



# Compressive sensing of ECG signals using plug-and-play regularization<sup>☆</sup>

V. S. Unni\*, Raturaj G. Gavaskar, Kunal N. Chaudhury

Department of Electrical Engineering, Indian Institute of Science, Bangalore, 560012, India

## ARTICLE INFO

### Article history:

Received 6 January 2022

Revised 6 July 2022

Accepted 17 August 2022

Available online 19 August 2022

### Keywords:

ECG signals

Compressive sensing

Proximal gradient descent

Plug-and-play regularization

GMM denoiser

## ABSTRACT

Compressive Sensing (CS) has recently attracted attention for ECG data compression. In CS, an ECG signal is projected onto a small set of random vectors. Recovering the original signal from such compressed measurements remains a challenging problem. Traditional recovery methods are based on solving a regularized minimization problem, where a sparsity-promoting prior is used. In this paper, we propose an alternative iterative recovery algorithm based on the Plug-and-Play (PnP) method, which has recently become popular for imaging problems. In PnP, a powerful denoiser is used to implicitly perform regularization, instead of using hand-crafted regularizers; this has been found to be more successful than traditional methods. In this work, we use a PnP version of the Proximal Gradient Descent (PGD) algorithm for ECG recovery. To ensure mathematical convergence of the PnP algorithm, the signal denoiser in question needs to satisfy some technical conditions. We use a high-quality ECG signal denoiser fulfilling this condition by learning a Bayesian prior for small-sized signal patches. This guarantees that the proposed algorithm converges to a fixed point irrespective of the initialization. Importantly, through extensive experiments, we show that the reconstruction quality of the proposed method is superior to that of state-of-the-art methods.

© 2022 Elsevier B.V. All rights reserved.

## 1. Introduction

Electrocardiogram (ECG) is widely used for diagnosis and monitoring cardiac conditions such as hypertension [1], heart failure [2], arrhythmia [3] etc. Essentially, an ECG signal is a representation of the electric activity in the heart over time. Extensive use of ECG in healthcare fosters a need for sophisticated signal processing approaches to efficiently compress, analyze, store and transmit ECG signals. For example, in wearable devices, there is a need to reduce energy consumption due to data transmission and to increase memory usage efficiency. Recently, several efforts have been made to develop wireless ECG sensors for continuous health monitoring, for which it is desirable to have devices with low power consumption or low complexity. However, continuous wireless transmission of long-term biomedical data consumes a significant amount of energy. Thus, compression of ECG signals would be helpful to achieve energy efficiency.

<sup>☆</sup> K. N. Chaudhury was partially supported by Core Research Grant CRG/2020/000527 and SERB-STAR Award STR/2021/000011 from the Department of Science and Technology, Government of India.

\* Corresponding author.

E-mail addresses: [unniv@iisc.ac.in](mailto:unniv@iisc.ac.in) (V. S. Unni), [raturajg@iisc.ac.in](mailto:raturajg@iisc.ac.in) (R.G. Gavaskar), [kunal@iisc.ac.in](mailto:kunal@iisc.ac.in) (K.N. Chaudhury).

Compressive Sensing (CS) is a possible solution for signal compression. It employs random linear projections that aim to preserve the structure of the signal. The signal can be reconstructed from its projections using nonlinear recovery methods. In fact, several works have explored the application of CS to biomedical signal processing, including ECG [4–6], EMG [7,8], EEG [9] signals and MRI images [10].

### 1.1. Compressive sensing of ECG signals

The data acquisition model in CS is given by Zhang et al. [5], Pant and Krishnan [6]

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{n} \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^N$  is the original ECG signal having length  $N$ ,  $\Phi \in \mathbb{R}^{M \times N}$  is a compression matrix with  $M \ll N$ , and  $\mathbf{n}$  denotes the noise in the acquisition system. Typically,  $\mathbf{n}$  is assumed to be white Gaussian noise with mean 0, whereas  $\Phi$  is taken to be a random Gaussian or binary matrix [4–6].

The original signal  $\mathbf{x}$  can be approximately recovered by solving the regularized inversion problem

$$\min_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}) + \lambda g(\mathbf{x}), \quad (2)$$

where the term  $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|^2$  forces consistency of the recovered signal w.r.t. the measurements, whereas  $g(\mathbf{x})$  (known as the regularizer) acts as a penalty function that forces the recovered signal to have some desirable properties such as smoothness. Here,  $\lambda$  is a positive scalar used to control the amount of regularization and  $\|\cdot\|$  denotes the  $\ell_2$  norm.

A good regularizer  $g(\mathbf{x})$  is necessary since recovering  $\mathbf{x}$  from  $\mathbf{y}$  is an ill-posed problem (as  $M \ll N$ ). Several regularizers have been explored for the ECG compressive recovery task, such as the weighted  $\ell_1$  norm [11], various other  $\ell_p$  norms [12], total variation (TV) [13], and second-order sparsity-promoting functions [6]. Moreover, efficient recovery algorithms have been derived by exploiting the temporal correlation between successive samples; examples include  $\ell_p^{1d}$ -regularized least-squares [14], Block Sparse Bayesian Learning Bound-Optimization (BSBL-BO) [15], and Block Sparse Bayesian Learning with Expectation Maximization (BSBL-EM) [16]. The latter two are considered to be state-of-the-art.

### 1.2. Classical regularization

It is well-known that the  $\ell_1$  norm promotes sparse solutions; moreover, natural signals such as ECG are known to be approximately sparse in suitably chosen domains [11,13]. Therefore, sparsity-promoting regularizers based on the  $\ell_1$  norm in the wavelet and gradient domains (TV) have traditionally been used for ECG reconstruction. The downside is that the  $\ell_1$  norm is not differentiable. Hence, the objective function in (2) as a whole is non-differentiable, even though  $f(\mathbf{x})$  is differentiable. A good choice of iterative numerical solvers to solve such problems is the class of *proximal algorithms* [17], such as ADMM and Proximal Gradient Descent (PGD). A proximal algorithm generally consists of smaller subproblems which individually involve only one of the two functions. Consequently, it is possible to take advantage of the differentiability of  $f(\mathbf{x})$ . In this paper, we focus on PGD since it is a particularly simple proximal algorithm. PGD is sometimes known as the Iterative Shrinkage-Thresholding Algorithm (ISTA) [18]. Starting from an initial point  $\mathbf{x}_0 \in \mathbb{R}^N$ , PGD creates a sequence of points  $\mathbf{x}_1, \mathbf{x}_2, \dots$  recursively using the rule

$$\mathbf{x}_{k+1} = \text{prox}_{\gamma \lambda g}(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)), \quad (3)$$

where  $\gamma > 0$  is a fixed parameter (known as the step size) and  $\text{prox}_{\gamma \lambda g}(\cdot)$  is a function known as the *proximal operator* of  $g$ :

$$\text{prox}_{\gamma \lambda g}(\mathbf{z}) = \underset{\mathbf{w} \in \mathbb{R}^N}{\text{argmin}} \left[ \frac{1}{2} \|\mathbf{w} - \mathbf{z}\|^2 + \gamma \lambda g(\mathbf{w}) \right]. \quad (4)$$

Note that if we put  $\Phi = \mathbf{I}$  in (2), then (2) reduces to (4). Thus, the proximal operator can be interpreted effectively as a Gaussian denoising operator. For regularizers such as the  $\ell_0$  and  $\ell_1$  norms, the proximal operator has a closed-form formula [17]. Hence, the PGD algorithm is easy to implement. As discussed in Section 2, the algorithm is guaranteed to converge under some mild conditions to a minimum of  $f(\mathbf{x}) + \lambda g(\mathbf{x})$ . Note that every PGD step can be seen as the composition of two operations: the first (computing  $\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)$ ) is effectively one step of the classical gradient descent algorithm and depends only on the function  $f$ , while the second (computing the proximal operator) depends only on  $g$ . This is why the algorithm is named as Proximal Gradient Descent.

