

Sparse Signal Reconstruction from Quantized Noisy Measurements via GEM Hard Thresholding

Kun Qiu and Aleksandar Dogandžić

Abstract

We develop a generalized expectation-maximization (GEM) algorithm for sparse signal reconstruction from quantized noisy measurements. The measurements follow an underdetermined linear model with sparse regression coefficients, corrupted by additive white Gaussian noise having unknown variance. These measurements are quantized into bins and only the bin indices are used for reconstruction. We treat the unquantized measurements as the missing data and propose a GEM iteration that aims at maximizing the likelihood function with respect to the unknown parameters. Under mild conditions, our GEM iteration yields a convergent monotonically non-decreasing likelihood function sequence and the Euclidean distance between two consecutive GEM signal iterates goes to zero as the number of iterations grows. We compare the proposed scheme with the state-of-the-art convex relaxation method for quantized compressed sensing via numerical simulations.

Index Terms

Compressed sensing, sparse signal reconstruction, quantization, generalized expectation-maximization (GEM) algorithm.

I. INTRODUCTION

In the past few years, compressed sensing [1]–[4] has attracted considerable attention spanning a wide variety of areas including applied mathematics, statistics and engineering. The compressed sensing theory asserts that, if the signal of interest is sparse or nearly sparse in some (e.g. Fourier, wavelet, etc.) domain, it is possible to reconstruct the underlying signal with high fidelity from only a few linear measurements whose dimension is much lower than that of the signal. The sparse signal reconstruction techniques can be roughly divided into three categories: convex relaxation, greedy pursuit, and probabilistic methods, see [5] for a survey. Digital storage and processing are integral parts of most modern systems, thereby necessitating quantization. There have recently been several efforts to incorporate the quantization effect into compressed sensing [6]–[9], see also the discussion on the state of the art in [9, Sec. I]. However, as observed in [9], references [6]–[8] focus only on the quantization effects and *do not* account for noise or approximately sparse signals. Most methods developed so far for quantized compressed sensing belong to the convex relaxation category, see [6]–[9]. In [9], Zymnis *et al.* consider a convex relaxation approach for signal reconstruction from quantized Gaussian-noise corrupted CS measurements using an ℓ_1 -norm regularization term and two convex cost functions: the negative log-likelihood function of the underlying sparse signal given the quantized data and a weighted least squares cost that employs virtual measurements constructed from centroids of the quantization bins. In both cases, the noise variance is assumed known and must be tuned to achieve good performance.

In this paper (see also [10]), we propose an *unrelaxed* probabilistic model with ℓ_0 -norm constrained signal space and derive a *generalized expectation-maximization* (GEM) algorithm for approximately computing the maximum likelihood

This work was supported by the National Science Foundation under Grant CCF-0545571.

Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

The authors are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (email: ald@iastate.edu; kqiu@iastate.edu).

(ML) estimates of the unknown sparse signal and noise variance parameter. As [9], we consider *both* the quantization and noise effects. However, in contrast to [9], our GEM algorithm estimates the noise variance from the data. We prove that, under certain mild conditions, the GEM iteration guarantees a convergent monotonically non-decreasing likelihood function sequence and diminishing Euclidean distance between consecutive signal parameter estimates as the number of iterations grows. The reconstruction performance of our method is studied via numerical examples and compared with the likelihood based convex relaxation approach from [9].

We introduce the notation used in this paper: $\mathcal{N}(\mathbf{y}; \boldsymbol{\mu}, \Sigma)$ denotes the multivariate probability density function (pdf) of a real-valued Gaussian random vector \mathbf{y} with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ ; $\phi(\cdot)$ and $\Phi(\cdot)$ denote the pdf and cdf of the standard normal random variable; I_n and $\mathbf{0}_{n \times 1}$ are the identity matrix of size n and the $n \times 1$ vector of zeros; $\|\cdot\|_p$ and “ T ” denote the ℓ_p norm and transpose, respectively; the hard thresholding operator $\mathcal{T}_r(\mathbf{s})$ keeps the r largest-magnitude elements of a vector \mathbf{s} intact and sets the rest to zero, e.g. $\mathcal{T}_2([0, 1, -5, 0, 3, 0]^T) = [0, 0, -5, 0, 3, 0]^T$.

II. MEASUREMENT MODEL

We model a $N \times 1$ real-valued measurement vector $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$ as

$$\mathbf{y} = H \mathbf{s} + \mathbf{e} \quad (1a)$$

where H is a known $N \times m$ full-rank *sensing matrix*, \mathbf{s} is an *unknown* $m \times 1$ sparse signal vector containing at most r nonzero elements ($r \leq m$), \mathbf{e} is an $N \times 1$ additive Gaussian noise vector with zero mean and covariance matrix $\sigma^2 I_N$; the noise variance σ^2 is *unknown*. The set of unknown parameters is

$$\boldsymbol{\theta} = (\mathbf{s}, \sigma^2) \in \Theta_r \quad (1b)$$

with the parameter space

$$\Theta_r = \mathcal{S}_r \times (0, +\infty) \quad (1c)$$

and

$$\mathcal{S}_r = \{\mathbf{s} \in \mathbb{R}^m : \|\mathbf{s}\|_0 \leq r\} \quad (1d)$$

is the sparse signal parameter space. We refer to r as the *sparsity level* of the signal and to the signal \mathbf{s} as being *r-sparse*. In this paper, we assume that the sparsity level r is *known*. In (1c), we impose strict positivity on the noise variance σ^2 , which is needed to ensure that the quantization is non-degenerate and the likelihood function of the unknown parameters is computable, see the following discussion. The elements of \mathbf{y} are *quantized* into codewords $\mathbf{b} = [b_1, b_2, \dots, b_N]^T$, where each b_i indexes the quantization interval (bin) that y_i falls in:

$$y_i \in \mathcal{D}(b_i) = [l(b_i), u(b_i)) \triangleq [l_i, u_i), \quad l_i < u_i, \quad i = 1, 2, \dots, N \quad (2)$$

where the real numbers l_i and u_i are the upper and lower boundaries of the quantization interval containing y_i .

Our goal is to estimate the parameters $\boldsymbol{\theta}$ from the quantized data \mathbf{b} . Since the unquantized measurements \mathbf{y} are not available for reconstruction, we refer to \mathbf{y} as the *unobserved (missing) data*; the concept of missing data is key for the development of our GEM algorithm, see Section III.

