

Bilinear Recovery Using Adaptive Vector-AMP

Subrata Sarkar¹, *Student Member, IEEE*, Alyson K. Fletcher¹, *Member, IEEE*,
Sundeep Rangan², *Fellow, IEEE*, and Philip Schniter³, *Fellow, IEEE*

Abstract—We consider the problem of jointly recovering the vector b and the matrix C from noisy measurements $Y = A(b)C + W$, where $A(\cdot)$ is a known affine linear function of b (i.e., $A(b) = A_0 + \sum_{i=1}^Q b_i A_i$ with known matrices A_i). This problem has applications in matrix completion, robust PCA, dictionary learning, self-calibration, blind deconvolution, joint-channel/symbol estimation, compressive sensing with matrix uncertainty, and many other tasks. To solve this bilinear recovery problem, we propose the Bilinear Adaptive Vector Approximate Message Passing (VAMP) algorithm. We demonstrate numerically that the proposed approach is competitive with other state-of-the-art approaches to bilinear recovery, including lifted VAMP and Bilinear Generalized Approximate Message Passing.

Index Terms—Approximate message passing, expectation propagation, expectation maximization, self-calibration, computed tomography, dictionary learning.

I. INTRODUCTION

A. Motivation

MANY problems of interest in science and engineering can be formulated as estimation of a structured matrix Z from noisy or incomplete measurements. The type of structure in Z determines the specific subproblem to be solved.

For example, when Z has a low-rank structure and only a subset of its entries are observed, the problem is known as *matrix completion* [1]. When $Z = L + S$ for low-rank L and sparse S , the problem of estimating L and S is known as *robust principle components analysis* (RPCA) [2]. When $Z = BC$ with sparse C , the problem of estimating B and C is known as *dictionary learning* [3]. When $Z = BC$ with

Manuscript received August 31, 2018; revised March 8, 2019; accepted May 1, 2019. Date of publication May 10, 2019; date of current version May 28, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Gwo Giun Lee. The work of S. Sarkar and P. Schniter was supported by the National Science Foundation under Grant 1716388. The work of A. K. Fletcher was supported in part by the National Science Foundation under Grant 1738285 and Grant 1738286, and in part by the Office of Naval Research under Grant N00014-15-1-2677. The work of S. Rangan was supported in part by the National Science Foundation under Grant 1116589, Grant 1302336, and Grant 1547332, and in part by the industrial affiliates of NYU WIRELESS. This paper was presented in part at the Information Theory and Applications Workshop, University of California San Diego, San Diego, CA, USA, February 2018. (Corresponding author: Philip Schniter.)

S. Sarkar and P. Schniter are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43202 USA (e-mail: sarkar.51@osu.edu; schniter.1@osu.edu).

A. K. Fletcher is with the Department of Statistics and Electrical Engineering, the University of California, Los Angeles, Los Angeles, CA 90095 USA (e-mail: akfletcher@ucla.edu).

S. Rangan is with the Department of Electrical and Computer Engineering, New York University, Brooklyn, NY 11201 USA (e-mail: srangan@nyu.edu).
Digital Object Identifier 10.1109/TSP.2019.2916100

nonnegative B and C , the problem is known as *nonnegative matrix factorization* (NMF) [4].

Sometimes Z has a more complicated structure. For example, the problems of *self-calibration* and *blind (circular) deconvolution* [5] can be formulated using $Z = \text{Diag}(Hb)\Psi C$, where H and Ψ are known and b and C are to be estimated.¹ The problem of *compressive sensing (CS) with matrix uncertainty* [6] can be formulated using $z = \sum_i b_i A_i c$, where $\{A_i\}$ are known and where b and sparse c are to be estimated. The latter covers the problem of *joint channel-symbol estimation* [7], in which case b_i are the data symbols, c contains (possibly sparse) channel coefficients, and the known $\{A_i\}$ are determined by the modulation scheme. The more general problem of *matrix CS* [8], [9] results from

$$z_m = \text{tr}\{A_m^\top (L + S)\} \text{ for } m = 1, \dots, M, \quad (1)$$

where $\{A_m\}$ are known and the goal is to estimate low-rank L and sparse S .

