

Compressed Sensing via Universal Denoising and Approximate Message Passing

Yanting Ma, Junan Zhu, and Dror Baron

Department of Electrical and Computer Engineering

North Carolina State University; Raleigh, NC 27695, USA

Email: {yma7, jzhu9, barondror}@ncsu.edu

Abstract—We study compressed sensing (CS) signal reconstruction problems where an input signal is measured via matrix multiplication under additive white Gaussian noise. Our signals are assumed to be stationary and ergodic, but the input statistics are unknown; the goal is to provide reconstruction algorithms that are universal to the input statistics. We present a novel algorithm that combines: (i) the approximate message passing (AMP) CS reconstruction framework, which converts the matrix channel recovery problem into scalar channel denoising; (ii) a universal denoising scheme based on context quantization, which partitions the stationary ergodic signal denoising into independent and identically distributed (i.i.d.) subsequence denoising; and (iii) a density estimation approach that approximates the probability distribution of an i.i.d. sequence by fitting a Gaussian mixture (GM) model. In addition to the algorithmic framework, we provide three contributions: (i) numerical results showing that state evolution holds for non-separable Bayesian sliding-window denoisers; (ii) a universal denoiser that does not require the input signal to be bounded; and (iii) we modify the GM learning algorithm, and extend it to an i.i.d. denoiser. Our universal CS recovery algorithm compares favorably with existing reconstruction algorithms in terms of both reconstruction quality and runtime, despite not knowing the input statistics of the stationary ergodic signal.

Index Terms—approximate message passing, compressed sensing, Gaussian mixture model, universal denoising.

I. INTRODUCTION

A. Motivation

Many scientific and engineering problems can be approximated as linear systems of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^N$ is the unknown input signal, $\mathbf{A} \in \mathbb{R}^{M \times N}$ is the matrix that characterizes the linear system, and $\mathbf{z} \in \mathbb{R}^M$ is measurement noise. The goal is to estimate \mathbf{x} from the measurements \mathbf{y} given \mathbf{A} and statistical information about \mathbf{z} . When $M \ll N$, the setup is known as compressed sensing (CS); by posing a sparsity or compressibility requirement on the signal, it is indeed possible to accurately recover \mathbf{x} from the ill-posed linear system [1, 2]. However, we might need $M > N$ when the signal is dense or the noise is substantial.

One popular scheme to solve CS recovery problems is LASSO [3] (also known as basis pursuit denoising [4]): $\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \gamma \|\mathbf{x}\|_1$, where $\|\cdot\|_p$ denotes

the ℓ_p -norm, and γ is a tuning parameter. This approach does not require statistical information about \mathbf{x} and \mathbf{z} , and can be conveniently solved via standard convex optimization tools or the approximate message passing (AMP) algorithm [5]. However, the reconstruction quality is often far from optimal.

Bayesian CS recovery algorithms based on message passing [6–8] usually achieve better reconstruction quality, but must know the prior for \mathbf{x} . For parametric signals with unknown parameters, one can infer the parameters and achieve the minimum mean square error (MMSE) in some settings; examples include EM-GM-AMP-MOS [9], turboGAMP [10], adaptive-GAMP [11], and AMP-MixD [12].

Unfortunately, possible uncertainty about the statistics of the signal may make it difficult to select a prior or model class for Bayesian algorithms. Therefore, it would be desirable to formulate universal algorithms to estimate \mathbf{x} that are agnostic to the particular statistics of the signal.

While approaches based on Kolmogorov complexity [13–16] are theoretically appealing for universal signal recovery, they are not computable in practice [17, 18]. Several algorithms based on Markov chain Monte Carlo (MCMC) [19–22] leverage the fact that for stationary ergodic signals, both the per-symbol empirical entropy and Kolmogorov complexity converge asymptotically almost surely to the entropy rate of the signal [17], and aim to minimize the empirical entropy. The best existing implementation of the MCMC approach [22] often achieves a mean square error (MSE) that is within 3 dB of the MMSE, which resembles a result by Donoho for universal denoising [13].

In this paper, we confine our attention to the system model defined in (1), where the input signal \mathbf{x} is generated by a stationary ergodic source, and merge concepts from AMP [5], a universal denoising algorithm for stationary ergodic signals [23, 24], and Gaussian mixture (GM) learning [25] for density estimation. The resulting universal CS recovery algorithm, which we call AMP-UD (AMP with a universal denoiser), compares favorably with existing approaches in terms of reconstruction quality and runtime.

B. Related work and main results

Approximate message passing: AMP is an iterative algorithm that solves a linear inverse problem by successively converting the matrix channel problem into scalar channel denoising problems with additive white Gaussian noise (AWGN). AMP has received considerable attention

This work was supported in part by the National Science Foundation under grant CCF-1217749 and in part by the U.S. Army Research Office under grant W911NF-14-1-0314.

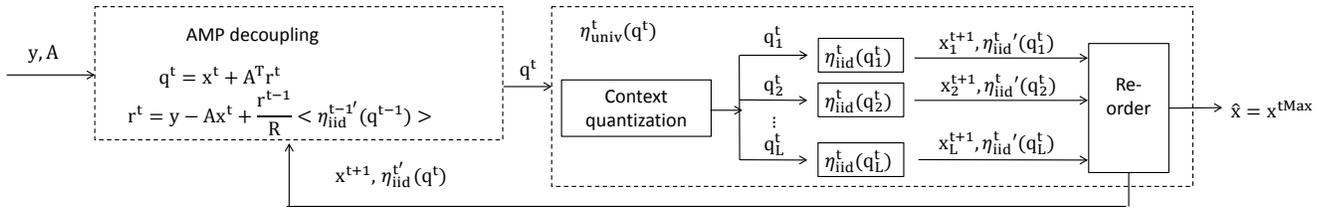


Fig. 1. Flow chart of AMP-UD. AMP decouples the linear inverse problem into scalar channel denoising problems. In the t -th iteration, the universal denoiser $\eta_{\text{univ}}^t(\cdot)$ converts stationary ergodic signal denoising into i.i.d. subsequences denoising. Each i.i.d. denoiser $\eta_{\text{iid}}^t(\cdot)$ (21) outputs the denoised subsequence x_i^{t+1} and the derivative of the denoiser $\eta_{\text{iid}}^t(\cdot)$ (22). The algorithm stops when the iteration index t reaches the predefined maximum t_{Max} , and outputs \hat{x}^{tMax} as the CS recovery result.

because of its fast convergence and the state evolution (SE) formalism [5, 26], which offers a precise characterization of the AWGN denoising problem in each iteration. AMP with separable denoisers has been rigorously proved to obey SE [26]. However, for non-i.i.d. signals we may want to explore non-separable denoisers. Donoho et al. [27] provide numerical results demonstrating that SE accurately predicts the phase transition of AMP when some well-behaved non-separable minimax denoisers are applied, and conjecture that SE holds for AMP with a broader class of denoisers. A compressive imaging algorithm that applies non-separable image denoisers within AMP appears in Tan et al. [28]. A potential challenge of implementing AMP is to obtain the Onsager correction term [5], which involves the calculation of the derivative of a denoiser. Metzler et al. [29] leverage a Monte Carlo technique to approximate the derivative of a denoiser when an explicit input-output relation of the denoiser is not available, and provide numerical results showing that SE holds for AMP with their approximation.