The joint distribution of the observed data \mathbf{b} and the unobserved data \mathbf{y} given the parameters $\boldsymbol{\theta}$ is

$$p_{\mathbf{y}, \mathbf{b} | \boldsymbol{\theta}}(\mathbf{y}, \mathbf{b} | \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y}; H\mathbf{s}, \sigma^2 I_N) \prod_{i=1}^N 1_{\mathcal{D}(b_i)}(y_i) = \prod_{i=1}^N \{\mathcal{N}(y_i; \mathbf{h}_i^T \mathbf{s}, \sigma^2) 1_{\mathcal{D}(b_i)}(y_i)\} \quad (3a)$$

where \mathbf{h}_i^T denotes the i th row of H and

$$1_A(y) = \begin{cases} 1, & y \in A, \\ 0, & \text{otherwise} \end{cases} \quad (3b)$$

is the indicator function. Consequently, the conditional pdf of \mathbf{y} given \mathbf{b} is

$$p_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}}(\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}) = \prod_{i=1}^N p_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}) = \prod_{i=1}^N \frac{\mathcal{N}(y_i; \mathbf{h}_i^T \mathbf{s}, \sigma^2) 1_{\mathcal{D}(b_i)}(y_i)}{\Phi((u_i - \mathbf{h}_i^T \mathbf{s})/\sigma) - \Phi((l_i - \mathbf{h}_i^T \mathbf{s})/\sigma)} \quad (3c)$$

where $\sigma = \sqrt{\sigma^2}$. We call the quantization *non-degenerate* if

$$\text{tr}[\text{cov}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}}(\mathbf{y} | \mathbf{b}, \boldsymbol{\theta})] = \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}) > 0 \quad (4)$$

for any $\boldsymbol{\theta} \in \Theta_r$, which ensures that there exists some uncertainty about \mathbf{y} given the quantized data \mathbf{b} . Since we assume in (1c) and (2) that $\sigma^2 > 0$ and $l_i < u_i$, $i = 1, 2, \dots, N$, our quantization is non-degenerate, see (3c).

The marginal log-likelihood function of $\boldsymbol{\theta}$ is

$$\mathcal{L}(\boldsymbol{\theta}) = \ln [p_{\mathbf{b} | \boldsymbol{\theta}}(\mathbf{b} | \boldsymbol{\theta})] = \sum_{i=1}^N \ln \left[\Phi\left(\frac{u_i - \mathbf{h}_i^T \mathbf{s}}{\sigma}\right) - \Phi\left(\frac{l_i - \mathbf{h}_i^T \mathbf{s}}{\sigma}\right) \right] \quad (5)$$

where the marginal likelihood $p_{\mathbf{b} | \boldsymbol{\theta}}(\mathbf{b} | \boldsymbol{\theta})$ is obtained by *integrating \mathbf{y} out* from the joint distribution of \mathbf{b} and \mathbf{y} in (3a). To compute $\mathcal{L}(\boldsymbol{\theta})$, we need the noise variance σ^2 to be strictly positive and restrict the parameter space accordingly, see (1c). Note that (5) has the same form as the log-likelihood function in [9, Sect. III-A]; however, unlike [9], which treats σ^2 as known, here the noise variance σ^2 is unknown and estimated from the data.

The ML estimate of $\boldsymbol{\theta}$ is

$$\hat{\boldsymbol{\theta}}_{\text{ML}} = (\hat{\mathbf{s}}_{\text{ML}}, \hat{\sigma}_{\text{ML}}^2) = \max_{\boldsymbol{\theta} \in \Theta_r} \mathcal{L}(\boldsymbol{\theta}). \quad (6)$$

Obtaining the exact ML estimate $\hat{\boldsymbol{\theta}}_{\text{ML}}$ in (6) requires a combinatorial search and is therefore infeasible in practice. In the following section, we develop a GEM algorithm that aims at maximizing (5) with respect to $\boldsymbol{\theta} \in \Theta_r$ and circumvents the combinatorial search.

III. THE GEM ALGORITHM

We now derive a GEM algorithm for estimating the parameters $\boldsymbol{\theta}$ by treating the unquantized measurements \mathbf{y} as the missing data. The observed data \mathbf{b} and missing data \mathbf{y} together make up the *complete data*. In our problem, the expectation (E) step can be readily computed, but the maximization (M) step is computational infeasible, see the following discussion. We propose an alternative computationally tractable generalized M step and prove that the generalized M parameter update improves the desired log-likelihood objective function (5).

Assume that the parameter estimate $\boldsymbol{\theta}^{(p)} = (\mathbf{s}^{(p)}, (\sigma^2)^{(p)})$ is available, where p denotes the iteration index. Then, the E step of the EM-type algorithms consists of evaluating the following expected complete-data log-likelihood function [see (3a)]:

$$\begin{aligned}
\mathcal{Q}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)}) &\triangleq \mathbb{E}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}} [\ln p_{\mathbf{y}, \mathbf{b}}(\mathbf{y}, \mathbf{b} | \boldsymbol{\theta}) | \mathbf{b}, \boldsymbol{\theta}^{(p)}] = -\frac{1}{2} N \ln(2\pi\sigma^2) - \mathbb{E}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}} [(\mathbf{y} - H\mathbf{s})^T (\mathbf{y} - H\mathbf{s}) | \mathbf{b}, \boldsymbol{\theta}^{(p)}] / (2\sigma^2) \\
&= -\frac{1}{2} N \ln(2\pi\sigma^2) - [\|\widehat{\mathbf{y}}^{(p)} - H\mathbf{s}\|_2^2 + \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}^{(p)})] / (2\sigma^2)
\end{aligned} \tag{7a}$$

where

$$\widehat{\mathbf{y}}^{(p)} = [\widehat{y}_1^{(p)}, \widehat{y}_2^{(p)}, \dots, \widehat{y}_N^{(p)}]^T = \mathbb{E}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}} [\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}^{(p)}] \tag{7b}$$

is the Bayesian minimum mean-square error (MMSE) estimate of the missing data \mathbf{y} for *known* $\boldsymbol{\theta}$ [11, Sec. 11.4] and we have used the fact that $\mathbb{E}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}} [\ln 1_{\mathcal{D}(b_i)}(y_i) | \mathbf{b}, \boldsymbol{\theta}^{(p)}] = 0$ for each $i = 1, 2, \dots, N$. From (7), the E step reduces to evaluating $\widehat{\mathbf{y}}^{(p)}$ and $\text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}^{(p)})$, $i = 1, 2, \dots, N$ using the expressions for the mean and variance of the truncated Gaussian pdf [12, eqs. (13.134) and (13.135)], see also (3c).

