

Expectation-Maximization Gaussian-Mixture Approximate Message Passing

Jeremy P. Vila, *Student Member, IEEE*, and Philip Schniter, *Senior Member, IEEE*

Abstract—When recovering a sparse signal from noisy compressive linear measurements, the distribution of the signal’s non-zero coefficients can have a profound effect on recovery mean-squared error (MSE). If this distribution was *a priori* known, then one could use computationally efficient approximate message passing (AMP) techniques for nearly minimum MSE (MMSE) recovery. In practice, however, the distribution is unknown, motivating the use of robust algorithms like LASSO—which is nearly minimax optimal—at the cost of significantly larger MSE for non-least-favorable distributions. As an alternative, we propose an empirical-Bayesian technique that simultaneously learns the signal distribution while MMSE-recovering the signal—according to the learned distribution—using AMP. In particular, we model the non-zero distribution as a Gaussian mixture and learn its parameters through expectation maximization, using AMP to implement the expectation step. Numerical experiments on a wide range of signal classes confirm the state-of-the-art performance of our approach, in both reconstruction error and runtime, in the high-dimensional regime, for most (but not all) sensing operators.

Index Terms—Compressed sensing, belief propagation, expectation maximization algorithms, Gaussian mixture model.

I. INTRODUCTION

WE consider estimating a K -sparse (or compressible) signal $\mathbf{x} \in \mathbb{R}^N$ from $M < N$ linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \in \mathbb{R}^M$, where \mathbf{A} is known and \mathbf{w} is additive white Gaussian noise (AWGN). For this problem, accurate (relative to the noise variance) signal recovery is known to be possible with polynomial-complexity algorithms when \mathbf{x} is sufficiently sparse and when \mathbf{A} satisfies certain restricted isometry properties [4], or when \mathbf{A} is large with i.i.d. zero-mean sub-Gaussian entries [5] as discussed below.

LASSO [6] (or, equivalently, Basis Pursuit Denoising [7]), is a well-known approach to the sparse-signal recovery problem that solves the convex problem

$$\hat{\mathbf{x}}_{\text{lasso}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda_{\text{lasso}} \|\mathbf{x}\|_1, \quad (1)$$

Manuscript received July 12, 2012; revised January 29, 2013 and June 14, 2013; accepted June 19, 2013. Date of publication July 10, 2013; date of current version August 28, 2013. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Namrata Vaswani. This work was supported in part by NSF-I/UCRC Grant IIP-0968910, by NSF Grant CCF-1018368, and by DARPA/ONR Grant N66001-10-1-4090. Portions of this work were presented at the Duke Workshop on Sensing and Analysis of High-Dimensional Data, July 2011 [1]; the Asilomar Conference on Signals, Systems, and Computers, November 2011 [2]; and the Conference on Information Science and Systems, March 2012 [3].

The authors are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: vila.2@osu.edu; schniter@ece.osu.edu).

Digital Object Identifier 10.1109/TSP.2013.2272287

with λ_{lasso} a tuning parameter that trades between the sparsity and measurement-fidelity of the solution. When \mathbf{A} is constructed from i.i.d. zero-mean sub-Gaussian entries, the performance of LASSO can be sharply characterized in the large system limit (i.e., as $K, M, N \rightarrow \infty$ with fixed undersampling ratio M/N and sparsity ratio K/M) using the so-called phase transition curve (PTC) [5], [8]. When the observations are noiseless, the PTC bisects the M/N -versus- K/M plane into the region where LASSO reconstructs the signal perfectly (with high probability) and the region where it does not. (See Figs. 3–5.) When the observations are noisy, the same PTC bisects the plane into the regions where LASSO’s noise sensitivity (i.e., the ratio of estimation-error power to measurement-noise power under the worst-case signal distribution) is either finite or infinite [9]. An important fact about LASSO’s noiseless PTC is that it is invariant to the distribution of the nonzero signal coefficients. In other words, if the vector \mathbf{x} is drawn i.i.d. from the pdf

$$p_X(x) = \lambda f_X(x) + (1 - \lambda)\delta(x), \quad (2)$$

where $\delta(\cdot)$ is the Dirac delta, $f_X(\cdot)$ is the active-coefficient pdf (with zero probability mass at $x = 0$), and $\lambda \triangleq K/N$, then the LASSO PTC is invariant to $f_X(\cdot)$. While this implies that LASSO is robust to “difficult” instances of $f_X(\cdot)$, it also implies that LASSO cannot benefit from the case that $f_X(\cdot)$ is an “easy” distribution. For example, when the signal is known a priori to be nonnegative, polynomial-complexity algorithms exist with PTCs that are better than LASSO’s [10].

At the other end of the spectrum is minimum mean-squared error (MMSE)-optimal signal recovery under *known* marginal pdfs of the form (2) and *known* noise variance. The PTC of MMSE recovery has been recently characterized [11] and shown to be well above that of LASSO. In particular, for *any* $f_X(\cdot)$, the PTC on the M/N -versus- K/M plane reduces to the line $K/M = 1$ in both the noiseless and noisy cases. Moreover, efficient algorithms for approximate MMSE-recovery have been proposed, such as the Bayesian version of Donoho, Maleki, and Montanari’s *approximate message passing* (AMP) algorithm from [12], which performs loopy belief-propagation on the underlying factor graph using central-limit-theorem approximations that become exact in the large-system limit under i.i.d. zero-mean sub-Gaussian \mathbf{A} . In fact, in this regime, AMP obeys [13] a state-evolution whose fixed points, when unique, are optimal. To handle arbitrary noise distributions and a wider class of matrices \mathbf{A} , Rangan proposed a *generalized AMP* (GAMP) [14] that forms the starting point of this work. (See Table I.) For more details and background on GAMP, we refer the reader to [14].

TABLE I
THE GAMP ALGORITHM FROM [14] WITH A STOPPING CONDITION IN (R10)
THAT USES THE NORMALIZED TOLERANCE PARAMETER τ_{GAMP}

inputs: $p_X(\cdot), p_{Y Z}(\cdot \cdot), \{A_{mn}\}, T_{\text{max}}, \tau_{\text{gamp}}$
definitions:
$p_{Z Y}(z_m \mathbf{y}; \hat{\mathbf{p}}_m, \mu_m^p, \mathbf{q}) \triangleq \frac{p_{Y Z}(y_m z_m; \mathbf{q}) \mathcal{N}(z_m; \hat{\mathbf{p}}_m, \mu_m^p)}{\int_z p_{Y Z}(y_m z; \mathbf{q}) \mathcal{N}(z; \hat{\mathbf{p}}_m, \mu_m^p)} \quad (\text{D1})$
$p_{X Y}(x_n \mathbf{y}; \hat{\mathbf{r}}_n, \mu_n^r, \mathbf{q}) \triangleq \frac{p_X(x_n; \mathbf{q}) \mathcal{N}(x_n; \hat{\mathbf{r}}_n, \mu_n^r)}{\int_x p_X(x; \mathbf{q}) \mathcal{N}(x; \hat{\mathbf{r}}_n, \mu_n^r)} \quad (\text{D2})$
initialize:
$\forall n: \hat{x}_n(1) = \int_x x p_X(x) \quad (\text{I1})$
$\forall n: \mu_n^x(1) = \int_x x - \hat{x}_n(1) ^2 p_X(x) \quad (\text{I2})$
$\forall m: \hat{s}_m(0) = 0 \quad (\text{I3})$
for $t = 1 : T_{\text{max}},$
$\forall m: \mu_m^p(t) = \sum_{n=1}^N A_{mn} ^2 \mu_n^x(t) \quad (\text{R1})$
$\forall m: \hat{\mathbf{p}}_m(t) = \sum_{n=1}^N A_{mn} \hat{x}_n(t) - \mu_m^p(t) \hat{s}_m(t-1) \quad (\text{R2})$
$\forall m: \mu_m^z(t) = \text{var}_{Z Y}\{z_m \mathbf{y}; \hat{\mathbf{p}}_m(t), \mu_m^p(t), \mathbf{q}\} \quad (\text{R3})$
$\forall m: \hat{z}_m(t) = \text{E}_{Z Y}\{z_m \mathbf{y}; \hat{\mathbf{p}}_m(t), \mu_m^p(t), \mathbf{q}\} \quad (\text{R4})$
$\forall m: \mu_m^s(t) = (1 - \mu_m^z(t)/\mu_m^p(t))/\mu_m^p(t) \quad (\text{R5})$
$\forall m: \hat{s}_m(t) = (\hat{z}_m(t) - \hat{\mathbf{p}}_m(t))/\mu_m^p(t) \quad (\text{R6})$
$\forall n: \mu_n^r(t) = (\sum_{m=1}^M A_{mn} ^2 \mu_m^s(t))^{-1} \quad (\text{R7})$
$\forall n: \hat{\mathbf{r}}_n(t) = \hat{x}_n(t) + \mu_n^r(t) \sum_{m=1}^M A_{mn}^* \hat{s}_m(t) \quad (\text{R8})$
$\forall n: \mu_n^x(t+1) = \text{var}_{X Y}\{x_n \mathbf{y}; \hat{\mathbf{r}}_n(t), \mu_n^r(t), \mathbf{q}\} \quad (\text{R9})$
$\forall n: \hat{x}_n(t+1) = \text{E}_{X Y}\{x_n \mathbf{y}; \hat{\mathbf{r}}_n(t), \mu_n^r(t), \mathbf{q}\} \quad (\text{R10})$
if $\sum_{n=1}^N \hat{x}_n(t+1) - \hat{x}_n(t) ^2 < \tau_{\text{gamp}} \sum_{n=1}^N \hat{x}_n(t) ^2$, break
end
outputs: $\{\hat{z}_m(t), \mu_m^z(t)\}, \{\hat{\mathbf{r}}_n(t), \mu_n^r(t)\}, \{\hat{x}_n(t+1), \mu_n^x(t+1)\}$

In practice, one ideally wants a recovery algorithm that does not need to know $p_X(\cdot)$ and the noise variance a priori, yet offers performance on par with MMSE recovery, which (by definition) requires knowing these prior statistics. Towards this goal, we propose a recovery scheme that aims to *learn* the prior signal distribution $p_X(\cdot)$, as well as the variance of the AWGN, while simultaneously recovering the signal vector \mathbf{x} from the noisy compressed measurements \mathbf{y} . To do so, we model the active component $f_X(\cdot)$ in (2) using a generic L -term Gaussian mixture (GM) and then learn the GM parameters and noise variance using the expectation-maximization (EM) algorithm [15]. As we will see, all of the quantities needed for the EM updates are already computed by the GAMP algorithm, making the overall process very computationally efficient. Moreover, GAMP provides approximately MMSE estimates of \mathbf{x} that suffice for signal recovery, as well as posterior activity probabilities that suffice for support recovery.

Since, in our approach, the prior pdf parameters are treated as deterministic unknowns, our proposed EM-GM-AMP algorithm can be classified as an “empirical-Bayesian” approach [16]. Compared with previously proposed empirical-Bayesian approaches to compressive sensing (e.g., [17]–[19]), ours has a more flexible signal model, and thus is able to better match a wide range of signal pdfs $p_X(\cdot)$, as we demonstrate through a detailed numerical study. In addition, the complexity scaling of our algorithm is superior to that in [17]–[19], implying lower complexity in the high dimensional regime, as we confirm numerically. Supplemental experiments demonstrate that our excellent results hold for a wide range of sensing operators \mathbf{A} , with some exceptions. Although this paper does not contain any convergence guarantees or a rigorous analysis/justification of the proposed EM-GM-AMP, Kamilov *et al.* showed (after the submission of this work) in [20] that a generalization of

EM-GM-AMP yields asymptotically (i.e., in the large system limit) consistent parameter estimates when \mathbf{A} is i.i.d. zero-mean Gaussian, when the parameterized signal and noise distributions match the true signal and noise distributions, and when those distributions satisfy certain identifiability conditions. We refer interested readers to [20] for more details.

Notation: For matrices, we use boldface capital letters like \mathbf{A} , and we use $\text{tr}(\mathbf{A})$ and $\|\mathbf{A}\|_F$ to denote the trace and Frobenius norm, respectively. Moreover, we use $(\cdot)^T$, $(\cdot)^*$, and $(\cdot)^H$ to denote transpose, conjugate, and conjugate transpose, respectively. For vectors, we use boldface small letters like \mathbf{x} , and we use $\|\mathbf{x}\|_p = (\sum_n |x_n|^p)^{1/p}$ to denote the ℓ_p norm, with x_n representing the n th element of \mathbf{x} . For a Gaussian random vector \mathbf{x} with mean \mathbf{m} and covariance matrix \mathbf{Q} , we denote the pdf by $\mathcal{N}(\mathbf{x}; \mathbf{m}, \mathbf{Q})$, and for its circular complex Gaussian counterpart, we use $\mathcal{CN}(\mathbf{x}; \mathbf{m}, \mathbf{Q})$. Finally, we use $\text{E}\{\cdot\}$, $\delta(\cdot)$, \mathbb{R} , and \mathbb{C} to denote the expectation operation, the Dirac delta, the real field, and the complex field, respectively.