Despite the encouraging results for using non-separable denoisers within AMP, a rigorous proof that SE holds for non-separable denoisers has yet to appear. Consequently, new evidence showing that AMP obeys SE may increase the community's confidence about using non-separable denoisers within AMP. Our first contribution is that we provide numerical results showing that SE holds for non-separable Bayesian sliding-window denoisers.

Universal denoising: Our proposed denoiser is inspired by an approach based on context quantization [24], where a universal denoiser for a stationary ergodic signal involves multiple separable denoisers for conditionally independent subsequences. Sivaramakrishnan and Weissman [24] have shown that their universal denoiser based on context quantization can achieve the MMSE for stationary ergodic signals with bounded components.

The boundedness condition of Sivaramakrishnan and Weissman [24] is due to their density estimation approach, in which the empirical distribution function is obtained by quantizing the bounded range of the signal. Such boundedness conditions may be undesirable in certain applications. We overcome this limitation by replacing their density estimation approach with GM model learning. Our second contribution is a universal denoiser that does not require the input signal to be bounded; we conjecture that our universal denoiser achieves the MMSE under some technical conditions.

Fitting Gaussian mixture models: Figueiredo and

Jain [25] propose an algorithm that fits a given data sequence with a GM model. The algorithm employs a cost function that resembles the minimum message length (MML) criterion, and the parameters are learned using a modified expectation-maximization (EM) algorithm.

Our GM fitting problem involves a sequence of observations \mathbf{q} corrupted by AWGN, and we estimate the probability density function (pdf) of \mathbf{x} based on the GM model for \mathbf{q} . Note that a GM convolved with Gaussian noise is still a GM in which the variance of each component is increased by the noise variance. Therefore, we send \mathbf{q} to the mixture model learning algorithm, and subtract the noise variance from each component of the estimated pdf $\hat{p}(q)$ to obtain $\hat{p}(x)$. Once $\hat{p}(x)$ is available, we denoise the subsequence by computing the conditional expectation of each entry of the subsequence of \mathbf{x} based on the estimated prior $\hat{p}(x)$.¹ Our third contribution is that we modify the GM learning algorithm, and extend it to an i.i.d. denoiser.

A flow chart of AMP-UD, which employs the AMP framework, along with our modified universal denoiser (η_{univ}) and the GM-based i.i.d. denoiser ($\eta_{\text{i.i.d.}}$), is shown in Fig. 1. Based on the numerical evidence that SE holds for AMP with the Bayesian sliding-window denoiser and the conjecture that our universal denoiser can achieve the MMSE, we further conjecture that AMP-UD achieves the MMSE under some technical conditions. The details of AMP-UD are developed in Sections II–V.

The remainder of the paper is arranged as follows. In Section II, we review AMP and provide new numerical evidence that AMP obeys SE with non-separable denoisers. In Section III, we extend the universal denoiser based on context quantization to overcome the boundedness condition. Section IV modifies the GM fitting algorithm, and extends it to an i.i.d. denoiser. Our proposed AMP-UD algorithm is described in detail in Section V. Numerical results are shown in Section VI, and we conclude the paper in Section VII.

II. APPROXIMATE MESSAGE PASSING FOR SLIDING-WINDOW DENOISERS

A. Review of AMP

Consider a linear inverse problem (1), where the empirical pdf of \mathbf{x} follows $p_X(x)$, the measurement matrix \mathbf{A} has i.i.d. Gaussian entries with unit-norm columns on average, and $\mathbf{z} \sim$

¹We remind the reader that MMSE-optimal estimators rely on conditional expectation.

$\mathcal{N}(\mathbf{z}; \mathbf{0}, \sigma_z^2 I)$, where $\mathcal{N}(\mathbf{u}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^n |\boldsymbol{\Sigma}|}} \exp(-\frac{1}{2}(\mathbf{u} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{u} - \boldsymbol{\mu}))$ denotes a multivariate Gaussian pdf of a random vector $\mathbf{u} \in \mathbb{R}^n$, $(\cdot)^T$ denotes the transpose, and I is the identity matrix.

Starting with $\mathbf{x}^0 = 0$, the AMP algorithm [5] proceeds iteratively according to

$$\mathbf{x}^{t+1} = \eta^t(\mathbf{A}^T \mathbf{r}^t + \mathbf{x}^t), \quad (2)$$

$$\mathbf{r}^t = \mathbf{y} - \mathbf{A} \mathbf{x}^t + \frac{1}{R} \mathbf{r}^{t-1} \langle \eta^{t-1'}(\mathbf{A}^T \mathbf{r}^{t-1} + \mathbf{x}^{t-1}) \rangle, \quad (3)$$

where $R = M/N$ represents the measurement rate, t represents the iteration index, $\eta^t(\cdot)$ is a denoising function, and $\langle \mathbf{u} \rangle = \frac{1}{N} \sum_{i=1}^N u_i$ for some vector $\mathbf{u} \in \mathbb{R}^N$. The denoising function $\eta^t(\cdot)$ is separable in the original derivation of AMP [5, 26, 30]. That is, $\eta^t(\mathbf{u}) = (\eta^t(u_1), \eta^t(u_2), \dots, \eta^t(u_N))$ and $\eta_t'(\mathbf{u}) = (\eta_t'(u_1), \eta_t'(u_2), \dots, \eta_t'(u_N))$, where $\eta_t'(\cdot)$ denotes the derivative of $\eta^t(\cdot)$. A useful property of AMP is that at each iteration, the vector $\mathbf{A}^T \mathbf{r}^t + \mathbf{x}^t \in \mathbb{R}^N$ in (2) is equivalent to the input signal \mathbf{x} corrupted by AWGN. The variance of the Gaussian noise $(\sigma^t)^2$ evolves following SE in the limit of large systems ($N \rightarrow \infty, M/N \rightarrow R$):

$$(\sigma^{t+1})^2 = \sigma_z^2 + \frac{1}{R} \mathbb{E}_{X,W} \left[(\eta^t(X + \sigma^t W) - X)^2 \right], \quad (4)$$