B. Prior Work

Many algorithms have been developed to solve the above problems. Some solve a convex relaxation of the original problem, while others attack non-convex formulations via alternating methods, greedy methods, variational methods, message-passing methods, and other techniques.

For matrix completion, well-known approaches include the nuclear-norm-based convex optimization method IALM [10], the non-convex successive over-relaxation approach LMAFit [11], the Grassmannian gradient-descent approach GROUSE [12], the greedy hard-thresholding approach Matrix-ALPS [13], and the variational-Bayes method VSBL [14]. For RPCA, there are also versions of IALM [10], LMAFit [11], and VSBL [14], as well as a robust cousin of GROUSE, called GRASTA [15]. For dictionary learning, there is the greedy K-SVD algorithm [16], the online SPAMS approach [17], and the ER-SpUD approach from [18]. A unified approach to matrix completion, RPCA, and dictionary learning was proposed in [19]–[21] using an extension of the approximate message-passing (AMP) methodology from [22], [23]. The resulting “bilinear generalized AMP” (BiGAMP) algorithm was compared to the aforementioned methods in [20] and found (empirically) to be competitive, if not superior, in phase transition and runtime. A related approach known as LowRAMP was proposed [24] and analyzed in [25], [26].

¹Here and in the sequel, we use lowercase bold notation for vectors and uppercase bold notation for matrices.

For self-calibration and blind deconvolution, well-known approaches include the convex relaxations from [5], [27], [28] and the alternating method from [29]. For CS with matrix uncertainty, there is the award-winning non-convex method [6]. For matrix CS, the well-known papers [8], [30], [31] proposed convex approaches and [9], [13] proposed greedy approaches. See the recent overview [32] for many other works. An AMP-based approach to self-calibration, blind deconvolution, CS with matrix uncertainty, and matrix CS was proposed in [33] and analyzed in [34]. This “parametric BiGAMP” (PBiGAMP) was compared to the above works in [33] and found to yield improved empirical phase transitions.

More recently, AMP methods for bilinear inference were proposed using the “lifting” approach (see, e.g., [5], [28], [32], [35] for seminal papers on lifting). To illustrate the idea, suppose that the measurement vector $y \in \mathbb{R}^M$ takes the form

$$y = \sum_{i=1}^Q \sum_{j=1}^N b_i a_{i,j} c_j + w, \quad (2)$$

where $a_{i,j} \in \mathbb{R}^M$ is known for all i, j and the goal is to recover $b = [b_1, \dots, b_Q]^\top$ and $c = [c_1, \dots, c_N]^\top$ in the presence of white noise w . Rewriting the measurements as

$$y = \sum_{i=1}^Q b_i \underbrace{[a_{i,1}, \dots, a_{i,N}]}_{\triangleq A_i} c + w \quad (3)$$

$$\begin{aligned} &= \underbrace{[A_1 \cdots A_Q]}_{\triangleq A} \begin{bmatrix} b_1 c \\ \vdots \\ b_Q c \end{bmatrix} + w \quad (4) \\ &= Ax + w \text{ for } x = b \otimes c = \text{vec}(cb^\top), \quad (5) \end{aligned}$$

we see that the noisy bilinear recovery problem (2) can be rewritten as the noisy linear recovery problem (5) with a rank-one structure on (the matrix form of) x . Thus, if this low-rank signal structure can be exploited by a linear inference algorithm, then bilinear inference can be accomplished. This is precisely what was proposed in [36], building on the non-separable-denoising version of the AMP algorithm from [37]. A rigorous analysis of “lifted AMP” was presented in [38].

The trouble with AMP is that its behavior is understood only in the case of large [39] or infinitely large, i.i.d. (sub) Gaussian A [40], [41]. Even small deviations from this scenario (e.g., mildly ill-conditioned and/or non-zero-mean A) can cause AMP to diverge [42]–[44]. To address this issue, an alternative called Vector AMP (VAMP) was proposed and analyzed in [45], with close connections to expectation propagation [46] (see also [47]–[49]). There it was established that, if A is an infinitely large right-rotationally invariant² random matrix and the denoising function used by VAMP is separable and Lipschitz, then VAMP’s performance can be exactly predicted by a scalar

state-evolution that also provides testable conditions for optimality. Since the class of right-rotationally invariant matrices is much larger than the class of i.i.d. Gaussian matrices, VAMP is much more robust than AMP with regards to the construction of A . For example, VAMP has no problem with ill-conditioned or mean-shifted matrices [45].