II. GAUSSIAN-MIXTURE GAMP

We first introduce Gaussian-mixture (GM) GAMP, a key component of our overall approach, where the coefficients in $\mathbf{x} = [x_1, \dots, x_N]^T$ are assumed to be i.i.d. with marginal pdf

$$p_X(x; \lambda, \boldsymbol{\omega}, \boldsymbol{\theta}, \boldsymbol{\phi}) = (1 - \lambda)\delta(x) + \lambda \sum_{\ell=1}^L \omega_\ell \mathcal{N}(x; \theta_\ell, \phi_\ell), \quad (3)$$

where $\delta(\cdot)$ is the Dirac delta, λ is the sparsity rate, and, for the k th GM component, ω_k , θ_k , and ϕ_k are the weight, mean, and variance, respectively. In the sequel, we use $\boldsymbol{\omega} \triangleq [\omega_1, \dots, \omega_L]^T$ and similar definitions for $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. By definition, $\sum_{\ell=1}^L \omega_\ell = 1$. The noise $\mathbf{w} = [w_1, \dots, w_M]^T$ is assumed to be i.i.d. Gaussian, with mean zero and variance ψ , i.e.,

$$p_W(w; \psi) = \mathcal{N}(w; 0, \psi), \quad (4)$$

and independent of \mathbf{x} . Although above and in the sequel we assume real-valued quantities, all expressions in the sequel can be converted to the circular-complex case by replacing \mathcal{N} with \mathcal{CN} and removing the $\frac{1}{2}$'s from (25), (44), and (58). We note that, from the perspective of GM-GAMP, the prior parameters $\mathbf{q} \triangleq [\lambda, \boldsymbol{\omega}, \boldsymbol{\theta}, \boldsymbol{\phi}, \psi]$ and the number of mixture components, L , are treated as fixed and known.

GAMP models the relationship between the m th observed output y_m and the corresponding noiseless output $z_m \triangleq \mathbf{a}_m^T \mathbf{x}$, where \mathbf{a}_m^T denotes the m th row of \mathbf{A} , using the conditional pdf $p_{Y|Z}(y_m|z_m; \mathbf{q})$. It then approximates the true marginal posterior $p(z_m|\mathbf{y}; \mathbf{q})$ by

$$p_{Z|Y}(z_m|\mathbf{y}; \hat{\mathbf{p}}_m, \mu_m^p, \mathbf{q}) \triangleq \frac{p_{Y|Z}(y_m|z_m; \mathbf{q}) \mathcal{N}(z_m; \hat{\mathbf{p}}_m, \mu_m^p)}{\int_z p_{Y|Z}(y_m|z; \mathbf{q}) \mathcal{N}(z; \hat{\mathbf{p}}_m, \mu_m^p)} \quad (5)$$

using quantities $\hat{\mathbf{p}}_m$ and μ_m^p that change with iteration t (see Table I), although here we suppress the t notation for brevity.

Under the AWGN assumption¹ (4) we have $p_{Y|Z}(y|z; \mathbf{q}) = \mathcal{N}(y; z, \psi)$, and thus the pdf (5) has moments [14]

$$E_{Z|Y}\{z_m | \mathbf{y}; \hat{p}_m, \mu_m^p, \mathbf{q}\} = \hat{p}_m + \frac{\mu_m^p}{\mu_m^p + \psi} (y_m - \hat{p}_m) \quad (6)$$

$$\text{var}_{Z|Y}\{z_m | \mathbf{y}; \hat{p}_m, \mu_m^p, \mathbf{q}\} = \frac{\mu_m^p \psi}{\mu_m^p + \psi}. \quad (7)$$

GAMP then approximates the true marginal posterior $p(x_n | \mathbf{y}; \mathbf{q})$ by

$$p_{X|Y}(x_n | \mathbf{y}; \hat{r}_n, \mu_n^r, \mathbf{q}) \triangleq \frac{p_X(x_n; \mathbf{q}) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r)}{\int_x p_X(x; \mathbf{q}) \mathcal{N}(x; \hat{r}_n, \mu_n^r)} \quad (8)$$

where again \hat{r}_n and μ_n^r vary with the GAMP iteration t .

Plugging the sparse GM prior (3) into (8) and simplifying, one can obtain² the GM-GAMP approximated posterior

$$p_{X|Y}(x_n | \mathbf{y}; \hat{r}_n, \mu_n^r, \mathbf{q}) = \left((1 - \lambda) \delta(x_n) + \lambda \sum_{\ell=1}^L \omega_\ell \mathcal{N}(x_n; \theta_\ell, \phi_\ell) \right) \frac{\mathcal{N}(x_n; \hat{r}_n, \mu_n^r)}{\zeta_n} \quad (9)$$

$$= (1 - \pi_n) \delta(x_n) + \pi_n \sum_{\ell=1}^L \bar{\beta}_{n,\ell} \mathcal{N}(x_n; \gamma_{n,\ell}, \nu_{n,\ell}) \quad (10)$$

with normalization factor

$$\zeta_n \triangleq \int_x p_X(x; \mathbf{q}) \mathcal{N}(x; \hat{r}_n, \mu_n^r) \quad (11)$$

$$= (1 - \lambda) \mathcal{N}(0; \hat{r}_n, \mu_n^r) + \lambda \sum_{\ell=1}^L \omega_\ell \mathcal{N}(0; \hat{r}_n - \theta_\ell, \mu_n^r + \phi_\ell) \quad (12)$$

and $(\hat{r}_n, \mu_n^r, \mathbf{q})$ -dependent quantities

$$\beta_{n,\ell} \triangleq \lambda \omega_\ell \mathcal{N}(\hat{r}_n; \theta_\ell, \phi_\ell + \mu_n^r) \quad (13)$$

$$\bar{\beta}_{n,\ell} \triangleq \frac{\beta_{n,\ell}}{\sum_{k=1}^L \beta_{n,k}} \quad (14)$$

$$\pi_n \triangleq \frac{1}{1 + \left(\frac{\sum_{\ell=1}^L \beta_{n,\ell}}{(1-\lambda) \mathcal{N}(0; \hat{r}_n, \mu_n^r)} \right)^{-1}} \quad (15)$$

$$\gamma_{n,\ell} \triangleq \frac{\hat{r}_n / \mu_n^r + \theta_\ell / \phi_\ell}{1 / \mu_n^r + 1 / \phi_\ell} \quad (16)$$

$$\nu_{n,\ell} \triangleq \frac{1}{1 / \mu_n^r + 1 / \phi_\ell}. \quad (17)$$

¹Because GAMP can handle an arbitrary $p_{Y|Z}(\cdot|\cdot)$, the extension of EM-GM-AMP to additive non-Gaussian noise, and even non-additive measurement channels (such as quantized outputs [21] or logistic regression [14]), is straightforward. Moreover, the parameters of the pdf $p_{Y|Z}(\cdot|\cdot)$ could be learned using a method similar to that which we propose for learning the AWGN variance ψ , as will be evident from the derivation in Section III.A. Finally, one could even model $p_{Y|Z}(\cdot|\cdot)$ as a Gaussian mixture and learn the corresponding parameters.

²Both (10) and (12) can be derived from (9) via the Gaussian-pdf multiplication rule: $\mathcal{N}(x; a, A) \mathcal{N}(x; b, B) = \mathcal{N}(x; \frac{a/A+b/B}{1/A+1/B}, \frac{1}{1/A+1/B}) \mathcal{N}(0; a-b, A+B)$.

The posterior mean and variance of $p_{X|Y}$ are given in steps (R9)–(R10) of Table I, and (10) makes it clear that π_n is GM-GAMP's approximation of the posterior support probability $\Pr\{x_n \neq 0 | \mathbf{y}; \mathbf{q}\}$.

In principle, one could specify GAMP for an arbitrary signal prior $p_X(\cdot)$. However, if the integrals in (R9)–(R10) are not computable in closed form (e.g., when $p_X(\cdot)$ is Student's-t), then they would need to be computed numerically, thereby drastically increasing the computational complexity of GAMP. In contrast, for GM signal models, we see above that all steps can be computed in closed form. Thus, a practical approach to the use of GAMP with an intractable signal prior $p_X(\cdot)$ is to *approximate* $p_X(\cdot)$ using an L -term GM, after which all GAMP steps can be easily implemented. The same approach could also be used to ease the implementation of intractable output priors $p_{Y|Z}(\cdot|\cdot)$.

III. EM LEARNING OF THE PRIOR PARAMETERS \mathbf{q}

We now propose an expectation-maximization (EM) algorithm [15] to learn the prior parameters $\mathbf{q} \triangleq [\lambda, \boldsymbol{\omega}, \boldsymbol{\theta}, \boldsymbol{\phi}, \psi]$. The EM algorithm is an iterative technique that increases a lower bound on the likelihood $p(\mathbf{y}; \mathbf{q})$ at each iteration, thus guaranteeing that the likelihood converges to a local maximum or at least a saddle point [22]. In our case, the EM algorithm manifests as follows. Writing, for arbitrary pdf $\hat{p}(\mathbf{x})$,

$$\ln p(\mathbf{y}; \mathbf{q}) = \int_{\mathbf{x}} \hat{p}(\mathbf{x}) \ln p(\mathbf{y}; \mathbf{q}) \quad (18)$$

$$= \int_{\mathbf{x}} \hat{p}(\mathbf{x}) \ln \left(\frac{p(\mathbf{x}, \mathbf{y}; \mathbf{q})}{\hat{p}(\mathbf{x})} \frac{\hat{p}(\mathbf{x})}{p(\mathbf{x}; \mathbf{y}; \mathbf{q})} \right) \quad (19)$$

$$= \underbrace{E_{\hat{p}(\mathbf{x})} \{ \ln p(\mathbf{x}, \mathbf{y}; \mathbf{q}) \}}_{\triangleq \mathcal{L}_{\hat{p}}(\mathbf{y}; \mathbf{q})} + \underbrace{H(\hat{p})}_{+ D(\hat{p} \| p_{X|Y}(\cdot | \mathbf{y}; \mathbf{q}))} \quad (20)$$

≥ 0

where $E_{\hat{p}(\mathbf{x})}\{\cdot\}$ denotes expectation over $\mathbf{x} \sim \hat{p}(\mathbf{x})$, $H(\hat{p})$ denotes the entropy of pdf \hat{p} , and $D(\hat{p} \| p)$ denotes the Kullback-Leibler (KL) divergence between \hat{p} and p . The non-negativity of the KL divergence implies that $\mathcal{L}_{\hat{p}}(\mathbf{y}; \mathbf{q})$ is a lower bound on $\ln p(\mathbf{y}; \mathbf{q})$, and thus the EM algorithm iterates over two steps: E) choosing \hat{p} to maximize the lower bound for fixed $\mathbf{q} = \mathbf{q}^i$, and M) choosing \mathbf{q} to maximize the lower bound for fixed $\hat{p} = \hat{p}^i$. For the E step, since $\mathcal{L}_{\hat{p}}(\mathbf{y}; \mathbf{q}^i) = \ln p(\mathbf{y}; \mathbf{q}^i) - D(\hat{p} \| p_{X|Y}(\cdot | \mathbf{y}; \mathbf{q}^i))$, the maximizing pdf would clearly be $\hat{p}^i(\mathbf{x}) = p_{X|Y}(\mathbf{x} | \mathbf{y}; \mathbf{q}^i)$, i.e., the true posterior under prior parameters \mathbf{q}^i . Then, for the M step, since $\mathcal{L}_{\hat{p}^i}(\mathbf{y}; \mathbf{q}) = E_{\hat{p}^i(\mathbf{x})} \{ \ln p(\mathbf{x}, \mathbf{y}; \mathbf{q}) \} + H(\hat{p}^i)$, the maximizing \mathbf{q} would clearly be $\mathbf{q}^{i+1} = \arg \max_{\mathbf{q}} E \{ \ln p(\mathbf{x}, \mathbf{y}; \mathbf{q}) | \mathbf{y}; \mathbf{q}^i \}$.

In our case, because the true posterior is very difficult to calculate, we instead construct our lower-bound $\mathcal{L}_{\hat{p}}(\mathbf{y}; \mathbf{q})$ using the GAMP approximated posteriors, i.e., we set $\hat{p}^i(\mathbf{x}) = \prod_n p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)$ for $p_{X|Y}$ defined in (8), resulting in

$$\mathbf{q}^{i+1} = \arg \max_{\mathbf{q}} \hat{E} \{ \ln p(\mathbf{x}, \mathbf{y}; \mathbf{q}) | \mathbf{y}; \mathbf{q}^i \}, \quad (21)$$

where “ \hat{E} ” indicates the use of the GAMP's posterior approximation. Moreover, since the joint optimization in (21) is dif-

difficult to perform, we update \mathbf{q} one component at a time (while holding the others fixed), which is the well known ‘‘incremental’’ variant on EM from [23]. In the sequel, we use ‘‘ $\mathbf{q}_{\setminus\lambda}^i$ ’’ to denote the vector \mathbf{q}^i with the element λ removed (and similar for the other parameters).