where $W \sim \mathcal{N}(w; 0, 1)$, $X \sim p_X(x)$, and $(\sigma^0)^2 = \sigma_z^2 + \frac{1}{R} \mathbb{E}[X^2]$. Formal statements for SE appear in [26, 30]. Additionally, it is convenient to use the following unbiased estimator for $(\sigma^t)^2$ [30]:

$$(\hat{\sigma}^t)^2 = \frac{1}{M} \|\mathbf{r}^t\|_2^2. \quad (5)$$

B. State evolution for Bayesian sliding-window denoisers

It has been conjectured by Donoho et al. [27] that AMP with a wide range of denoisers obeys SE, including many non-separable denoisers. We now provide new evidence to support this conjecture by constructing non-separable Bayesian denoisers within the sliding-window denoising scheme for two Markov signal models, and showing that SE accurately predicts the performance of AMP with this class of denoisers for large signal dimension N . Our rationale for examining the SE performance of sliding-window denoisers is that the context quantization approach of Sivaramakrishnan and Weissman [24] resembles a sliding-window denoiser.

The mathematical model for an AWGN channel denoising problem is defined as

$$\mathbf{q} = \mathbf{x} + \mathbf{v}, \quad (6)$$

where $\mathbf{x} \in \mathbb{R}^N$ is the input signal, $\mathbf{v} \in \mathbb{R}^N$ is AWGN that follows the distribution $\mathcal{N}(\mathbf{v}; \mathbf{0}, \sigma_v^2 I)$, and $\mathbf{q} \in \mathbb{R}^N$ is a sequence of noisy observations.

In a separable denoiser, x_j is estimated only from its noisy observation q_j . The separable Bayesian denoiser that minimizes the MSE is the point-wise conditional expectation,

$$\hat{x}_j = \mathbb{E}[x_j | q_j] = \int x_j p(x_j | q_j) dx_j, \quad (7)$$

where Bayes' rule yields $p(x_j | q_j) = \frac{p(y_j | x_j) p(x_j)}{p(q_j)}$. If entries of the input signal \mathbf{x} are independent and x_j is drawn from $p(x_j)$, then (7) achieves the MMSE.

When there are statistical dependencies among the entries of \mathbf{x} , a sliding-window denoising scheme can be applied to improve performance. We consider two Markov sources as examples that contain statistical dependencies, and highlight that our true motivation is the richer class of stationary ergodic sources.

Example source 1: Consider a two-state Markov state machine that contains states s_0 (zero state) and s_1 (nonzero state). The transition probabilities are $p_{10} = p(s_0 | s_1)$ and $p_{01} = p(s_1 | s_0)$. In the steady state, the marginal probability of state s_1 is $p(s_1) = \frac{p_{01}}{p_{01} + p_{10}}$. We call our first example source Markov-constant (MConst for short); it is generated by the two-state Markov machine with $p_{01} = \frac{3}{970}$ and $p_{10} = 0.10$, and in the nonzero state the signal value is the constant 1. These state transition parameters yield 3% nonzero entries in a MConst signal on average.

Example source 2: Our second example is a four-state Markov switching signal (M4 for short) that follows the pattern $+1, +1, -1, -1, +1, +1, -1, -1, \dots$ with 3% error probability in state transitions, resulting in the signal switching from -1 to $+1$ or vice versa either too early or too late; the four states $s_1 = [-1 \ -1]$, $s_2 = [-1 \ +1]$, $s_3 = [+1 \ -1]$, and $s_4 = [+1 \ +1]$ have equal marginal probabilities 0.25 in the steady state.

Bayesian sliding-window denoiser: Denote a block $(q_s, q_{s+1}, \dots, q_t)$ of a sequence \mathbf{q} by \mathbf{q}_s^t for $s < t$. The $(2k+1)$ -Bayesian sliding-window denoiser η_{MConst} for the MConst signal is defined as

$$\eta_{\text{MConst},j}(\mathbf{q}_{j-k}^{j+k}) = \mathbb{E}[x_j | \mathbf{q}_{j-k}^{j+k}] = \frac{p(x_j = 1, \mathbf{q}_{j-k}^{j+k})}{p(\mathbf{q}_{j-k}^{j+k})}, \quad (8)$$

and the MSE of η_{MConst} can be shown to be

$$\begin{aligned} \text{MSE}^{\text{MConst}} &= \mathbb{E} \left[\left(x_j - \eta_{\text{MConst},j}(\mathbf{q}_{j-k}^{j+k}) \right)^2 \right] \\ &= \int \frac{p(x_j = 0, \mathbf{q}_{j-k}^{j+k}) p(x_j = 1, \mathbf{q}_{j-k}^{j+k})}{p(\mathbf{q}_{j-k}^{j+k})} d\mathbf{q}_{j-k}^{j+k}, \end{aligned} \quad (9)$$

where

$$p(x_j = s, \mathbf{q}_{j-k}^{j+k}) = \sum_{\mathbf{x}_{j-k}^{j+k} \setminus x_j} \mathcal{N}(\mathbf{q}_{j-k}^{j+k}; \mathbf{x}_{j-k}^{j+k}, \sigma_v^2 I) p(\mathbf{x}_{j-k}^{j+k}) \Big|_{x_j=s}, \quad (10)$$

$\mathbf{x}_{j-k}^{j+k} \setminus x_j$ denotes the sequence \mathbf{x}_{j-k}^{j+k} not including the middle symbol x_j ,

$$p(\mathbf{q}_{j-k}^{j+k}) = \sum_{\mathbf{x}_{j-k}^{j+k}} \mathcal{N}(\mathbf{q}_{j-k}^{j+k}; \mathbf{x}_{j-k}^{j+k}, \sigma_v^2 I) p(\mathbf{x}_{j-k}^{j+k}), \quad (11)$$

$$p(\mathbf{x}_{j-k}^{j+k}) = p(x_{j-k}) \prod_{t=-k}^{k-1} p(x_{j+t+1} | x_{j+t}). \quad (12)$$

To obtain the Onsager correction term $\frac{1}{R} \mathbf{r}^{t-1} \langle \eta^{t-1'}(\mathbf{A}^T \mathbf{r}^{t-1} + \mathbf{x}^{t-1}) \rangle$ in (3), we need to calculate

