

GENERALIZED APPROXIMATE MESSAGE PASSING FOR COSPARSE ANALYSIS COMPRESSIVE SENSING

Mark Borgerding, Philip Schniter, Jeremy Vila

Dept. ECE, The Ohio State University
Columbus, OH 43210

Sundeep Rangan

Dept. of ECE, NYU Polytechnic Institute
Brooklyn, NY 11201

ABSTRACT

In cosparse analysis compressive sensing (CS), one seeks to estimate a non-sparse signal vector from noisy sub-Nyquist linear measurements by exploiting the knowledge that a given linear transform of the signal is cosparse, i.e., has sufficiently many zeros. We propose a novel approach to cosparse analysis CS based on the generalized approximate message passing (GAMP) algorithm. Unlike other AMP-based approaches to this problem, ours works with a wide range of analysis operators and regularizers. In addition, we propose a novel ℓ_0 -like soft-thresholder based on MMSE denoising for a spike-and-slab distribution with an infinite-variance slab. Numerical demonstrations on synthetic and practical datasets demonstrate advantages over existing AMP-based, greedy, and reweighted- ℓ_1 approaches.

Index Terms— Approximate message passing, belief propagation, compressed sensing.

1. INTRODUCTION

We consider the problem of recovering a signal $\mathbf{x} \in \mathbb{R}^N$ (e.g., an N -pixel image) from the possibly noisy linear measurements

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{w} \in \mathbb{R}^M, \quad (1)$$

where Φ represents a known linear measurement operator \mathbf{w} represents noise, and $M \ll N$. We focus on the *analysis compressive sensing* (CS) problem [1, 2] where, for a given analysis operator Ω ,

$$\mathbf{u} \triangleq \Omega \mathbf{x} \in \mathbb{R}^D \quad (2)$$

is assumed to be *cosparse* (i.e., contain sufficiently many zero-valued coefficients). This differs from the *synthesis CS* problem, where \mathbf{x} itself is assumed to be *sparse* (i.e., contain sufficiently few non-zero coefficients). Although the two problems become interchangeable when Ω is invertible, we are mainly interested in non-invertible Ω , e.g., the “overcomplete” case where $D > N$. For notational simplicity, we assume real-valued quantities in the sequel. However, our methods are easily generalized to complex-valued quantities, as we demonstrate in the numerical experiments.

The analysis CS problem is typically formulated as a regularized loss minimization problem of the form

$$\hat{\mathbf{x}}_{\text{rlm}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|_2^2 + h(\Omega \mathbf{x}), \quad (3)$$

with separable regularizer $h(\mathbf{u}) = \sum_{d=1}^D h_d(u_d)$. One of the most famous instances of $h(\mathbf{u})$ is that of total-variation (TV) regularization [3], where $h(\mathbf{u}) = \lambda \|\mathbf{u}\|_1$ and Ω computes variation across

neighboring pixels. In the anisotropic case, this variation is measured by finite difference operators, e.g., $\Omega = [\mathbf{D}_h^H, \mathbf{D}_v^H]^H$, where \mathbf{D}_h computes horizontal differences and \mathbf{D}_v computes vertical differences. Of course, ℓ_1 regularization can be used with generic Ω , with the desirable property that it always renders (3) convex. The resulting problem, sometimes referred to as the *generalized LASSO* (GrLASSO) [4], is amenable to a wide range of efficient optimization techniques like Douglas-Rachford splitting [5] and NESTA [6].

Despite the elegance of the ℓ_1 norm, several studies have shown improvements from the use of ℓ_0 -like norms for $h(\mathbf{u})$, especially for highly overcomplete Ω (i.e., $D \gg N$). For example, the use of iteratively reweighted ℓ_1 [7] has demonstrated significant improvements over ℓ_1 regularization in the context of analysis CS [8, 9]. Likewise, greedy approaches to locate the zero-valued elements in \mathbf{u} have also demonstrated significant improvements over ℓ_1 . Examples include greedy analysis pursuit (GAP), analysis iterative hard thresholding (AIHT), analysis hard thresholding pursuit (AHTP), analysis CoSaMP (ACoSaMP), and analysis subspace pursuit (ASP) [2, 10].

In this paper, we propose a Bayesian approach to analysis CS that leverages recent advances in approximate message passing (AMP) algorithms [11, 12], and in particular the generalized AMP (GAMP) algorithm from [13].

While other AMP-based approaches have been recently proposed for the special case where Ω is a 1D finite difference operator, i.e., the TV-AMP from [14] and the ssAMP from [15], our approach works with a *generic* analysis operator Ω and a much broader range of signal priors and likelihoods. Furthermore, our approach facilitates both MAP and (approximate) MMSE estimation of \mathbf{x} in a computationally efficient manner. We also note that a different Bayesian approach to cosparse analysis CS, based on multivariate Gauss-mixture priors, was recently presented in [16]. The MAP and MMSE estimation methods proposed in [16], which employ greedy pursuit and Gibbs sampling, respectively, have computational complexities that scale as $O(D^2 N)$ and $O(D^2 N^3)$ (assuming $M \leq D$). In contrast, ours scales like $O(DN)$ for generic Ω , or $O(N \log N)$ when Φ and Ω have fast implementations, which is often the case in imaging applications.