Very recently, [50] performed a rigorous analysis of VAMP under *non*-separable Lipschitz denoisers, showing that—here too—VAMP’s behavior is exactly predicted by a scalar state-evolution when A is infinitely large and right-rotationally invariant. Furthermore, [50] demonstrated the success of lifted VAMP on bilinear problems such as self-calibration and CS with matrix uncertainty. In addition, [50] gave evidence that, like AMP, the PBiGAMP algorithm is sensitive to deviations from the i.i.d. assumptions used in its derivation [33] and analysis [34]. For this reason, lifted VAMP significantly outperformed PBiGAMP in some cases [50].

Despite its good performance and rigorous analyses under infinitely large right-rotationally invariant random A , lifted VAMP suffers from computational issues brought on by the lifting itself: The $N + Q$ unknowns $[b, c]$ in the bilinear problem (2) manifest as NQ unknowns x after lifting to (5). This is a serious problem when N and Q are both large. As a concrete example, consider the application of lifting to (square) dictionary learning, where the goal is to recover $B \in \mathbb{R}^{N \times N}$ and sparse $C \in \mathbb{R}^{N \times L}$ from noisy measurements $Y = BC + W$. This bilinear relationship can be lifted via

$$Y = \sum_{ij} b_{i,j} A_{i,j} C + W \quad (6)$$

$$= \underbrace{[A_{1,1} \cdots A_{N,N}]}_{\triangleq A \in \mathbb{R}^{N \times N^3}} \underbrace{(b \otimes C)}_{\triangleq X \in \mathbb{R}^{N^3 \times L}} + W, \quad (7)$$

where $A_{i,j} \in \mathbb{R}^{N \times N}$ is constructed with a 1 in the (i, j) th position and zeros elsewhere, and where $b = [b_{1,1}, \dots, b_{N,N}]^\top \in \mathbb{R}^{N^2}$. Even at the relatively small patch size of 8×8 (i.e., $N = 64$), the matrix A has dimension $64 \times 262\,144$, and the unknown matrix X has dimension $262\,144 \times L$. The rule-of-thumb $L = 5N \ln N$ [18] then gives $L = 1331$, in which case X contains 3.5×10^8 entries, which leads to difficulties with computation and memory.

C. Contributions

In this paper, we present a novel VAMP-based approach to bilinear recovery. With the aim of computational efficiency, we avoid lifting and instead build on the recently proposed Adaptive VAMP framework from [51]. However, different from [51], which focused on noisy *linear* recovery, we focus on noisy *bilinear* recovery.

In particular, we focus on recovering $\{b_i\}$ and C from noisy measurements $Y \in \mathbb{R}^{M \times L}$ of the form

$$Y = \sum_{i=1}^Q b_i A_i C + W, \quad (8)$$

²If A is right-rotationally invariant then its singular value decomposition $A = USV^\top$ has Haar distributed V , i.e., V is uniformly distributed over the group of orthogonal matrices.

where $\{A_i\}$ are known and W contains white noise. Note that (8) is a multiple-measurement vector (MMV) extension of (3), and that it covers all of the motivating problems discussed in Section I-A. For example, in self-calibration, where we estimate b and C from $Y = \text{Diag}(Hb)\Psi C + W$, we can set $A_i = \text{Diag}(h_i)\Psi$, where h_i is the i th column of H . Or, in dictionary learning, where we estimate B and C from $Y = BC + W$, we can write $B = \sum_{i=1}^{MN} b_i A_i$ for $A_i = e_{\langle i-1 \rangle_M} e_{\lfloor (i-1)/M \rfloor}^\top$, where $\langle i \rangle_M$ denotes i -modulo- M , $\lfloor \cdot \rfloor$ denotes floor, and $\{e_i\}$ is the standard basis.