### 1.3. Plug-and-play regularization

Plug-and-play (PnP) regularization is a novel regularization technique developed in the image processing community [19,20]. The main step in PnP is to replace the proximal operator by a powerful signal denoiser. As discussed in Section 2, this is due to the similarity of the proximal operator with a denoising operation. In the context of PGD, the function  $\text{prox}_{\gamma \lambda g}(\cdot)$  is replaced by

a signal denoiser  $D(\cdot)$ , so that the  $k$ th step now becomes  $\mathbf{x}_{k+1} = D(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k))$ . This algorithm is known as PnP-PGD. Essentially, this amounts to taking one step of gradient descent, followed by denoising. Note that we no longer need to choose a regularizer  $g(\mathbf{x})$  since  $g$  does not appear in the modified algorithm; the regularization is performed implicitly by the denoiser.

PnP has yielded state-of-the-art results in many imaging problems. However, since there is no regularizer  $g$  involved, the aforementioned convergence to a minimum of  $f(\mathbf{x}) + \lambda g(\mathbf{x})$  does not apply.

### 1.4. Contribution

The contributions of this work are as follows.

1. We introduce the PnP framework for reconstructing ECG signals from CS measurements. To the best of our knowledge, PnP has never been used for this application. Through extensive experiments, we show that the proposed method, based on PnP-PGD, outperforms the current state-of-the-art CS recovery methods for ECG signals.
2. Even though convergence of PGD to a minimum of  $f + \lambda g$  is not applicable to PnP-PGD, we show that a different form of convergence, known as *fixed-point convergence*, can be guaranteed if the ECG denoiser  $D(\cdot)$  satisfies a condition known as *contractivity*. Thus, the challenge lies in designing a contractive ECG denoiser.
3. We derive a high-quality contractive ECG denoiser  $D(\cdot)$  by modeling small patches as random vectors following a Gaussian Mixture Model (GMM). We experimentally show that the denoising performance of the GMM denoiser is comparable or better than existing state-of-the-art ECG denoisers.

The rest of this paper is organized as follows. In Section 2, we give an overview of the PnP-PGD algorithm and explain the motivation behind its development. We derive the GMM denoiser in Section 3 and compare its denoising quality with existing ECG signal denoisers. In Section 4, we discuss how the this denoiser can be used in a way that guarantees convergence of PnP-PGD for CS recovery. Numerical experiments on CS recovery of ECG signals are shown in Section 5, and we conclude the paper in Section 6. Some of the mathematical proofs are given in the Appendix.

## 2. Plug-and-Play PGD

We first state a standard convergence result for PGD and then move on to discuss some convergence-related aspects of PnP. What makes PGD a simple yet powerful algorithm is its guarantee of convergence to a minimum of  $f(\mathbf{x}) + \lambda g(\mathbf{x})$ . In the following theorem (and the rest of the paper), we denote the largest singular value of a matrix by  $\sigma_{\max}(\cdot)$ .

**Theorem 1** [18]. *Consider the PGD algorithm for minimizing the function  $f(\mathbf{x}) + \lambda g(\mathbf{x})$ , where  $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|^2$ . Suppose  $g$  is continuous and convex, and that  $0 < \gamma < 2/\sigma_{\max}(\Phi^T \Phi)$ . Then the sequence  $f(\mathbf{x}_k) + \gamma g(\mathbf{x}_k)$  converges to the minimum of  $f(\mathbf{x}) + \lambda g(\mathbf{x})$  as  $k \rightarrow \infty$ .*

We now turn our attention to the PnP framework. Consider the definition of  $\text{prox}_{\gamma \lambda g}(\mathbf{z})$  in (4). Note that the minimization problem in (4) is similar to (2) if we put  $\Phi = \mathbf{I}$ , the identity matrix. Hence,  $\text{prox}_{\gamma \lambda g}(\mathbf{z})$  is essentially a regularized inverse corresponding to the additive noise model  $\mathbf{z} = \mathbf{w} + \mathbf{n}$ , where  $\mathbf{n}$  is Gaussian noise. Thus, the proximal operator is simply an additive Gaussian denoising operator.

It is well-known in the image processing community that specially designed denoisers such as nonlocal means (NLM) [21] and

BM3D [22] are superior to traditional denoisers based on regularization, e.g.  $\ell_1$  or TV-regularized denoising. Motivated by this observation, the work in [19] explored how the performance of a proximal algorithm for image recovery problems is affected if we replace  $\text{prox}_{\gamma\lambda g}(\cdot)$  by some arbitrary Gaussian denoiser  $D(\cdot)$ , such as NLM or BM3D. This scheme was named as plug-and-play, since the denoiser  $D$  serves as a pluggable module that replaces the proximal operator in an already existing numerical solver. In the original work [19], the PnP scheme was explored for a different proximal algorithm – ADMM – but it was adapted to PGD in [20]. The PnP-PGD algorithm is thus recursively defined by

$$\mathbf{x}_{k+1} = D(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)). \quad (5)$$

Note that the same denoiser  $D$  can be utilized for several kinds of image recovery problems using the PnP framework, since only the function  $f$  changes from problem to problem. For this reason, in the past few years, PnP has gained a lot of interest in the imaging community. However, the use of PnP for recovering one-dimensional signals such as ECG signals has remained an unexplored territory.

An immediate question that arises from the PnP scheme is as follows: Does the sequence  $(\mathbf{x}_k)$  converge to some  $\mathbf{x}^*$ ? And if so, is  $\mathbf{x}^*$  optimal in some sense? The latter question can be resolved if  $D(\cdot)$  is expressible as the proximal operator of some function  $g$ . In general, however, an arbitrary  $D$  cannot be expressed in this way. As a result, the PnP-PGD algorithm cannot be interpreted as minimizing an objective function of the form  $f + \lambda g$ , and the convergence result in Theorem 1 is not generally applicable. Therefore, we are left with trying to determine whether at least the sequence  $(\mathbf{x}_k)$  converges. It turns out that such a guarantee can indeed be given under a technical condition on  $D$ .

**Definition 1.** The denoiser  $D : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is said to be contractive if there exists  $\delta \in [0, 1)$  such that for all points  $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^N$ ,

$$\|D(\mathbf{z}_1) - D(\mathbf{z}_2)\| \leq \delta \|\mathbf{z}_1 - \mathbf{z}_2\|$$

We now state a theorem that guarantees the convergence of PnP-PGD using a contractive denoiser.

**Theorem 2.** Consider the PnP-PGD algorithm,  $\mathbf{x}_{k+1} = D(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k))$ , where  $f(\mathbf{x}) = (1/2) \|\Phi \mathbf{x} - \mathbf{y}\|^2$ . Suppose  $0 < \gamma \leq 2/\sigma_{\max}(\Phi^T \Phi)$ . Moreover, suppose the denoiser  $D$  is contractive. Then, as  $k \rightarrow \infty$ , the sequence  $\mathbf{x}_1, \mathbf{x}_2, \dots$  converges linearly (at an exponential rate) to a unique fixed point  $\mathbf{x}^*$  that does not depend on the initialization  $\mathbf{x}_0$ .

While Theorem 2 is proved in the Appendix, we mention here that the proof uses the Banach Fixed Point Theorem [23, Th. 9.23], from which the linear rate of convergence follows.