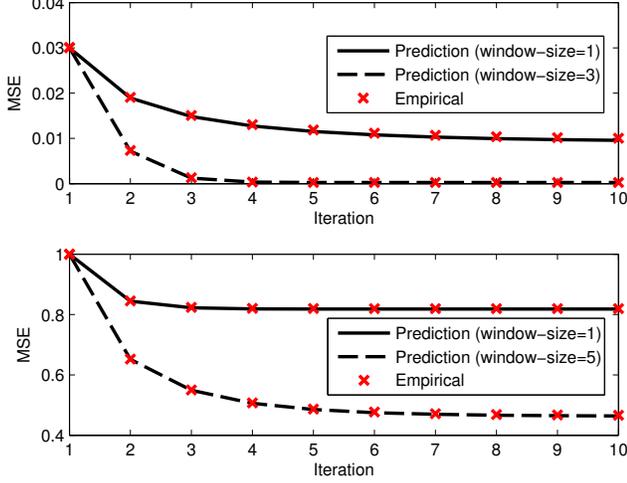


Fig. 2. Top: Numerical verification of SE for AMP with η_{MConst} (8) when the input is a MConst signal. ($N = 20,000$, $R = 0.2$, $\text{SNR} = 5$ dB.) Bottom: Numerical verification of SE for AMP with η_{M4} (14) when the input is a M4 signal. ($N = 20,000$, $R = 0.2$, $\text{SNR} = 10$ dB.)

the derivative of $\eta_{\text{MConst},j}$. It can be shown that

$$\frac{\partial}{\partial q_j} \eta_{\text{MConst},j}(\mathbf{q}_{j-k}^{j+k}) = \frac{p(x_j = 0, \mathbf{q}_{j-k}^{j+k})p(x_j = 1, \mathbf{q}_{j-k}^{j+k})}{(\sigma_v p(\mathbf{q}_{j-k}^{j+k}))^2}. \quad (13)$$

Similarly, the $(2k+1)$ -Bayesian sliding-window denoiser η_{M4} for the M4 signal is defined as

$$\begin{aligned} \eta_{\text{M4},j}(\mathbf{q}_{j-k}^{j+k}) &= \mathbb{E}[x_j | \mathbf{q}_{j-k}^{j+k}] \\ &= \frac{p(x_j = 1, \mathbf{q}_{j-k}^{j+k}) - p(x_j = -1, \mathbf{q}_{j-k}^{j+k})}{p(\mathbf{q}_{j-k}^{j+k})}, \end{aligned} \quad (14)$$

where $p(x_j = s, \mathbf{q}_{j-k}^{j+k})$ for $s \in \{-1, 1\}$ is defined in (10), $p(\mathbf{q}_{j-k}^{j+k})$ is defined in (11), and (12) becomes

$$p(\mathbf{x}_{j-k}^{j+k}) = p(x_{j-k+1}, x_{j-k}) \prod_{t=-k}^{k-2} p(x_{j+t+2} | x_{j+t+1}, x_{j+t}).$$

It can be shown that

$$\begin{aligned} \text{MSE}^{\text{M4}} &= \mathbb{E} \left[\left(x_j - \eta_{\text{M4},j}(\mathbf{q}_{j-k}^{j+k}) \right)^2 \right] \\ &= \int \frac{4p(x_j = 0, \mathbf{q}_{j-k}^{j+k})p(x_j = 1, \mathbf{q}_{j-k}^{j+k})}{p(\mathbf{q}_{j-k}^{j+k})} d\mathbf{q}_{j-k}^{j+k}, \end{aligned} \quad (15)$$

$$\frac{\partial}{\partial q_j} \eta_{\text{M4},j}(\mathbf{q}_{j-k}^{j+k}) = \frac{4p(x_j = 0, \mathbf{q}_{j-k}^{j+k})p(x_j = 1, \mathbf{q}_{j-k}^{j+k})}{(\sigma_z p(\mathbf{q}_{j-k}^{j+k}))^2}. \quad (16)$$

If AMP with η_{MConst} or η_{M4} obeys SE, then the noise variance $(\sigma^t)^2$ at each iteration should evolve according to (4). As a consequence, the reconstruction error at iteration t can be predicted by evaluating (9) or (15) with σ_v^2 in (10) and (11) being replaced by $(\sigma^t)^2$.

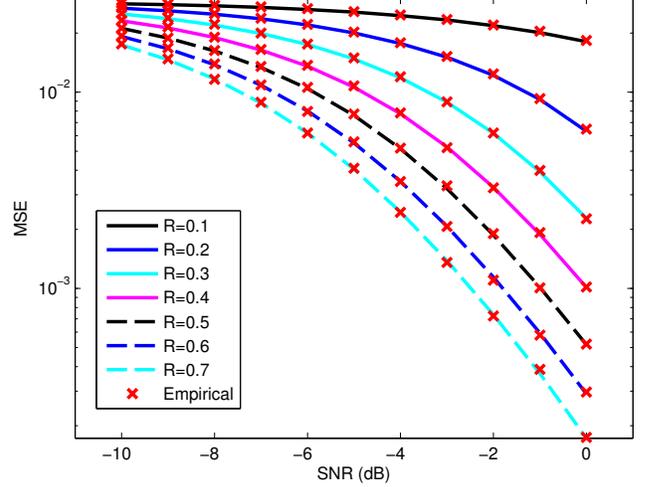


Fig. 3. Numerical verification of SE for AMP with η_{MConst} (8) at various measurement rates (R) and noise levels (SNR). The lines are predicted by SE, whereas the markers represent the empirical MSE. ($N = 20,000$.)

C. Numerical evidence

We apply η_{MConst} (8) within AMP for MConst signals, and η_{M4} (14) within AMP for M4 signals. The window size $2k+1$ is chosen to be 1 or 3 for η_{MConst} , and 1 or 5 for η_{M4} . Note that when the window size is 1, η_{MConst} and η_{M4} become separable denoisers. The MSE predicted by SE is compared to the empirical MSE at each iteration where the input signal to noise ratio ($\text{SNR} = 10 \log_{10}[(N\mathbb{E}[x^2])/(M\sigma_z^2)]$) is 5 dB or 10 dB. It is shown in Fig. 2 for AMP with η_{MConst} and η_{M4} that the markers representing the empirical MSE track the lines predicted by SE, and that side-information from neighboring entries helps improve the MSE performance.

Results for AMP with η_{MConst} at various measurement rates and noise levels are shown in Fig. 3. The markers that represent the empirical MSE lie on the lines predicted by SE, which further verifies the correctness of SE.

Our SE results for the two Markov signals increase our confidence that applying non-separable denoisers within AMP for non-i.i.d. signals will track SE. This confidence motivates us to apply a universal denoiser within AMP for CS reconstruction of stationary ergodic signals with unknown input statistics. Indeed, the numerical results in Section VI show that AMP with a universal denoiser leads to a promising universal CS recovery algorithm.