When deriving³ the proposed method, we treat $\{b_i\}$ as deterministic unknowns and the entries of C as random variables. The prior distribution on C is assumed to be known up to some (possibly) unknown hyperparameters, which are learned jointly with $\{b_i\}$ and C . Also, W is treated as additive white Gaussian noise (AWGN) with an unknown variance that is also learned. More details are provided in the sequel.

We show (empirically) that the proposed Bilinear Adaptive VAMP (BAd-VAMP) method performs as well as the EM-PBiGAMP algorithm from [33], with regard to accuracy and computational complexity, when the underlying matrices are i.i.d., as assumed for the derivation of PBiGAMP. However, we will show that BAd-VAMP outperforms EM-PBiGAMP when the underlying matrices become ill-conditioned. In the ill-conditioned case, we show that BAd-VAMP performs as well as, and sometimes significantly better than, lifted VAMP. However, BAd-VAMP is much more computationally efficient due to its avoidance of lifting. In this sense, the proposed BAd-VAMP is shown to be accurate, robust, and computationally efficient.

Notation: In this paper, we use boldface uppercase letters to denote matrices (e.g., X), boldface lowercase letters for vectors (e.g., x) and non-bold letters for scalars (e.g., x). Given a matrix X , we use x_l to denote the l th column and x_{nl} to denote the element in the n th row and l th column. We use $\mathbb{E}[f(x)|b]$ to denote the expectation of $f(x)$ w.r.t. the density b , i.e., $\mathbb{E}[f(x)|b] = \int f(x)b(x) dx$, and we use $\text{var}[f(x)|b]$ for the corresponding variance. We use $\text{Diag}(x)$ to denote the diagonal matrix created from vector x , and $\text{diag}(X)$ to denote the vector of elements on the diagonal of the matrix X .

II. PROPOSED FRAMEWORK

In an effort to make our algorithmic development more consistent with the VAMP papers [45], [50], [51], we now make some minor notational changes relative to (8). First, we will use the notation $A(b) \triangleq \sum_i b_i A_i$ to be concise. Second, the quantities b_i and C in (8) will be changed to $\theta_{A,i}$ and X , respectively.

The problem of interest can thus be stated as follows: estimate the matrix $X \in \mathbb{R}^{N \times L}$ and learn the parameters $\Theta \triangleq \{\theta_A, \theta_x, \gamma_w\}$ in the statistical model

$$Y = A(\theta_A)X + W \quad (9a)$$

$$X \sim p_{\mathbf{X}}(\cdot; \theta_x), \quad w_{ml} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \gamma_w^{-1}), \quad (9b)$$

³Although the derivation treats the entries of C as random variables and the associated denoiser as Bayesian, the final algorithm is more general in that it only requires the denoiser to be Lipschitz.

where $A(\cdot)$ is a known matrix-valued linear function, and $p_{\mathbf{X}}(\cdot; \theta_x)$ is a density on X parameterized by the vector θ_x . Here, γ_w is the noise precision, i.e., the inverse noise variance.

More precisely, we aim to compute the maximum-likelihood (ML) estimate of Θ and, under that estimate, compute the minimum mean-squared error (MMSE) estimate of X , i.e.,

$$\hat{\Theta}_{\text{ML}} = \arg \max_{\Theta} p_{\mathbf{Y}}(Y; \Theta) \quad (10)$$

$$\hat{X}_{\text{MMSE}} = \mathbb{E}[X|Y; \hat{\Theta}_{\text{ML}}]. \quad (11)$$

In (10), $p_{\mathbf{Y}}(Y; \Theta)$ is the likelihood function of Θ , which can be written as

$$p_{\mathbf{Y}}(Y; \Theta) = \int p_{\mathbf{X}}(X; \Theta) p_{\mathbf{Y}|\mathbf{X}}(Y|X; \Theta) dX. \quad (12)$$

In (11), the expectation is taken over the posterior density

$$p_{\mathbf{X}|\mathbf{Y}}(X|Y; \hat{\Theta}_{\text{ML}}) = \frac{p_{\mathbf{X}}(X; \hat{\Theta}_{\text{ML}}) p_{\mathbf{Y}|\mathbf{X}}(Y|X; \hat{\Theta}_{\text{ML}})}{p_{\mathbf{Y}}(Y; \hat{\Theta}_{\text{ML}})}. \quad (13)$$

