subband, has a large magnitude, which suggests that the wavelet coefficient with index 127 has a high *a priori* likelihood of belonging to the signal support.

The weight matrix of the RBM model reveals coefficient dependencies. Let $h_\star$ denote the hidden unit with the largest bias term of the RBM model using the learned overcomplete dictionary. Fig. 7(a) shows the weights associated with $h_\star$. This hidden unit is connected to a group of visible units via large positive weights, which tend to be
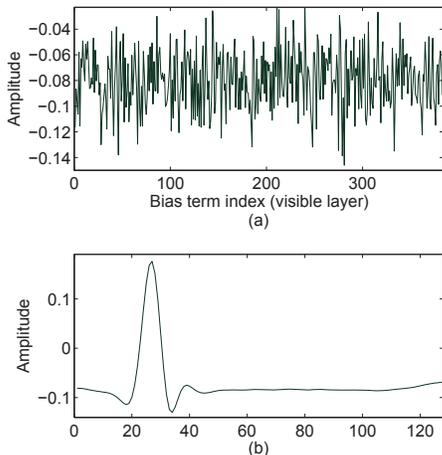
Figure 5: (a) Visible layer bias terms of the RBM model using learned overcomplete dictionaries. (b) Dictionary atom sharing the same index as the largest visible bias term.

Figure 6: (a) Visible layer bias terms of the RBM model using wavelets. (b) Wavelet representation of an ECG segment of 128 samples.

active simultaneously with $h_\star$, in order to lower the system energy. Figs. 7(b) and (c) show the dictionary atoms sharing the same indexes as the 2 visible units with the most positive connections to $h_\star$. Contrarily, Fig. 7(d) corresponds to the dictionary atom that shares the same index as the visible unit connected to $h_\star$ via the most negative weight. Note that both the patterns in Figs. 7(b) and (d) resemble the QRS complex, and therefore, are unlikely to happen simultaneously in an ECG segment of 128 samples (sampling frequency of 360 Hz).

### 3.1.3. QRS Detection Performance

The transmission and reconstruction of ECG signals is typically followed by other processing tasks, *e.g.* arrhythmia classification, whose first step is usually the detection of the QRS complexes. Therefore, it is worth evaluating how the performance of a QRS detection algorithm varies with the number of measurements. To achieve this goal, the threshold-independent QRS detector proposed in [23] is used. The ground truth for the R peaks is determined by the QRS detection provided by the algorithm when using the original ECG signal as input. The QRS detection performance of the reconstructed signals is measured using precision and recall. An R-peak detection in the reconstructed signal is classified as true positive when it is within ±4 samples from an R-peak detected in the origi-
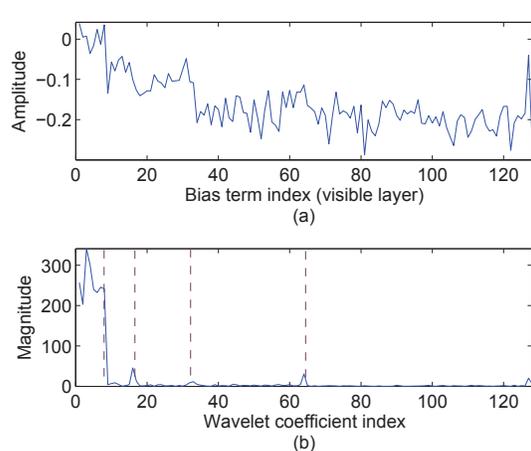
nal signal. Otherwise, it is classified as false positive. For example, if an R-peak is detected in the original signal at sample 200, a true positive detection in the reconstructed signal will be obtained if an R-peak is detected anywhere in the sample range [196, 204]. If none R-peaks in the reconstructed signal are detected within ±4 samples from an R-peak detected in the original signal, it will lead to a false negative result.