2. GENERALIZED AMP FOR ANALYSIS CS

2.1. The proposed Bayesian model

Our approach is Bayesian in that it treats the true signal \mathbf{x} as a realization of a random vector $\mathbf{x} \in \mathbb{R}^N$ with prior pdf $p_{\mathbf{x}}(\mathbf{x})$ and likelihood function $p_{\mathbf{y}|\mathbf{q}}(\mathbf{y}|\Phi \mathbf{x})$, where \mathbf{y} are the observed noisy measurements and $\mathbf{q} \triangleq \Phi \mathbf{x}$ are akin to hidden noiseless measurements. (For clarity, we write random quantities using sans-serif fonts and deterministic ones using serif fonts.) Furthermore, we assume that the

This work was supported by NSF grants CCF-1018368, CCF-1218754, and an allocation of computing time from the Ohio Supercomputer Center.

prior and likelihood have the forms

$$p_{\mathbf{y}|\mathbf{q}}(\mathbf{y}|\Phi\mathbf{x}) \propto \prod_{m=1}^M \exp(-l_m([\Phi\mathbf{x}]_m)) \quad (4)$$

$$p_{\mathbf{x}}(\mathbf{x}) \propto \prod_{d=1}^D \exp(-h_d([\Omega\mathbf{x}]_d)) \prod_{n=1}^N \exp(-g_n(x_n)) \quad (5)$$

with scalar functions $l_m(\cdot)$, $h_d(\cdot)$, and $g_n(\cdot)$. Note that the function $l_m(\cdot)$ is dependent on the measured value y_m . We discuss the design of these functions in the sequel.

Given the form of (4) and (5), the MAP estimate $\hat{\mathbf{x}}_{\text{MAP}} \triangleq \arg \max_{\mathbf{x}} p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y})$ can be written (using Bayes rule) as

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \min_{\mathbf{x}} \{l(\Phi\mathbf{x}) + h(\Omega\mathbf{x}) + g(\mathbf{x})\} \quad (6)$$

with separable loss function $l(\mathbf{q}) = \sum_{m=1}^M l_m(q_m)$ and separable regularizers $g(\mathbf{x}) = \sum_{n=1}^N g_n(x_n)$ and $h(\mathbf{u}) = \sum_{d=1}^D h_d(u_d)$. Note that, with trivial $g(\mathbf{x}) = 0$ and quadratic loss $l(\mathbf{q}) = \frac{1}{2} \|\mathbf{q} - \mathbf{y}\|_2^2$, the MAP estimation problem (6) reduces to the regularized loss minimization problem (3). But clearly (6) is more general.

As for the MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}} \triangleq \int \mathbf{x} p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y}) d\mathbf{x}$, exact evaluation requires the computation of a high dimensional integral that is intractable for most problem sizes of interest. In the sequel, we present a computationally efficient approach to approximate MMSE estimation that is based on loopy belief propagation and, in particular, the GAMP algorithm from [13].

2.2. Background on GAMP

The GAMP algorithm [13] aims to estimate the signal \mathbf{x} , which is assumed to be a realization of random vector $\mathbf{x} \in \mathbb{R}^N$ with known prior $p_{\mathbf{x}}(\mathbf{x})$, from the observations \mathbf{y} generated in accordance with the likelihood function $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{A}\mathbf{x})$. Here, the prior and likelihood are assumed to be separable in the sense that

$$p_{\mathbf{x}}(\mathbf{x}) \propto \prod_{n=1}^N \exp(-g_n(x_n)), \quad p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z}) \propto \prod_{i=1}^I \exp(-f_i(z_i)), \quad (7)$$

where $\mathbf{z} \triangleq \mathbf{A}\mathbf{x} \in \mathbb{R}^I$ can be interpreted as hidden transform outputs. The MAP version of GAMP aims to compute $\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y})$, i.e., solve the optimization problem

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \min_{\mathbf{x}} \sum_{i=1}^I f_i([\mathbf{A}\mathbf{x}]_i) + \sum_{n=1}^N g_n(x_n), \quad (8)$$

while the MMSE version of GAMP aims to compute the MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}} \triangleq \int \mathbf{x} p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y}) d\mathbf{x}$, in both cases by iterating simple, scalar optimizations. MAP-GAMP can be considered as the extension of the AMP algorithm [11] from the quadratic loss $f(\mathbf{z}) = \|\mathbf{y} - \mathbf{z}\|_2^2$ to generic separable losses of the form $f(\mathbf{z}) = \sum_{i=1}^I f_i(z_i)$. Likewise, MMSE-GAMP can be considered as a similar extension of the Bayesian-AMP algorithm [12] from additive white Gaussian noise (AWGN) $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z})$ to generic $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z})$ of the form in (7).