The statistical model (9) implies that⁴

$$p_{\mathbf{Y}|\mathbf{X}}(Y|X; \Theta) = \prod_{l=1}^L p_{\mathbf{y}|x}(y_l|x_l; \Theta) \quad (14)$$

where

$$p_{\mathbf{y}|x}(y|x; \Theta) = \mathcal{N}(y; A(\theta_A)x, I/\gamma_w). \quad (15)$$

For simplicity, we treat $\{x_l\}_{l=1}^L$ as i.i.d. in the sequel, so that

$$p_{\mathbf{X}}(X; \Theta) = \prod_{l=1}^L p_{\mathbf{x}}(x_l; \theta_x) \quad (16)$$

for some density $p_{\mathbf{x}}(\cdot; \theta_x)$ parameterized by θ_x . In this case, the posterior density decouples as

$$p_{\mathbf{X}|\mathbf{Y}}(X|Y; \Theta) \propto \prod_{l=1}^L p_{\mathbf{x}}(x_l; \Theta) p_{\mathbf{y}|x}(y_l|x_l; \Theta). \quad (17)$$

III. BACKGROUND

A. Background on VAMP

Recalling (11), we are interested in computing the MMSE estimate of X from the noisy measurements Y . This problem is solved (under certain conditions) by the VAMP approach from [45]. We now review VAMP at the detail needed for further development of BAd-VAMP.

From (11), the MMSE estimate of X equals the mean of the posterior pdf $p_{\mathbf{X}|\mathbf{Y}}$. Because $p_{\mathbf{X}|\mathbf{Y}}$ decouples across the columns of X , as in (17), it suffices to consider a single column and drop the l notation for simplicity. Also, for now, we will assume that Θ are the true parameters used to generate (Y, X) and drop the Θ notation for simplicity; we will revisit the estimation of Θ in Section III-B. With these simplifications, (9) reduces to

$$y = Ax + \mathcal{N}(0, I/\gamma_w), \quad x \sim p_{\mathbf{x}}. \quad (18)$$

⁴In (14)–(16), to promote notational simplicity, the left side of the equation is written using Θ even though the right side depends on a subset of Θ .

Recall that the MMSE estimate of x equals

$$\mathbb{E}[x|y] = \int x p_{\mathbf{x}|y}(x|y) dx \quad (19)$$

for the posterior density

$$p_{\mathbf{x}|y}(x|y) = Z(y)^{-1} p_{\mathbf{x}}(x) p_{y|\mathbf{x}}(y|x), \quad (20)$$

where $Z(y)$ is the normalizing constant

$$Z(y) \triangleq \int p_{\mathbf{x}}(x) p_{y|\mathbf{x}}(y|x) dx. \quad (21)$$

For high-dimensional x , the integrals in (19) and (21) are difficult to compute directly. Thus other methods must be used.

Variational inference (VI) [52] can be used to bypass the computation of $Z(y)$. For example, notice that the true posterior $p_{\mathbf{x}|y}$ can be recovered by solving the variational optimization (over densities)

$$\hat{q} = \arg \min_q D_{\text{KL}}(q \| p_{\mathbf{x}|y}), \quad (22)$$

where $D_{\text{KL}}(q \| p)$ denotes the KL divergence from p to q , i.e.,

$$D_{\text{KL}}(q \| p) \triangleq \int q(x) \ln \frac{q(x)}{p(x)} dx. \quad (23)$$

Plugging (20) into (23), we see that

$$\begin{aligned} D_{\text{KL}}(q \| p_{\mathbf{x}|y}) &= D_{\text{KL}}(q \| p_{\mathbf{x}}) + D_{\text{KL}}(q \| p_{y|\mathbf{x}}) + H(q) \\ &\quad + \ln Z(y) \end{aligned} \quad (24)$$

where $H(q) \triangleq - \int q(x) \ln q(x) dx$ is the differential entropy of $x \sim q$. Thus it follows from (22) and (24) that

$$\hat{q} = \arg \min_q \{ D_{\text{KL}}(q \| p_{\mathbf{x}}) + D_{\text{KL}}(q \| p_{y|\mathbf{x}}) + H(q) \}, \quad (25)$$

which bypasses $Z(y)$. Still, solving (25) is difficult in most cases of interest. The typical response is to impose constraints on q , but doing so compromises \hat{q} and its mean.