The same RBM model trained in Section 3.1.1 is employed for reconstruction in this experiment. Measurements are artificially contaminated with Gaussian noise of variance $\sigma_n^2 = 0.1$. The Daubechies-4 wavelet transform, using a decomposition level $L = 4$, and a sparsity threshold $K = 0.1N$ is employed as well. Records 103, 105, 106, 108, 112, 113 and 116 are selected to evaluate the performance of the QRS detector since they were not employed to train the RBM in Section 3.1.1. The selected ECG records are first sampled using a certain number of measurements and a sliding window of size $N = 128$, then reconstructed using the RBM-OMP-like algorithm. Finally, the reconstructed segments of length $N = 128$ are concatenated to reconstruct the original signal and the R-peaks are detected. Precision and recall results as a function of the number of measurements are presented in Fig. 8. High precision values, above 0.8, are obtained even for low number of measurements $M = 0.25N$ for all the selected ECG records. Similarly, high recall values, above
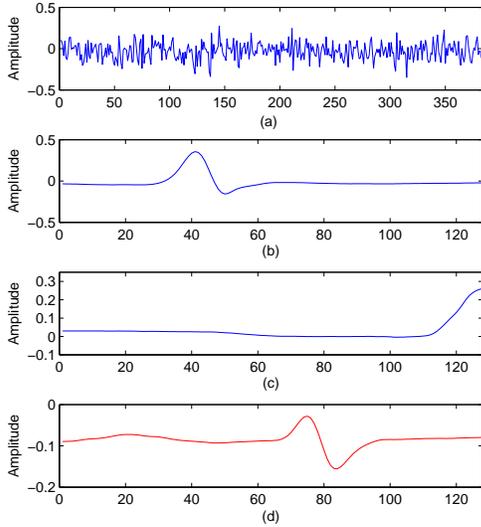
7

Figure 7: (a) Weights associated with the hidden unit with the largest bias term. (b-c) Dictionary atoms sharing the same indexes as the visible units with the 2 most positive weights. (d) Dictionary atom sharing the same index as the visible unit with the most negative weight.



Figure 8: Precision and recall performance for the detection of QRS complexes.

0.8, are obtained for $M \geq 0.2N$.

As a baseline for comparison, QRS detection is also performed with signals reconstructed using the BPDN algorithm. The same subset of ECG records, noise variance, RBM model, wavelet transform settings, sliding window size, and sparsity threshold used in the previous experiment are employed for the comparison. Precision and recall values are averaged across the selected ECG records for both the proposed algorithm and the BPDN algorithm. The results, shown in Table 1, indicate that the RBM-OMP-like algorithm significantly outperforms the BPDN algorithm in terms of QRS detection performance.

To the best of our knowledge, only the works in [14] and [20] address the problem of QRS detection using signals reconstructed with compressed sensing methods. However, the authors in [14] only reported QRS detection performance as a function of the number of measurements, but did not report the length of the signals. Without that information, it is not possible to estimate the $M/N$ ratio and compare [14] with our proposed metho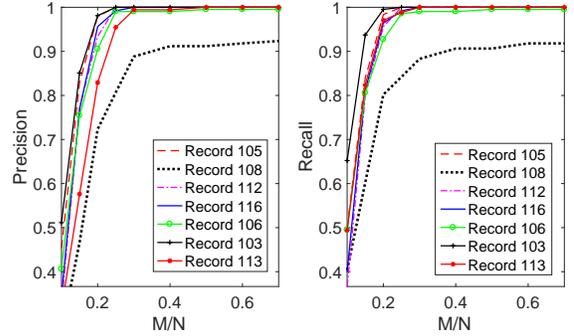d. The work in [20] reported PSim values for precision and re-call above 62% for a compression ratio of $N/M = 25$, using a sliding window of size 1024 and overcomplete patient-agnostic dictionaries. By using the proposed algorithm with overcomplete dictionaries, the same sliding window as in [20], and a compression ratio of $N/M = 30$, PSim values for precision and recall of 53% and 59% are obtained. The percentage of similarity metric depends on the employed QRS detection algorithm. A QRS detector different from that used in [20] is employed in this paper, which may have contributed to the difference in performance. In addition, the results in [20] were obtained with all the records from the MIT-BIH Arrhythmia Database, while our results correspond to the selected subset of ECG records.