In the large-system limit (i.e., $I, N \rightarrow \infty$ with I/N converging to a finite positive constant) under i.i.d sub-Gaussian \mathbf{A} , GAMP is characterized by a state evolution whose fixed points, when unique, are Bayes optimal [17, 18]. For generic \mathbf{A} , it has been shown [19] that MAP-GAMP's fixed points coincide with the critical points of the cost function (8) and that MMSE-GAMPs fixed points coincide with those of a certain variational cost that was connected to the

definitions for MMSE-GAMP:

$$F_i(\hat{p}, \nu^p) \triangleq \frac{\int z \exp(-f_i(z)) \mathcal{N}(z; \hat{p}, \nu^p) dz}{\int \exp(-f_i(z)) \mathcal{N}(z; \hat{p}, \nu^p) dz} \quad (D1)$$

$$G_n(\hat{r}, \nu^r) \triangleq \frac{\int x \exp(-g_n(x)) \mathcal{N}(x; \hat{r}, \nu^r) dx}{\int \exp(-g_n(x)) \mathcal{N}(x; \hat{r}, \nu^r) dx} \quad (D2)$$

definitions for MAP-GAMP:

$$F_i(\hat{p}, \nu^p) \triangleq \arg \min_z f_i(z) + \frac{1}{2\nu^p} |z - \hat{p}|^2 \quad (D3)$$

$$G_n(\hat{r}, \nu^r) \triangleq \arg \min_x g_n(x) + \frac{1}{2\nu^r} |x - \hat{r}|^2 \quad (D4)$$

inputs:

$$\forall i, n: F_i, G_n, \hat{x}_n(1), \nu_n^x(1), a_{in}, T_{\max} \geq 1, \epsilon \geq 0, \beta_0 \in (0, 1]$$

initialize:

$$\forall i: \hat{s}_i(0) = 0, t = 1$$

for $t = 1, \dots, T_{\max}$,

$$\text{if } t = 1, \text{ then } \beta = 1, \text{ else } \beta = \beta_0 \quad (R1)$$

$$\forall i: \nu_i^p(t) = \beta \sum_{n=1}^N |a_{in}|^2 \nu_n^x(t) + (1-\beta) \nu_i^p(t-1) \quad (R2)$$

$$\forall i: \hat{p}_i(t) = \sum_{n=1}^N a_{in} \hat{x}_n(t) - \nu_i^p(t) \hat{s}_i(t-1) \quad (R3)$$

$$\forall i: \nu_i^z(t) = \nu_i^p(t) F'_i(\hat{p}_i(t), \nu_i^p(t)) \quad (R4)$$

$$\forall i: \hat{z}_i(t) = F_i(\hat{p}_i(t), \nu_i^p(t)) \quad (R5)$$

$$\forall i: \nu_i^s(t) = \beta \left(1 - \frac{\nu_i^z(t)}{\nu_i^p(t)} \right) \frac{1}{\nu_i^p(t)} + (1-\beta) \nu_i^s(t-1) \quad (R6)$$

$$\forall i: \hat{s}_i(t) = \beta \frac{\hat{z}_i(t) - \hat{p}_i(t)}{\nu_i^p(t)} + (1-\beta) \hat{s}_i(t-1) \quad (R7)$$

$$\forall n: \tilde{x}_n(t) = \beta \hat{x}_n(t) + (1-\beta) \tilde{x}_n(t-1) \quad (R8)$$

$$\forall n: \nu_n^r(t) = \beta \left(\frac{1}{\sum_{i=1}^I |a_{in}|^2 \nu_i^s(t)} \right) + (1-\beta) \nu_n^r(t-1) \quad (R9)$$

$$\forall n: \hat{x}_n(t) = \tilde{x}_n(t) + \nu_n^r(t) \sum_{i=1}^I a_{in}^* \hat{s}_i(t) \quad (R10)$$

$$\forall n: \nu_n^x(t+1) = \nu_n^r(t) G'_n(\hat{x}_n(t), \nu_n^r(t)) \quad (R11)$$

$$\forall n: \hat{x}_n(t+1) = G_n(\hat{x}_n(t), \nu_n^r(t)) \quad (R12)$$

$$\text{if } \|\hat{\mathbf{x}}(t) - \hat{\mathbf{x}}(t+1)\| / \|\hat{\mathbf{x}}(t+1)\| < \epsilon, \text{ then stop} \quad (R13)$$

end

outputs: $\forall n: \hat{x}_n(t+1)$

Table 1. The damped GAMP algorithm. In (R4) and (R11), F'_i and G'_n denote the derivatives of F_i and G_n w.r.t their first arguments.