We take a different approach. Using the “Gibbs free energy”

$$J(q_1, q_2, q_3) \triangleq D_{\text{KL}}(q_1 \| p_{\mathbf{x}}) + D_{\text{KL}}(q_2 \| p_{y|\mathbf{x}}) + H(q_3), \quad (26)$$

one can rewrite (25) as⁵

$$\arg \min_{q_1} \min_{q_2} \max_{q_3} J(q_1, q_2, q_3) \quad (27a)$$

$$\text{s.t. } q_1 = q_2 = q_3. \quad (27b)$$

But, as discussed earlier, (27) is difficult to solve. In the *expectation consistent approximate inference* (EC) scheme proposed by Opper and Winther in [46], the density constraint (27b) is relaxed to moment-matching constraints, i.e.,

$$\arg \min_{q_1} \min_{q_2} \max_{q_3} J(q_1, q_2, q_3) \quad (28a)$$

$$\text{s.t. } \mathbb{E}[x|q_1] = \mathbb{E}[x|q_2] = \mathbb{E}[x|q_3] \quad (28b)$$

$$\text{tr}\{\text{Cov}[x|q_1]\} = \text{tr}\{\text{Cov}[x|q_2]\} = \text{tr}\{\text{Cov}[x|q_3]\}, \quad (28c)$$

⁵We minimize over q_1 and q_2 because $D_{\text{KL}}(q_1 \| p_{\mathbf{x}})$ and $D_{\text{KL}}(q_2 \| p_{y|\mathbf{x}})$ are convex, while we maximize over q_3 because $H(q_3)$ is concave.

Algorithm 1: VAMP algorithm [45].

```

1: initialize:
    $r_1^0, \gamma_1^0$ 
2: for  $t = 0, \dots, T_{\text{max}}$  do
3:    $x_1^t = g_1(r_1^t, \gamma_1^t)$ 
4:    $1/\eta_1^t = \langle g_1'(r_1^t, \gamma_1^t) \rangle / \gamma_1^t$ 
5:    $\gamma_2^t = \eta_1^t - \gamma_1^t$ 
6:    $r_2^t = (\eta_1^t x_1^t - \gamma_1^t r_1^t) / \gamma_2^t$ 
7:    $x_2^t = g_2(r_2^t, \gamma_2^t)$ 
8:    $1/\eta_2^t = \langle g_2'(r_2^t, \gamma_2^t) \rangle / \gamma_2^t$ 
9:    $\gamma_1^{t+1} = \eta_2^t - \gamma_2^t$ 
10:   $r_1^{t+1} = (\eta_2^t x_2^t - \gamma_2^t r_2^t) / (\eta_2^t - \gamma_2^t)$ 
11: end for

```

where $\mathbb{E}[x|q^t]$ denotes $\mathbb{E}[x]$ under $x \sim q^t$. This yields stationary points of the form

$$q_1(x) \propto p_{\mathbf{x}}(x) \exp\left(-\frac{\gamma_1}{2} \|x - r_1\|_2^2\right) \quad (29a)$$

$$q_2(x) \propto p_{y|\mathbf{x}}(y|x) \exp\left(-\frac{\gamma_2}{2} \|x - r_2\|_2^2\right) \quad (29b)$$

$$q_3(x) \propto \exp\left(-\frac{\eta}{2} \|x - \hat{x}\|_2^2\right), \quad (29c)$$

for $\{r_1, \gamma_1, r_2, \gamma_2, \hat{x}, \eta\}$ that lead to satisfaction of (28b)–(28c). Various approaches can be used to solve for $\{r_1, \gamma_1, r_2, \gamma_2, \hat{x}, \eta\}$. One is to alternate the update of $\{(r_1, \gamma_1), (\hat{x}, \eta)\}$ and $\{(r_2, \gamma_2), (\hat{x}, \eta)\}$ such that, at each iteration, the moments of q_3 are consistent with either q_1 or q_2 . This approach is summarized in Algorithm 1 using⁶