The work in [24] evaluates QRS detection on reconstructed signals after using compression methods. A comparison with [24] is not entirely fair since compressed sensing is not a traditional compression method. Compressed sensing needs to use the compressed measurements to estimate not only the magnitude of the most significant coefficients, but also their positions. As a reference, it is reported that the work in [24] detects the QRS of reconstructed signals, which were previously compressed with a compression ratio of 4.5, with recall and precision values of 99.78% and 99.92%, respectively. They used all the 48 records from the MIT-BIH Arrhythmia Database to validate performance. Using our proposed method with a ratio $N/M = 4.5$ leads to averaged recall and precision values of 96.34% and 93.92%, respectively, across records 103, 105, 106, 108, 112, 113 and 116. The met-

Table 1: Comparison of the RBM-OMP-like algorithm with the BPDN algorithm in terms of QRS detection performance.

| Algorithm | Metric | M/N | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | 0.1 | 0.15 | 0.2 | 0.25 | 0.3 | 0.4 | 0.5 | 0.6 |
| RBM-OMP-like | Precision | 0.382 | 0.746 | 0.9 | 0.976 | 0.982 | 0.989 | 0.988 | 0.988 |
| | Recall | 0.465 | 0.833 | 0.937 | 0.981 | 0.983 | 0.989 | 0.988 | 0.988 |
| BPDN | Precision | 0.275 | 0.607 | 0.866 | 0.937 | 0.974 | 0.984 | 0.988 | 0.988 |
| | Recall | 0.338 | 0.68 | 0.911 | 0.967 | 0.977 | 0.984 | 0.987 | 0.987 |

ric values of the proposed method are smaller because, as indicated before, compressed sensing is not a traditional compression method, and it has been shown that it does not compare favorably with compression methods when only compression ratio is considered [25]. The value of compressed sensing is that it allows sub-Nyquist sampling. In addition, unlike the work in [24], the complete set of records from the MIT-BIH Arrhythmia Database was not used for testing in this experiment because some of the records had already been used for training.

### 3.1.4. Robustness to Noise

The third experiment evaluates the robustness of the RBM-OMP-like algorithm to measurement noise. An example of measurement noise is quantization noise. Even though it was assumed that the measurement noise was Gaussian in Section 3.1.1, it is assumed to be uniformly distributed in this section to simplify the calculations. Using this assumption, the noise variance of an uniform quantizer takes the form $\sigma_n^2 = (\Delta f)^2/(12 \times 2^{2m})$, where $\Delta f$ is the difference between the maximum and minimum values of the signal to be quantized and $m$ is the number of bits. Since analog-to-digital converters of 10 bits are often employed in WBANs [26], the number of bits $m$ is set to 10. Since $\Delta f$ for the compressed measurements calculated in Section 3.1.1 ranges from 37.5 to 8600, the noise variance varies approximately from $1 \times 10^{-4}$ to 5.88. Therefore, for this experiment, all the parameters are set to the same values as in the first experiment, except the measurement noise variance, which is varied in the range $[1 \times 10^{-4}, 5.88]$, while the Gaussian sensing matrix is kept fixed. Results are shown in Fig. 9. The successful performance of the RBM-OMP-like algorithm in the presence of sampling noise is expected as the noise variance is explicitly taken into account in the algorithm.
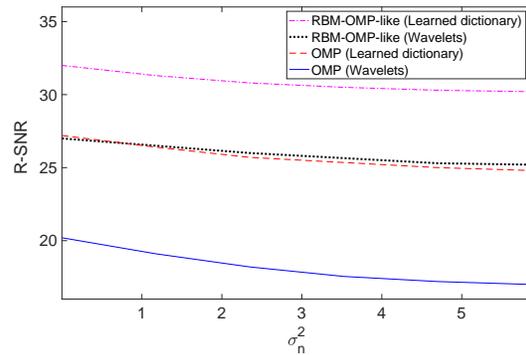


Figure 9: Performance of the RBM-OMP-like and OMP algorithms in the presence of measurement noise.

### 3.1.5. Parameter Evaluation

We evaluate the sensitivity of the RBM-OMP-like algorithm using overcomplete dictionaries to other parameters, such as the dictionary size and the sparsity threshold. In each experiment, we only changed the values of one parameter and kept the other parameters fixed. The default values are set the same as in the first experiment. Fig. 10 illustrates the performance of the RBM-OMP-like algorithm for dictionaries of size 2N, 3N and 4N. The increase in the number of atoms improves the capability of capturing details and thus leads to better representations.