Bethe free entropy in [20]. However, with general \mathbf{A} (e.g., non-zero-mean \mathbf{A} [21] or ill-conditioned \mathbf{A} [22]) GAMP may not converge to its fixed points, i.e., it may diverge. In an attempt to prevent divergence with generic \mathbf{A} , damped [22], adaptively damped [23], sequential [20], and double-loop [24] versions of GAMP have been proposed.

A damped version of the GAMP algorithm is summarized in Table 1. There, smaller values of the damping parameter β_0 make GAMP more robust to difficult \mathbf{A} at the expense of convergence speed, and $\beta_0 = 1$ recovers the original GAMP algorithm from [13]. Note that the only difference between MAP-GAMP and MMSE-GAMP is the definition of the scalar denoisers in (D1)-(D4). Denoisers of the type in (D3)-(D4) are often referred to “proximal operators” in the optimization literature. In fact, as noted in [19] and [22], max-sum GAMP is closely related to primal-dual algorithms from convex optimization, such as the classical Arrow-Hurwicz and recent Chambolle-Pock and primal-dual hybrid gradient algorithms [25–27]. The primary difference between MAP-GAMP and those algorithms is that the primal and dual stepsizes (i.e., $\nu_n^r(t)$ and $1/\nu_i^p(t)$ in Table 1) are adapted, rather than fixed or scheduled.

2.3. GAMP Enables Analysis CS

If we configure GAMP's transform \mathbf{A} and loss function $f(\cdot)$ as

$$\mathbf{A} = \begin{bmatrix} \Phi \\ \Omega \end{bmatrix}, \quad f_i(\cdot) = \begin{cases} l_i(\cdot) & i \in \{1, \dots, M\} \\ h_{i-M}(\cdot) & i \in \{M+1, \dots, M+D\} \end{cases} \quad (9)$$

where Φ and Ω are the measurement and analysis operators from Sec. 2.1, and $l_i(\cdot)$ and $h_d(\cdot)$ are the loss and regularization functions from Sec. 2.1, then MAP-GAMP's optimization problem (8) coincides with the MAP optimization (6), which (as discussed earlier) is a generalization of the analysis-CS problem (3). Likewise, MMSE-GAMP will return an approximation of the MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ under the statistical model (4)-(5).

In the sequel, we refer to GAMP under (9) (with suitable choices of f_q , g_n , and h_d) as “Generalized AMP for Analysis CS,” or GrAMPA. Despite the simplicity of this idea and its importance to, e.g., image recovery, it has (to our knowledge) not been proposed before, outside of our preprint [28].

2.4. Choice of loss and regularization

One of the strengths of GrAMPA is the freedom to choose the loss function $l_i(\cdot)$ and the regularizations $g_n(\cdot)$ and $h_d(\cdot)$.

The quadratic loss $l_m(q) = |y_m - q|^2$, as used in (3), is appropriate for many applications. GrAMPA, however, also supports non-quadratic losses, as needed for 1-bit compressed sensing [29], phase retrieval [30], and Poisson-based photon-limited imaging [31].

The pixel regularization $g_n(\cdot)$ could be used to enforce known positivity in x_n (via $g_n(x) = -\ln \mathbb{1}_{x \geq 0}$), real-valuedness in x_n despite complex-valued measurements (via $g_n(x) = -\ln \mathbb{1}_{x \in \mathbb{R}} \forall n$), or zero-valuedness in x_n (via $g_n(x) = -\ln \mathbb{1}_{x=0} \forall n$). Here, we use $\mathbb{1}_A \in \{0, 1\}$ to denote the indicator of the event A .

As for the analysis regularization $h_d(\cdot)$, the use of $h_d(u) = \lambda|u|$ with MAP-GrAMPA would allow it to tackle the GrLASSO and anisotropic TV problems defined in Sec. 1. With MMSE-GrAMPA, a first instinct might be to use the Bernoulli-Gaussian (BG) prior commonly used for synthesis CS, i.e., $h_d(u) = -\ln((1-\beta)\delta(u) + \beta\mathcal{N}(u; 0, \sigma^2))$, where $\delta(\cdot)$ is the Dirac delta pdf and the parameters β and σ^2 control sparsity and variance, respectively. But the need to tune two parameters is inconvenient, and bias effects from the use of finite σ^2 can degrade performance, especially when $D \gg N$.