A. EM Update of the Gaussian Noise Variance ψ

We first derive the EM update for the noise variance ψ given a previous parameter estimate \mathbf{q}^i . For this, we write $p(\mathbf{x}, \mathbf{y}; \mathbf{q}) = Cp(\mathbf{y}|\mathbf{x}; \psi) = C \prod_{m=1}^M p_{Y|Z}(y_m | \mathbf{a}_m^T \mathbf{x}; \psi)$ for a ψ -invariant constant C , so that

$$\psi^{i+1} = \arg \max_{\psi > 0} \sum_{m=1}^M \hat{\mathbb{E}} \{ \ln p_{Y|Z}(y_m | \mathbf{a}_m^T \mathbf{x}; \psi) | \mathbf{y}; \mathbf{q}^i \} \quad (22)$$

$$= \arg \max_{\psi > 0} \sum_{m=1}^M \int_{z_m} p_{Z|\mathbf{Y}}(z_m | \mathbf{y}; \mathbf{q}^i) \ln p_{Y|Z}(y_m | z_m; \psi) \quad (23)$$

since $z_m = \mathbf{a}_m^T \mathbf{x}$. The maximizing value of ψ in (23) is necessarily a value of ψ that zeroes the derivative of the sum, i.e., that satisfies³

$$\sum_{m=1}^M \int_{z_m} p_{Z|\mathbf{Y}}(z_m | \mathbf{y}; \mathbf{q}^i) \frac{d}{d\psi} \ln p_{Y|Z}(y_m | z_m; \psi) = 0. \quad (24)$$

Because $p_{Y|Z}(y_m | z_m; \psi) = \mathcal{N}(y_m; z_m, \psi)$, we can obtain

$$\frac{d}{d\psi} \ln p_{Y|Z}(y_m | z_m; \psi) = \frac{1}{2} \left(\frac{|y_m - z_m|^2}{\psi^2} - \frac{1}{\psi} \right), \quad (25)$$

which, when plugged into (24), yields the unique solution

$$\psi^{i+1} = \frac{1}{M} \sum_{m=1}^M \int_{z_m} p_{Z|\mathbf{Y}}(z_m | \mathbf{y}; \mathbf{q}^i) |y_m - z_m|^2 \quad (26)$$

$$= \frac{1}{M} \sum_{m=1}^M (|y_m - \hat{z}_m|^2 + \mu_m^z), \quad (27)$$

where the use of \hat{z}_m and μ_m^z follows from (R3)–(R4) in Table I.

B. EM Updates of the Signal Parameters: BG Case

Suppose that the signal distribution $p_X(\cdot)$ is modeled using an $L = 1$ -term GM, i.e., a Bernoulli-Gaussian (BG) pdf. In this case, the marginal signal prior in (3) reduces to

$$p_X(x; \lambda, \omega, \theta, \phi) = (1 - \lambda)\delta(x) + \lambda\mathcal{N}(x; \theta, \phi). \quad (28)$$

Note that, in the BG case, the mixture weight ω is, by definition, unity and does not need to be learned.

We now derive the EM update for λ given previous parameters $\mathbf{q}^i \triangleq [\lambda^i, \theta^i, \phi^i, \psi^i]$. Because we can write $p(\mathbf{x}, \mathbf{y}; \mathbf{q}) = C \prod_{n=1}^N p_X(x_n; \lambda, \theta, \phi)$ for a λ -invariant constant C ,

$$\lambda^{i+1} = \arg \max_{\lambda \in (0,1)} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \lambda, \mathbf{q}_{\setminus\lambda}^i) | \mathbf{y}; \mathbf{q}^i \right\}. \quad (29)$$

³The continuity of both the integrand and its partial derivative with respect to ψ allow the use of Leibniz’s integral rule to exchange differentiation and integration.

The maximizing value of λ in (29) is necessarily a value of λ that zeroes the derivative of the sum, i.e., that satisfies⁴

$$\sum_{n=1}^N \int_{x_n} p_{X|\mathbf{Y}}(x_n | \mathbf{y}; \mathbf{q}^i) \frac{d}{d\lambda} \ln p_X(x_n; \lambda, \mathbf{q}_{\setminus\lambda}^i) = 0. \quad (30)$$

For the BG $p_X(x_n; \lambda, \theta, \phi)$ in (28), it is readily seen that

$$\frac{d}{d\lambda} \ln p_X(x_n; \lambda, \mathbf{q}_{\setminus\lambda}^i) = \frac{\mathcal{N}(x_n; \theta^i, \phi^i) - \delta(x_n)}{p_X(x_n; \lambda, \mathbf{q}_{\setminus\lambda}^i)} \quad (31)$$

$$= \begin{cases} \frac{1}{\lambda} & x_n \neq 0 \\ \frac{-1}{1-\lambda} & x_n = 0. \end{cases} \quad (32)$$

Plugging (32) and (9) into (30), it becomes evident that the neighborhood around the point $x_n = 0$ should be treated differently than the remainder of \mathbb{R} . Thus, we define the closed ball $\mathcal{B}_\epsilon \triangleq [-\epsilon, \epsilon]$ and its complement $\bar{\mathcal{B}}_\epsilon \triangleq \mathbb{R} \setminus \mathcal{B}_\epsilon$, and note that, in the limit $\epsilon \rightarrow 0$, the following is equivalent to (30):

$$\begin{aligned} \sum_{n=1}^N \underbrace{\int_{x_n \in \bar{\mathcal{B}}_\epsilon} p_{X|\mathbf{Y}}(x_n | \mathbf{y}; \mathbf{q}^i)}_{\stackrel{\epsilon \rightarrow 0}{=} \pi_n} &= \frac{\lambda}{1-\lambda} \sum_{n=1}^N \underbrace{\int_{x_n \in \mathcal{B}_\epsilon} p_{X|\mathbf{Y}}(x_n | \mathbf{y}; \mathbf{q}^i)}_{\stackrel{\epsilon \rightarrow 0}{=} 1 - \pi_n} \quad (33) \end{aligned}$$

where the values taken by the integrals are evident from (10). Finally, the EM update for λ is the unique value satisfying (33) as $\epsilon \rightarrow 0$, which is readily shown to be

$$\lambda^{i+1} = \frac{1}{N} \sum_{n=1}^N \pi_n. \quad (34)$$

Conveniently, the posterior support probabilities $\{\pi_n\}_{n=1}^N$ are easily calculated from the GM-GAMP outputs via (15).

Similar to (29), the EM update for θ can be written as

$$\theta^{i+1} = \arg \max_{\theta \in \mathbb{R}} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \theta, \mathbf{q}_{\setminus\theta}^i) | \mathbf{y}; \mathbf{q}^i \right\}. \quad (35)$$

The maximizing value of θ in (35) is again a necessarily a value of θ that zeroes the derivative, i.e., that satisfies

$$\sum_{n=1}^N \int_{x_n} p_{X|\mathbf{Y}}(x_n | \mathbf{y}; \mathbf{q}^i) \frac{d}{d\theta} \ln p_X(x_n; \theta, \mathbf{q}_{\setminus\theta}^i) = 0. \quad (36)$$

For the BG $p_X(x_n; \lambda, \theta, \phi)$ given in (28),

$$\frac{d}{d\theta} \ln p_X(x_n; \lambda^i, \theta, \phi^i) = \frac{(x_n - \theta) \lambda^i \mathcal{N}(x_n; \theta, \phi^i)}{\phi^i p_X(x_n; \theta, \mathbf{q}_{\setminus\theta}^i)} \quad (37)$$

$$= \begin{cases} \frac{x_n - \theta}{\phi^i} & x_n \neq 0 \\ 0 & x_n = 0. \end{cases} \quad (38)$$

⁴To justify the exchange of differentiation and integration via Leibniz’s integral rule here, one could employ the Dirac approximation $\delta(x) = \mathcal{N}(x; 0, \epsilon)$ for fixed arbitrarily small $\epsilon > 0$, after which the integrand and its derivative w.r.t. λ become continuous. The same comment applies in to all exchanges of differentiation and integration in the sequel.

Splitting the domain of integration in (36) into \mathcal{B}_ϵ and $\bar{\mathcal{B}}_\epsilon$ as before, and then plugging in (38), we find that the following is equivalent to (36) in the limit of $\epsilon \rightarrow 0$:

$$\sum_{n=1}^N \int_{x_n \in \bar{\mathcal{B}}_\epsilon} (x_n - \theta) p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i) = 0. \quad (39)$$

The unique value of θ satisfying (39) as $\epsilon \rightarrow 0$ is then

$$\theta^{i+1} = \frac{\sum_{n=1}^N \lim_{\epsilon \rightarrow 0} \int_{x_n \in \bar{\mathcal{B}}_\epsilon} x_n p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)}{\sum_{n=1}^N \lim_{\epsilon \rightarrow 0} \int_{x_n \in \bar{\mathcal{B}}_\epsilon} p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)} \quad (40)$$

$$= \frac{1}{\lambda^{i+1} N} \sum_{n=1}^N \pi_n \gamma_{n,1} \quad (41)$$

where $\{\gamma_{n,1}\}_{n=1}^N$ defined in (16) are easily computed from the GM-GAMP outputs. The equality in (41) can be verified by plugging the GAMP posterior expression (10) into (40).

Similar to (29), the EM update for ϕ can be written as

$$\hat{\phi}^{i+1} = \arg \max_{\phi > 0} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \phi, \mathbf{q}_{\phi}^i) \mid \mathbf{y}; \mathbf{q}^i \right\}. \quad (42)$$

The maximizing value of ϕ in (42) is again necessarily a value of ϕ that zeroes the derivative, i.e., that satisfies

$$\sum_{n=1}^N \int_{x_n} p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i) \frac{d}{d\phi} \ln p_X(x_n; \phi, \mathbf{q}_{\phi}^i) = 0. \quad (43)$$

For the $p_X(x_n; \lambda, \theta, \phi)$ given in (28), it is readily seen that

$$\begin{aligned} & \frac{d}{d\phi} \ln p_X(x_n; \lambda^i, \theta^i, \phi) \\ &= \frac{1}{2} \left(\frac{|x_n - \theta^i|^2}{(\phi)^2} - \frac{1}{\phi} \right) \frac{\lambda^i \mathcal{N}(x_n; \theta^i, \phi)}{p_X(x_n; \phi, \mathbf{q}_{\phi}^i)} \\ &= \begin{cases} \frac{1}{2} \left(\frac{|x_n - \theta^i|^2}{(\phi)^2} - \frac{1}{\phi} \right) & x_n \neq 0 \\ 0 & x_n = 0 \end{cases} \end{aligned} \quad (44)$$

Splitting the domain of integration in (43) into \mathcal{B}_ϵ and $\bar{\mathcal{B}}_\epsilon$ as before, and then plugging in (44), we find that the following is equivalent to (43) in the limit of $\epsilon \rightarrow 0$:

$$\sum_{n=1}^N \int_{x_n \in \bar{\mathcal{B}}_\epsilon} (|x_n - \theta^i|^2 - \phi) p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i) = 0. \quad (45)$$

The unique value of ϕ satisfying (45) as $\epsilon \rightarrow 0$ is then

$$\phi^{i+1} = \frac{\sum_{n=1}^N \lim_{\epsilon \rightarrow 0} \int_{x_n \in \bar{\mathcal{B}}_\epsilon} |x_n - \theta^i|^2 p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)}{\sum_{n=1}^N \lim_{\epsilon \rightarrow 0} \int_{x_n \in \bar{\mathcal{B}}_\epsilon} p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)}. \quad (46)$$

Finally, we expand $|x_n - \theta^i|^2 = |x_n|^2 - 2\text{Re}(x_n^* \theta^i) + |\theta^i|^2$ which gives

$$\phi^{i+1} = \frac{1}{\lambda^{i+1} N} \sum_{n=1}^N \pi_n (|\theta^i - \gamma_{n,1}|^2 + \nu_{n,1}) \quad (47)$$

where $\{\nu_{n,1}\}_{n=1}^N$ from (17) are easily computed from the GAMP outputs. The equality in (47) can be readily verified by plugging (10) into (46).

C. EM Updates of the Signal Parameters: GM Case

We now generalize the EM updates derived in Section III.B to the GM prior given in (3) for $L \geq 1$. As we shall see, it is not possible to write the exact EM updates in closed-form when $L > 1$, and so some approximations will be made.

We begin by deriving the EM update for λ given the previous parameters $\mathbf{q}^i \triangleq [\lambda^i, \boldsymbol{\omega}^i, \boldsymbol{\theta}^i, \boldsymbol{\phi}^i, \boldsymbol{\psi}^i]$. The first two steps are identical to the steps (29) and (30) presented for the BG case, and for brevity we do not repeat them here. In the third step, use of the GM prior (3) yields

$$\begin{aligned} \frac{d}{d\lambda} \ln p_X(x_n; \lambda, \mathbf{q}_{\lambda}^i) &= \frac{\sum_{\ell=1}^L \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i) - \delta(x_n)}{p_X(x_n; \lambda, \mathbf{q}_{\lambda}^i)} \\ &= \begin{cases} \frac{1}{\lambda} & x_n \neq 0 \\ \frac{-1}{1-\lambda} & x_n = 0 \end{cases}, \end{aligned} \quad (48)$$

which coincides with the BG expression (32). The remaining steps also coincide with those in the BG case, and so the final EM update for λ , in the case of a GM,⁵ is given by (34).