Fig. 11 illustrates how the sparsity threshold affects the performance of the RBM-OMP-like algorithm when overcomplete dictionaries are employed as the sparsifying transform. For $M > 0.3N$, the R-SNR decreases rapidly as the sparsity threshold decreases from 0.2N to 0.05N. The reason is that useful information is lost when the sparsity threshold decreases. On the other hand, for $M < 0.2N$ the R-SNR tends to be slightly higher for smaller sparsity thresholds since sparser signals require fewer mea-
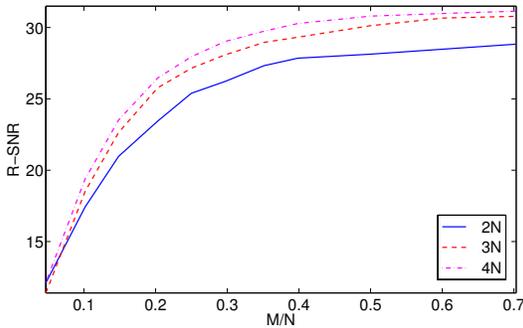
9

Figure 10: Reconstruction performance of the RBM-OMP-like algorithm using an overcomplete dictionary as a function of the dictionary size.
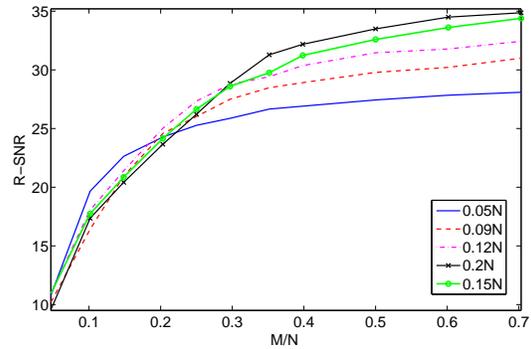


Figure 11: Reconstruction performance of the RBM-OMP-like algorithm using an overcomplete dictionary as a function of the sparsity threshold.



Figure 12: Evaluation of the reconstruction of the selected records from the European ST-T Database using RBM models with different number of hidden units and the wavelet transform as the sparsifying transform.

surements to yield stable reconstruction [5].

These results demonstrate that the RBM-OMP-like algorithm is not extremely sensitive to variations of the selected parameters. Even though we are aware that optimizing the parameters may lead to reconstruction improvement, how to perform the optimization is still an open problem that is beyond the scope of this paper.

### 3.2. Experiments on the European ST-T Database

The European ST-T Database consists of 90 annotated excerpts of ambulatory ECG recordings from 79 subjects. Each record is two hours long and contains two signals, each sampled at 250 samples per second with 12-bit resolution over a nominal 20 millivolt input range. Single leads from records 103, 104, 106, 107, 108, 110, 112, 118, 123, 129, 136, and 147 are employed for the experiments in this section. This data set consists of a variety of signals with changes in ST segment and T-wave morphology, rhythm, and signal quality.