Note the difference between the types of convergence addressed in Theorems 1 and 2: Theorem 1 claims the convergence of the sequence of objective function values  $f(\mathbf{x}_k) + \lambda g(\mathbf{x}_k)$ , whereas Theorem 2 is concerned with the sequence of variables  $\mathbf{x}_k$ . Theorem 2 essentially claims that the PnP-PGD algorithm eventually stabilizes, in the sense that two consecutive iterates  $\mathbf{x}_k$  and  $\mathbf{x}_{k+1}$  are close to each other. This property is known as fixed-point convergence [24], and is desirable for any recovery algorithm. Thus, by Theorem 2, it is sufficient for the denoiser  $D$  to be contractive, in order to have fixed-point convergence.

It is useful to compare our result with a similar result in a recent work [25]. In [25], the authors proved fixed-point convergence of PnP-PGD under a different set of assumptions than ours. The convergence result in [25] is applicable to the case where the loss function  $f(\mathbf{x})$  is strongly convex and  $D - I$  is Lipschitz continuous, where  $I$  is the identity operator. In contrast, we require  $D$  to be contractive and we do not require strong convexity of  $f$ . In fact, in

our case,  $f$  is not strongly convex since the sensing matrix  $\Phi$  has a non-trivial null space.

Various methods for Gaussian denoising of ECG signals have been explored in the literature; see [26] for a review. The state-of-the-art techniques are optimization-based, e.g. TV, multiresolution analysis methods such as wavelets, empirical mode decomposition methods etc. A combination of these methods is sometimes used [27]. Further, nonlocal means (NLM) denoising has also been found to be promising [28]. However, to the best of our knowledge, there is no work that determines whether any of these denoisers are contractive. Can we design a high-quality contractive ECG signal denoiser? In the next section, we show that this can indeed be done. Specifically, we design a Gaussian denoiser which takes the form  $D(\mathbf{z}) = \mathbf{W}\mathbf{z}$ , where  $\mathbf{W}$  is a symmetric matrix. The resulting PnP-PGD algorithm outperforms state-of-the-art methods for ECG.

### 3. GMM Denoiser

Our ECG denoiser is inspired from an observation that was made in the context of images [29,30]: A small patch of some fixed size belonging to a clean (i.e. noiseless) image can be well-modeled as a random vector having a Gaussian Mixture Model (GMM) as its density. Such a density can be learned by fitting a GMM to a large collection of patches extracted from a set of clean images, usually belonging to a common class (e.g. face images). We apply this idea to model patches belonging to ECG signals. Specifically, we extract a large collection of patches of length  $P \ll N$  from a set of noiseless ECG signals as our training data set. This is used to fit a GMM density (with a pre-determined number of components  $K$ ) using the expectation-maximization (EM) algorithm. Essentially, we model clean ECG patches of length  $P$  as random vectors  $\mathbf{v} \in \mathbb{R}^P$  drawn from this learned GMM density, which we denote by  $p(\mathbf{v})$ . For  $j = 1, \dots, K$ , let  $\alpha_j \geq 0$  be the mixture coefficient of the  $j$ th component,  $\boldsymbol{\mu}_j \in \mathbb{R}^P$  be the mean and  $\boldsymbol{\Sigma}_j$  be the positive definite covariance matrix. Then  $p(\mathbf{v})$  is given by


$$p(\mathbf{v}) = \sum_{j=1}^K \alpha_j \mathcal{N}(\mathbf{v}; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j), \quad (6)$$

where  $\mathcal{N}$  denotes a Gaussian density function.

How can this model be used for denoising an ECG signal corrupted with Gaussian noise? Again, we borrow from a patch-based denoising framework that is quite popular in image processing [29,30]. At a high level, this framework consists of the following steps:

1. Extract all possible patches of length  $P$  from the noisy signal; if the signal length is  $N$ , then there are  $N$  such patches (we apply circular padding to the signal). If  $\mathbf{z}$  denotes the noisy signal, the collection of patches is given by  $\mathbf{P}_1 \mathbf{z}, \dots, \mathbf{P}_N \mathbf{z}$ , where  $\mathbf{P}_i : \mathbb{R}^N \rightarrow \mathbb{R}^P$  is the linear operator that extracts the patch starting at the  $i$ th location. This is defined as the segment  $(z_i, z_{i+1}, \dots, z_{i+P-1})$ .
2. Denoise each patch independently by computing a Bayesian estimate of its corresponding clean patch under an additive Gaussian noise model, using  $p(\mathbf{v})$  as the prior distribution of clean patches. Letting  $G$  denote the denoising operator, the collection of denoised patches is given by  $G(\mathbf{P}_1 \mathbf{z}), \dots, G(\mathbf{P}_N \mathbf{z})$ .
3. Place each denoised patch back into its corresponding location in the signal. Each sample location  $i \in \{1, \dots, N\}$  belongs to  $P$  overlapping denoised patches; take the average of the  $P$  values at this location as the estimate of the  $i$ th sample of the denoised signal. This completes the overall denoising process, which is given by

$$D(\mathbf{z}) = \frac{1}{P} \sum_{i=1}^N \mathbf{P}_i^T G(\mathbf{P}_i \mathbf{z}). \quad (7)$$



**Fig. 1.** Comparison of the denoising performance using various state-of-the-art signal denoisers on an ECG signal of length  $N = 200$ . The clean signal (green) is overlaid in every plot for comparison. Note that the denoised signal using GMM has more structural similarity with the clean signal. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The patch denoiser  $G$  forms the core of this framework, and the overall denoising performance depends on the performance of  $G$ . One approach to incorporate the Bayesian prior  $p(\cdot)$  in the patch denoising is to take  $G(\cdot)$  as the maximum a-posteriori (MAP) estimator of the clean patch under a Gaussian noise model. That is, for a noisy patch  $\mathbf{u} = \mathbf{v} + \mathbf{n}$  (where  $\mathbf{n}$  is Gaussian noise), we can define  $G(\mathbf{u})$  to be the mode of the conditional density  $p(\mathbf{v}|\mathbf{u})$ . However, it is known that this cannot be computed in closed form when the prior  $p(\mathbf{v})$  is a GMM [29]. Instead, motivated by Teodoro et al. [30], we choose  $G$  to be the minimum mean-squared-error (MMSE) estimator of the clean patch:

$$G(\mathbf{u}) = \mathbb{E}[\mathbf{v}|\mathbf{u}]. \quad (8)$$

The theorem below gives a closed-form expression for  $\mathbb{E}[\mathbf{v}|\mathbf{u}]$ .

**Theorem 3 [30].** Consider the additive noise model  $\mathbf{u} = \mathbf{v} + \mathbf{n}$ , where  $\mathbf{n}$  is zero-mean Gaussian noise having variance  $\sigma^2$ . Suppose  $\mathbf{v}$  has the GMM density given by (6). Then,

$$G(\mathbf{u}) = \left( \sum_{j=1}^K \beta_j(\mathbf{u}) \mathbf{C}_j \right) \mathbf{u},$$

where  $\mathbf{C}_j = \boldsymbol{\Sigma}_j (\boldsymbol{\Sigma}_j + \sigma^2 \mathbf{I})^{-1}$  and

$$\beta_j(\mathbf{u}) = \frac{\alpha_j \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j + \sigma^2 \mathbf{I})}{\sum_{l=1}^K \alpha_l \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_l, \boldsymbol{\Sigma}_l + \sigma^2 \mathbf{I})}.$$

To summarize, the overall GMM denoiser  $D$  is given by

$$D(\mathbf{z}) = \frac{1}{P} \sum_{i=1}^N \mathbf{P}_i^T \left[ \left( \sum_{j=1}^K \beta_j(\mathbf{P}_i \mathbf{z}) \mathbf{C}_j \right) \mathbf{P}_i \mathbf{z} \right]. \quad (9)$$

In order to gauge the quality of the GMM denoiser, we perform a denoising experiment on a noisy ECG signal. Specifically, we compare its performance with the following ECG signal denoising schemes: TV [31], NLM [28], and wavelet- $\ell_1$  regularization. In Figure 1, we show a segment of signal #115 (assumed noiseless) from the Physionet MIT-BIH Arrhythmia Database [32–34]. The segment has length  $N = 200$ . We add white Gaussian noise such that the signal-to-noise ratio (SNR) of the noisy signal is 30 dB. The SNR for an estimated signal  $\hat{\mathbf{x}}$  (here, the denoised signal) with respect to a reference signal  $\mathbf{x}$  (here, the clean signal) is defined as

$$\text{SNR} = 10 \log_{10} \left( \frac{\|\mathbf{x}\|^2}{\|\mathbf{x} - \hat{\mathbf{x}}\|^2} \right).$$

**Table 1**

Denoising performance (SNR in dB). The signal length is  $N = 200$ .