We next derive the EM updates for the GM parameters $\boldsymbol{\omega}$, $\boldsymbol{\theta}$, and $\boldsymbol{\phi}$. For each $k = 1, \dots, L$, we incrementally update θ_k , then ϕ_k , and then the entire vector $\boldsymbol{\omega}$, while holding all other parameters fixed. The EM updates are thus

$$\theta_k^{i+1} = \arg \max_{\theta_k \in \mathbb{R}} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \theta_k, \mathbf{q}_{\theta_k}^i) \mid \mathbf{y}; \mathbf{q}^i \right\}, \quad (49)$$

$$\phi_k^{i+1} = \arg \max_{\phi_k > 0} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \phi_k, \mathbf{q}_{\phi_k}^i) \mid \mathbf{y}; \mathbf{q}^i \right\} \quad (50)$$

$$\boldsymbol{\omega}^{i+1} = \arg \max_{\boldsymbol{\omega} > 0: \sum_k \omega_k = 1} \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \boldsymbol{\omega}, \mathbf{q}_{\boldsymbol{\omega}}^i) \mid \mathbf{y}; \mathbf{q}^i \right\}. \quad (51)$$

Following (36), the maximizing value of θ_k in (49) is again necessarily a value of θ_k that zeros the derivative, i.e.,

$$\sum_{n=1}^N \int_{x_n} p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i) \frac{d}{d\theta_k} \ln p_X(x_n; \theta_k, \mathbf{q}_{\theta_k}^i) = 0, \quad (52)$$

Plugging in the derivative [see (53), shown at the bottom of the page] and the version of $p_{X|Y}(x_n | \mathbf{y}; \mathbf{q}^i)$ from (9), integrating (52) separately over \mathcal{B}_ϵ and $\bar{\mathcal{B}}_\epsilon$ as in (33), and taking $\epsilon \rightarrow 0$, we find that the \mathcal{B}_ϵ portion vanishes, giving the necessary condition

$$\sum_{n=1}^N \int_{x_n} \frac{p(x_n | x_n \neq 0, \mathbf{y}; \mathbf{q}^i) \lambda^i \omega_k^i \mathcal{N}(x_n; \theta_k, \phi_k^i) (x_n - \theta_k)}{\zeta_n(\omega_k^i \mathcal{N}(x_n; \theta_k, \phi_k^i) + \sum_{\ell \neq k} \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i))} = 0. \quad (54)$$

⁵The arguments in this section reveal that, under signal priors of the form $p_X(x) = (1 - \lambda)\delta(x) + \lambda f_X(x)$, where $f_X(\cdot)$ can be arbitrary, the EM update for λ is that given in (34).

$$\frac{d}{d\theta_k} \ln p_X(x_n; \theta_k, \mathbf{q}_{\theta_k}^i) = \left(\frac{x_n - \theta_k}{\phi_k^i} \right) \frac{\lambda^i \omega_k^i \mathcal{N}(x_n; \theta_k, \phi_k^i)}{(1 - \lambda^i)\delta(x_n) + \lambda^i (\omega_k^i \mathcal{N}(x_n; \theta_k, \phi_k^i) + \sum_{\ell \neq k} \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i))} \quad (53)$$

Since this integral cannot be evaluated in closed form, we apply the approximation $\mathcal{N}(x_n; \theta_k, \phi_k^i) \approx \mathcal{N}(x_n; \theta_k^i, \phi_k^i)$ in both the numerator and denominator, and subsequently exploit the fact that $p(x_n|x_n \neq 0, \mathbf{y}; \mathbf{q}^i) = \mathcal{N}(x_n; \hat{r}_n, \mu_n^r) \sum_{\ell} \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i)$ from (9) to cancel terms, and so obtain the (approximated) necessary condition

$$\sum_{n=1}^N \int_{x_n} \frac{\lambda^i \omega_k^i \mathcal{N}(x_n; \hat{r}_n, \mu_n^r) \mathcal{N}(x_n; \theta_k^i, \phi_k^i)}{\zeta_n} (x_n - \theta_k) = 0. \quad (55)$$

We then simplify (55) using the Gaussian-pdf multiplication rule, and set θ_k^{i+1} equal to the value of θ_k that satisfies (55), which can be found to be

$$\theta_k^{i+1} = \frac{\sum_{n=1}^N \pi_n \bar{\beta}_{n,k} \gamma_{n,k}}{\sum_{n=1}^N \pi_n \bar{\beta}_{n,k}} \quad (56)$$

Note from (10) that $\pi_n \bar{\beta}_{n,k}$ can be interpreted as the probability that x_n originated from the k th mixture component.

For sparse signals \mathbf{x} , we find that learning the GM means $\{\theta_k\}$ using the above EM procedure yields excellent recovery MSE. However, for “heavy-tailed” signals (i.e., whose pdfs have tails that are not exponentially bounded, such as Student’s-t), our experience indicates that the EM-learned values of $\{\theta_k\}$ tend to gravitate towards the outliers in $\{x_n\}_{n=1}^N$, resulting in an overfitting of $p_X(\cdot)$ and thus poor reconstruction MSE. For such heavy-tailed signals, we find that better reconstruction performance is obtained by fixing the means at zero (i.e., $\theta_k^i = 0 \forall k, i$). Thus, in the remainder of the paper, we consider two modes of operation: a “sparse” mode where θ is learned via the above EM procedure, and a “heavy-tailed” mode that fixes $\theta = \mathbf{0}$.

Following (52), the maximizing value of ϕ_k in (50) is necessarily a value of ϕ_k that zeroes the derivative, i.e.,

$$\sum_{n=1}^N \int_{x_n} p_{X|Y}(x_n|\mathbf{y}; \mathbf{q}^i) \frac{d}{d\phi_k} \ln p_X(x_n; \phi_k, \mathbf{q}_{\setminus \phi_k}^i) = 0. \quad (57)$$

As for the derivative in the previous expression, we find (58), shown at the bottom of the page, integrating (57) separately over \mathcal{B}_{ϵ} and $\bar{\mathcal{B}}_{\epsilon}$, as in (33), and taking $\epsilon \rightarrow 0$, we find that the \mathcal{B}_{ϵ} portion vanishes, giving

$$\sum_{n=1}^N \int_{x_n} \frac{p(x_n|x_n \neq 0, \mathbf{y}; \mathbf{q}^i) \lambda^i \omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k) / \zeta_n}{\omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k) + \sum_{\ell \neq k} \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i)} \times \left(\frac{|x_n - \theta_k^i|^2}{\phi_k} - 1 \right) \quad (59)$$

Similar to (54), this integral is difficult to evaluate, and so we again apply the approximation $\mathcal{N}(x_n; \theta_k^i, \phi_k) \approx \mathcal{N}(x_n; \theta_k^i, \phi_k^i)$

in the numerator and denominator, after which several terms cancel, yielding the necessary condition

$$\sum_{n=1}^N \int_{x_n} \frac{\mathcal{N}(x_n; \hat{r}_n, \mu_n^r) \lambda^i \omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k^i)}{\zeta_n} \times \left(\frac{|x_n - \theta_k^i|^2}{\phi_k} - 1 \right) = 0. \quad (60)$$

To find the value of ϕ_k satisfying (60), we expand $|x_n - \theta_k^i|^2 = |x_n|^2 - 2\text{Re}(x_n^* \theta_k^i) + |\theta_k^i|^2$ and apply the Gaussian-pdf multiplication rule, which gives

$$\phi_k^{i+1} = \frac{\sum_{n=1}^N \pi_n \bar{\beta}_{n,k} (|\theta_k^i - \gamma_{n,k}|^2 + \nu_{n,k})}{\sum_{n=1}^N \pi_n \bar{\beta}_{n,k}}. \quad (61)$$

Finally, the value of the positive ω maximizing (51) under the pmf constraint $\sum_{k=1}^L \omega_k = 1$ can be found by solving the unconstrained optimization problem $\max_{\omega, \xi} J(\omega, \xi)$, where ξ is a Lagrange multiplier and

$$\begin{aligned} J(\omega, \xi) &\triangleq \sum_{n=1}^N \hat{\mathbb{E}} \left\{ \ln p_X(x_n; \omega, \mathbf{q}^i \omega) | \mathbf{y}; \mathbf{q}^i \right\} \\ &\quad - \xi \left(\sum_{\ell=1}^L \omega_{\ell} - 1 \right) \\ &= \sum_{n=1}^N \int_{x_n} p_{X|Y}(x_n|\mathbf{y}; \mathbf{q}^i) \ln p_X(x_n; \omega, \mathbf{q}^i \omega) \\ &\quad - \xi \left(\sum_{\ell=1}^L \omega_{\ell} - 1 \right). \end{aligned} \quad (62)$$

We start by setting $\frac{d}{d\omega_k} J(\omega, \xi) = 0$, which yields

$$\begin{aligned} &\sum_{n=1}^N \int_{x_n} \frac{p_X(x_n; \mathbf{q}^i) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r)}{\zeta_n} \\ &\quad \times \frac{d}{d\omega_k} \ln p_X(x_n; \omega, \mathbf{q}^i \omega) = \xi. \end{aligned} \quad (63)$$

$$\Leftrightarrow \sum_{n=1}^N \int_{x_n} \frac{p_X(x_n; \mathbf{q}^i) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r) \lambda^i \mathcal{N}(x_n; \theta_k^i, \phi_k^i)}{\zeta_n p_X(x_n; \omega, \mathbf{q}^i \omega)} = \xi. \quad (64)$$

Like in (54) and (59), the above integral is difficult to evaluate, and so we approximate $\omega \approx \omega^i$, which reduces the previous equation to

$$\xi = \sum_{n=1}^N \int_{x_n} \frac{\lambda^i \mathcal{N}(x_n; \theta_k^i, \phi_k^i) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r)}{\zeta_n}. \quad (65)$$

$$\frac{d}{d\phi_k} \ln p_X(x_n; \phi_k, \mathbf{q}_{\setminus \phi_k}^i) = \frac{1}{2} \left(\frac{|x_n - \theta_k^i|^2}{\phi_k^2} - \frac{1}{\phi_k} \right) \frac{\lambda^i \omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k)}{(1 - \lambda^i) \delta(x_n) + \lambda^i (\omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k) + \sum_{\ell \neq k} \omega_{\ell}^i \mathcal{N}(x_n; \theta_{\ell}^i, \phi_{\ell}^i))}. \quad (58)$$

Multiplying both sides by ω_k^i for $k = 1, \dots, L$, summing over k , employing the fact $1 = \sum_k \omega_k^i$, and simplifying, we obtain the equivalent condition

$$\xi = \sum_{n=1}^N \int_{x_n} \frac{\lambda^i \sum_{k=1}^L \omega_k^i \mathcal{N}(x_n; \theta_k^i, \phi_k^i) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r)}{\zeta_n} \quad (66)$$

$$= \sum_{n=1}^N \pi_n. \quad (67)$$

Plugging (67) into (65) and multiplying both sides by ω_k , the derivative-zeroing value of ω_k is seen to be

$$\omega_k = \frac{\sum_{n=1}^N \lambda^i \omega_k \mathcal{N}(x_n; \theta_k^i, \phi_k^i) \mathcal{N}(x_n; \hat{r}_n, \mu_n^r) / \zeta_n}{\sum_{n=1}^N \pi_n}, \quad (68)$$

where, if we use $\omega_k \approx \omega_k^i$ on the right of (68), then we obtain

$$\omega_k^{i+1} = \frac{\sum_{n=1}^N \pi_n \bar{\beta}_{n,k}}{\sum_{n=1}^N \pi_n}. \quad (69)$$

Although, for the case of GM priors, approximations were used in the derivation of the EM updates (56), (61), and (69), it is interesting to note that, in the case of $L = 1$ mixture components, these approximate EM-GM updates coincide with the *exact* EM-BG updates derived in Section III.B. In particular, the approximate-EM update of the GM parameter θ_1 in (56) coincides with the exact-EM update of the BG parameter θ in (41), the approximate-EM update of the GM parameter ϕ_1 in (61) coincides with the exact-EM update of the BG parameter ϕ in (47), and the approximate-EM update of the GM parameter ω_1 in (69) reduces to the fixed value 1. Thus, one can safely use the GM updates above in the BG setting without any loss of optimality.

D. EM Initialization

Since the EM algorithm may converge to a local maximum or at least a saddle point of the likelihood function, proper initialization of the unknown parameters \mathbf{q} is essential. Here, we propose initialization strategies for both the “sparse” and “heavy-tailed” modes of operation, for a given value of L . Regarding the value of L , we prescribe a method to learn it in Section III.F. However, the fixed choices $L = 3$ for “sparse” mode and $L = 4$ for “heavy tailed” mode usually perform well, as shown in Section IV.