E step: Compute

$$\widehat{\mathbf{y}}^{(p)} = H\mathbf{s}^{(p)} + \sigma^{(p)} \boldsymbol{\delta}^{(p)} \tag{8a}$$

$$\text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}^{(p)}) = (\sigma^2)^{(p)} (1 - \xi_i^{(p)}) \tag{8b}$$

where $\boldsymbol{\delta}^{(p)} = [\delta_1^{(p)}, \delta_2^{(p)}, \dots, \delta_N^{(p)}]^T$, $\sigma^{(p)} = [(\sigma^2)^{(p)}]^{1/2}$, and

$$v_i^{(p)} = (u_i - \mathbf{h}_i^T \mathbf{s}^{(p)}) / \sigma^{(p)} \tag{8c}$$

$$\lambda_i^{(p)} = (l_i - \mathbf{h}_i^T \mathbf{s}^{(p)}) / \sigma^{(p)} \tag{8d}$$

$$\delta_i^{(p)} = -\frac{\phi(v_i^{(p)}) - \phi(\lambda_i^{(p)})}{\Phi(v_i^{(p)}) - \Phi(\lambda_i^{(p)})} \tag{8e}$$

$$\xi_i^{(p)} = (\delta_i^{(p)})^2 + \frac{v_i^{(p)} \phi(v_i^{(p)}) - \lambda_i^{(p)} \phi(\lambda_i^{(p)})}{\Phi(v_i^{(p)}) - \Phi(\lambda_i^{(p)})} \tag{8f}$$

for $i = 1, 2, \dots, N$. When u_i or l_i is infinite, so is v_i or λ_i , and, in this case, $v_i^{(p)} \phi(v_i^{(p)})$ or $\lambda_i^{(p)} \phi(\lambda_i^{(p)})$ in (8f) becomes zero.

The standard M step of an EM algorithm for the above model requires the maximization of $\mathcal{Q}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)})$ in (7a) with respect to $\boldsymbol{\theta} \in \Theta_r$. For any given \mathbf{s} , $\mathcal{Q}((\mathbf{s}, \sigma^2) | \boldsymbol{\theta}^{(p)})$ is uniquely maximized with respect to σ^2 at

$$\widehat{\sigma}^2(\mathbf{s}, \boldsymbol{\theta}^{(p)}) = [\|\widehat{\mathbf{y}}^{(p)} - H\mathbf{s}\|_2^2 + \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | \mathbf{b}, \boldsymbol{\theta}^{(p)})] / N \tag{9a}$$

implying

$$\mathcal{Q}((\mathbf{s}, \widehat{\sigma}^2(\mathbf{s}, \boldsymbol{\theta}^{(p)})) | \boldsymbol{\theta}^{(p)}) \geq \mathcal{Q}((\mathbf{s}, \sigma^2) | \boldsymbol{\theta}^{(p)}) \tag{9b}$$

where the equality in (9b) holds only if $\sigma^2 = \widehat{\sigma}^2(\mathbf{s}, \boldsymbol{\theta}^{(p)})$. The exact M step for updating the estimates of \mathbf{s} hence reduces to the maximization of the *concentrated* expected complete-data log-likelihood function $\mathcal{Q}((\mathbf{s}, \widehat{\sigma}^2(\mathbf{s}, \boldsymbol{\theta}^{(p)})) | \boldsymbol{\theta}^{(p)})$ with respect to \mathbf{s} or, equivalently, to solving the following optimization problem:

$$\min_{\mathbf{s} \in \mathcal{S}_r} \|\widehat{\mathbf{y}}^{(p)} - H\mathbf{s}\|_2^2 \tag{10}$$

which requires combinatorial search and is therefore infeasible in practice. We now propose a generalized M step that increases (rather than maximizes) the expected complete-data log-likelihood function (7a).

Generalized M step: Compute

$$\mathbf{s}^{(p+1)} = \mathcal{T}_r(\mathbf{s}^{(p)} + \frac{1}{c^2} H^T (\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)})) = \mathcal{T}_r(\mathbf{s}^{(p)} + \frac{\sigma^{(p)}}{c^2} H^T \boldsymbol{\delta}^{(p)}) \quad (11a)$$

$$\begin{aligned} (\sigma^2)^{(p+1)} &= [\|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2 + \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}^{(p)})] / N \\ &= \|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2 / N + (\sigma^2)^{(p)} (1 - \sum_{i=1}^N \xi_i^{(p)} / N) \end{aligned} \quad (11b)$$

where $\boldsymbol{\delta}^{(p)} = [\delta_1^{(p)}, \delta_2^{(p)}, \dots, \delta_N^{(p)}]^T$ and $\xi_1^{(p)}, \xi_2^{(p)}, \dots, \xi_N^{(p)}$ are computed using (8). Construct the new parameter estimate as

$$\boldsymbol{\theta}^{(p+1)} = (\mathbf{s}^{(p+1)}, (\sigma^2)^{(p+1)}). \quad (12)$$

Iterate between the above E and generalized M steps until two consecutive sparse signal estimates $\mathbf{s}^{(p)}$ and $\mathbf{s}^{(p+1)}$ do not differ significantly.

In (11a), c is a step size coefficient chosen to satisfy the following inequality:

$$c \geq \rho_H \quad (13)$$

where ρ_H denotes the largest singular value of H , also known as the spectral norm of H . Interestingly, (11a) is closely related to the hard-thresholded gradient-search step for maximizing the marginal log likelihood in (5):

$$\mathbf{s}^{(p+1)} = \mathcal{T}_r\left(\mathbf{s}^{(p)} + \tau \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \mathbf{s}} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(p)}}\right). \quad (14)$$

Indeed, choosing the step size in (14) adaptively as $\tau = (\sigma^2)^{(p)} / c^2$ leads to our GEM update (11a). Observe that (11b) follows by substituting $\mathbf{s} = \mathbf{s}^{(p+1)}$ into (9a). Because the quantization is non-degenerate, i.e. (4) holds, we have

$$\frac{1}{N} \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | b_i, \boldsymbol{\theta}^{(p)}) = (\sigma^2)^{(p)} (1 - \sum_{i=1}^N \xi_i^{(p)} / N) > 0 \quad (15)$$

and, consequently, $(\sigma^2)^{(p+1)} > 0$ for all indices $p = 0, 1, \dots$ as long as $(\sigma^2)^{(0)} > 0$, see (11b).