III. UNIVERSAL DENOISING

A. Background

Sivaramakrishnan and Weissman [24] propose to quantize the noisy symbols to generate quantized contexts that are used to partition the unquantized symbols into subsequences. That is, given the noisy observations $\mathbf{q} \in \mathbb{R}^N$, define the context of q_j as $\mathbf{c}_j = [\mathbf{q}_{j-k}^{j-1}; \mathbf{q}_{j+1}^{j+k}]$ for $j = 1+k, \dots, N-k$, where $\mathbf{c}_j \in \mathbb{R}^{2k}$, and $[\alpha; \beta]$ denotes the concatenation of the sequences α and β . Vector quantization can be applied to the context set $\mathcal{C} = \{\mathbf{c}_j : j = 1+k, \dots, N-k\}$, and each \mathbf{c}_j is assigned a label $l_j \in \{1, \dots, L\}$ that represents the cluster that \mathbf{c}_j belongs to. Finally, the L subsequences that consist of

symbols from \mathbf{q} with the same label are obtained by taking $\mathbf{q}_l = \{q_j : l_j = l\}$, for $l = 1, \dots, L$.

The symbols in each subsequence \mathbf{q}_l are regarded as approximately conditionally independent given the common quantized contexts [24]. The rationale underlying this concept is that a sliding-window denoiser uses information from the contexts to estimate the current symbol, and symbols with similar contexts in the noisy output of the scalar channel have similar contexts in the original signal. Therefore, symbols with similar contexts can be grouped together and denoised using the same denoiser.

Sivaramakrishnan and Weissman [24] estimate the pdf of \mathbf{x}_l , which is the clean subsequence corresponding to \mathbf{q}_l , by first obtaining an estimate of $p_l(q)$, which is the empirical pdf of \mathbf{q}_l , via quantization followed by kernel density estimation, and then quantizing $p_l(x)$ in both the support domain and the probability domain to find a $\hat{p}_l(x)$ that matches q_l well. Once $\hat{p}_l(x)$ is obtained, the conditional expectation of the symbols in the l -th subsequence can be calculated (7).

For any well-defined error metric, Sivaramakrishnan and Weissman [24] have proved for stationary ergodic signals with bounded components that their universal denoiser asymptotically achieves the optimal estimation error among all sliding-window denoising schemes despite not knowing the prior for the signal. When the error metric is square error, the optimal error is the MMSE.

B. Extension to unbounded signals

Sivaramakrishnan and Weissman [24] have shown that one can denoise a stationary ergodic signal by (i) grouping together symbols with similar contexts; and (ii) applying a separable denoiser to each group. Such a scheme is optimal in the limit of large signal dimension N . However, their denoiser assumes a bounded input, which might make it inapplicable to some real-world settings.

We modify step (i), while employing an entirely different approach for step (ii) to construct our proposed universal denoiser, which does not require the boundedness condition.

Step (i): The context set \mathcal{C} is acquired in the same way as described in Section III-A, while we add weights to the contexts before clustering. That is, for each $\mathbf{c}_j \in \mathcal{C}$ of length $2k$, the weighted context is defined as

$$\mathbf{c}'_j = \mathbf{c}_j \odot \mathbf{w}, \quad (17)$$

where \odot denotes a point-wise product, and the weights take values,

$$w_{k_i} = \begin{cases} e^{-\beta(k-k_i)}, & k_i = 1, \dots, k \\ e^{-\beta(k_i-k-1)}, & k_i = k+1, \dots, 2k \end{cases}, \quad (18)$$

for some $\beta \in (0, 1)$, so that the symbols in \mathbf{c}_j that are closer in index to q_j have larger weights than the ones that are located farther away. The exponential decay rate β is made adaptive to the noise level:

$$\beta = b_1 \log_{10}(\sigma_v^2 / (\|\mathbf{q}\|_2^2 / N - \sigma_v^2)) + b_2, \quad (19)$$

where b_1 and b_2 can be determined numerically.

The weighted context set $\mathcal{C}' = \{\mathbf{c}'_j : j = 1+k, \dots, N-k\}$ is then sent to a k -means algorithm [31], and \mathbf{q}_l 's are obtained

according to the labels determined via clustering. A post-processing step is added to ensure that the empirical pdf of \mathbf{q}_l is learned from no less than T symbols. That is, if the size of \mathbf{q}_l , which is denoted by B , is less than T , then $T-B$ symbols in other clusters whose contexts are closest to the centroid of the current cluster are included to estimate the empirical pdf of \mathbf{q}_l , while after the pdf is learned, the extra symbols are removed, and only \mathbf{q}_l is denoised with the currently learned pdf. The rationale for requiring at least T symbols is that using a small amount of data to estimate the pdf may result in a poor estimate.

Step (ii): In order to overcome the limitation of boundedness, we fit $p(q)$ with a GM model. In the case of Gaussian noise channel, we can then estimate $p(x)$ by subtracting the noise variance from each Gaussian component in $\hat{p}(q)$. Details are provided in Section IV.

IV. I.I.D. DENOISING VIA GAUSSIAN MIXTURE FITTING

A. Background

The pdf of a GM has the form:

$$p(x) = \sum_{s=1}^S \alpha_s \mathcal{N}(x; \mu_s, \sigma_s^2), \quad (20)$$

where S is the number of Gaussian components, and $\sum_{s=1}^S \alpha_s = 1$, so that $p(x)$ is a proper pdf.

Figueiredo and Jain [25] propose to fit a GM model for a given data sequence by starting with some arbitrarily large S , and inferring the structure of the mixture by letting the mixing probabilities α_s of some components be zero. This approach resembles the concept underlying the minimum message length (MML) criterion that selects the best overall model from the entire model space, which differs from model class selection based on the best model within each class.² A component-wise EM algorithm that updates $(\alpha_s, \mu_s, \sigma_s^2)$ sequentially in s is used to implement the MML-based approach. The main feature of the algorithm is that if α_s is estimated as 0, then the s -th component is immediately removed, and all the estimates in the expectation step are recomputed before moving to the maximization step that learns the parameters of the $(s+1)$ -th component.

B. Extension to denoising

Consider the scalar channel denoising problem defined in (6) with an i.i.d. input. We propose to estimate \mathbf{x} from its Gaussian noise corrupted observations \mathbf{q} by posing a GM prior on \mathbf{x} , and learning the parameters of the GM model with a modified version of the algorithm proposed by Figueiredo and Jain [25].

Initialization of EM: The EM algorithm must be initialized for each parameter, $\{\alpha_s, \mu_s, \sigma_s^2\}$, $s = 1, \dots, S$. One may choose to initialize the Gaussian components with equal mixing probabilities and equal variances, and the initial value of the means are randomly sampled from the input data sequence [25]. However, in CS recovery problems, the input

²All models with the same number of components belong to one model class, and different models within a model class have different parameters for each component.

signal is often sparse, and it becomes difficult to correct the initial value if the initialized values are far from the truth. To see why a poor initialization might be problematic, consider the following scenario: a sparse binary signal that contains a few ones and is corrupted by Gaussian noise is sent to the algorithm. If the initialization levels of the μ_s 's are all around zero, then the algorithm is likely to fit a Gaussian component with near-zero mean and large variance rather than two narrow Gaussian components, one of which has mean close to zero while the other has mean close to one.