For the “sparse” mode, we set the initial sparsity rate λ^0 equal to the theoretical noiseless LASSO PTC, i.e., $\lambda^0 = \frac{M}{N} \rho_{SE}(\frac{M}{N})$, where [10]

$$\rho_{SE}\left(\frac{M}{N}\right) = \max_{c>0} \frac{1 - \frac{2N}{M} [(1+c^2)\Phi(-c) - c\phi(c)]}{1 + c^2 - 2[(1+c^2)\Phi(-c) - c\phi(c)]} \quad (70)$$

describes the maximum value of $\frac{K}{M}$ supported by LASSO for a given $\frac{M}{N}$, and where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cdf and pdf of the $\mathcal{N}(0, 1)$ distribution, respectively. Using the energies $\|\mathbf{y}\|_2^2$ and $\|\mathbf{A}\|_F^2$ and an assumed value of SNR^0 , we initialize the noise and signal variances, respectively, as

$$\psi^0 = \frac{\|\mathbf{y}\|_2^2}{(\text{SNR}^0 + 1)M}, \quad \varphi^0 = \frac{\|\mathbf{y}\|_2^2 - M\psi^0}{\|\mathbf{A}\|_F^2 \lambda^0}, \quad (71)$$

TABLE II
THE EM-GM-AMP ALGORITHM (FIXED- L CASE)

```

Initialize  $L$  and  $\mathbf{q}^0$  as described in Section III-D.
Initialize  $\hat{\mathbf{x}}^0 = \mathbf{0}$ .
for  $i = 1$  to  $I_{\max}$  do
  Generate  $\hat{\mathbf{x}}^i, \hat{\mathbf{z}}^i, (\boldsymbol{\mu}^z)^i, \boldsymbol{\pi}^i, \{\beta_k^i, \gamma_k^i, \nu_k^i\}_{k=1}^L$  using GM-GAMP
  with  $\mathbf{q}^{i-1}$  (see Table I).
  if  $\|\hat{\mathbf{x}}^i - \hat{\mathbf{x}}^{i-1}\|_2^2 < \tau_{\text{em}} \|\hat{\mathbf{x}}^{i-1}\|_2^2$  then
    break.
  end if
  Compute  $\lambda^i$  from  $\boldsymbol{\pi}^{i-1}$  as described in (34).
  for  $k = 1$  to  $L$  do
    if sparse mode enabled then
      Compute  $\theta_k^i$  from  $\boldsymbol{\pi}^{i-1}, \gamma_k^{i-1}, \{\beta_l^{i-1}\}_{l=1}^L$  as described in
      (56).
    else if heavy-tailed mode enabled then
      Set  $\theta_k^i = 0$ .
    end if
    Compute  $\phi_k^i$  from  $\theta_k^{i-1}, \boldsymbol{\pi}^{i-1}, \gamma_k^{i-1}, \nu_k^{i-1}, \{\beta_l^{i-1}\}_{l=1}^L$  as
    described in (61).
    Compute  $\omega^i$  from  $\boldsymbol{\pi}^{i-1}$  and  $\{\beta_l^{i-1}\}_{l=1}^L$  as described in (69).
  end for
  Compute  $\psi^i$  from  $\hat{\mathbf{z}}^i$  and  $(\boldsymbol{\mu}^z)^i$  as in (27).
end for

```

where, in the absence of (user provided) knowledge about the true $\text{SNR} \triangleq \|\mathbf{A}\mathbf{x}\|_2^2 / \|\mathbf{w}\|_2^2$, we suggest $\text{SNR}^0 = 100$, because in our experience this value works well over a wide range of true SNR. Then, we uniformly space the initial GM means $\boldsymbol{\theta}^0$ over $[\frac{-L+1}{2L}, \frac{L-1}{2L}]$, and subsequently fit the mixture weights $\boldsymbol{\omega}^0$ and variances $\boldsymbol{\phi}^0$ to the uniform pdf supported on $[-0.5, 0.5]$ (which can be done offline using the standard approach to EM-fitting of GM parameters, e.g., ([24], p. 435)). Finally, we multiply $\boldsymbol{\theta}^0$ by $\sqrt{12\varphi^0}$ and $\boldsymbol{\phi}^0$ by $12\varphi^0$ to ensure that the resulting signal variance equals φ^0 .

For the “heavy-tailed” mode, we initialize λ^0 and ψ^0 as above and set, for $k = 1, \dots, L$,

$$\omega_k^0 = \frac{1}{L}, \quad \phi_k^0 = \frac{k}{\sqrt{L}} \frac{(\|\mathbf{y}\|_2^2 - M\psi^0)}{\|\mathbf{A}\|_F^2 \lambda^0}, \quad \text{and } \theta_k^0 = 0. \quad (72)$$

E. EM-GM-AMP Summary and Demonstration

The fixed- L EM-GM-AMP⁶ algorithm developed in the previous sections is summarized in Table II. For EM-BG-AMP (as previously described in [2]), one would simply run EM-GM-AMP with $L = 1$.

To demonstrate EM-GM-AMP’s ability to learn the underlying signal distribution, Fig. 1 shows examples of the GM-modeled signal distributions learned by EM-GM-AMP in both “sparse” and “heavy-tailed” modes. To create the figure, we first constructed the true signal vector $\mathbf{x} \in \mathbb{R}^N$ using $N = 2000$ independent draws of the true distribution $p_X(\cdot)$ shown in each of the subplots. Then, we constructed measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}$ by drawing $\mathbf{A} \in \mathbb{R}^{M \times N}$ with i.i.d. $\mathcal{N}(0, M^{-1})$ elements and $\mathbf{w} \in \mathbb{R}^M$ with i.i.d. $\mathcal{N}(0, \sigma^2)$ elements, with $M = 1000$ and σ^2 chosen to achieve $\text{SNR} = 25$ dB. Finally, we ran EM-GM-AMP according to Table II, and plotted the GM approximation $p_X(x; \mathbf{q}^i)$ from (3) using the learned pdf parameters $\mathbf{q}^i = [\lambda^i, \boldsymbol{\omega}^i, \boldsymbol{\theta}^i, \boldsymbol{\phi}^i, \psi^i]$. Fig. 1 confirms that EM-GM-AMP is successful in learning a reasonable approximation of the unknown true pdf $p_X(\cdot)$ from the noisy

⁶Matlab code at <http://www.ece.osu.edu/~schmitter/EMturboGAMP>.

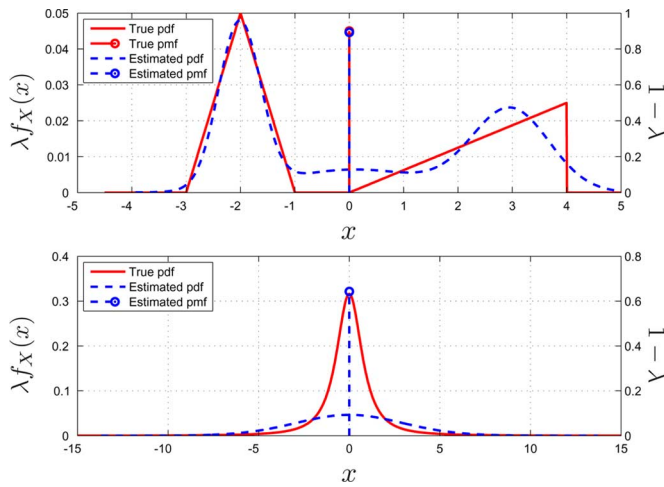


Fig. 1. True and EM-GM-AMP-learned versions of the signal distribution $p_X(x) = \lambda f_X(x) + (1 - \lambda)\delta(x)$. The top subplot shows “sparse” mode EM-GM-AMP run using GM-order $L = 3$ on a sparse signal whose non-zero components were generated according to a triangular mixture, whereas the bottom subplot shows “heavy-tailed” EM-GM-AMP run using $L = 4$ on a Student’s-t signal with rate parameter $q = 1.67$ (defined in (82)). The density of the continuous component $\lambda f_X(x)$ is marked on the left axis, while the mass of the discrete component $(1 - \lambda)\delta(x)$ is marked on the right axis.

compressed observations \mathbf{y} , in both sparse and heavy-tailed modes.

F. Selection of GM Model Order L

We now propose a method to learn the number of GM components, L , based on standard maximum likelihood (ML)-based model-order-selection methodology [25], i.e.,

$$\arg \max_{L \in \mathbb{Z}^+} \ln p(\mathbf{y}; \hat{\mathbf{q}}_L) - \eta(L), \tag{73}$$

where $\hat{\mathbf{q}}_L$ is the ML estimate of \mathbf{q} under the hypothesis L and $\eta(L)$ is a penalty term. For $\eta(L)$, there are several possibilities, but we focus on the Bayesian information criterion (BIC) [25]:

$$\eta_{\text{BIC}}(L) = |\hat{\mathbf{q}}_L| \ln U, \tag{74}$$

where $|\hat{\mathbf{q}}_L|$ denotes the number⁷ of real-valued parameters affected by L , and U is the sample size (see below).

Because $\ln p(\mathbf{y}; \hat{\mathbf{q}}_L)$ is difficult to evaluate, we work with the lower bound (where for now L^j , $\hat{\mathbf{q}}_{L^j}$, and $\hat{\mathbf{q}}_{L^{j+1}}$ are arbitrary)

$$\ln p(\mathbf{y}; \hat{\mathbf{q}}_L) = \ln \int_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}; \hat{\mathbf{q}}_{L^j}) \frac{p(\mathbf{x}, \mathbf{y}; \hat{\mathbf{q}}_L)}{p(\mathbf{x}|\mathbf{y}; \hat{\mathbf{q}}_{L^j})} \tag{75}$$

$$\geq \int_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}; \hat{\mathbf{q}}_{L^j}) \ln \frac{p(\mathbf{x}, \mathbf{y}; \hat{\mathbf{q}}_L)}{p(\mathbf{x}|\mathbf{y}; \hat{\mathbf{q}}_{L^j})} \tag{76}$$

$$= \int_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}; \hat{\mathbf{q}}_{L^j}) \ln p(\mathbf{x}, \mathbf{y}; \hat{\mathbf{q}}_L) + \text{const} \tag{77}$$

$$= \sum_{n=1}^N \int_{x_n} p(x_n|\mathbf{y}; \hat{\mathbf{q}}_{L^j}) \ln p_X(x_n; \hat{\mathbf{q}}_L) + \text{const} \tag{78}$$

⁷In our case, the parameters affected by L are the GM means, variances, and weights, so that, for real-valued signals, we use $|\hat{\mathbf{q}}_L| = 3L - 1$ in “sparse” mode and $|\hat{\mathbf{q}}_L| = 2L - 1$ in heavy-tailed mode, and for complex-valued signals, we use $|\hat{\mathbf{q}}_L| = 4L - 1$ in “sparse” mode and $|\hat{\mathbf{q}}_L| = 2L - 1$ in heavy-tailed mode.

$$= \sum_{n=1}^N \underbrace{\int_{x_n \neq 0} p(x_n|\mathbf{y}; \hat{\mathbf{q}}_{L^j}) \ln f_X(x_n; \hat{\mathbf{q}}_L)}_{\triangleq \mathcal{L}_{L^j}(\mathbf{y}; \hat{\mathbf{q}}_L)} + \text{const}, \tag{79}$$

where (76) applies Jensen’s inequality, “const” denotes a constant term w.r.t. L , and (78) holds because $\ln p(\mathbf{x}, \mathbf{y}; \hat{\mathbf{q}}_L) = \ln p(\mathbf{x}; \hat{\mathbf{q}}_L) + \ln p(\mathbf{y}|\mathbf{x}; \hat{\psi}) = \sum_{n=1}^N \ln p_X(x_n; \hat{\mathbf{q}}_L) + \text{const}$. Equation (79) can then be obtained integrating (78) separately over \mathcal{B}_ϵ and $\bar{\mathcal{B}}_\epsilon$ and taking $\epsilon \rightarrow 0$, as done several times in Section III.B. Using this lower bound in place of $\ln p(\mathbf{y}; \hat{\mathbf{q}}_L)$ in (73), we obtain the BIC-inspired model order estimate (where now $\hat{\mathbf{q}}_L$ is specifically the ML estimate of \mathbf{q}_L)

$$L^{j+1} \triangleq \arg \max_{L \in \mathbb{Z}^+} \mathcal{L}_{L^j}(\mathbf{y}; \hat{\mathbf{q}}_L) - \eta_{\text{BIC}}(L). \tag{80}$$

We in fact propose to perform (80) iteratively, with $j = 0, 1, 2, \dots$ denoting the iteration index. Notice that (80) can be interpreted as a “penalized” EM update for L ; if we neglect the penalty term $\eta(L)$, then (75)–(79) becomes a standard derivation for the EM-update of L (recall, e.g., the EM derivation in Section III). The penalty term is essential, though, because the unpenalized log-likelihood lower bound $\mathcal{L}_{L^j}(\mathbf{y}; \hat{\mathbf{q}}_L)$ is non-decreasing⁸ in L .

We now discuss several practical aspects of our procedure. First, we are forced to approximate the integral in (79). To start, we use GM-GAMP’s approximation of the posterior $p(x_n|\mathbf{y}; \hat{\mathbf{q}}_{L^j})$ from (9), and the EM approximations of the ML-estimates $\hat{\mathbf{q}}_{L^j}$ and $\hat{\mathbf{q}}_L$ outlined in Section III.C. In this case, the integral in (79) takes the form

$$\int_{x_n} \pi_n \sum_{l=1}^{L^j} \bar{\beta}_{n,l} \mathcal{N}(x_n; \gamma_{n,l}, \nu_{n,l}) \ln \sum_{k=1}^L \omega_k \mathcal{N}(x_n; \theta_k, \phi_k) \tag{81}$$

which is still difficult due to the log term. Hence, we evaluate (81) using the point-mass approximation $\mathcal{N}(x_n; \gamma_{n,l}, \nu_{n,l}) \approx \delta(x_n - \gamma_{n,l})$. Second, for the BIC penalty (74), we use the sample size $U = \sum_{n=1}^N \pi_n$, which is the effective number of terms in the sum in (79). Third, when maximizing L over \mathbb{Z}^+ in (80), we start with $L = 1$ and increment L in steps of one until the penalized metric decreases. Fourth, for the initial model order L^0 , we recommend using $L^0 = 3$ in “sparse” mode and $L^0 = 4$ in “heavy-tailed” mode, i.e., the fixed- L defaults from Section III.D. Finally, (80) is iterated until either $L^{j+1} = L^j$ or a predetermined maximum number of allowed model-order iterations J_{max} has been reached.