Thus, for the MMSE case, we propose a *sparse non-informative parameter estimator* (SNIPE) that can be understood as the MMSE denoiser for a “spike-and-slab” prior with an infinite-variance slab. In particular, SNIPE computes the MMSE estimate of random variable u_d from a $\mathcal{N}(0, \nu_d^q)$ -corrupted observation \hat{q}_d under the prior

$$p_{u_d}(u) = \beta_d p_0(u/\sigma)/\sigma + (1 - \beta_d)\delta(u), \quad (10)$$

in the limiting case that $\sigma \rightarrow \infty$. Here, $\beta_d \in (0, 1]$ is the prior probability that $u_d \neq 0$ and the “slab” pdf $p_0(u)$ is continuous, finite, and non-zero at $u = 0$, but otherwise *arbitrary*. Note that, for fixed σ and β_d , the MMSE estimator can be stated as

$$\begin{aligned} F_d(\hat{q}_d; \nu_d^q) &\triangleq \mathbb{E}\{u_d | \hat{q}_d; \nu_d^q\} = \frac{\int u p_{u_d}(u) \mathcal{N}(u; \hat{q}_d, \nu_d^q) du}{\int p_{u_d}(u) \mathcal{N}(u; \hat{q}_d, \nu_d^q) du} \\ &= \frac{\int u p_0(u/\sigma) \mathcal{N}(u; \hat{q}_d, \nu_d^q) du}{\int p_0(u/\sigma) \mathcal{N}(u; \hat{q}_d, \nu_d^q) du + \sigma \frac{1-\beta_d}{\beta_d} \mathcal{N}(0; \hat{q}_d, \nu_d^q)}. \end{aligned} \quad (11)$$

Since, with any fixed sparsity $\beta_d < 1$, the estimator (11) trivializes to $F_d(\hat{q}_d; \nu_d^q) = 0 \forall \hat{q}_d$ as $\sigma \rightarrow \infty$, we scale the sparsity with σ as $\beta_d = \sigma / (\sigma + p_0(0) \sqrt{2\pi\nu_d^q} \exp(\omega))$ for a tunable parameter $\omega \in \mathbb{R}$, in which case it can be shown that

$$F_d(\hat{q}_d; \nu_d^q, \omega) \xrightarrow{\sigma \rightarrow \infty} \frac{\hat{q}_d}{1 + \exp(\omega - \frac{1}{2}|\hat{q}_d|^2/\nu_d^q)}. \quad (12)$$

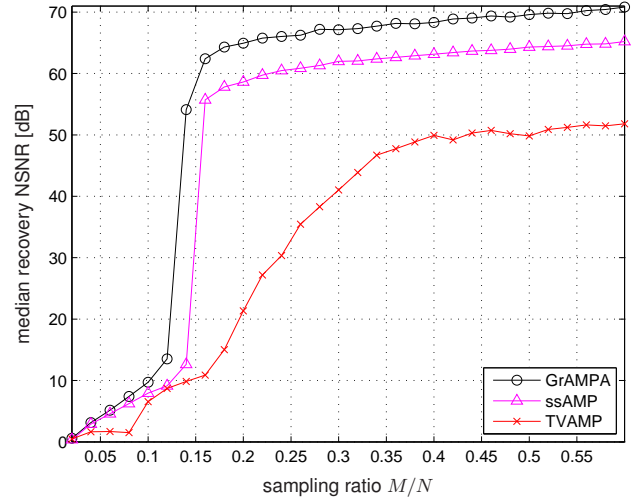


Fig. 1. Recovery of 0.05-sparse Bernoulli-Gaussian finite-difference signals from AWGN-corrupted measurements at SNR = 60 dB.

3. NUMERICAL RESULTS

We now provide numerical results that compare GrAMPA with SNIPE denoising to several existing algorithms for cosparsity analysis CS. In all cases, recovery performance was quantified using $\text{NSNR} \triangleq \|\mathbf{x}\|^2 / \|\hat{\mathbf{x}} - \mathbf{x}\|^2$. Each algorithm was given perfect knowledge of relevant statistical parameters (e.g., noise variance) or in cases where an algorithmic parameter needed to be tuned (e.g., GrLASSO λ or SNIPE ω), the NSNR-maximizing value was used.

3.1. Comparison to ssAMP and TV-AMP

We first replicate an experiment from the ssAMP paper [15]. Using the demonstration code for [15], we generated signal realizations $\mathbf{x} \in \mathbb{R}^N$ that yield BG 1D-finite-difference sequences $\Omega\mathbf{x}$ with sparsity rate 0.05. Then we attempted to recover those signals from AWGN-corrupted observations $\mathbf{y} = \Phi\mathbf{x} + \mathbf{w} \in \mathbb{R}^M$, at an $\text{SNR} \triangleq \|\Phi\mathbf{x}\|_2^2 / \|\mathbf{w}\|_2^2$ of 60 dB, generated with i.i.d Gaussian measurement matrices Φ .

Figure 1 shows median NMSE versus sampling ratio M/N for ssAMP, TV-AMP, and GrAMPA, over 100 problem realizations. There we see GrAMPA uniformly outperforming ssAMP, which uniformly outperforms TV-AMP. We attribute the performance differences to choice of regularization: GrAMPA's SNIPE regularization is closer to ℓ_0 than ssAMP's BG-based regularization, which is closer to ℓ_0 than TV-AMP's ℓ_1 regularization. We note the performance of GrAMPA in Fig. 1 is much better than that reported in [15] due to the misconfiguration of GrAMPA in [15].