To address this issue, we modify the initialization to examine the maximal distance between each symbol of the input data sequence and the current initialization of the μ_s 's. If the distance is greater than $0.1\sigma_{\text{init}}$, then we add a Gaussian component whose mean is initialized as the value of the symbol being examined, where σ_{init}^2 is the initialization level of the variance. We found in our simulations that the modified initialization improves the accuracy of the density estimation, and speeds up the convergence of the EM algorithm; the details of the simulation are omitted for brevity.

Leverage side-information about the noise: Because we know that each Gaussian component in $\hat{p}(q)$ should have variance no less than the noise variance σ_v^2 , during the parameter learning process, if a component has variance that is significantly less than σ_v^2 , we assume that this low-variance component is spurious, and set the corresponding α_s to zero. However, if the variance of the component is only slightly less than the noise variance σ_v^2 , then we allow the algorithm to keep tracking this component, and set the component variance to be equal to σ_v^2 at the end of the parameter learning process. That said, when subtracting the noise variance σ_v^2 from the Gaussian components of $\hat{p}(q)$ to obtain the components of $\hat{p}(x)$, we could have components with zero-valued variance, which leads to deltas in $\hat{p}(x)$.

Denoising: Once the parameters in (20) are estimated, we define a denoiser for i.i.d. signals as conditional expectation:

$$\begin{aligned} \eta_{\text{hid}}(q) &= \mathbb{E}[x|q] \\ &= \frac{\int xp(q|x)\hat{p}(x)dx}{\hat{p}(q)} \\ &= \frac{\sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2) \left(\frac{\sigma_s^2}{\sigma_s^2 + \sigma_v^2} (q - \mu_s) + \mu_s \right)}{\sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2)}. \end{aligned} \quad (21)$$

V. PROPOSED UNIVERSAL CS RECOVERY ALGORITHM

Combining the three components that have been discussed in Sections II–IV, we are now ready to introduce our proposed universal CS recovery algorithm AMP-UD.

Consider a linear inverse problem (1), where the input signal \mathbf{x} is generated by a stationary ergodic source.³ To estimate \mathbf{x} from \mathbf{y} given \mathbf{A} , we apply AMP as defined in (2) and (3). In each iteration, observations corrupted by AWGN $\mathbf{q}^t = \mathbf{x}^t + \mathbf{A}^T \mathbf{r}^t$ are obtained. We utilize the k -means clustering algorithm [31] to perform vector quantization over the weighted contexts (17). Once we obtain L i.i.d.

³For some measurement matrices such as zero mean Gaussian, AMP tends to follow SE. For other matrices, convergence of AMP might be problematic.

subsequences $\{\mathbf{q}_1, \dots, \mathbf{q}_L\}$, we fit the empirical pdf of each \mathbf{q}_l with a GM model using the approach described in Section IV, and denoise the corresponding subsequence via (21).

To obtain the Onsager correction term in (3), we need to calculate the derivative of η_{hid} (21). Denoting

$$\begin{aligned} f(q) &= \sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2) \left(\frac{\sigma_s^2}{\sigma_s^2 + \sigma_v^2} (q - \mu_s) + \mu_s \right), \\ g(q) &= \sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2), \end{aligned}$$

we have that

$$\begin{aligned} f'(q) &= \sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2) \\ &\quad \cdot \left(\frac{\sigma_s^2 + \mu_s^2 - q\mu_s}{\sigma_s^2 + \sigma_v^2} - \left(\frac{\sigma_s(q - \mu_s)}{\sigma_s^2 + \sigma_v^2} \right)^2 \right), \\ g'(q) &= \sum_{s=1}^S \alpha_s \mathcal{N}(q; \mu_s, \sigma_s^2 + \sigma_v^2) \left(-\frac{q - \mu_s}{\sigma_s^2 + \sigma_v^2} \right). \end{aligned}$$

Therefore,

$$\eta'_{\text{hid}}(q) = \frac{d}{dq} \eta_{\text{hid}}(q) = \frac{f'(q)g(q) - f(q)g'(q)}{(g(q))^2}. \quad (22)$$

Conjecture 1. Consider a linear inverse problem (1) where we want to estimate the input signal \mathbf{x} from its noisy measurements \mathbf{y} and the measurement matrix \mathbf{A} . Under some technical conditions, the AMP-UD algorithm, where the proposed universal denoiser based on context quantization and GM-based i.i.d. subsequence denoising (21), is applied within AMP iterations (2, 3), achieves the MMSE.

VI. NUMERICAL RESULTS

We implement AMP-UD in Matlab (we are in the process of posting our code online) on a Dell OPTIPLEX 9010 running an Intel(R) Core™ i7-3770 with 16GB RAM, and test it utilizing different types of signals at various measurement rates and SNR levels, where SNR is defined as $\text{SNR} = 10 \log_{10}[(N\mathbb{E}[x^2])/(M\sigma_z^2)]$. The input signal length N is 10,000. The context quantization is implemented via the k -means algorithm [31], and the initial number of clusters is set to be 10. That is, L is initialized as 10, and may become smaller if empty clusters occur. The lower bound T on the number of symbols required to learn the GM parameters is 256. The context size $2k$ is chosen to be 12, and the contexts are weighted according to (18) and (19). In order to avoid possible divergence of AMP-UD, we employ a damping technique [32] to slow down the evolution. Specifically, damping is an extra step in the AMP iteration (3), instead of updating the value of \mathbf{x}^{t+1} by the output of the denoiser $\eta^t(\mathbf{A}^T \mathbf{r}^t + \mathbf{x}^t)$, a weighted sum of $\eta^t(\mathbf{A}^T \mathbf{r}^t + \mathbf{x}^t)$ and \mathbf{x}^t is taken as follows,

$$\mathbf{x}^{t+1} = \lambda \eta^t(\mathbf{A}^T \mathbf{r}^t + \mathbf{x}^t) + (1 - \lambda) \mathbf{x}^t, \quad (23)$$

for some constant $0 < \lambda \leq 1$. In our simulation, λ is set to be 0.1, and we run 100 AMP iterations. The recovery performance is evaluated by signal to distortion ratio (SDR =

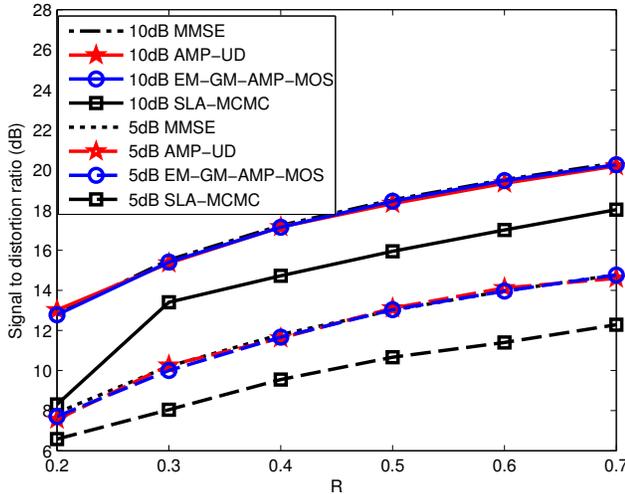


Fig. 4. AMP-UD, SLA-MCMC, and EM-GM-AMP-MOS reconstruction results for an i.i.d. sparse Laplace signal as a function of measurement rate. ($N = 10,000$, SNR = 5 dB or 10 dB.)