As a demonstration of the proposed model-order selection procedure, we estimated a realization of \mathbf{x} with $N = 1000$ coefficients drawn i.i.d. from the triangular mixture pdf shown in Fig. 1 (top, red) with $\lambda = 0.1$, from the $M = 500$ noisy measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}$, where \mathbf{A} was i.i.d. $\mathcal{N}(0, M^{-1})$, and \mathbf{w} was AWGN such that SNR = 20 dB. For illustrative purposes, we set the initial model order at $L^0 = 1$. Iteration $j = 1$ yielded the metric $\mathcal{L}_{L^j}(\mathbf{y}; \hat{\mathbf{q}}_L) - \eta_{\text{BIC}}(L)$ shown at the top of Fig. 2, which was maximized by $L = 3 \triangleq L^1$. The metric resulting from iteration $j = 2$ is shown in the middle of Fig. 2, which was maximized by $L = 2 \triangleq L^2$. At iteration $j = 3$, we

⁸Note that $\mathcal{L}_{L^j}(\mathbf{y}; \hat{\mathbf{q}}_L)$ can be written as a constant plus a scaled value of the negative KL divergence between $p(\mathbf{x}|\mathbf{x} \neq 0, \mathbf{y}; \hat{\mathbf{q}}_{L^j})$ and the GMM $f_X(\mathbf{x}; \hat{\mathbf{q}}_L)$, where the KL divergence is clearly non-increasing in L .

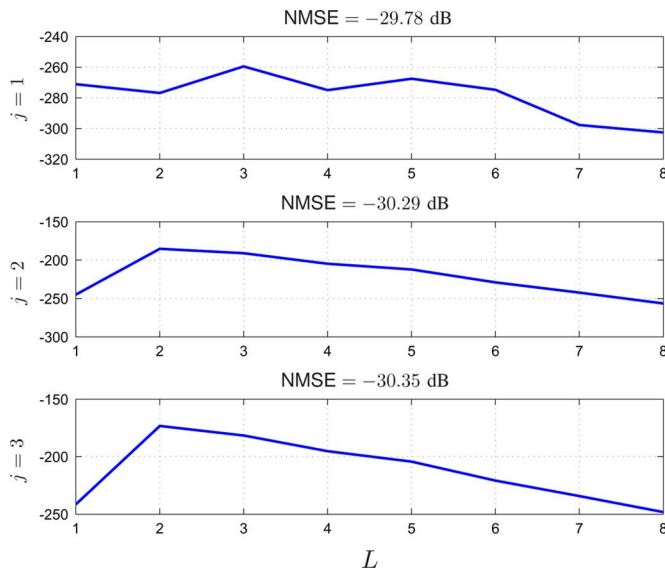


Fig. 2. An example of the model-order metric in (80) over several iterations $j = 1, 2, 3$ using initial model-order $L^j|_{j=0} = 1$, together with the NMSE of the resulting estimates.

obtained the metric at the bottom of Fig. 2, which is also maximized by $L = 2 \triangleq L^3$. Since $L^3 = L^2$, the algorithm terminates with final model order estimate $L = 2$. Fig. 2 also indicates the per-iteration MSE, which is best at the final model order.

IV. NUMERICAL RESULTS

In this section we report the results of a detailed numerical study that investigate the performance of EM-GM-AMP under both noiseless and noisy settings. For all experiments, we set the GM-GAMP tolerance to $\tau_{\text{GAMP}} = 10^{-5}$ and the maximum GAMP-iterations to $T_{\text{max}} = 20$ (recall Table I), and we set the EM tolerance to $\tau_{\text{em}} = 10^{-5}$ and the maximum EM-iterations to $I_{\text{max}} = 20$ (recall Table II). For fixed- L EM-GM-AMP, we set $L = 3$ in “sparse” and $L = 4$ in “heavy-tailed” modes.

A. Noiseless Phase Transitions

We first describe the results of experiments that computed noiseless empirical phase transition curves (PTCs) under three sparse-signal distributions. To evaluate each empirical PTC, we fixed $N = 1000$ and constructed a 30×30 grid where (M, K) were chosen to yield a uniform sampling of oversampling ratios $\frac{M}{N} \in [0.05, 0.95]$ and sparsity ratios $\frac{K}{M} \in [0.05, 0.95]$. At each grid point, we generated $R = 100$ independent realizations of a K -sparse signal \mathbf{x} from a specified distribution and an $M \times N$ measurement matrix \mathbf{A} with i.i.d. $\mathcal{N}(0, M^{-1})$ entries. From the noiseless measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$, we recovered the signal \mathbf{x} using several algorithms. A recovery $\hat{\mathbf{x}}$ from realization $r \in \{1, \dots, R\}$ was defined a success if the NMSE $\triangleq \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 / \|\mathbf{x}\|_2^2 < 10^{-6}$, and the average success rate was defined as $\bar{S} \triangleq \frac{1}{R} \sum_{r=1}^R S_r$, where $S_r = 1$ for a success and $S_r = 0$ otherwise. The empirical PTC was then plotted, using Matlab’s `contour` command, as the $\bar{S} = 0.5$ contour over the sparsity-undersampling grid.

Figs. 3–5 show the empirical PTCs for five recovery algorithms: the proposed EM-GM-AMP algorithm (in “sparse” mode) for both L fixed and L learned through model-order




Fig. 3. Empirical PTCs and LASSO theoretical PTC for noiseless recovery of Bernoulli-Gaussian signals.

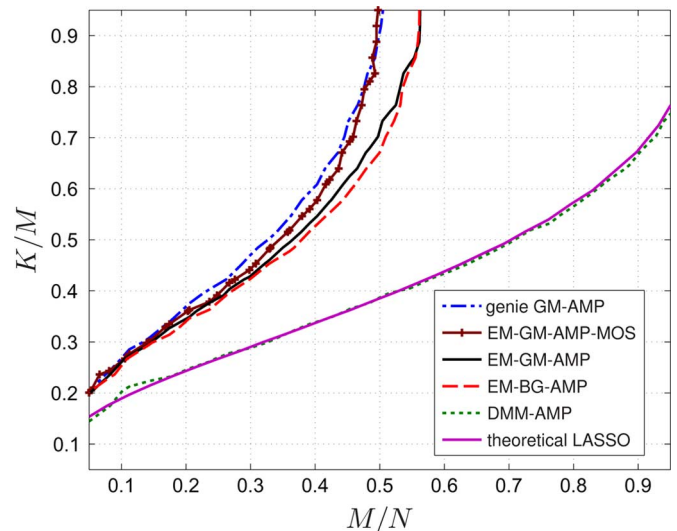


Fig. 4. Empirical PTCs and LASSO theoretical PTC for noiseless recovery of Bernoulli signals.

selection (MOS), the proposed EM-BG-AMP algorithm, a genie-tuned⁹ GM-AMP that uses the true parameters $\mathbf{q} = [\lambda, \boldsymbol{\omega}, \boldsymbol{\theta}, \boldsymbol{\phi}, \psi]$, and the Donoho/Maleki/Montanari (DMM) LASSO-style AMP from [10]. For comparison, Figs. 3–5 also display the theoretical LASSO PTC (70). The signals were generated as Bernoulli-Gaussian (BG) in Fig. 3 (using mean $\boldsymbol{\theta} = 0$ and variance $\boldsymbol{\phi} = 1$ for the Gaussian component), as Bernoulli in Fig. 4 (i.e., all non-zero coefficients set equal to 1), and as Bernoulli-Rademacher (BR) in Fig. 5.

For all three signal types, Figs. 3–5 show that the empirical PTC of EM-GM-AMP significantly improves on the empirical PTC of DMM-AMP as well as the theoretical PTC of LASSO. (The latter two are known to converge in the large system limit [10].) For BG signals, Fig. 3 shows that EM-GM-AMP-MOS,

⁹For genie-tuned GM-AMP, for numerical reasons, we set the noise variance at $\psi = 10^{-6}$ and, with Bernoulli and BR signals, the mixture variances at $\boldsymbol{\phi}_k = 10^{-2}$.




Fig. 5. Empirical PTCs and LASSO theoretical PTC for noiseless recovery of Bernoulli-Rademacher signals.

EM-GM-AMP, and EM-BG-AMP all yield PTCs that are nearly identical to that of genie-GM-AMP, suggesting that our EM-learning procedures are working well. For Bernoulli signals, Fig. 4 shows EM-GM-AMP-MOS performing very close to genie-GM-AMP, and both EM-GM-AMP and EM-BG-AMP performing slightly worse but far better than DMM-AMP. Finally, for BR signals, Fig. 5 shows EM-GM-AMP performing significantly better than EM-BG-AMP, since the former is able to accurately model the BR distribution (with $L \geq 2$ mixture components) whereas the latter (with a single mixture component) is not, and on par with genie-GM-AMP, whereas EM-GM-AMP-MOS performs noticeably better than genie-GM-AMP. The latter is due to EM-GM-AMP-MOS doing per-realization parameter tuning, while genie-GM-AMP employs the best set of *fixed* parameters over all realizations.

To better understand the performance of EM-GM-AMP when $\frac{M}{N} \ll 1$, we fixed $N = 8192$ and constructed a 12×9 grid of (M, K) values spaced uniformly in the log domain. At each grid point, we generated $R = 100$ independent realizations of a K -sparse BG signal and an i.i.d. $\mathcal{N}(0, M^{-1})$ matrix \mathbf{A} . We then recovered \mathbf{x} from the noiseless measurements using EM-GM-AMP-MOS, EM-GM-AMP, EM-BG-AMP, genie-GM-AMP, and the Lasso-solver¹⁰ FISTA¹¹ [26]. Fig. 6 shows that the PTCs of EM-GM-AMP-MOS and EM-GM-AMP are nearly identical, slightly better than those of EM-BG-AMP and genie-GM-AMP (especially at very small M), and much better than FISTA's.

Next, we studied the effect of the measurement matrix construction on the performance of EM-GM-AMP in “sparse” mode with fixed $L = 3$. For this, we plotted EM-GM-AMP empirical PTCs for noiseless recovery of a length- $N = 1000$ BG signal under several types of measurement matrix \mathbf{A} : i.i.d. $\mathcal{N}(0, 1)$, i.i.d. Uniform $[-\frac{1}{2}, \frac{1}{2}]$, i.i.d. centered Cauchy with

¹⁰For this experiment, we also tried DMM-AMP but found that it had convergence problems, and we tried SPGL1 but found performance degradations at small M .

¹¹For FISTA, we used the regularization parameter $\lambda_{\text{FISTA}} = 10^{-5}$, which is consistent with the values used for the noiseless experiments in [26].




Fig. 7. Empirical PTCs for EM-GM-AMP noiseless recovery of Bernoulli-Gaussian signals under various \mathbf{A} : i.i.d. $\mathcal{N}(0, 1)$, i.i.d. Uniform $[-\frac{1}{2}, \frac{1}{2}]$, i.i.d. Bernoulli with $\lambda_A \triangleq \text{Pr}\{a_{mn} \neq 0\} = 0.15$, i.i.d. zero-mean Bernoulli-Rademacher with $\lambda_A \in \{0.05, 0.15, 1\}$, i.i.d. Cauchy, and randomly row-sampled DCT.




Fig. 6. Empirical PTCs for noiseless recovery of Bernoulli-Gaussian signals of length $N = 8192$ when $M \ll N$.

scale 1, i.i.d. Bernoulli¹² (i.e., $a_{mn} \in \{0, 1\}$) with $\lambda_A \triangleq \text{Pr}\{a_{mn} \neq 0\} = 0.15$, i.i.d. zero-mean BR (i.e., $a_{mn} \in \{0, 1, -1\}$) with $\lambda_A \in \{0.05, 0.15, 1\}$, and randomly row-sampled Discrete Cosine Transform (DCT). Fig. 7 shows that the EM-GM-AMP PTC with i.i.d. $\mathcal{N}(0, 1)$ matrices also holds with the other i.i.d. zero-mean sub-Gaussian examples (i.e., Uniform and BR with $\lambda_A = 1$). This is not surprising given that AMP itself has rigorous guarantees for i.i.d. zero-mean sub-Gaussian matrices [5]. Fig. 7 shows that the i.i.d. \mathcal{N} PTC is also preserved with randomly row-sampled DCT matrices, which is not surprising given AMP's excellent empirical performance with many types of deterministic \mathbf{A} [27] even in the absence of theoretical guarantees. Fig. 7

¹²For the Bernoulli and BR matrices, we ensured that no two columns of a given realization \mathbf{A} were identical.




Fig. 8. NMSE versus undersampling ratio M/N for noisy recovery of Bernoulli-Gaussian signals.




Fig. 9. NMSE versus undersampling ratio M/N for noisy recovery of Bernoulli signals.

shows, however, that EM-GM-AMP's PTC can degrade with non-zero-mean i.i.d. matrices (as in the Bernoulli example) or with super-Gaussian i.i.d. matrices (as in the BR example with sparsity rate $\lambda_A = 0.05$ and the Cauchy example). Surprisingly, the i.i.d.- \mathcal{N} PTC is preserved by i.i.d.-BR matrices with sparsity rate $\lambda_A = 0.15$, even though $\lambda_A > \frac{1}{3}$ is required for a BR matrix to be sub-Gaussian [28].

B. Noisy Sparse Signal Recovery

Figs. 8–10 show NMSE for noisy recovery of BG, Bernoulli, and BR signals, respectively. To construct these plots, we fixed $N = 1000$, $K = 100$, $\text{SNR} = 25$ dB, and varied M . Each data point represents NMSE averaged over $R = 500$ realizations, where in each realization we drew an \mathbf{A} with i.i.d. $\mathcal{N}(0, M^{-1})$ elements, an AWGN noise vector, and a random signal vector. For comparison, we show the performance of the proposed EM-GM-AMP (in “sparse” mode) for both MOS and $L = 3$




Fig. 10. NMSE versus undersampling ratio M/N for noisy recovery of Bernoulli-Rademacher signals.

versions, EM-BG-AMP, genie-tuned¹³ Orthogonal Matching Pursuit (OMP) [29], genie-tuned¹³ Subspace Pursuit (SP) [30], Bayesian Compressive Sensing (BCS) [19], Sparse Bayesian Learning [18] (via the more robust T-MSBL [31]), de-biased genie-tuned¹⁴ LASSO (via SPGL1 [32]), and Smoothed- ℓ_0 (SL0) [33]. All algorithms were run under the suggested defaults, with $\text{noise} = \text{small}$ in T-MSBL.