$$g_1(r_1, \gamma_1) \triangleq \frac{\int x p_{\mathbf{x}}(x) \mathcal{N}(x; r_1, I/\gamma_1) dx}{\int p_{\mathbf{x}}(x) \mathcal{N}(x; r_1, I/\gamma_1) dx} \quad (30)$$

$$g_2(r_2, \gamma_2) \triangleq \frac{\int x p_{y|\mathbf{x}}(y|x) \mathcal{N}(x; r_2, I/\gamma_2) dx}{\int p_{y|\mathbf{x}}(y|x) \mathcal{N}(x; r_2, I/\gamma_2) dx}, \quad (31)$$

which, under these definitions of g_1 and g_2 , can be recognized as an instance of *expectation propagation* (EP) [53, Section 3.2], [54], [55]. In lines 4 and 8, $g_i'(r_i, \gamma_i) \in \mathbb{R}^N$ denotes the diagonal of the Jacobian matrix of $g_i(\cdot, \gamma_i)$ at r_i , i.e.,

$$g_i'(r_i, \gamma_i) \triangleq \text{diag}\left(\frac{\partial g_i(r_i, \gamma_i)}{\partial r_i}\right), \quad (32)$$

and $\langle x \rangle$ denotes the average coefficient value, i.e., $\langle x \rangle \triangleq \frac{1}{N} \sum_{i=1}^N x_i$ for $x \in \mathbb{R}^N$. Due to the form of $p_{y|\mathbf{x}}$ in (15), it can be shown that

$$g_2(r_2, \gamma_2) = (\gamma_2 I + \gamma_w A^T A)^{-1} (\gamma_2 r_2 + \gamma_w A^T y) \quad (33)$$

$$\langle g_2'(r_2, \gamma_2) \rangle = \gamma_2 \text{tr}\{(\gamma_2 I + \gamma_w A^T A)^{-1}\} / N. \quad (34)$$

Meanwhile, the form of $g_1(\cdot)$ depends on $p_{\mathbf{x}}$ through (30).

Based on the description above, one might wonder whether the EC stationary point $\hat{x} = \mathbb{E}[x|q_1] = \mathbb{E}[x|q_2] = \mathbb{E}[x|q_3]$ is a good approximation of the true conditional mean $\mathbb{E}[x|y]$, and additionally one might question whether Algorithm 1 converges to this \hat{x} . Both of these concerns were resolved in the VAMP paper [45]. In particular, [45] showed that, when

⁶In [47], different g_1 and g_2 were proposed so that the EP algorithm accomplishes joint MAP estimation of x from y , i.e., $\hat{x} = \arg \max_x p(x|y)$.

A is right rotationally invariant and asymptotically large, the per-iteration behavior of Algorithm 1 with $g_2(\cdot)$ from (33) and Lipschitz⁷ $g_1(\cdot)$ is exactly predicted by a scalar state evolution. Furthermore, in the case where $g_1(\cdot)$ is matched to generating p_X from (18) as in (30), and where $g_2(\cdot)$ uses the true AWGN precision $\gamma_w < \infty$ as in (31), the MSE of the fixed point \hat{x} of Algorithm 1 was shown in [45] to match the MMSE predicted by the replica method [56]. This replica prediction is conjectured to be correct [57], in which case the \hat{x} generated by Algorithm 1 under (30) and (33) will be MMSE for infinitely large, right-rotationally invariant A when the state evolution has unique fixed points. Note that, for infinitely large i.i.d. A , the replica prediction has been proven to be correct [58], [59].

In the sequel, we will refer to Algorithm 1 with generic Lipschitz $g_1(\cdot)$ as the VAMP algorithm, noting that it coincides with EP in the special case of Bayesian $g_1(\cdot)$ from (30). VAMP is more general than EP because it can be used with denoisers $g_1(\cdot)$ that have no probabilistic interpretation and still lead to precisely predictable behavior under infinitely large, right-rotationally A [45], [50]. We note that, when VAMP is applied to the MMV model (9a), a separate copy of $\{r_1, \gamma_1, r_2, \gamma_2, \hat{x}, \eta\}$ must be tracked for each column of Y .