$10 \log_{10}(\mathbb{E}[x^2]/\text{MSE})$), where the MSE is averaged over 50 random draws of \mathbf{x} , \mathbf{A} , and \mathbf{z} .

We compare the performance of AMP-UD to (i) the universal CS recovery algorithm SLA-MCMC [22]; and (ii) the Bayesian message passing approaches EM-GM-AMP-MOS [9] for i.i.d. inputs and turboGAMP [10] for non-i.i.d. inputs. Note that EM-GM-AMP-MOS assumes during recovery that the input is i.i.d., whereas turboGAMP is designed for non-i.i.d. inputs with a known statistical model. We do not include results for other well-known CS algorithms such as compressive sensing matching pursuit (CoSaMP) [33] or gradient projection for sparse reconstruction (GPSR) [34], because their SDR performance is consistently weaker than the three algorithms being compared.

i.i.d. sparse Laplace signal: An i.i.d. sparse Laplace signal follows the distribution $p(x) = 0.03\mathcal{L}(0, 1) + 0.97\delta(x)$, where $\mathcal{L}(0, 1)$ denotes a Laplacian distribution with mean zero and variance one, and $\delta(\cdot)$ is the delta function [35]. It is shown in Fig. 4 that AMP-UD and EM-GM-AMP-MOS achieve the MMSE, whereas SLA-MCMC has weaker performance, because the MCMC approach is expected to sample from the posterior and its MSE is twice the MMSE [13, 22].

Markov-uniform signal: Consider the two-state Markov state machine defined in Section II-B with $p_{01} = \frac{3}{970}$ and $p_{10} = 0.10$. A Markov-uniform signal (MUnif for short) follows a uniform distribution $U[0, 1]$ at the nonzero state s_1 . These parameters lead to 3% nonzero entries in a MUnif signal on average. It is shown in Fig. 5 that at low SNR, AMP-UD and SLA-MCMC have similar SDR performance and are both better than turboGAMP despite not knowing the Markovian structure of the signal. At high SNR, the three algorithms are comparable in SDR.

Dense Markov Rademacher signal: Consider the two-state Markov state machine defined in Section II-B with $p_{01} = \frac{3}{70}$ and $p_{10} = 0.10$. A dense Markov Rademacher signal (MRad for short) takes values from $\{-1, +1\}$ with equal probability at s_1 . These parameters lead to 30% nonzero entries in an MRad signal on average. Because

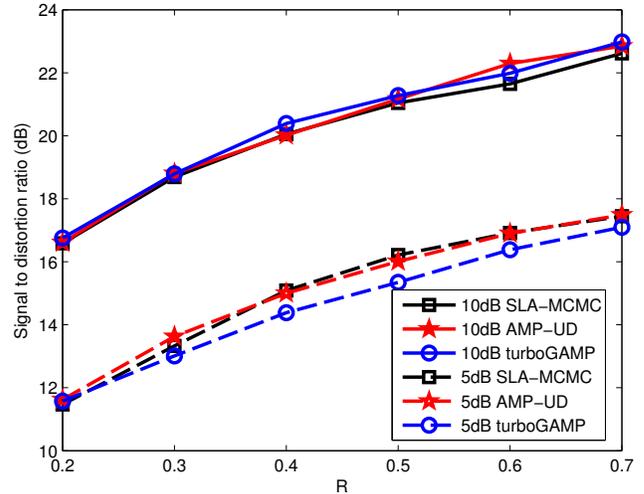


Fig. 5. AMP-UD, SLA-MCMC, and turboGAMP reconstruction results for a two-state Markov signal with nonzero entries drawn from a uniform distribution $U[0, 1]$ as a function of measurement rate. ($N = 10,000$, SNR = 5 dB or 10 dB.)

the MRad signal is dense (non-sparse), we must measure it with somewhat larger measurement rates and SNRs than before. It is shown in Fig. 6 that AMP-UD and SLA-MCMC achieve better overall performance than turboGAMP despite not knowing the Markovian signal structure. AMP-UD outperforms SLA-MCMC except for the lowest tested measurement rate at low SNR. That said, further tests suggest that AMP-UD can be improved in this configuration.

Runtime: The runtime of AMP-UD for MUnif and MRad is typically under 10 minutes, but somewhat more for signals such as sparse Laplace that require a large number of Gaussian components to be fit. For comparison, the runtime of SLA-MCMC is typically an hour, whereas typical runtimes of EM-GM-AMP-MOS and turboGAMP are 30 minutes.

VII. CONCLUSION

In this paper, we introduced a universal CS recovery algorithm AMP-UD that applies our proposed universal denoiser (UD) within approximate message passing (AMP). AMP-UD is designed to reconstruct stationary ergodic inputs from noisy linear measurements. Although AMP-UD is universal and does not know the input statistics, it achieves favorable signal to distortion ratios compared to existing algorithms, and its runtime is typically faster.

The algorithm combines three existing schemes: (i) AMP [5]; (ii) universal denoising [24]; and (iii) a density estimation approach based on Gaussian mixture fitting [25]. In addition to the algorithmic framework, we provided three contributions. First, we provided numerical results showing that SE holds for non-separable Bayesian sliding-window denoisers. Second, we designed a universal denoiser that does not require the input signal to be bounded. Third, we modified the GM learning algorithm, and extended it to an i.i.d. denoiser.