Noisy	TV	NLM	Wavelet	GMM
15	25.721	26.257	23.216	<b>27.492</b>
20	26.901	27.023	25.861	<b>28.373</b>
25	29.252	29.415	29.157	<b>29.646</b>
30	33.081	33.139	31.396	<b>33.819</b>
35	35.322	35.892	32.699	<b>36.276</b>
40	40.391	40.455	33.681	<b>41.262</b>

A higher SNR value indicates a better estimation quality. The denoised signals obtained using the aforementioned denoising schemes are shown in Fig. 1. Observe that the GMM denoiser yields the highest SNR of all the methods; the visual quality is considerably better compared to TV and wavelet, and comparable to NLM.

For a more extensive comparison, we repeat this experiment for different SNR values of the noisy signal. The SNR values of the denoised signal are noted in Table 1. Again, it is observed that the GMM denoiser outperforms the other denoisers, while NLM is the second-best method. In fact, for high noise levels (SNR of 15 and 20 dB), the gap in performance between GMM and NLM is quite high. A possible explanation is that when the noise level is high, reliable computation of weights for NLM is difficult and can result in spikes in the denoised signal.

#### 4. Convergence analysis

Recall that we would like  $D$  to be contractive; however, due to the complexity of the expression in (9), it is difficult to determine whether this is the case. Fortunately, while using it as part of the larger PnP-PGD framework, we can modify the denoiser to make it contractive using a simple trick. Note that the coefficients  $\beta_j(\mathbf{P}_i \mathbf{z})$  in (9) are nonnegative and sum to 1 for each  $i$ . Consider the situation where we replace the  $\beta_j(\mathbf{P}_i \mathbf{z})$ 's by some fixed universal constants  $b_{ji}$  that do not depend on  $\mathbf{z}$ , but have the same properties:  $b_{ji} \geq 0$  for all  $i, j$ , and  $\sum_j b_{ji} = 1$  for all  $i$ . Then  $D(\mathbf{z})$  becomes a linear function of  $\mathbf{z}$ . In fact, we can write  $D(\mathbf{z}) = \mathbf{W}\mathbf{z}$ , where  $\mathbf{W} \in \mathbb{R}^{N \times N}$  is given by

$$\mathbf{W} = \frac{1}{P} \sum_{i=1}^N \mathbf{P}_i^T \left( \sum_{j=1}^K b_{ji} \mathbf{C}_j \right) \mathbf{P}_i. \quad (10)$$

**Theorem 4.** Let  $\mathbf{W}$  be defined as in (10), where  $b_{ji} \geq 0$  for all  $i, j$ , and  $\sum_{j=1}^K b_{ji} = 1$  for all  $i$ . Suppose  $N$  is a multiple of  $P$ . Then the largest eigenvalue of  $\mathbf{W}$ ,  $\lambda_{\max}(\mathbf{W})$ , is  $< 1$ . Consequently, the denoiser  $D(\mathbf{z}) = \mathbf{W}\mathbf{z}$  is contractive, with the constant  $\delta$  being  $\lambda_{\max}(\mathbf{W})$ .

The proof is given in the appendix. Note that the requirement for  $N$  to be a multiple of  $P$  is not too restrictive, since we can pad the signal if it is not. We only need to find suitable coefficients  $b_{ij}$  to replace  $\beta_j(\mathbf{P}_i\mathbf{z})$  in (9). This can be done as follows. We first run a small number  $T$  (say,  $T = 10$  or  $20$ ) of PnP-PGD iterations using the coefficients  $\beta_j(\mathbf{P}_i\mathbf{z})$  in the denoiser, to get an intermediate estimate  $\mathbf{x}_T$ . We then set  $b_{ij} = \beta_j(\mathbf{P}_i\mathbf{x}_T)$  for all  $i$  and  $j$ . That is, we fix the  $b_{ij}$ 's as the coefficients obtained from the intermediate point  $\mathbf{x}_T$ . The subsequent PnP-PGD iterations are run using the denoiser in (10) that uses these fixed coefficients. Since  $\mathbf{W}$  is contractive, it follows from Theorem 2 that the sequence  $\mathbf{x}_{T+1}, \mathbf{x}_{T+2}, \dots$  converges to some fixed point  $\mathbf{x}^*$ . Consequently, the PnP-PGD iterations  $\mathbf{x}_1, \mathbf{x}_2, \dots$  converge to some fixed point  $\mathbf{x}^*$ .

The idea behind fixing the coefficients after  $T$  iterations is that as  $k$  increases,  $\mathbf{x}_k$  is expected to become more refined (in the sense of looking similar to the unknown signal  $\mathbf{x}$ ); therefore, the coefficients  $\beta_j(\mathbf{P}_i\mathbf{x}_T)$  after  $T$  iterations would be good enough to use for all subsequent iterates as well. In fact, this scheme has been used in PnP algorithms for image restoration [24,35].

We note that although the patch denoiser  $G$  in (8) is an MMSE estimator, the overall image denoiser  $D$  is not. Therefore existing convergence results for PnP with MMSE denoisers, e.g. [36], do not apply to our case. We make an important remark on the similarity and differences of the proposed GMM denoiser with the denoiser in [30,37]. Indeed, the idea of our GMM denoiser is inspired by that in [30,37]. However, there are a couple of subtle differences:

- The denoiser in [30,37] is *scene-adapted*, in the sense that the GMM distribution is tailored for the specific scene being reconstructed. This is possible because the application considered there is hyperspectral image sharpening, in which a complementary image of the same scene is available to obtain training data tailored for that scene. In contrast, in this work, we learn just one GMM distribution that is kept common for all the signals being reconstructed.
- The method used to replace  $\beta_j(\mathbf{P}_i\mathbf{z})$  by a fixed coefficient  $b_{ij}$  is different in our paper as compared to [2]. This is because the approach in [2] fundamentally relies on the availability of a complementary image, and thus cannot be applied to compressive sensing. In particular, in our paper, we take the  $b_{ij}$ 's to be the coefficients acquired from a surrogate signal that is obtained by running a few iterations of the PnP-PGD algorithm; this idea was inspired by [Sreehari]. On the other hand, in [2], the  $b_{ij}$ 's are taken to be the coefficients obtained from the complementary image (multispectral or panchromatic image).

## 5. Experimental results

**Database:** To validate the proposed PnP-PGD method for ECG CS recovery, we use a subset of the data from the Physionet MIT-BIH Arrhythmia Database [32–34]. Every file in the database consists of two lead recordings sampled at 360Hz with 11 bits per sample of resolution. It contains 48 half-hour excerpts of two-channel ambulatory ECG recordings, obtained from 47 subjects studied by the BIH Arrhythmia Laboratory.