3.2. Comparison to GAP: Synthetic cosparsity recovery

We now compare GrAMPA with SNIPE denoising to Greedy Analysis Pursuit (GAP) [2] using an experiment from [2] that constructed $\Omega^T \in \mathbb{R}^{N \times D}$ as a random, almost-uniform, almost-tight frame and \mathbf{x} as an exactly L -cosparsity vector. The objective was then to recover \mathbf{x} from noiseless measurements $\mathbf{y} = \Phi\mathbf{x}$ using analysis operator Ω and i.i.d Gaussian Φ . For this experiment, we used $N = 200$.

Figure 2 shows the empirical phase-transition curves (PTCs) for GAP and GrAMPA versus sampling ratio $\delta = M/N$ and uncer-

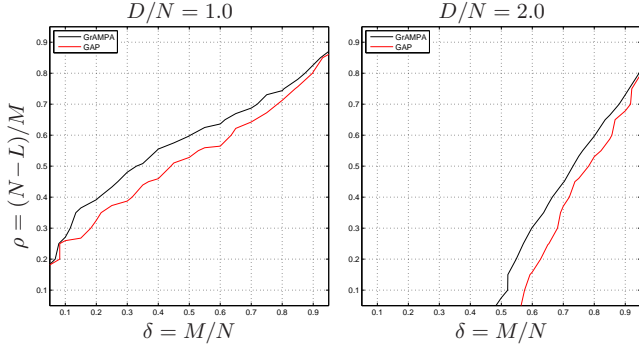


Fig. 2. Phase transition curves for recovery of L -cosparse N -length signals from M noiseless measurements under i.i.d Gaussian Φ and an $N \times D$ random, almost-uniform, almost-tight frame Ω^T .

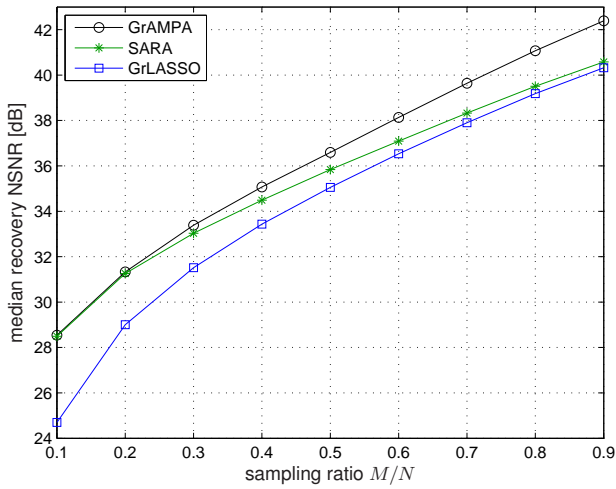


Fig. 3. Recovery of the 512×512 Lena image under a measurement SNR of 40 dB, spread-spectrum Φ , and Db1-8 concatenated Ω .

tainty ratio $\rho = (N - L)/M$. For points below the PTC, recovery was successful with high probability, while for points above the PTC, recovery was unsuccessful with high probability. Here, we defined “success” as $\text{NSNR} \geq 10^6$. Figure 2 shows that the PTC of GrAMPA is uniformly better than that of GAP. It also shows that, for both algorithms, the PTC approaches the feasibility boundary (i.e., $\rho = 1$) as $M/N \rightarrow 1$ but that, as the analysis operator becomes more overcomplete (i.e., D/N increases), the PTC progressively weakens.

3.3. Compressive image recovery via sparsity averaging

Next, we repeat an experiment from [8], where the $N = 512 \times 512$ Lena image \mathbf{x} was recovered from M noisy complex-valued measurements $\mathbf{y} = \Phi \mathbf{x} + \mathbf{w}$ at $\text{SNR} = 40$ dB. The measurements were of the “spread spectrum” form: $\Phi = \mathbf{MFC}$, where \mathbf{C} was diagonal with random ± 1 entries, \mathbf{F} was an N -FFT, and $\mathbf{M} \in \{0, 1\}^{M \times N}$ contained rows of \mathbf{I}_N selected uniformly at random. An overcomplete dictionary $\Psi \in \mathbb{R}^{N \times 8N}$ was constructed from a horizontal concatenation of the first 8 Daubechies orthogonal DWT matrices, yielding the analysis operator $\Omega = \Psi^T$. The use of highly overcomplete concatenated dictionaries is dubbed “sparsity averaging” in [8].