There are numerous directions for future work. First, can we provide theoretical guarantees that our denoiser asymptotically achieves the MMSE for unknown stationary ergodic

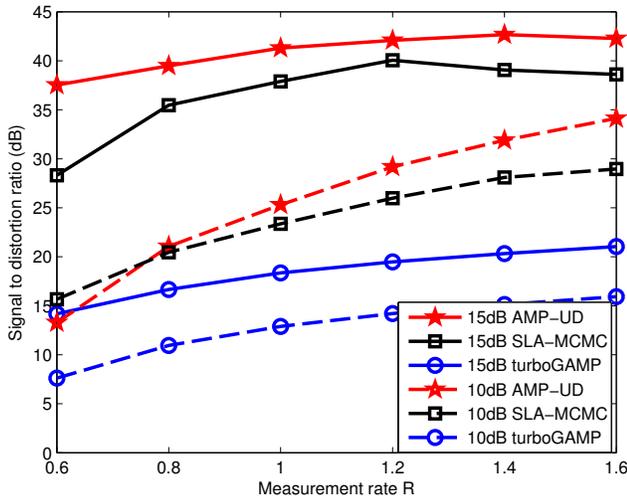


Fig. 6. AMP-UD, SLA-MCMC, and turboGAMP reconstruction results for a dense two-state Markov signal with nonzero entries drawn from a Rademacher (± 1) distribution as a function of measurement rate. ($N = 10,000$, $\text{SNR} = 10$ dB or 15 dB.)

signals? Such guarantees were proven by Sivaramkrishnan and Weissman [24], but we have extended their approach to unbounded signals. Second, it is not clear whether our Gaussian mixture fit can deal with outliers, and we aim to develop a robust density estimation scheme. Third, our current algorithm was designed to minimize the square error, and the denoiser could be modified to minimize other error metrics [36]. Finally, AMP-UD was designed to reconstruct one-dimensional signals. In order to support applications that process multi-dimensional signals such as images, it might be instructive to employ universal image denoisers within AMP.

ACKNOWLEDGEMENTS

We thank Mario Figueiredo and Tsachy Weissman for informative discussions.

REFERENCES

- [1] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [2] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [3] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Royal Stat. Soc. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [4] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.
- [5] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.
- [6] D. Baron, S. Sarvotham, and R. G. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Trans. Signal Process.*, vol. 58, pp. 269–280, Jan. 2010.
- [7] D. L. Donoho, A. Maleki, and A. Montanari, "Message Passing Algorithms for Compressed Sensing: I. Motivation and Construction," in *IEEE Inf. Theory Workshop*, Jan. 2010.
- [8] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, July 2011, pp. 2168–2172.
- [9] J. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Trans. Signal Process.*, vol. 61, no. 19, pp. 4658–4672, Oct. 2013.

- [10] J. Ziniel, S. Rangan, and P. Schniter, "A generalized framework for learning and recovery of structured sparse signals," in *Proc. IEEE Stat. Sig. Proc. Workshop (SSP)*, Aug. 2012, pp. 325–328.
- [11] U. Kamilov, S. Rangan, A. K. Fletcher, and M. Unser, "Approximate message passing with consistent parameter estimation and applications to sparse learning," in *Workshop Neural Info. Proc. Sys. (NIPS)*, Dec. 2012, pp. 2447–2455.
- [12] Y. Ma, J. Tan, N. Krishnan, and D. Baron, "Empirical Bayes and full Bayes for signal estimation," *Arxiv preprint arxiv:1405.2113v1*, May 2014.
- [13] D. L. Donoho, "The Kolmogorov sampler," Stanford University, Stanford, CA, Department of Statistics Technical Report 2002-4, Jan. 2002.
- [14] D. Donoho, H. Kakavand, and J. Mammen, "The simplest solution to an underdetermined system of linear equations," in *Proc. Int. Symp. Inf. Theory (ISIT)*, July 2006, pp. 1924–1928.
- [15] S. Jalali and A. Maleki, "Minimum complexity pursuit," in *Proc. Allerton Conf. Commun., Control, Computing*, Sept. 2011, pp. 1764–1770.
- [16] S. Jalali, A. Maleki, and R. G. Baraniuk, "Minimum complexity pursuit for universal compressed sensing," *IEEE Trans. Inf. Theory*, vol. 60, no. 4, pp. 2253–2268, Apr. 2014.
- [17] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience, 2006.
- [18] M. Li and P. M. B. Vitanyi, *An introduction to Kolmogorov complexity and its applications*. Springer-Verlag, New York, 2008.
- [19] D. Baron, "Information complexity and estimation," in *Workshop Inf. Theoretic Methods Science Eng. (WITMSE)*, Aug. 2011.
- [20] D. Baron and M. F. Duarte, "Universal MAP estimation in compressed sensing," in *Proc. Allerton Conf. Comm., Control, Computing*, Sept. 2011, pp. 768–775.
- [21] J. Zhu, D. Baron, and M. F. Duarte, "Complexity-adaptive universal signal estimation for compressed sensing," in *Proc. IEEE Stat. Sig. Proc. Workshop (SSP)*, June 2014, pp. 416–419.
- [22] —, "Complexity-matching universal signal estimation in compressed sensing," July 2014, submitted.
- [23] K. Sivaramkrishnan and T. Weissman, "Universal denoising of discrete-time continuous-amplitude signals," *IEEE Trans. Inf. Theory*, vol. 54, no. 12, pp. 5632–5660, Dec. 2008.
- [24] —, "A context quantization approach to universal denoising," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2110–2129, June 2009.
- [25] M. Figueiredo and A. Jain, "Unsupervised learning of finite mixture models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, pp. 381–396, Mar. 2002.
- [26] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [27] D. Donoho, I. Johnstone, and A. Montanari, "Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising," *IEEE Trans. Inf. Theory*, vol. 59, no. 6, pp. 3396–3433, June 2013.
- [28] J. Tan, Y. Ma, and D. Baron, "Compressive imaging via approximate message passing with image denoising," *Arxiv preprint arxiv:1405.4429*, May 2014.
- [29] C. Metzler, A. Maleki, and R. G. Baraniuk, "From denoising to compressed sensing," *Arxiv preprint arxiv:1406.4175v2*, June 2014.
- [30] A. Montanari, "Graphical models concepts in compressed sensing," *Compressed Sensing: Theory and Applications*, pp. 394–438, 2012.
- [31] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proc. 5th Berkeley Symp. Math. Stat. & Prob.*, vol. 1, no. 14, 1967, pp. 281–297.
- [32] S. Rangan, P. Schniter, and A. Fletcher, "On the convergence of approximate message passing with arbitrary matrices," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, July 2014, pp. 236–240.
- [33] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harm. Anal.*, vol. 26, no. 3, pp. 301–321, May 2009.
- [34] M. Figueiredo, R. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Top. Sign. Proces.*, vol. 1, pp. 586–597, Dec. 2007.
- [35] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. McGraw Hill Book Co., 1991.
- [36] J. Tan, D. Carmon, and D. Baron, "Signal estimation with additive error metrics in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 60, no. 1, pp. 150–158, Jan. 2014.