In the limiting case where all quantization bins are infinitely small, our GEM iteration reduces to the iterative hard thresholding method in [13] for *unquantized* measurements \mathbf{y}

$$\mathbf{s}^{(p+1)} = \mathcal{T}_r(\mathbf{s}^{(p)} + \mu H^T (\mathbf{y} - H \mathbf{s}^{(p)})) \quad (16)$$

with $\mu = 1/c^2$. We refer to the resulting method as GEM_∞ . Clearly, GEM_∞ is the benchmark for the reconstruction performance achievable by our GEM algorithm.

In Lemma 1, we verify that, for the choice of c in (13), the above scheme is indeed a GEM iteration by proving that the parameter update (11) guarantees monotonically non-decreasing expected complete-data log-likelihood function (7).

Lemma 1: Assuming that the parameter estimate in the p -th iteration $\boldsymbol{\theta}^{(p)} = (\mathbf{s}^{(p)}, (\sigma^2)^{(p)})$ belongs to the parameter space Θ_r , (13) holds, and the quantization is non-degenerate [i.e. (4) holds], the sparse signal update $\boldsymbol{\theta}^{(p+1)} = (\mathbf{s}^{(p+1)}, (\sigma^2)^{(p+1)})$ in (11) also belongs to Θ_r and satisfies

$$\mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) \geq \mathcal{Q}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}). \quad (17)$$

Proof: See the Appendix. □

A. Initialization and Termination of the GEM Algorithm

We initialize the proposed GEM iteration as follows:

$$\mathbf{s}^{(0)} = \mathbf{0}_{m \times 1}, \quad (\sigma^2)^{(0)} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i^{(-1)})^2 \quad (18a)$$

where

$$\hat{y}_i^{(-1)} = \begin{cases} (l_i + u_i)/2, & \text{if } l_i > -\infty \text{ and } u_i < +\infty \\ l_i, & \text{if } u_i = +\infty \\ u_i, & \text{if } l_i = -\infty \end{cases}. \quad (18b)$$

Denote by $\boldsymbol{\theta}^{(+\infty)} = (\mathbf{s}^{(+\infty)}, (\sigma^2)^{(+\infty)})$ and $\hat{\mathbf{y}}^{(+\infty)}$ the estimates of the unknown parameter set and missing data upon convergence of the GEM iteration. To estimate the signal \mathbf{s} , we propose to use either the sparse estimate $\mathbf{s}^{(+\infty)}$ or the non-sparse estimate *before thresholding* [see (11a)]:

$$\tilde{\mathbf{s}} = \mathbf{s}^{(+\infty)} + \frac{1}{c^2} H^T (\hat{\mathbf{y}}^{(+\infty)} - H \mathbf{s}^{(+\infty)}) \quad (19)$$

which is appealing for recovering nearly sparse signals with many small nonzero elements. If the sensing matrix H has orthonormal rows, i.e. $H H^T = I_N$, then the marginal likelihood function under our measurement model in Section II *coincides with* the marginal likelihood function under the random signal model (where the random signal is the sum of the sparse signal component and additive signal ‘noise’ that models approximate signal sparsity) in [14] and [15]; consequently, the MMSE estimates $E_{\mathbf{y}|\mathbf{b},\boldsymbol{\theta}}[\mathbf{y}|\mathbf{b},\boldsymbol{\theta}]$ coincide as well for the two models. Under the random signal model and for $c = \rho_H = 1$, (19) is closely related to the *empirical Bayesian (EB) estimate* of the random signal in [15, eq. (7.7)]: It follows by replacing the unobserved data vector \mathbf{y} with its empirical Bayesian MMSE estimate $\hat{\mathbf{y}}^{(+\infty)} = E_{\mathbf{y}|\mathbf{b},\boldsymbol{\theta}}[\mathbf{y}|\mathbf{b},\boldsymbol{\theta}^{(+\infty)}]$ in [15, eq. (7.7)], see also (7b).

B. Convergence Analysis

Theorem 1 establishes convergence properties of our GEM algorithm.

Theorem 1: Under the conditions of Lemma 1, the GEM iteration guarantees monotonically nondecreasing marginal log-likelihood function (5):

$$\mathcal{L}(\boldsymbol{\theta}^{(p+1)}) \geq \mathcal{L}(\boldsymbol{\theta}^{(p)}) \quad (20)$$

and the log-likelihood sequence $\mathcal{L}(\boldsymbol{\theta}^{(p)})$ converges to a limit as the iteration index p goes to infinity.

Furthermore, if

$$c > \rho_H \quad (21)$$

and if there exists at least one $i \in \{1, 2, \dots, N\}$ such that both the upper and lower boundaries u_i and l_i of the i th quantization interval are finite, then the Euclidean distances between two consecutive GEM signal and variance parameter iterates $\|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2$ and $\|(\sigma^2)^{(p+1)} - (\sigma^2)^{(p)}\|_2$ go to zero as the iteration index p goes to infinity.

Proof: See the Appendix. □

The additional convergence condition in the second part of Theorem 1 requiring the existence of finite upper and lower quantization boundaries is related to the parameter identifiability under our model. For example, consider the 1-bit

quantization scheme in [6] where the quantization threshold is zero for all $i = 1, 2, \dots, N$, i.e. only the sign information of the unquantized measurements is recorded. In this case, for any index i , one of the quantization boundaries u_i and l_i is infinite and the parameter sets (s, σ^2) and $(as, a^2\sigma^2)$ yield the same marginal distribution $p_{b_i|\theta}(\mathbf{b}|\boldsymbol{\theta})$ for any positive constant a , implying that the model is not identifiable; it is also impossible to determine the magnitude of the signal s unless additional information about the signal magnitude is provided, see [6].

IV. NUMERICAL EXAMPLES

We consider reconstructions of one- and two-dimensional signals from quantized compressive samples and compare the performances of

- our GEM algorithm in Section III with the step size coefficient set to $c = \rho_H$ and
- the fixed point continuation algorithm (labeled FPC) in [9] for solving the ℓ_1 -regularized ML optimization problem:

$$\min_{\mathbf{s} \in \mathbb{R}^m} [-\mathcal{L}(\mathbf{s}, \sigma^2) + \lambda_{\text{reg}} \|\mathbf{s}\|_1] \quad (22)$$

where the noise variance σ^2 is assumed *known* and the regularization parameter λ_{reg} controls the sparsity of the output.

The unquantized measurements are partitioned into B bins, where the quantization thresholds are chosen so that the bins contain approximately equal numbers of measurements on average. In the following examples, GEM_B and FPC_B denote the GEM and FPC algorithms that use B bins for quantization.