**Metrics:** We quantify the performance of the proposal using following metrics: SNR, which is defined in Section 3, and mean-squared error (MSE), which is defined below.

$$\text{MSE} = \frac{1}{N} \|\mathbf{x} - \hat{\mathbf{x}}\|^2,$$

where  $\hat{\mathbf{x}}$  and  $\mathbf{x}$  are the reconstructed and original ECG signals, respectively. Note that here we are assuming the Physionet signals are the true signals  $\mathbf{x}$ ; in reality these signals also contain noise, which the metrics above neglect, though at most SNRs the (simulated) additive noise dominates [28].

**Compared methods:** We compare with the following state-of-the-art methods: BSBL-BO [15], BSBL-EM [16], sparse prior on the ECG wavelet representation [11], and TV regularization [13]. We tuned the parameters of all the methods so that maximum SNR is obtained. The codes are used from the publicly available sites [5,15,16]. All simulations were performed using MATLAB (R2021a) on a Quad core, 3.80GHz machine with 32GB RAM.

The sensing matrix  $\Phi$  is constructed by randomly drawing each entry from the standard normal distribution  $\mathcal{N}(0, 1)$ , and applying an orthonormalization step to ensure that the rows of  $\Phi$  are orthonormal [12]. We trained the GMM on the set of all possible overlapping patches of size  $P = 30$  extracted from signal #104 in the dataset [32], which is of length 10,800. The number of GMM components  $K$  is set to be 10. The training time is found to be 2.86s. In all the experiments on CS reconstruction, we terminate the PnP-PGD algorithm after 150 iterations since we observed that the algorithm stabilizes by then.

In addition to the denoising results reported in Section 3, the results in the subsequent sections support the claim that GMM is a good prior for modeling ECG signal patches. Note that the entire training process can be done offline.

### 5.1. Goodness of GMM modeling

In this experiment, we evaluate the goodness of GMM modeling by visualizing some of the eigenvectors of the covariance matrices from the learned GMM distribution. This type of visualization is commonly utilized in papers on image processing, e.g. [29]. In particular, in [29] it is observed that the eigenvectors corresponding to the largest few eigenvalues (of each covariance matrix) are relatively smooth and capture the large-scale structure of the patches, whereas the eigenvalues corresponding to the smallest few eigenvectors contain many fluctuations and thus capture the local structure. If the GMM is a good model, we expect to see a similar trend for ECG signal patches.

In our case, we fit a GMM with  $K = 10$  components on ECG patches of size 30. In Fig. 3, we plot 10 randomly selected eigenvectors (having unit norm) corresponding to eigenvalue indices  $\geq 28$  (i.e. largest few eigenvalues) from the fitted GMM model. In Fig. 3, we show similar plots but for eigenvalue indices  $\leq 3$  (i.e. smallest few eigenvalues). It is evident that the expected trend described in the previous paragraph holds true in practice: the richness of textures, fluctuations and other local structures is captured by the signals in Fig. 3, while most of the large-scale details are captured by the signals in Fig. 2.

### 5.2. Recovery from noiseless measurements

We study the signal reconstruction performance of our method (under zero noise), especially when the number of measurements  $M$  is much lesser than  $N$ . In Fig. 4, we show a segment of the original ECG signal #100 from Moody and Mark [33] and its reconstruction obtained using our proposed method with 10% measurements. We can see that the recovered signal using GMM as the denoiser produces visually similar result on comparison with the other methods and also has better performance in terms of SNR.

### 5.3. Study of SNR for different $M$

We next perform an exhaustive experiment where we vary the number of measurements ( $M$ ) while fixing the length of the sig-

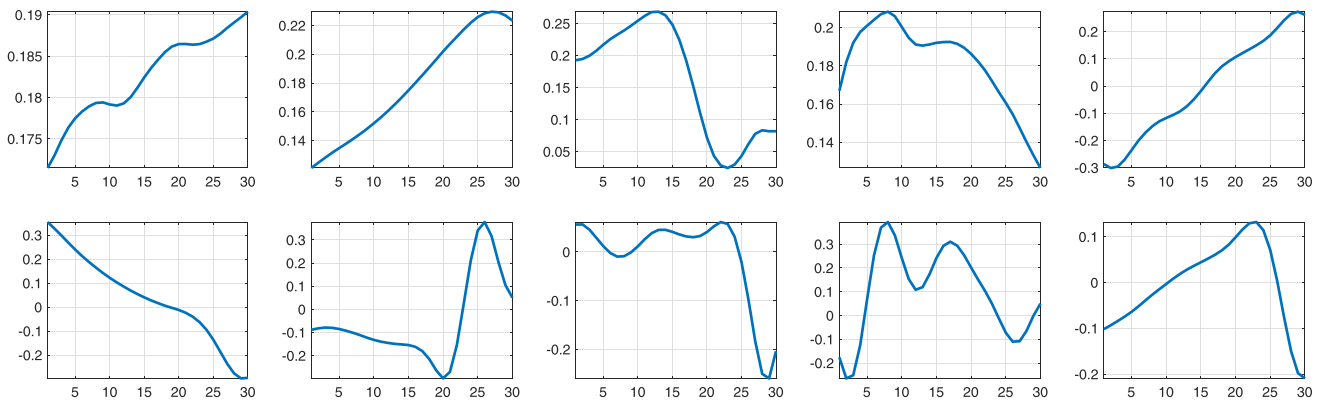


Fig. 2. Ten randomly selected eigenvectors corresponding to the largest few eigenvalues of the covariance matrices from the learned GMM model.




Fig. 3. Ten randomly selected eigenvectors corresponding to the smallest few eigenvalues of the covariance matrices from the learned GMM model.




Fig. 4. Visual comparison of a CS recovered ECG signal with 10% measurements and no additive noise. The length of the original signal is  $N = 512$ . The proposed method produces more structurally similar (with the clean signal) output even under lower number of observations.

nal. The signal #105 of length  $N = 512$  is used for the experiment. The results are reported in Fig. 6. Each instance of the SNR values shown is obtained by averaging of 500 independent trials. We note that our proposed method achieves the best recovery among all the methods.

#### 5.4. Study of average SNR over different signals

In this section, we study the performance of the proposed method with that of various ECG signals from MIT-BIH database [33]. We consider a larger dataset of 10 signals (#100 to #109),

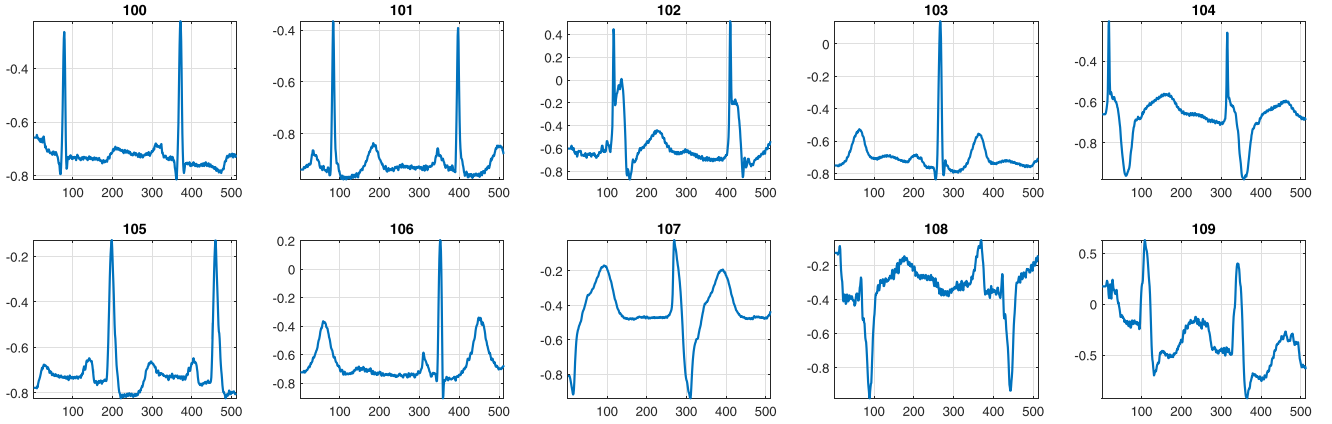


Fig. 5. Different test signals from the MIT-BIH Arrhythmia database.