Experiments on the European ST-T Database evaluate the performance of the RBM-OMP-like algorithm by varying the number of hidden units of the RBM employed to model the probability distribution of the sparsity pattern. The training and testing data sets for this experiment consist of 68000 and 16372 segments of size $N = 256$, extracted from the selected ECG records from the European ST-T Database, respectively. The same number of segments is extracted from each record. Compressed measurements are artificially contaminated with Gaussian noise of variance $\sigma_n^2 = 0.4$.
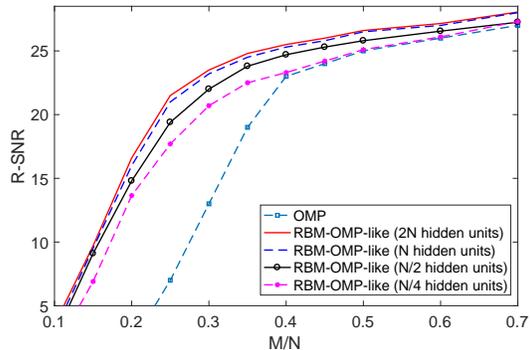
Figures 12 and 13 illustrate the performance of the RBM-OMP-like algorithm while varying the number of hidden units of the RBM model and using wavelets and learned overcomplete dictionaries as the sparsifying transform, respectively. The Daubechies-4 wavelet transform, using a decomposition level $L = 4$ and a sparsity threshold $K = 0.1N$, was employed to generate the results in Fig. 12. Similarly, a learned overcomplete dictionary with $3N$ dictionary atoms and a sparsity level $K = 0.12N$ is employed to generate the results in Fig. 13. The mean R-SNR across samples of the testing dataset increases as the number of hidden units changes from $P = N/4$ to $P = N$, respectively. Setting the number of hidden units to
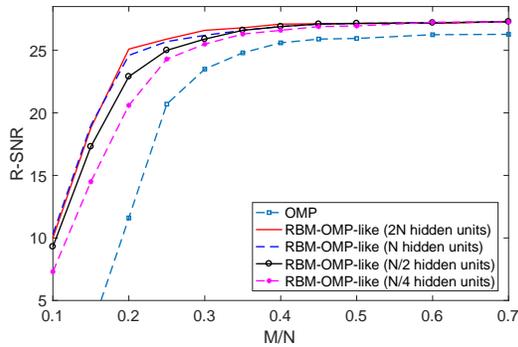
10

Figure 13: Evaluation of the reconstruction of the selected records from the European ST-T Database using RBM models with different number of hidden units and a learned overcomplete dictionary as the sparsifying transform.

$P = 2N$ improves the reconstruction performance slightly compared to the case when $P = N$.

## 4. Conclusions

The potential of CS to lower energy consumption in WBAN-enabled ECG monitors lies in its ability to reduce the number of samples at the sensor node. With the goal of further reducing the number of samples, we proposed to use the RBM-CS scheme for ECG reconstruction. This scheme fully exploits signal structure in a statistical fashion, and subsequently, leads to stable reconstruction with low number of measurements. Even though other CS algorithms that exploit signal structure have been applied to ECG reconstruction, they all assume a specific type of structure, *e.g.*, tree-structured sparsity model and block-sparsity model, which is also tied to a specific sparsifying transform. Instead, the RBM-OMP like algorithm uses a general model that accounts for a large set of structures and sparsifying transforms. It was shown that significant performance gains in the low-measurement regime can be obtained by using the RBM-OMP-like algorithm relative to state-of-the-art algorithms for ECG reconstruction. It was also shown that the algorithm is robust to quantization measurement noise and that the detection of QRS complexes on the reconstructed signals remains accurate for $M/N \geq 0.25$. Simulations using real ECG data from the MIT-BIH Arrhythmia database and the the European ST-T database revealed that the performance of the algorithm is better when using learned overcomplete dictionaries than when using wavelets.

## References

[1] H. Cao, V. Leung, C. Chow, H. Chan, Enabling technologies for wireless body area networks: A survey and outlook, IEEE Communications Magazine 47 (12) (2009) 84–93.

[2] M. Seyedi, B. Kibret, D. Lai, M. Faulkner, A survey on intrabody communications for body area network applications, IEEE Transactions on Biomedical Engineering, 60 (8) (2013) 2067–2079.

[3] H. Mamaghanian, N. Khaled, D. Atienza, P. Vandergheynst, Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes, IEEE Transactions on Biomedical Engineering 58 (9) (2011) 2456–2466.

[4] D. Donoho, Compressed sensing, IEEE Transactions on Information Theory 52 (4) (2006) 1289–1306.

[5] E. Candès, M. Wakin, An introduction to compressive sampling, IEEE Signal Processing Magazine 25 (2) (2008) 21–30.

[6] A. Dixon, E. Allstot, D. Gangopadhyay, D. Allstot, Compressed sensing system considerations for ECG and EMG wireless biosensors, IEEE Transactions on Biomedical Circuits and Systems 6 (2) (2012) 156 –166.

[7] Z. Zhang, T. Jung, S. Makeig, B. Rao, Compressed sensing for energy-efficient wireless telemonitoring of noninvasive fetal ECG via block sparse Bayesian learning, IEEE Transactions on Biomedical Engineering 60 (2) (2013) 300–309.