The main step of the FPC iteration in [9] can be obtained by replacing the hard thresholding operator in (14) with a soft thresholding operator. Observe that the FPC method in [9] requires tuning of several quantities, whereas our GEM algorithm requires only the knowledge of the signal sparsity level r . Upon tuning the noise variance σ^2 , we set the step-size parameter of FPC to

$$\tau = \sigma^2 / \rho_H^2 \quad (23)$$

which results in better performance and numerical stability than the suggested value $1/\rho_H^2$ in [9, Sec. IV], see also the following discussion. We set the shrinkage parameter β of FPC algorithm to the suggested value 0.5, see [9, Sec. IV].

We employ the following convergence criterion:

$$\|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2 / m < \epsilon \quad (24)$$

where, for the FPC method, we apply (24) in the inner loop of the FPC iteration, see also [9, Sec. IV].

A. One-dimensional Signal Reconstruction

We generated sparse signals \mathbf{s} of length $m = 500$ containing 20 *randomly located* nonzero elements, see also the simulation examples in [2] and [16]. The nonzero components of \mathbf{s} are independent, identically distributed (i.i.d.) random variables that are either -1 or $+1$ with equal probability. The sensing matrices H are constructed by creating an $N \times m$ matrix containing i.i.d. samples from the zero-mean Gaussian distribution with variance $1/m$; therefore,

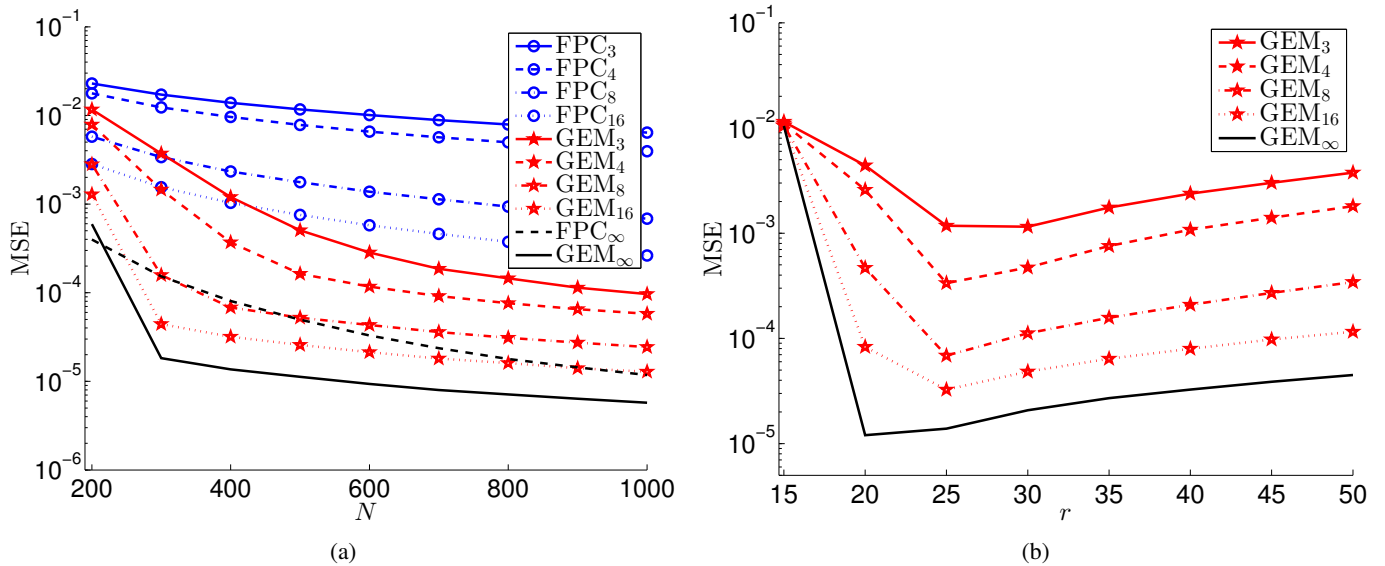


Fig. 1. (a) Average MSEs of various estimators of \mathbf{s} as functions of the number of measurements N , and (b) average MSEs of the GEM estimators of \mathbf{s} as functions of the sparsity level r with the number of measurements N fixed at 400.

each row of H approximately has a unit magnitude. The $N \times 1$ unquantized measurement vector \mathbf{y} is generated using (1a) with noise variance $\sigma^2 = 10^{-4}$.

Our performance metric is the *average* mean-square error (MSE) of a signal estimate $\hat{\mathbf{s}}$:

$$\text{MSE}\{\hat{\mathbf{s}}\} = \text{E}[\|\hat{\mathbf{s}} - \mathbf{s}\|_2^2]/m \quad (25)$$

computed using 1000 Monte Carlo trials, where *averaging* is performed over the random sensing matrices H , the sparse signal \mathbf{s} , and the noise \mathbf{e} . We selected the convergence threshold $\epsilon = 10^{-13}$ in (24).

Fig. 1(a) shows the average MSEs of the compared methods as we vary the number of measurements N and the number of quantization bins B . Since the true signal \mathbf{s} is sparse, we use $\mathbf{s}^{(+\infty)}$ upon convergence of the GEM iteration as our GEM signal estimate. GEM_∞ and FPC_∞ are the limiting cases for the GEM and FPC algorithms, where the main steps of GEM_∞ and FPC_∞ are the hard thresholding step (16) and its soft thresholding counterpart, respectively. The sparsity level of the GEM algorithm is set to 25, slightly higher than the true signal support size 20. To implement the FPC algorithm, we chose the true value of the noise variance $\sigma^2 = 10^{-4}$ and the regularization parameter $\lambda_{\text{reg}} = 400$, which yields sparse signal estimates with approximately 20 to 25 nonzero signal elements. We use the FPC step-size parameter in (23): In this example, FPC does not always converge if we apply $\tau = 1/\rho_H^2$ suggested in [9, Sec. IV]. From Fig. 1(a), our GEM algorithm achieves consistently lower MSEs than the convex FPC method over wide ranges of N and B . When $N > 400$, GEM with only 3 quantization bins outperforms the FPC method with 16 quantization bins. The performance of the GEM method with 16 quantization bins is quite close to that of the limiting GEM_∞ algorithm.

We now study the sensitivity of our GEM algorithm to the choice of the signal sparsity level r . Fig. 1(b) shows the average MSEs achieved by the GEM scheme as we vary r for $N = 400$ measurements. Setting $r = 25$ yields the smallest MSE for the GEM method with $B = 3, 4, 8$, and 16 bins. The MSE increases for $r \geq 25$, but the rate of MSE degradation is quite mild.