Figure 3 shows median NSNR (over 30 Monte-Carlo trials) versus sampling ratio M/N for GrAMPA with SNIPE denoising; for

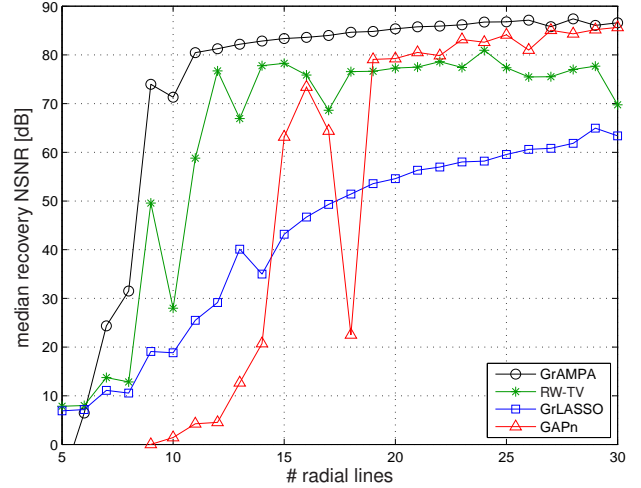


Fig. 4. Recovery of the 64×64 Shepp-Logan phantom from 2D FFT Φ and 2D horizontal, vertical, and diagonal finite-difference Ω .

“SARA” from [8], which employs iteratively-reweighted- ℓ_1 [7]; and for GrLASSO implemented via the “SOPT” Matlab code that accompanies [8], which employs Douglas-Rachford splitting [5]. All algorithms enforced non-negativity in the estimate. Figure 3 shows GrAMPA outperforming the other algorithms in NSNR at all sampling ratios M/N . Averaging over trials where all algorithms gave recovery $\text{NSNR} \geq 30$ dB, the runtimes of GrAMPA, GrLASSO, and SARA were 220, 255, and 2687 seconds, respectively.

3.4. Shepp-Logan phantom recovery via 2D finite-differences

Finally, we investigated the recovery of the $N = 64 \times 64$ Shepp-Logan Phantom image from 2D Fourier radial-line measurements $\mathbf{y} = \Phi \mathbf{x} + \mathbf{w}$ at $\text{SNR} = 80$ dB, using an analysis operator Ω composed of horizontal, vertical, diagonal, and anti-diagonal 2D finite differences, as described in the noise-tolerant GAP paper [32].

Figure 4 plots median recovery NSNR (over 11 Monte-Carlo trials) versus number of radial lines for GrAMPA with SNIPE denoising, GAPn [32], the “RW-TV” approach from [8], which employs iteratively-weighted- ℓ_1 [7], and GrLASSO, implemented using the Douglas-Rachford based “SOPT” Matlab code from [8]. The figure shows that GrAMPA achieved the best phase transition and also the best NSNR (for all numbers of radial lines above 6). Averaging over trials where all algorithms gave recovery $\text{NSNR} \geq 30$ dB, the runtimes of GrAMPA, GrLASSO, RW-TV, and GAP were 0.28, 1.8, 9.7, and 30.1 seconds, respectively.

4. CONCLUSIONS

In this work, we proposed the “Generalized AMP for Analysis CS” (GrAMPA) algorithm, a new AMP-based approach to analysis CS that can be used with a wide range of loss functions, regularization terms, and analysis operators. In addition, we proposed the “Sparse Non-informative Parameter Estimator” (SNIPE), an ℓ_0 -like soft thresholder that corresponds to the MMSE denoiser for a spike-and-slab distribution with an infinite-variance slab. Numerical experiments comparing GrAMPA with SNIPE to several other recently proposed analysis-CS algorithms show improved recovery performance and excellent runtime. Online tuning of the SNIPE parameter ω will be considered in future work.