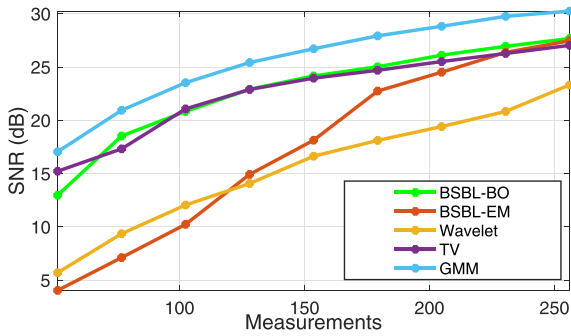


Fig. 6. Average SNR vs. measurements ( $M$ ) over 500 runs, with no additive noise. The length of the original ECG signal is  $N = 512$ .




Fig. 7. Average SNR vs. measurements ( $M$ ) averaged over the 10 ECG signals in Fig. 5. The measurements do not contain any additive noise.

shown in Fig. 5 for performance evaluation. All the signals are of length  $N = 512$ . We also compare with a more recent method, ECGLet [38], in this section. We measure the average SNR (over all the 10 test signals) of the reconstructed signal for different values of  $M$ . The results are reported in Fig. 7. Note that the proposed method yields the highest SNR among all the methods considered, for every  $M$ .

### 5.5. Effect of compression ratio

In this section we investigate the effect of compression ratio (CR) on the quality of the reconstructed ECG signals. The CR is defined as:

$$CR = \frac{N - M}{N} \times 100 \quad (11)$$

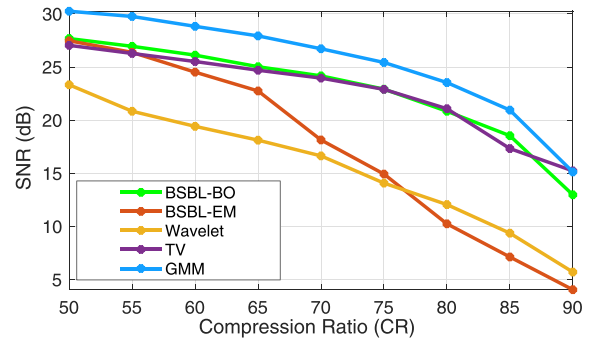


Fig. 8. Average SNR vs. CR over 500 runs. The length of the original ECG signal is  $N = 512$ .





Fig. 9. Average SNR vs. MSE over 500 runs. The length of the original ECG signal is  $N = 512$ .

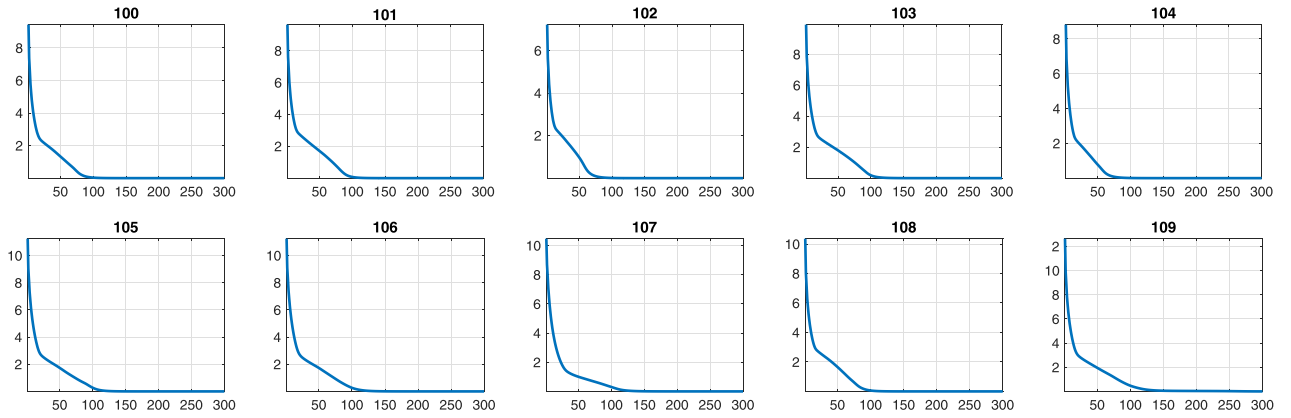
where  $N$  is the length of the original signal and  $M$  is the length of the compressed signal. For each value of CR, we repeated the experiment 500 times, and in each time, the sensing matrix was randomly generated [5]. Figure 8, shows the variation of SNR with CR. It is worth noting that we obtain superior performance over the whole range of CR values. The input signal is a segment of ECG signal #105 from the MIT-BIH database [33].

### 5.6. Recovery from noisy measurements

From previous experiments we notice that the proposed method performs well when compared with the other methods in noiseless scenario, i.e. when  $\mathbf{n} = \mathbf{0}$  in (1). Now we examine the recovery performance of our method from noisy compressed measurements. The signal #103 of length  $N = 512$  and  $M = 256$  is used for this experiment. In Fig. 9, we plot the MSE of the recovered



**Fig. 10.** Visual comparison of a CS recovered ECG signal with 50% measurements and additive Gaussian noise. The length of the original signal is  $N = 512$ . The proposed method is able to capture the minute variations in the signal faithfully.




**Fig. 11.** Plots of  $\|\mathbf{x}_{k+1} - \mathbf{x}_k\|$  vs  $k$  in the proposed algorithm (with  $M = 0.5N$ ) for the test signals in Fig. 5.

signal as a function of the noise level in the input  $\mathbf{y}$  (specified in terms of the SNR of the input). The values reported are averaged over 500 independent trials. To simulate the noisy measurements, we followed the approach in [28]. It is evident that the proposed method produces quality reconstructions under noisy measurements. Finally we show a reconstruction result from noisy measurements in Fig. 10. In the experiment, we add random Gaussian noise to the compressed measurements. On comparison, none of the three methods in Fig. 10 are able to completely mitigate the effect of noise. However, our method performs the best by a significant margin, resulting in an SNR of 24.24 dB in the recovered signal, and is visually similar to the original signal. We observed that in low SNR scenarios, TV acts as a better denoiser than GMM, which might explain why the CS reconstruction performance is higher for TV as compared to the proposed method when the input SNR is low.

### 5.7. Numerical convergence

In this section, we numerically verify the convergence of the proposed PnP-PGD algorithm. For all the signals in Fig. 5, we use PnP-PGD for reconstruction from  $M = 0.5N$  measurements. As explained in Section 4, we take the surrogate signal as the signal obtained after running the algorithm for  $T = 10$  iterations. The plots of  $\|\mathbf{x}_{k+1} - \mathbf{x}_k\|$  as a function of  $k$  are shown in Fig. 11 for each of



**Fig. 12.** Variation of SNR with trials for BSBL-BO (blue) and the proposed method (red). The length of the original ECG signal  $N = 512$ . The proposed method exhibits stable recovery of the ECG signal even for a low compression ratio. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the signals. Note that  $\|\mathbf{x}_{k+1} - \mathbf{x}_k\|$  decays to 0 as  $k$  increases, which is a necessary condition for convergence.

### 5.8. Comparison with deep learning

For completeness, we compare the proposed method with a deep learning method for CS reconstruction, known as SDAE [39]. In this experiment, we use the non-invasive Fetal ECG dataset from



**Table 2**

Performance comparison with SDAE on ECG data, collected from Physionet database (patient id: ecgca154) by varying the compression ratio.