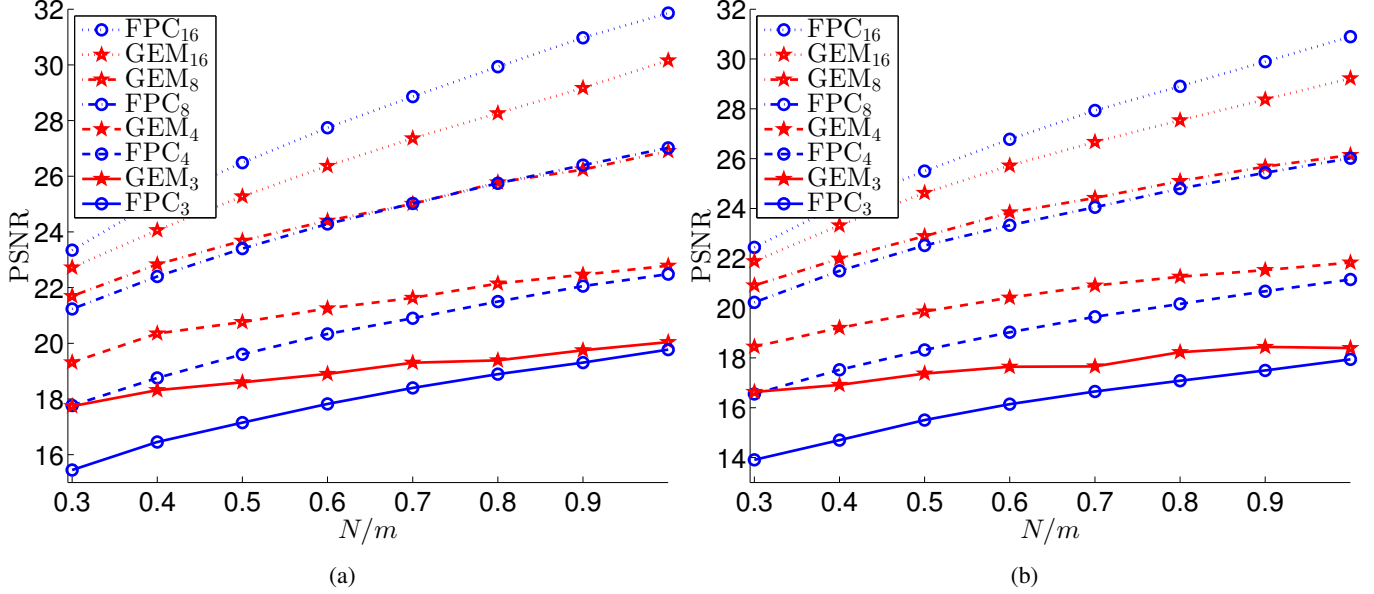


Fig. 2. PSNR curves for GEM and FPC reconstructions of (a) the 256×256 Lena image and (b) the 256×256 Cameraman image as functions of the normalized number of measurements N/m .

B. Two-dimensional Image Reconstruction

In this example, we reconstruct the standard Lena and Cameraman test images, both of size $m = 256^2$. Here, the sensing matrix H has the structure $H = \Phi \Psi$, where Φ is the $N \times m$ *structurally random sampling matrix* [17] and Ψ is the $m \times m$ inverse Daubechies-8 discrete wavelet transform (DWT) matrix. Under these choices of Φ and Ψ , the rows of H are orthonormal, i.e. $H H^T = I_N$ and, consequently, $\rho_H = 1$. The signal vectors s consisting of the wavelet coefficients of the Lena and Cameraman images are not strictly sparse and contain many small but nonzero elements. No noise is added, and therefore the unquantized measurements satisfy $\mathbf{y} = H \mathbf{s}$. Here, the role of the noise variance parameter σ^2 is to account for the fact that the wavelet coefficients of the test images are not strictly sparse (i.e. for the presence of the signal ‘noise’).

Our performance metric is the peak signal-to-noise ratio (PSNR) of a signal estimate $\hat{\mathbf{s}}$ [18, eq. (3.7)]:

$$\text{PSNR (dB)} = 10 \log_{10} \left\{ \frac{[(\Psi \mathbf{s})_{\text{MAX}} - (\Psi \mathbf{s})_{\text{MIN}}]^2}{\|\hat{\mathbf{s}} - \mathbf{s}\|_2^2 / m} \right\} \quad (26)$$

where $(\Psi \mathbf{s})_{\text{MIN}}$ and $(\Psi \mathbf{s})_{\text{MAX}}$ denote the smallest and largest elements of the image $\Psi \mathbf{s}$. We selected the convergence threshold $\epsilon = 10^{-10}$ in (24). The sparsity level of the GEM algorithm is set to $r = 4000 N/m$. For the FPC algorithm, we tuned manually the regularization and noise variance parameters to achieve good performance, yielding $\lambda_{\text{reg}} = 0.1$ and $\sigma^2 = 10$, respectively. Upon the convergence of the GEM algorithm, we use the empirical Bayesian signal estimate (19) to reconstruct the Lena and Cameraman images. Similarly, instead of the sparse signal estimate $\hat{\mathbf{s}}_{\text{FPC}}$ obtained upon convergence of the FPC iteration, we chose the approximately sparse signal estimate $\hat{\mathbf{s}}_{\text{FPC}} + \tau \partial \mathcal{L}(\boldsymbol{\theta}) / \partial \mathbf{s} |_{\boldsymbol{\theta} = (\hat{\mathbf{s}}_{\text{FPC}}, \sigma^2)} = \hat{\mathbf{s}}_{\text{FPC}} - \tau H^T \nabla f_{\text{ml}}(H \hat{\mathbf{s}}_{\text{FPC}})$ in [9, Sec. IV]. In this example, these signal estimates lead to better reconstructions compared with the corresponding purely sparse estimates.

Fig. 2 shows the PSNR performances of the GEM and FPC methods as function of subsampling factor N/m for

various numbers of quantization bins B . For both the Lena and Cameraman images, the GEM algorithm outperforms the convex FPC method for coarser quantization ($B = 3$ and 4), where the performance gap increases as the number of measurements N decreases. For 8 quantization bins, the reconstruction performance of GEM is similar to FPC for Lena reconstruction and slightly better than FPC for Cameraman reconstruction, see Figs. 2(a) and 2(b). When $B = 16$, FPC exhibits better performance than GEM. Note that, in addition to the regularization parameter λ_{reg} , the FPC method requires tuning of the noise variance parameter σ^2 , and we have found that its performance is sensitive to the choice of σ^2 . In contrast, our GEM achieves good performance with automatic estimation of the noise variance parameter.