5. REFERENCES

- [1] Michael Elad, Peyman Milanfar, and Ron Rubinstein, "Analysis versus synthesis in signal priors," *Inverse Problems*, vol. 23, pp. 947–968, 2007.
- [2] S. Nam, M. E. Davies, M. Elad, and R. Gribonval, "The cospase analysis model and algorithms," *Appl. Computational Harmonic Anal.*, vol. 34, no. 1, pp. 30–56, Jan. 2013.
- [3] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, pp. 259–268, 1992.
- [4] Ryan J. Tibshirani, "Solution path of the generalized lasso," *Ann. Statist.*, vol. 39, no. 3, pp. 1335–1371, 2011.
- [5] P. L. Combettes and J.-C. Pesquet, "A Douglas-Rachford splitting approach to nonsmooth convex variational signal recovery," *IEEE J. Sel. Topics Signal Process.*, vol. 1, no. 4, pp. 6564–574, Dec. 2007.
- [6] S. Becker, J. Bobin, and E. J. Candès, "NESTA: A fast and accurate first-order method for sparse recovery," *SIAM J. Imag. Sci.*, vol. 4, no. 1, pp. 1–39, 2011.
- [7] E. J. Candès, M. B. Wakin, and S. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. App.*, vol. 14, no. 5, pp. 877–905, Dec. 2008.
- [8] Rafael E. Carrillo, Jason D. McEwen, Dimitri Van De Ville, Jean-Philippe Thiran, and Yves Wiaux, "Sparsity averaging for compressive imaging," *IEEE Signal Process. Lett.*, vol. 20, no. 6, pp. 591–594, 2013.
- [9] R. Chartrand, E. Y. Sidky, and X. Pan, "Nonconvex compressive sensing for X-ray CT: An algorithm comparison," in *Proc. Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, Nov. 2013, pp. 665–669.
- [10] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. E. Davies, "Greedy-like algorithms for the cospase analysis model," *arXiv:1207.2456*, July 2012.
- [11] 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.
- [12] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inform. Theory Workshop*, Cairo, Egypt, Jan. 2010, pp. 1–5.
- [13] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inform. Thy.*, Aug. 2011, pp. 2168–2172, (full version at *arXiv:1010.5141*).
- [14] D. L. Donoho, I. M. Johnstone, and A. Montanari, "Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising," *IEEE Trans. Inform. Theory*, vol. 59, no. 6, June 2013.
- [15] Jaewook Kang, Hyoyoung Jung, Heung-No Lee, and Kiseon Kim, "Spike-and-slab approximate message-passing for high-dimensional piecewise-constant recovery," *arXiv:1408.3930v1*, Aug. 2014, (Matlab codes at https://sites.google.com/site/jwkang10/ssAMP_algorithm_ver1_Aug2014.zip).
- [16] J. S. Turek, I. Yavneh, and M. Elad, "On MAP and MMSE estimators for the co-sparse analysis model," *Digital Signal Process.*, vol. 28, pp. 57–74, 2014.
- [17] Adel Javanmard and Andrea Montanari, "State evolution for general approximate message passing algorithms, with applications to spatial coupling," *Inform. Inference*, vol. 2, no. 2, pp. 115–144, 2013.
- [18] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and iterative algorithms," in *Proc. IEEE Int. Symp. Inform. Thy.*, Boston, MA, June 2012, pp. 1643–1647, (full paper at *arXiv:1207.7321*).
- [19] S. Rangan, P. Schniter, E. Riegler, A. Fletcher, and V. Cevher, "Fixed points of generalized approximate message passing with arbitrary matrices," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2013, pp. 664–668, (full version at *arXiv:1301.6295*).
- [20] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová, "Variational free energies for compressed sensing," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 1499–1503.
- [21] F. Caltagirone, F. Krzakala, and L. Zdeborová, "On convergence of approximate message passing," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 1812–1816.
- [22] S. Rangan, P. Schniter, and A. Fletcher, "On the convergence of generalized approximate message passing with arbitrary matrices," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 236–240, (full version at *arXiv:1402.3210*).
- [23] Jeremy Vila, Philip Schniter, Sundeeep Rangan, Florent Krzakala, and Lenka Zdeborová, "Adaptive damping and mean removal for the generalized approximate message passing algorithm," in *IEEE ICASSP*, 2015, to appear.
- [24] S. Rangan, A. K. Fletcher, P. Schniter, and U. Kamilov, "Inference for generalized linear models via alternating directions and Bethe free energy minimization," *arXiv:1501.01797*, Jan. 2015.
- [25] E. Esser, X. Zhang, and T. F. Chan, "A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science," *SIAM J. Imaging Sci.*, vol. 3, no. 4, pp. 1015–1046, 2010.
- [26] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imaging Vis.*, vol. 40, pp. 120–145, 2011.
- [27] B. He and X. Yuan, "Convergence analysis of primal-dual algorithms for a saddle-point problem: From contraction perspective," *SIAM J. Imaging Sci.*, vol. 5, pp. 119–149, 2012.
- [28] M. Borgerding and P. Schniter, "Generalized approximate message passing for the cospase analysis model," *arXiv:1312.3968*, 2013, (Matlab codes at <http://www2.ece.ohio-state.edu/~schniter/GrAMPA>).
- [29] U. S. Kamilov, V. K. Goyal, and S. Rangan, "Message-passing de-quantization with applications to compressed sensing," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6270–6281, Dec. 2012.
- [30] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Proc. Allerton Conf. Commun. Control Comput.*, Monticello, IL, Oct. 2012, pp. 815–822.
- [31] D. L. Snyder, A. M. Hammond, and R. L. White, "Image recovery from data acquired with a charge-coupled-device camera," *J. Am. Statist. Assoc.*, vol. 10, pp. 1014–1023, 1993.
- [32] S. Nam, M. E. Davies, M. Elad, and R. Gribonval, "Recovery of cospase signals with greedy analysis pursuit in the presence of noise," in *Proc. IEEE Workshop Comp. Adv. Multi-Sensor Adaptive Process.*, Puerto Rico, Dec. 2011.