For BG signals, Fig. 8 shows that EM-GM-AMP-MOS, EM-GM-AMP, and EM-BG-AMP together exhibit the best performance among the tested algorithms, reducing the M/N breakpoint (i.e., the location of the knee in the NMSE curve, which represents a sort of phase transition) from 0.3 down to 0.26, but also improving NMSE by ≈ 1 dB relative to the next best algorithm, which was BCS. Relative to the other EM-AMP variants, MOS resulted in a slight degradation of performance for $\frac{M}{N}$ between 0.26 and 0.31, but was otherwise identical. For Bernoulli signals, Fig. 9 shows much more significant gains for EM-GM-AMP-MOS, EM-GM-AMP and EM-BG-AMP over the other algorithms: the M/N breakpoint was reduced from 0.4 down to 0.32 (and even 0.3 with MOS), and the NMSE was reduced by ≈ 8 dB relative to the next best algorithm, which was T-MSBL in this case. Finally, for BR signals, Fig. 10 shows a distinct advantage for EM-GM-AMP and EM-GM-AMP-MOS over the other algorithms, including EM-BG-AMP, due to the formers' ability to accurately model the BR signal prior. In particular, for $M/N \geq 0.36$, EM-GM-AMP-MOS reduces the NMSE by 10 dB relative to the best of the other algorithms (which was either EM-BG-AMP or T-MSBL depending on the value of M/N) and reduces the M/N breakpoint from 0.38 down to 0.35.

To investigate each algorithm's robustness to AWGN, we plotted the NMSE attained in the recovery of BR signals with

¹³We ran both OMP (using the implementation from <http://sparselab.stanford.edu/OptimalTuning/code.htm>) and SP under 10 different sparsity assumptions, spaced uniformly from 1 to $2K$, and reported the lowest NMSE among the results.

¹⁴We ran SPGL1 in ‘BPDN’ mode: $\min_{\mathbf{x}} \|\mathbf{x}\|_1$ s.t. $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \sigma$, for hypothesized tolerances $\sigma^2 \in \{0.1, 0.2, \dots, 1.5\} \times M\psi$, and reported the lowest NMSE among the results.




Fig. 11. NMSE versus SNR for noisy recovery of Bernoulli-Rademacher signals.

$N = 1000$, $M = 500$, and $K = 100$ as a function of SNR in Fig. 11, where each point represents an average over $R = 100$ problem realizations, where in each realization we drew an \mathbf{A} with i.i.d. $\mathcal{N}(0, M^{-1})$ elements, an AWGN noise vector, and a random signal vector. All algorithms were under the same conditions as those reported previously, except that T-MSBL used `noise = small` when $\text{SNR} > 22$ dB and `noise = mild` when $\text{SNR} \leq 22$ dB, as recommended in [34]. From Fig. 11, we see that the essential behavior observed in the fixed-SNR BR plot Fig. 10 holds over a wide range of SNRs. In particular, Fig. 11 shows that EM-GM-AMP and EM-GM-AMP-MOS yield significantly lower NMSE than all other algorithms over the full SNR range, while EM-BG-AMP and T-MSBL yield the second lowest NMSE (also matched by BCS for SNRs between 30 and 40 dB). Note, however, that T-MSBL must be given some knowledge about the true noise variance in order to perform well [34], unlike the proposed algorithms.

C. Heavy-Tailed Signal Recovery

In many applications of compressive sensing, the signal to be recovered is not perfectly sparse, but instead contains a few large coefficients and many small ones. While the literature often refers to such signals as “compressible,” there are many real-world signals that do not satisfy the technical definition of compressibility (see, e.g., [35]), and so we refer to such signals more generally as “heavy tailed.”

To investigate algorithm performance for these signals, we first consider an i.i.d. Student’s-t signal, with prior pdf

$$p_X(x; q) \triangleq \frac{\Gamma((q+1)/2)}{\sqrt{\pi}\Gamma(q/2)}(1+x^2)^{-(q+1)/2} \quad (82)$$

under the (non-compressible) rate $q = 1.67$, which has been shown to be an excellent model for wavelet coefficients of natural images [35]. For such signals, Fig. 12 plots NMSE versus the number of measurements M for fixed $N = 1000$, $\text{SNR} = 25$ dB, and an average of $R = 500$ realizations, where in each realization we drew an \mathbf{A} with i.i.d. $\mathcal{N}(0, M^{-1})$ elements, an AWGN noise vector, and a random signal vector. Fig. 12 shows both variants of EM-GM-AMP (here run in




Fig. 12. NMSE versus undersampling ratio M/N for noisy recovery of Student-t signals with rate parameter 1.67.

“heavy-tailed” mode) outperforming all other algorithms under test.¹⁵ We have also verified (in experiments not shown here) that “heavy-tailed” EM-GM-AMP exhibits similarly good performance with other values of the Student’s-t rate parameter q , as well as for i.i.d. centered Cauchy signals.

To investigate the performance for positive heavy-tailed signals, we conducted a similar experiment using i.i.d. log-normal \mathbf{x} , generated using the distribution

$$p_X(x; \mu, \sigma^2) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\ln x - \mu)}{2\sigma^2}\right) \quad (83)$$

with location parameter $\mu = 0$ and scale parameter $\sigma^2 = 1$. Fig. 13 confirms the excellent performance of EM-GM-AMP-MOS, EM-GM-AMP, and EM-BG-AMP over all tested undersampling ratios M/N . We postulate that, for signals known a priori to be positive, EM-GM-AMP’s performance could be further improved through the use of a prior p_X with support restricted to the positive reals, via a mixture of positively truncated Gaussians.

It may be interesting to notice that, with the perfectly sparse signals examined in Figs. 8–10, SL0 and SPGL1 performed relatively poorly, the relevance-vector-machine (RVM)-based approaches (i.e., BCS, T-MSBL) performed relatively well, and the greedy approaches (OMP and SP) performed in-between. With the heavy-tailed signals in Figs. 12–13, it is more difficult to see a consistent pattern. For example, with the Student’s-t signal, the greedy approaches performed the worse, the RVM approaches were in the middle, and SL0 and SPGL1 performed very well. But with the log-normal signal, the situation was very different: the greedy approaches performed very well, SPGL1 performed moderately well, but SL0 and the RVM approaches performed very poorly.

In conclusion, for *all* of the many signal types tested above, the best recovery performance came from EM-GM-AMP and its MOS variant. We attribute this behavior to EM-GM-AMP’s

¹⁵In this experiment, we ran both OMP and SP under 10 different sparsity hypotheses, spaced uniformly from 1 to $K_{\text{lasso}} = M\rho_{\text{SE}}(\frac{M}{N})$, and reported the lowest NMSE among the results.




Fig. 13. NMSE versus undersampling ratio M/N for noisy recovery of log-normal signals with location parameter 0 and scale parameter 1.

ability to tune itself to the signal (and in fact the realization) at hand.

D. Runtime and Complexity Scaling With N

Next we investigated how complexity scales with signal length N by evaluating the runtime of each algorithm on a typical personal computer. For this, we fixed $K/N = 0.1$, $M/N = 0.5$, $\text{SNR} = 25$ dB and varied the signal length N . Fig. 14 shows the runtimes for noisy recovery of a Bernoulli-Rademacher signal, while Fig. 15 shows the corresponding NMSEs. In these plots, each datapoint represents an average over $R = 50$ realizations. The algorithms that we tested are the same ones that we described earlier. However, to fairly evaluate runtime, we configured some a bit differently than before. In particular, for genie-tuned SPGL1, in order to yield a better runtime-vs-NMSE tradeoff, we reduced the tolerance grid (recall footnote 14) to $\sigma^2 \in \{0.6, 0.8, \dots, 1.4\} \times M\psi$ and turned off debiasing. For OMP and SP, we used the fixed support size $K_{\text{lasso}} = M\rho_{\text{SE}}(\frac{M}{N})$ rather than searching for the size that minimizes NMSE over a grid of 10 hypotheses, as before. Otherwise, all algorithms were run under the suggested defaults, with T-MSBL run under `noise = small` and EM-GM-AMP run in “sparse” mode.

The complexities of the proposed EM-GM-AMP methods are dominated by one matrix multiplication by \mathbf{A} and \mathbf{A}^T per iteration. Thus, when these matrix multiplications are explicitly implemented and \mathbf{A} is dense, the total complexity of EM-GM-AMP should scale as $\mathcal{O}(MN)$. This scaling is indeed visible in the runtime curves of Fig. 14. There, $\mathcal{O}(MN)$ becomes $\mathcal{O}(N^2)$ since the ratio M/N was fixed, and the horizontal axis plots N on a logarithmic scale, so that this complexity scaling manifests, at sufficiently large values of N , as a line with slope 2. Fig. 14 confirms that genie-tuned SPGL1 also has the same complexity scaling, albeit with longer overall runtimes. Meanwhile, Fig. 14 shows T-MSBL, BCS, SL0, OMP, and SP exhibiting a complexity scaling of $\mathcal{O}(N^3)$ (under fixed K/N and M/N), which results in orders-of-magnitude larger runtimes for long signals (e.g., $N \geq 10^4$). With short signals (e.g., $N < 1300$), though, OMP, SP, SL0, and SPGL1




Fig. 14. Runtime versus signal length N for noisy recovery of Bernoulli-Rademacher signals.

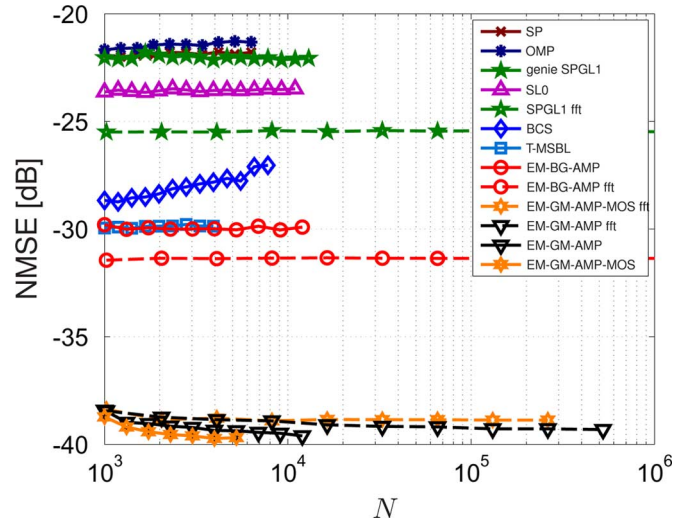


Fig. 15. NMSE versus signal length N for noisy recovery of Bernoulli-Rademacher signals.

are faster than EM-GM-AMP. Finally, Fig. 15 verifies that, for most of the algorithms, the NMSEs are relatively insensitive to signal length N when the undersampling ratio M/N and sparsity ratio K/M are both fixed, although the performance of EM-GM-AMP improves with N (which is not surprising in light of AMP’s large-system-limit optimality properties [13]) and the performance of BCS degrades with N .

Both the proposed EM-GM-AMP methods and SPGL1 can exploit the case where multiplication by \mathbf{A} and \mathbf{A}^T is implemented using a fast algorithm like the fast Fourier transform (FFT)¹⁶, which reduces the complexity to $\mathcal{O}(N \log N)$, and avoids the need to store \mathbf{A} in memory—a potentially serious problem when MN is large. The dashed lines in Figs. 14–15 (labeled “fft”) show the average runtime and NMSE of the proposed algorithms and SPGL1 in case that \mathbf{A} was a randomly row-sampled FFT. As expected, the runtimes are dramatically reduced. While EM-BG-AMP retains its place as the fastest

¹⁶For our FFT-based experiments, we used the complex-valued versions of EM-BG-AMP, EM-GM-AMP, and SPGL1.

algorithm, SPGL1 now runs $1.5\times$ faster than EM-GM-AMP (at the cost of 14 dB higher NMSE). The MOS version of EM-GM-AMP yields slightly better NMSE, but takes ≈ 2.5 times as long to run as the fixed- L version.