V. CONCLUDING REMARKS

We developed a generalized expectation-maximization (GEM) hard thresholding reconstruction algorithm for sparse signal reconstruction from quantized Gaussian-noise corrupted compressed sensing measurements. We showed that the likelihood of our GEM iteration is monotonically non-decreasing and converges to a limit and the Euclidean distance between two GEM iterates goes to zero. The major advantage of our proposed method over the existing FPC method in [9] is that our method automates the estimation of the noise parameter from the quantized measurements whereas the FPC method requires tuning of the noise variance. The numerical examples showed good reconstruction performance of our GEM algorithm. Specifically, in the one dimensional sparse signal simulation, our GEM methods consistently outperforms the FPC method. We also find that in this experiment, the performance of our GEM algorithm is quite stable when the sparsity level r is larger than or equal to the true signal support size. For the two dimensional image reconstruction method, our method performs better when the quantization is coarser whereas FPC achieves higher PSNR when the quantization bins shrink. Further research will include developing theoretical analysis of the reconstruction accuracy of the proposed algorithm and automating the GEM method by estimating the signal sparsity level from the data.

APPENDIX

Proof of Lemma 1: Without loss of generality, we assume $\mathbf{s}^{(p+1)} \neq \mathbf{s}^{(p)}$ (Lemma 1 holds trivially when $\mathbf{s}^{(p+1)} = \mathbf{s}^{(p)}$). The claim that $\boldsymbol{\theta}^{(p+1)} = (\mathbf{s}^{(p+1)}, (\sigma^2)^{(p+1)}) \in \Theta_r$ is an immediate consequence of the GEM update (11) and the non-degenerate quantization assumption (4). Now, consider the following inequality:

$$\|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2 = c^2 \|(\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)})/c\|_2^2 \quad (\text{A1a})$$

$$\leq c^2 (\|(\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)})/c\|_2^2 + \|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2 - \|H(\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)})/c\|_2^2) \quad (\text{A1b})$$

$$\leq c^2 [\|(\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)})/c\|_2^2 + \|\mathbf{s}^{(p)} - \mathbf{s}^{(p)}\|_2^2 - \|H(\mathbf{s}^{(p)} - \mathbf{s}^{(p)})/c\|_2^2] \quad (\text{A1c})$$

$$= \|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)}\|_2^2 \quad (\text{A1d})$$

where (A1b) follows by using (13) and the Rayleigh-quotient property [19, Theorem 21.5.6]

$$\frac{\|H(\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)})\|_2^2}{\|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2} \leq \rho_H^2 \leq c^2 \quad (\text{A2})$$

and (A1c) follows by the fact that $\mathbf{s}^{(p+1)}$ in (11a) minimizes the following function of \mathbf{s} over all $\mathbf{s} \in \mathcal{S}_r$:

$$\|(\widehat{\mathbf{y}}^{(p)} - H \mathbf{s})/c\|_2^2 + \|\mathbf{s} - \mathbf{s}^{(p)}\|_2^2 - \|H(\mathbf{s} - \mathbf{s}^{(p)})/c\|_2^2. \quad (\text{A3a})$$

To see this, observe that (A3a) can be rewritten as

$$\|\mathbf{s} - \mathbf{s}^{(p)} - H^T (\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)})/c^2\|_2^2 \quad (\text{A3b})$$

up to an additive constant that is not a function of \mathbf{s} .

Now, (17) follows easily:

$$\mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) \geq \mathcal{Q}((\mathbf{s}^{(p+1)}, (\sigma^2)^{(p)}) | \boldsymbol{\theta}^{(p)}) \quad (\text{A4a})$$

$$\geq \mathcal{Q}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) \quad (\text{A4b})$$

where (A4a) follows by setting $\mathbf{s} = \mathbf{s}^{(p+1)}$ and $\sigma^2 = (\sigma^2)^{(p)}$ in (9b) and noting that $\widehat{\sigma}^2(\mathbf{s}^{(p+1)}, \boldsymbol{\theta}^{(p)}) = (\sigma^2)^{(p+1)}$;

(A4b) follows by using (A1). \square

Proof of Theorem 1: The first claim in (20) is a direct consequence of Lemma 1 and the following property of the EM-type algorithms [20] (see e.g. [20, eq. (3.2)]):

$$\mathcal{L}(\boldsymbol{\theta}) = \mathcal{Q}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)}) - \mathcal{H}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)}) \quad (\text{A5})$$

where

$$\mathcal{H}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)}) \triangleq \mathbb{E}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}} [\ln p_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}}(\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}) | \mathbf{b}, \boldsymbol{\theta}^{(p)}] \quad (\text{A6})$$

which is maximized at $\boldsymbol{\theta} = \boldsymbol{\theta}^{(p)}$, i.e.

$$\mathcal{H}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(p)}) \leq \mathcal{H}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) \quad (\text{A7})$$

see [20, Lemma 1]. The first claim follows by combining the result of Lemma 1 in (17) with (A5) and (A7):

$$\mathcal{L}(\boldsymbol{\theta}^{(p+1)}) = \mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) - \mathcal{H}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) \geq \mathcal{Q}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) - \mathcal{H}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) = \mathcal{L}(\boldsymbol{\theta}^{(p)}). \quad (\text{A8})$$

Now, we move on to prove the second part of Theorem 1. From (A5), we have

$$\mathcal{L}(\boldsymbol{\theta}^{(p+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(p)}) = \mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) - \mathcal{Q}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) + \mathcal{H}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) - \mathcal{H}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) \quad (\text{A9a})$$

$$\geq \mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) - \mathcal{Q}(\boldsymbol{\theta}^{(p)} | \boldsymbol{\theta}^{(p)}) \quad (\text{A9b})$$

$$= \mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) - \mathcal{Q}((\mathbf{s}^{(p+1)}, (\sigma^2)^{(p)}) | \boldsymbol{\theta}^{(p)}) \quad (\text{A9c})$$

$$+ \frac{1}{2(\sigma^2)^{(p)}} (\|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)}\|_2^2 - \|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2)$$

$$\geq \frac{1}{2(\sigma^2)^{(p)}} (\|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)}\|_2^2 - \|\widehat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2) \quad (\text{A9d})$$

$$\geq 0 \quad (\text{A9e})$$

where (A9b)–(A9e) follow by using (A7), (7), (A4a) and Lemma 1, respectively. Since the sequence $\mathcal{L}(\boldsymbol{\theta}^{(p)})$ is monotonically nondecreasing and upper-bounded by zero, it converges to a limit. Consequently, $\mathcal{L}(\boldsymbol{\theta}^{(p+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(p)})$ converges to zero. Therefore, the quantity in (A9d) [sandwiched by $\mathcal{L}(\boldsymbol{\theta}^{(p+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(p)})$ and zero via inequalities] converges to zero as well.