[8] H. Mamaghanian, N. Khaled, D. Atienza, P. Vandergheynst, Structured sparsity models for compressively sensed electrocardiogram signals: A comparative study, in: 2011 IEEE Biomedical Circuits and Systems Conference, 2011, pp. 125–128.

[9] L. Polania, R. Carrillo, M. Blanco-Velasco, K. Barner, Compressed sensing based method for ECG compression, in: Proc., IEEE ICASSP, 2011, pp. 761 –764.

[10] J. Pant, S. Krishnan, Compressive sensing of electrocardiogram signals by promoting sparsity on the second-order difference and by using dictionary learning, IEEE Transactions on Biomedical Circuits and Systems 8 (2) (2014) 293–302.

[11] J. Zhang, Z. Gu, Z. Yu, Y. Li, Energy-efficient ECG compression on wireless biosensors via minimal coherence sensing and weighted minimization reconstruction, IEEE Journal of Biomedical and Health Informatics 19 (2) (2015) 520–528.

[12] L. Polania, R. Carrillo, M. Blanco-Velasco, K. Barner, Exploiting prior knowledge in compressed sensing wireless ECG systems, IEEE Journal of Biomedical and Health Informatics 19 (2) (2015) 508–519.

[13] L. Polanía, K. Barner, Exploiting restricted Boltzmann machines and deep belief networks in compressed sensing, IEEE Transactions on Signal Processing 65 (17) (2017) 4538–4550.

[14] M. Duarte, Y. Eldar, Structured compressed sensing: From theory to applications, IEEE Transactions on Signal Processing 59 (9) (2011) 4053–4085.

[15] A. Goldberger, et al., Physiobank, Physiotoolkit, and Physionet: Components of a new research resource for complex physiologic signals, Circulation 101 (23) (2000) e215–e220.

[16] A. Taddei, G. Distante, M. Emdin, P. Pisani, G. Moody, C. Zeelenberg, C. Marchesi, The European ST-T database: standard for evaluating systems for the analysis of ST-T changes in ambulatory electrocardiography, European heart journal 13 (9) (1992) 1164–1172.

[17] M. Aharon, M. Elad, A. Bruckstein, K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation, IEEE Transactions on Signal Processing 54 (11) (2006) 4311–4322.

[18] G. Hinton, Training products of experts by minimizing contrastive divergence, Neural Comput. 14 (8) (2002) 1771–1800.

[19] F. Chen, A. Chandrakasan, V. Stojanovic, Design and analysis of a hardware-efficient compressed sensing architecture for data compression in wireless sensors, IEEE Journal of Solid-State Circuits 47 (3) (2012) 744–756.

[20] D. Craven, B. McGinley, L. Kilmartin, M. Glavin, E. Jones, Impact of compressed sensing on clinically relevant metrics for ambulatory ECG monitoring, Electronics Letters 51 (4) (2015) 323–325.

[21] R. Chávez, A. Khaleghi, I. Balasingham, T. Ramstad, Architecture of an ultra wideband wireless body area network for medical applications, in: International Symposium on Applied Sciences in Biomedical and Communication Technologies, 2009, pp. 1–6.

[22] L. Sharma, Coding ECG beats using multiscale compressed sensing based processing, Computers & Electrical Engineering 45 (2015) 211–221.

[23] A. Ramakrishnan, A. Prathosh, T. Ananthapadmanabha, Threshold-independent QRS detection using the dynamic plosion index, IEEE Signal Processing Letters 21 (5) (2014) 554–558.

[24] M. Elgendi, A. Mohamed, R. Ward, Efficient ECG compression and QRS detection for e-health applications, Scientific reports 7 (1) (2017) 459.

[25] D. Craven, B. McGinley, L. Kilmartin, M. Glavin, E. Jones, Compressed sensing for bioelectric signals: A review, IEEE journal of biomedical and health informatics 19 (2) (2015) 529–540.

[26] K. Thotahewa, J. Redouté, M. Yuce, System implementation and evaluation of an energy efficient UWB-based MAC protocol for wireless body area networks, in: Ultra Wideband Wireless Body Area Networks, 2014, pp. 117–144.