E. Example: Compressive Recovery of Audio

As a practical example, we experimented with the recovery of an audio signal from compressed measurements. The full length-81920 audio signal was first partitioned into T blocks $\{\mathbf{u}_t\}_{t=1}^T$ of length N . Noiseless compressed measurements $\mathbf{y}_t = \Phi \mathbf{u}_t \in \mathbb{R}^M$ were then collected using $M = N/2$ samples per block. Rather than reconstructing \mathbf{u}_t directly from \mathbf{y}_t , we first reconstructed¹⁷ the transform coefficients $\mathbf{x}_t = \Psi^T \mathbf{u}_t$, using the (orthogonal) discrete cosine transform (DCT) $\Psi \in \mathbb{R}^{N \times N}$, and later reconstructed \mathbf{u}_t via $\mathbf{u}_t = \Psi \mathbf{x}_t$. Our effective sparse-signal model can thus be written as $\mathbf{y}_t = \mathbf{A} \mathbf{x}_t$ with $\mathbf{A} = \Phi \Psi$. We experimented with two types of measurement matrix Φ : i.i.d. zero-mean Gaussian and random selection (i.e., containing rows of the identity matrix selected uniformly at random), noting that the latter allows a fast implementation of \mathbf{A} and \mathbf{A}^T . Table III shows the resulting time-averaged NMSE, i.e., $\text{TNMSE} \triangleq \frac{1}{T} \sum_{t=1}^T \|\mathbf{u}_t - \hat{\mathbf{u}}_t\|^2 / \|\mathbf{u}_t\|^2$, and total runtime achieved by the previously described algorithms at block lengths $N = 1024, 2048, 4096, 8192$, which correspond to $T = 80, 40, 20, 10$ blocks, respectively. The numbers reported in the table represent an average over 50 realizations of Φ . For these experiments, we configured the algorithms as described in Section IV.C for the heavy-tailed experiment except that, for genie-SPGL1, rather than using $\psi = 0$, we used $\psi = 10^{-6}$ for the tolerance grid (recall footnote 14) because we found that this value minimized TNMSE and, for T-MSBL, we used the setting `prune_gamma` = 10^{-12} as recommended in a personal correspondence with the author. For certain combinations of algorithm and blocklength, excessive runtimes prevented us from carrying out the experiment, and thus no result appears in the table.

Table III shows that, for this audio experiment, the EM-GM-AMP methods and SLO performed best in terms of TNMSE. As in the synthetic examples presented earlier, we attribute EM-GM-AMP's excellent TNMSE to its ability to tune itself to whatever signal is at hand. As for SLO's excellent TNMSE, we reason that it had the good fortune of being particularly well-tuned to this audio signal, given that it performed relatively poorly with the signal types used for Figs. 8–11 and Fig. 13. From the runtimes reported in Table III, we see that, with i.i.d. Gaussian Φ and the shortest block length ($N = 1024$), genie-OMP is by far the fastest, whereas the EM-GM-AMP methods are the slowest. But, as the block length grows, the EM-GM-AMP methods achieve better and better runtimes as a consequence of their excellent complexity scaling, and eventually EM-BG-AMP and fixed- L EM-GM-AMP become the two fastest algorithms under test (as shown with i.i.d. Gaussian Φ at $N = 8192$). For this audio example, the large-block regime may be the more important, because that is where all algorithms give their smallest

¹⁷Although one could exploit additional structure among the multiple-timestep coefficients $\{\mathbf{x}_t\}_{t=1}^T$ for improved recovery (e.g., sparsity clustering in the time and/or frequency dimensions, as well as amplitude correlation in those dimensions) as demonstrated in [36], such techniques are outside the scope of this paper.

TABLE III
AVERAGE TNMSE (IN dB) AND TOTAL RUNTIME (IN SECONDS) FOR COMPRESSIVE AUDIO RECOVERY

		$N = 1024$		$N = 2048$		$N = 4096$		$N = 8192$	
		TNMSE	time	TNMSE	time	TNMSE	time	TNMSE	time
i.i.d. Gaussian Φ	EM-GM-AMP-MOS	-17.3	468.9	-18.3	487.2	-21.0	967.9	-21.8	2543
	EM-GM-AMP	-16.9	159.2	-18.0	213.2	-20.7	434.0	-21.4	1129
	EM-BG-AMP	-15.9	115.2	-17.0	174.1	-19.4	430.2	-20.0	1116
	SLO	-16.8	41.6	-17.9	128.5	-20.6	629.0	-21.3	2739
	genie SPGL1	-14.3	90.9	-16.2	200.6	-18.6	514.3	-19.5	1568
	BCS	-15.0	67.5	-15.8	149.1	-18.4	428.0	-18.8	2295
	T-MSBL	-16.3	1.2e4	–	–	–	–	–	–
	genie OMP	-13.9	20.1	-14.9	109.9	-17.6	527.0	–	–
	genie SP	-14.5	87.7	-15.5	305.9	-18.0	1331	–	–
	random selection Φ	EM-GM-AMP-MOS	-16.6	233.0	-17.5	136.1	-20.5	109.6	-21.6
	EM-GM-AMP	-16.7	56.1	-17.7	43.7	-20.5	38.0	-21.5	37.8
	EM-BG-AMP	-16.2	29.6	-17.2	22.3	-19.7	19.4	-20.5	18.0
	SLO	-16.7	35.7	-17.6	119.5	-20.4	597.8	-21.2	2739
	genie SPGL1	-14.0	34.4	-15.9	24.5	-18.4	21.7	-19.7	19.6
	BCS	-15.5	60.5	-16.1	126.2	-19.4	373.8	-20.2	2295
	T-MSBL	-15.5	1.2e4	–	–	–	–	–	–
	genie OMP	-15.1	20.1	-15.7	106.8	-18.9	506.0	–	–
	genie SP	-15.2	104.5	-16.1	395.3	-18.7	1808	–	–

TNMSE. Next, looking at the runtimes under random-selection Φ , we see dramatic speed improvements for the EM-GM-AMP methods and SPGL1, which were all able to leverage Matlab's fast DCT. In fact, the total runtimes of these four algorithms *decrease* as N is increased from 1024 to 8192. We conclude by noting that EM-BG-AMP (at $N = 8192$ with random selection Φ) achieves the fastest runtime in the entire table while yielding a TNMSE that is within 1.3 dB of the best value in the entire table. Meanwhile, fixed- L EM-GM-AMP (at $N = 8192$ with random selection Φ) gives TNMSE only 0.3 dB away from the best in the entire table with a runtime of only about twice the best in the entire table. Finally, the best TNMSEs in the entire table are achieved by EM-GM-AMP-MOS (at $N = 8192$), which takes ≈ 2.5 times as long to run as its fixed- L counterpart.

V. CONCLUSION

Those interested in practical compressive sensing face the daunting task of choosing among literally hundreds of signal reconstruction algorithms (see, e.g., [37]). In testing these algorithms, they are likely to find that some work very well with particular signal classes, but not with others. They are also likely to get frustrated by those algorithms that require the tuning of many parameters. Finally, they are likely to find that some of the algorithms that are commonly regarded as “very fast” are actually very slow in high-dimensional problems. Meanwhile, those familiar with the theory of compressive sensing know that the workhorse LASSO is nearly minimax optimal, and that its phase transition curve is robust to the nonzero-coefficient distribution of sparse signals. However, they also know that, for most signal classes, there is a large gap between the MSE performance of LASSO and that of the MMSE estimator derived under full knowledge of the signal and noise statistics [11]. Thus, they may wonder whether there is a way to close this gap by designing a signal reconstruction algorithm that *both learns and exploits* the signal and noise statistics.

With these considerations in mind, we proposed an empirical Bayesian approach to compressive signal recovery that merges two powerful inference frameworks: expectation maximization (EM) and approximate message passing (AMP). We then

demonstrated—through a detailed numerical study—that our approach, when used with a flexible Gaussian-mixture signal prior, achieves a state-of-the-art combination of reconstruction error and runtime on a very wide range of signal and matrix types in the high-dimensional regime. However, certain non-zero-mean and super-Gaussian sensing matrices give our AMP-based method trouble. Making AMP robust to these matrices remains a topic of importance for future research.

REFERENCES

- [1] J. P. Vila and P. Schniter, "An empirical-Bayes approach to compressive sensing via approximate message passing," presented at the Duke Workshop Sens. Anal. High-Dimensional Data, Durham, NC, USA, Jul. 2011.
- [2] J. P. Vila and P. Schniter, "Expectation-maximization Bernoulli-Gaussian approximate message passing," in *Proc. Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, USA, Nov. 2011, pp. 799–803.
- [3] J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," in *Proc. Conf. Inf. Sci. Syst.*, Princeton, NJ, USA, Mar. 2012, pp. 1–6.
- [4] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. New York, NY, USA: Cambridge Univ. Press, 2012.
- [5] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and iterative algorithms," in *Proc. IEEE Int. Symp. Inf. Theory*, Boston, MA, USA, Jun. 2012, pp. 1–5.
- [6] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [7] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientif. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [8] D. L. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing," *Philos. Trans. Roy. Soc. A*, vol. 367, no. 1906, pp. 4273–4293, 2009.
- [9] D. L. Donoho, A. Maleki, and A. Montanari, "The noise-sensitivity phase transition in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6920–6941, Oct. 2011.
- [10] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," in *Proc. Natl. Acad. Sci.*, Nov. 2009, vol. 106, pp. 18914–18919.
- [11] Y. Wu and S. Verdú, "Optimal phase transitions in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, pp. 6241–6263, Oct. 2012.
- [12] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inf. Theory Workshop*, Cairo, Egypt, Jan. 2010, pp. 1–5.
- [13] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, pp. 764–785, Feb. 2011.
- [14] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Saint Petersburg, Russia, Aug. 2011, pp. 2168–2172.
- [15] A. Dempster, N. M. Laird, and D. B. Rubin, "Maximum-likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc.*, vol. 39, pp. 1–17, 1977.
- [16] B. Efron, *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. New York, NY, USA: Cambridge Univ. Press, 2010.
- [17] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, 2001.
- [18] D. P. Wipf and B. D. Rao, "Sparse Bayesian learning for basis selection," *IEEE Trans. Signal Process.*, vol. 52, pp. 2153–2164, Aug. 2004.
- [19] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 56, pp. 2346–2356, Jun. 2008.
- [20] U. S. Kamilov, S. Rangan, A. K. Fletcher, and M. Unser, "Approximate message passing with consistent parameter estimation and applications to sparse learning," in *Proc. Neural Inf. Process. Syst. Conf. (NIPS)*, Lake Tahoe, NV, USA, Dec. 2012, pp. 2447–2455.
- [21] U. S. Kamilov, A. Bourquard, A. Amini, and M. Unser, "One-bit measurements with adaptive thresholds," *IEEE Signal Process. Lett.*, vol. 19, pp. 607–610, Oct. 2012.
- [22] C. F. J. Wu, "On the convergence properties of the EM algorithm," *Ann. Statist.*, vol. 11, no. 1, pp. 95–103, 1983.
- [23] R. Neal and G. Hinton, M. I. Jordan, Ed., "A view of the EM algorithm that justifies incremental, sparse, and other variants," in *Learning in Graphical Models*. Cambridge, MA, USA: MIT Press, 1999, pp. 355–368.
- [24] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2007.
- [25] P. Stoica and Y. Selén, "Model-order selection: A review of information criterion rules," *IEEE Signal Process. Mag.*, vol. 21, pp. 36–47, Jul. 2004.
- [26] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [27] H. Monajemi, S. Jafarpour, M. Gavish, and D. L. Donoho, "Deterministic matrices matching the compressed sensing phase transitions of Gaussian random matrices," in *Proc. Natl. Acad. Sci.*, Jan. 2013, vol. 110, pp. 1181–1186, Stat 330/CME 362 Collaboration.
- [28] V. V. Buldygin and Y. V. Kozachenko, *Metric Characterization of Random Variables and Random Processes*. Providence, RI, USA: Amer. Math. Soc., 2000.
- [29] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, USA, 1993, pp. 40–44.
- [30] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, pp. 2230–2249, Mar. 2009.
- [31] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE J. Sel. Topics Signal Process.*, vol. 5, pp. 912–926, Sep. 2011.
- [32] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Scientif. Comput.*, vol. 31, no. 2, pp. 890–912, 2008.
- [33] H. Mohimani, M. Babaie-Zadeh, and C. Jutten, "A fast approach for overcomplete sparse decomposition based on smoothed norm," *IEEE Trans. Signal Process.*, vol. 57, pp. 289–301, Jan. 2009.
- [34] Z. Zhang, "Master the usage of T-MSBL in 3 minutes," Univ. of California, San Diego, CA, USA, Tech. Rep., Nov. 2011.
- [35] V. Cevher, "Learning with compressible priors," in *Proc. Neural Inf. Process. Syst. Conf.*, Vancouver, BC, Canada, Dec. 2009, pp. 261–269.
- [36] J. Ziniel, S. Rangan, and P. Schniter, "A generalized framework for learning and recovery of structured sparse signals," in *Proc. IEEE Workshop Statist. Signal Process.*, Ann Arbor, MI, USA, Aug. 2012, pp. 325–328.
- [37] Compressive Sensing Resources: References and Software [Online]. Available: <http://dsp.rice.edu/cs>



Jeremy P. Vila (S'11) received the B.S. degree in electrical and computer engineering from the Ohio State University in 2010. He is currently a Ph.D. student in the Information Processing Systems Lab in the Department of Electrical and Computer Engineering at OSU. His primary research interests include compressive sensing, statistical signal processing, and machine learning.



Philip Schniter (SM'05) received the B.S. and M.S. degrees in electrical engineering from the University of Illinois at Urbana-Champaign in 1992 and 1993, respectively, and the Ph.D. degree in electrical engineering from Cornell University in Ithaca, NY, in 2000.

From 1993 to 1996 he was employed by Tektronix Inc. in Beaverton, OR as a systems engineer. After receiving the Ph.D. degree, he joined the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, where he is currently a Professor and a member of the Information Processing Systems (IPS) Lab. In 2008–2009 he was a visiting professor at Eurecom, Sophia Antipolis, France, and Supélec, Gif-sur-Yvette, France.

In 2003, Dr. Schniter received the National Science Foundation CAREER Award. His areas of interest currently include statistical signal processing, wireless communications and networks, and machine learning.