Now, we have

$$\frac{\|\hat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p)}\|_2^2 - \|\hat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2}{2(\sigma^2)^{(p)}} \geq \frac{c^2 \|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2 - \|H(\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)})\|_2^2}{2(\sigma^2)^{(p)}} \quad (\text{A10a})$$

$$\geq \frac{c^2 - \rho_H^2}{2(\sigma^2)^{(p)}} \|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2 \quad (\text{A10b})$$

where (A10a) follows by (A1b)–(A1d) and (A10b) results from (A2). Since the quantization is non-degenerate, $(\sigma^2)^{(p)} > 0$. Further, since there exists an index i such that both u_i and l_i are finite, we claim $(\sigma^2)^{(p)} < \infty$ for all $p > 0$. If $(\sigma^2)^{(p)}$ grows to infinity, then the i -th term in the summation of the marginal likelihood (5) goes to $-\infty$. Note that all summands in (5) are upper bounded by zero. This implies that the marginal likelihood $\mathcal{L}(\boldsymbol{\theta}^{(p)})$ goes to negative infinity if $(\sigma^2)^{(p)}$ grows to infinity, which is certainly less than $\mathcal{L}(\boldsymbol{\theta}^{(0)})$ for any reasonable initialization. This then contradicts with (20). Note also that the step size coefficient satisfies $c > \rho_H$, see (21). Therefore, the term $(c^2 - \rho_H^2)/[2(\sigma^2)^{(p)}]$ in (A10b) is positive and bounded away from zero, which implies that $\|\mathbf{s}^{(p+1)} - \mathbf{s}^{(p)}\|_2^2$ goes to zero.

Finally, from (A9), we also conclude that the sequence $\mathcal{Q}(\boldsymbol{\theta}^{(p+1)} | \boldsymbol{\theta}^{(p)}) - \mathcal{Q}((\mathbf{s}^{(p+1)}, (\sigma^2)^{(p)}) | \boldsymbol{\theta}^{(p)})$ converges to zero. Since the quantization is non-degenerate, $\text{tr}[\text{cov}_{\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}}(\mathbf{y} | \mathbf{b}, \boldsymbol{\theta}^{(p)})] = \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | \mathbf{b}, \boldsymbol{\theta}^{(p)}) > 0$ [see (4)], the function

$$\mathcal{Q}((\mathbf{s}^{(p+1)}, \sigma^2) | \boldsymbol{\theta}^{(p)}) = -\frac{1}{2}N \ln(2\pi\sigma^2) - \frac{\|\hat{\mathbf{y}}^{(p)} - H \mathbf{s}^{(p+1)}\|_2^2 + \sum_{i=1}^N \text{var}_{y_i | b_i, \boldsymbol{\theta}}(y_i | \mathbf{b}, \boldsymbol{\theta}^{(p)})}{2\sigma^2}$$

is a continuous and unimodal function of σ^2 , with the unique maximum achieved at $\sigma^2 = (\sigma^2)^{(p+1)}$, see also (9a). We conclude that $(\sigma^2)^{(p)} - (\sigma^2)^{(p+1)}$ must go to zero. The second claim of Theorem 1 follows. □

REFERENCES

- [1] *IEEE Signal Process. Mag. (Special Issue on Sensing, Sampling, and Compression)*, Mar. 2008.
- [2] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inf. Theory*, vol. 51, pp. 4203–4215, Dec. 2005.
- [3] E. J. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [4] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.
- [5] J. A. Tropp and S. J. Wright, “Computational methods for sparse solution of linear inverse problems,” *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, 2010.
- [6] P. Boufounos and R. Baraniuk, “1-bit compressive sensing,” in *Proc. 42rd Annu. Conf. Inform. Sci. Syst.*, Princeton, NJ, Mar. 2008, pp. 16–21.
- [7] W. Dai, H. V. Pham, and O. Milenkovic, “Distortion-rate functions for quantized compressive sensing,” in *IEEE Inf. Theory Workshop Networking Inf. Theory*, Volos, Greece, Jun. 2009, pp. 171–175.
- [8] L. Jacques, D. Hammond, and M. Fadili, “Dequantizing compressed sensing: When oversampling and non-Gaussian constraints combine,” *IEEE Trans. Inf. Theory*, vol. 57, no. 1, 2011.
- [9] A. Zymnis, S. Boyd, and E. Candès, “Compressed sensing with quantized measurements,” *IEEE Signal Process. Lett.*, vol. 17, pp. 149–152, Feb. 2010.
- [10] K. Qiu and A. Dogandžić, “A GEM hard thresholding method for reconstructing sparse signals from quantized noisy measurements,” in *Proc. 4th IEEE Int. Workshop Comput. Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, San Juan, PR, Dec. 2011, pp. 349–352.
- [11] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [12] N. L. Johnson and S. Kotz, *Continuous Univariate Distributions*, 2nd ed. New York: Wiley, 1994, vol. 1.
- [13] T. Blumensath and M. E. Davies, “Iterative hard thresholding for compressed sensing,” *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.
- [14] K. Qiu and A. Dogandžić, “Double overrelaxation thresholding methods for sparse signal reconstruction,” in *Proc. 44th Annu. Conf. Inform. Sci. Syst.*, Princeton, NJ, 2010.
- [15] K. Qiu and A. Dogandžić, “ECME thresholding methods for sparse signal reconstruction,” Iowa State University, Tech. Rep., Apr. 2010. [Online]. Available: <http://arxiv.org/abs/1004.4880>
- [16] K. Qiu and A. Dogandžić, “Variance-component based sparse signal reconstruction and model selection,” *IEEE Trans. Signal Process.*, vol. 58, pp. 2935–2952, Jun. 2010.

- [17] T. T. Do, T. D. Tran, and L. Gan, "Fast compressive sampling with structurally random matrices," in *Proc. Int. Conf. Acoust., Speech, Signal Processing*, Las Vegas, NV, Apr. 2008, pp. 3369–3372.
- [18] J. Starck, F. Murtagh, and J. Fadili, *Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity*. New York: Cambridge Univ. Press, 2010.
- [19] D. A. Harville, *Matrix Algebra From a Statistician's Perspective*. New York: Springer-Verlag, 1997.
- [20] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. Ser. B*, vol. 39, no. 1, pp. 1–38, 1977, with discussion.