CR	25	37.5	50	62.5	75	87.5	93.75
SDAE	37.69	36.73	35.39	32.40	26.78	<b>23.56</b>	<b>22.29</b>
GMM	<b>42.85</b>	<b>39.56</b>	<b>37.18</b>	<b>33.29</b>	<b>27.16</b>	20.57	11.44

Physionet, which was used in [39]. This database contains a series of 55 multichannel abdominal non-invasive fetal electrocardiogram (ECG) recordings, taken from a single subject between 21 to 40 weeks of pregnancy. We conducted an experiment on the signal with patient id ecgca154. Table 2 shows the variation of SNR with CR for the proposed method and SDAE. We note that the proposed method is able to outperform SDAE in all cases except when the CR is very high.

### 5.9. Stable recovery

We show that the proposed method produces very stable reconstructions. For this experiment, we considered the signal #105 with 10% measurements ( $N = 512$ ). We ran 500 trials of the method, so that its stability can be observed for different realizations of the sensing matrix  $\Phi$ . In Fig. 12, we show the SNR variation for BSBL-BO (blue) and the proposed method (red). We noted a standard deviation of 6.90 dB in SNR for BSBL-BO and 2.70 dB for the proposed method. Thus, the proposed method is more stable as compared to BSBL-BO. In fact, we observed that the contrast in stability is more pronounced for smaller  $M$ .

## 6. Conclusion

We introduced a novel framework for recovering ECG signals from compressively sensed measurements. Our method is based on the plug-and-play (PnP) paradigm that has recently become popular for image restoration problems. Essentially, the recovery method consists of repeating two main steps – inverting the forward model, and denoising – until stability is attained. We designed a high-quality ECG signal denoiser to be used in the denoising step. Moreover, we proved that the recovery algorithm is guaranteed to converge. Importantly, we showed via numerical experiments that our proposed method is superior to current state-of-the-art methods used for ECG CS recovery.

### Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### CRediT authorship contribution statement

**V. S. Unni:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Ruturaj G. Gavaskar:** Methodology, Software, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Kunal N. Chaudhury:** Methodology, Formal analysis, Funding acquisition, Project administration, Writing – original draft, Writing – review & editing.

## Appendix A

### A1. Proof of Theorem 2

Since  $\nabla f(\mathbf{x}) = \Phi^\top(\Phi\mathbf{x} - \mathbf{y})$ , we can write the PnP-PGD algorithm as  $\mathbf{x}_{k+1} = S(\mathbf{x}_k)$ , where

$$S(\mathbf{x}) = D(\mathbf{x} - \gamma\Phi^\top\Phi\mathbf{x} + \gamma\Phi^\top\mathbf{y}).$$

It is enough to prove that the function  $S(\cdot)$  is contractive, since the convergence of  $(\mathbf{x}_k)$  to some unique fixed point  $\mathbf{x}^*$  at a linear rate would then follow by the Banach Fixed Point Theorem [23, Th. 9.23].

Let  $\delta < 1$  be the constant in Definition 1. Then for any  $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^N$ , we have,

$$\begin{aligned} \|S(\mathbf{z}_1) - S(\mathbf{z}_2)\| &\leq \delta \|(\mathbf{z}_1 - \gamma\Phi^\top\Phi\mathbf{z}_1) - (\mathbf{z}_2 - \gamma\Phi^\top\Phi\mathbf{z}_2)\| \\ &= \delta \|(\mathbf{I} - \gamma\Phi^\top\Phi)(\mathbf{z}_1 - \mathbf{z}_2)\| \\ &\leq \delta \cdot \sigma_{\max}(\mathbf{I} - \gamma\Phi^\top\Phi) \cdot \|\mathbf{z}_1 - \mathbf{z}_2\|. \end{aligned}$$

Let  $L = \sigma_{\max}(\Phi^\top\Phi)$ . Since  $\Phi^\top\Phi$  is positive semidefinite, its singular values are also its eigenvalues. In particular, its eigenvalues lie in  $[0, L]$ . Since  $0 < \gamma \leq 2/L$ , the eigenvalues of  $\mathbf{I} - \gamma\Phi^\top\Phi$  lie in  $[-1, 1]$ . Therefore,  $\sigma_{\max}(\mathbf{I} - \gamma\Phi^\top\Phi) \leq 1$ . Thus, for all  $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^N$  we have,

$$\|S(\mathbf{z}_1) - S(\mathbf{z}_2)\| \leq \delta \|\mathbf{z}_1 - \mathbf{z}_2\|.$$

Since  $\delta < 1$ , the function  $S(\cdot)$  is contractive.

### A2. Proof of Theorem 4

A proof can be found in [30, Appendix B]; for completeness, here we give a different and more concise proof. Note that each  $\mathbf{C}_j$  is symmetric positive semidefinite (p.s.d.); hence, for each  $i$ , the matrix  $\mathbf{B}_i := \sum_j b_{ji}\mathbf{C}_j$  is p.s.d. (as a convex combination of p.s.d. matrices). By the same logic, we get that  $\mathbf{W}$  is p.s.d. Thus, to show  $\lambda_{\max}(\mathbf{W}) < 1$ , we only need to show that  $\mathbf{z}^\top\mathbf{W}\mathbf{z} < \|\mathbf{z}\|^2$  for all  $\mathbf{z} \in \mathbb{R}^N$ .

To prove this, first note that  $\lambda_{\max}(\mathbf{C}_j) < 1$  for all  $j$ . Since each  $\mathbf{B}_i$  is a convex combination of all  $\mathbf{C}_j$ 's and since  $\lambda_{\max}(\cdot)$  is a convex function on the set of symmetric matrices, we have  $\lambda_{\max}(\mathbf{B}_i) < 1$  for all  $i = 1, \dots, N$ . Let

$$\delta = \max(\lambda_{\max}(\mathbf{B}_1), \dots, \lambda_{\max}(\mathbf{B}_N)).$$

Clearly,  $\delta < 1$ . Note that for each  $i$ , we have  $\mathbf{u}^\top\mathbf{B}_i\mathbf{u} \leq \lambda_{\max}(\mathbf{B}_i)\|\mathbf{u}\|^2 \leq \delta\|\mathbf{u}\|^2$ . Therefore, for any  $\mathbf{z} \in \mathbb{R}^N$ ,

$$\mathbf{z}^\top\mathbf{W}\mathbf{z} = \frac{1}{P} \sum_{i=1}^N (\mathbf{P}_i\mathbf{z})^\top\mathbf{B}_i(\mathbf{P}_i\mathbf{z}) \leq \frac{\delta}{P} \sum_{i=1}^N \|\mathbf{P}_i\mathbf{z}\|^2.$$

Since  $N$  is a multiple of  $P$ ,  $\sum_{i=1}^N \|\mathbf{P}_i\mathbf{z}\|^2$ , which is the sum of all patches of length  $P$  extracted from  $\mathbf{z}$  (using circular padding), is simply equal to  $P\|\mathbf{z}\|^2$ . Thus, we get that  $\mathbf{z}^\top\mathbf{W}\mathbf{z} \leq \delta\|\mathbf{z}\|^2 < \|\mathbf{z}\|^2$  for all  $\mathbf{z} \in \mathbb{R}^N$ .

## References

- [1] S. Ahmad, S. Chen, K. Soueidan, I. Batkin, M. Bolic, H. Dajani, V. Groza, Electrocardiogram-assisted blood pressure estimation, *IEEE Trans. Biomed. Eng.* 59 (3) (2012) 608–618.
- [2] L. Pecchia, P. Mellillo, M. Sansone, M. Bracale, Discrimination power of short-term heart rate variability measures for CHF assessment, *IEEE Trans. Inf. Technol. Biomed. Eng.* 15 (1) (2010) 40–46.
- [3] M.I. Owis, A.H. Abou-Zied, A.-B. Youssef, Y.M. Kadam, Study of features based on nonlinear dynamical modeling in ECG arrhythmia detection and classification, *IEEE Trans. Biomed. Eng.* 49 (7) (2002) 733–736.
- [4] H. Mamaghanian, N. Khaled, D. Atienza, P. Vandergheynst, Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes, *IEEE Trans. Biomed. Eng.* 58 (9) (2011) 2456–2466.
- [5] Z. Zhang, T.P. Jung, S. Makeig, B.D. Rao, Compressed sensing for energy-efficient wireless telemonitoring of noninvasive fetal ECG via block sparse Bayesian learning, *IEEE Trans. Biomed. Eng.* 60 (2) (2012) 300–309.

- [6] J.K. Pant, S. Krishnan, Compressive sensing of electrocardiogram signals by promoting sparsity on the second-order difference and by using dictionary learning, *IEEE Trans. Biomed. Circuits Syst.* 8 (2) (2013) 293–302.
- [7] A. Salman, E.G. Allstot, A.Y. Chen, A.M. Dixon, D. Gangopadhyay, D.J. Allstot, Compressive sampling of EMG bio-signals, in: *Proc. IEEE International Symposium of Circuits and Systems*, 2011, pp. 2095–2098.
- [8] A.M. Dixon, E.G. Allstot, D. Gangopadhyay, D.J. Allstot, Compressed sensing system considerations for ECG and EMG wireless biosensors, *IEEE Trans. Biomed. Circuits Syst.* 6 (2) (2012) 156–166.
- [9] S. Aviyente, Compressed sensing framework for EEG compression, in: *Proc. IEEE/SP 14th Workshop on Statistical Signal Processing*, 2007, pp. 181–184.
- [10] M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: the application of compressed sensing for rapid MR imaging, *Magn. Reson. Med.* 58 (6) (2007) 1182–1195.
- [11] L.F. Polania, K.E. Barner, A weighted  $\ell_1$  minimization algorithm for compressed sensing ECG, in: *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 4413–4417.
- [12] J.K. Pant, W.-S. Lu, A. Antoniou, New improved algorithms for compressive sensing based on  $\ell_p$  norm, *IEEE Trans. Circuits Syst. II* 61 (3) (2014) 198–202.
- [13] Y. Liu, M. De Vos, I. Gligorijevic, V. Matic, Y. Li, S. Van Huffel, Multi-structural signal recovery for biomedical compressive sensing, *IEEE Trans. Biomed. Eng.* 60 (10) (2013) 2794–2805.
- [14] J.K. Pant, S. Krishnan, Reconstruction of ECG signals for compressive sensing by promoting sparsity on the gradient, in: *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 993–997.
- [15] Z. Zhang, B.D. Rao, Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning, *IEEE J. Sel. Top. Signal Process.* 5 (5) (2011) 912–926.
- [16] Z. Zhang, B.D. Rao, Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation, *IEEE Trans. Signal Process.* 61 (8) (2013) 2009–2015.
- [17] A. Beck, *First-Order Methods in Optimization*, SIAM, Philadelphia, PA, USA, 2017.
- [18] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* 2 (1) (2009) 183–202.
- [19] S.V. Venkatakrishnan, C.A. Bouman, B. Wohlberg, Plug-and-play priors for model based reconstruction, in: *Proc. IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945–948.
- [20] U.S. Kamilov, H. Mansour, B. Wohlberg, A plug-and-play priors approach for solving nonlinear imaging inverse problems, *IEEE Signal Process. Lett.* 24 (12) (2017) 1872–1876.
- [21] A. Buades, B. Coll, J.M. Morel, A non-local algorithm for image denoising, *Proc. IEEE Comput. Vis. Pattern Recognit.* 2 (2005) 60–65.
- [22] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, Image denoising by sparse 3-D transform-domain collaborative filtering, *IEEE Trans. Image Process.* 16 (8) (2007) 2080–2095.
- [23] W. Rudin, *Principles of Mathematical Analysis*, McGraw-Hill, New York, 1976.
- [24] P. Nair, R.G. Gavaskar, K.N. Chaudhury, Fixed-point and objective convergence of plug-and-play algorithms, *IEEE Trans. Comput. Imaging* (7) (2021) 337–348.
- [25] E. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, W. Yin, Plug-and-play methods provably converge with properly trained denoisers, in: *International Conference on Machine Learning*, 2019, pp. 5546–5557.
- [26] S. Chatterjee, R.S. Thakur, R.N. Yadav, L. Gupta, D.K. Raghuvanshi, Review of noise removal techniques in ECG signals, *IET Signal Proc.* 14 (9) (2020) 569–590.
- [27] S. Kumar, D. Panigrahy, P.K. Sahu, Denoising of electrocardiogram (ECG) signal by using empirical mode decomposition (EMD) with non-local mean (NLM) technique, *Biocybern. Biomed. Eng.* 38 (2) (2018) 297–312.
- [28] B.H. Tracey, E.L. Miller, Nonlocal means denoising of ECG signals, *IEEE Trans. Biomed. Eng.* 59 (9) (2012) 2383–2386.
- [29] D. Zoran, Y. Weiss, From learning models of natural image patches to whole image restoration, in: *Proc. International Conference on Computer Vision*, 2011, pp. 479–486.
- [30] A.M. Teodoro, J.M. Bioucas-Dias, M.A. Figueiredo, A convergent image fusion algorithm using scene-adapted Gaussian-mixture-based denoising, *IEEE Trans. Image Process.* 28 (1) (2018) 451–463.
- [31] L. Condat, A direct algorithm for 1-D total variation denoising, *IEEE Signal Process. Lett.* 101 (23) (2013) 1054–1057.
- [32] A.L. Goldberger, L.A. Amaral, L. Glass, J.M. Hausdorff, P.C. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C.K. Peng, H.E. Stanley, Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals, *Circulation* 101 (23) (2000) e215–e220.
- [33] G.B. Moody, R.G. Mark, The impact of the MIT-BIH arrhythmia database, *IEEE Eng. Med. Biol. Mag.* 20 (3) (2001) 45–50.
- [34] Z. Lu, D.Y. Kim, W.A. Pearlman, Wavelet compression of ECG signals by the set partitioning in hierarchical trees algorithm, *IEEE Trans. Biomed. Eng.* 47 (7) (2000) 849–856.
- [35] S. Sreehari, S.V. Venkatakrishnan, B. Wohlberg, G.T. Buzzard, L.F. Drummy, J.P. Simmons, C.A. Bouman, Plug-and-play priors for bright field electron tomography and sparse interpolation, *IEEE Trans. Comput. Imaging* 2 (4) (2016) 408–423.
- [36] X. Xu, Y. Sun, J. Liu, B. Wohlberg, U.S. Kamilov, Provable convergence of plug-and-play priors with MMSE denoisers, *IEEE Signal Process. Lett.* 27 (2020) 1280–1284.
- [37] A.M. Teodoro, J.M. Bioucas-Dias, M.A. Figueiredo, Scene-adapted plug-and-play algorithm with convergence guarantees, in: *IEEE International Workshop on Machine Learning for Signal Processing*, 2017, pp. 1–6.
- [38] N. Ansari, A. Gupta, WNC-ECGlet: weighted non-convex minimization based reconstruction of compressively transmitted ECG using ECglet, *Biomed. Signal Process. Control* 49 (2019) 1–13.
- [39] P.R. Muduli, R.R. Gunukula, A. Mukherjee, A deep learning approach to fetal-ecg signal reconstruction, in: *National Conference on Communications*, 2016, pp. 1–6.