B. Background on Expectation Maximization

We now return to the case where Θ is unknown and the goal is to compute its ML estimate, $\hat{\Theta}_{\text{ML}}$. From (10) and (12), we have

$$\hat{\Theta}_{\text{ML}} = \arg \min_{\Theta} -\ln \int p_X(X; \Theta) p_{Y|X}(Y|X; \Theta) dX, \quad (35)$$

but (35) is impractical to optimize directly due to the high-dimensional integral.

Expectation-maximization (EM) [60] is a well known iterative approach to ML that alternates between i) minimizing an upper-bound of the negative log-likelihood and ii) tightening the upper-bound. The EM algorithm is usually written as

$$Q(\Theta; \hat{\Theta}^t) \triangleq -\mathbb{E} [\ln p_{X,Y}(X, Y; \Theta) | Y; \hat{\Theta}^t] \quad (36a)$$

$$\hat{\Theta}^{t+1} = \arg \min_{\Theta} Q(\Theta; \hat{\Theta}^t). \quad (36b)$$

Letting $q^t = p_{X|Y}(\cdot | Y; \hat{\Theta}^t)$, we can write

$$\begin{aligned} Q(\Theta; \hat{\Theta}^t) &= -\mathbb{E} [\ln p_X(X; \Theta) | Y; \hat{\Theta}^t] \\ &\quad - \mathbb{E} [\ln p_{Y|X}(Y|X; \Theta) | Y; \hat{\Theta}^t] \end{aligned} \quad (37a)$$

$$= -\mathbb{E} [\ln p_X(X; \Theta) | q^t] \quad (37b)$$

$$\begin{aligned} &\quad - \mathbb{E} [\ln p_{Y|X}(Y|X; \Theta) | q^t] \\ &= J(q^t, q^t, q^t; \Theta) + \text{const.} \end{aligned} \quad (37c)$$

where J , also known as the Gibbs free energy, is defined as

$$\begin{aligned} J(q_1, q_2, q_3; \Theta) &\triangleq D_{\text{KL}}(q_1 || p_X(\cdot; \Theta)) \\ &\quad + D_{\text{KL}}(q_2 || p_{Y|X}(Y | \cdot; \Theta)) + H(q_3). \end{aligned} \quad (38)$$

⁷While the original VAMP paper [45] focused on separable Lipschitz $g_1(\cdot)$, [50] extended the results to non-separable Lipschitz $g_1(\cdot)$.

Thus, (36) can also be written as [61]

$$q^t = p_{X|Y}(\cdot | Y; \hat{\Theta}^t) \quad (39a)$$

$$\hat{\Theta}^{t+1} = \arg \min_{\Theta} J(q^t, q^t, q^t; \Theta). \quad (39b)$$

Note that $J(q^t, q^t, q^t; \Theta)$ is an upper bound on $-\ln p_Y(Y; \Theta)$ for any q^t since

$$J(q^t, q^t, q^t; \Theta) = -\ln p_Y(Y; \Theta) + D_{\text{KL}}(q^t || p_{Y|X}(Y | \cdot; \Theta)), \quad (40)$$

where $D_{\text{KL}} \geq 0$ by construction. Thus, while the specific choice of q^t in (39a) yields a tight upper bound in that

$$J(q^t, q^t, q^t; \hat{\Theta}^t) = -\ln p_Y(Y; \hat{\Theta}^t), \quad (41)$$

other choices of bounding q^t can also be used in EM [61].

IV. BILINEAR ADAPTIVE VAMP

We now propose an algorithm that approximates the quantities in (10)–(11), i.e., the ML estimate of $\Theta = \{\theta_A, \theta_x, \gamma_w\}$ and the MMSE estimate of X under the statistical model (9). We start by developing Bilinear EM-VAMP and then add “variance auto-tuning” to obtain Bilinear Adaptive VAMP (BAd-VAMP).

A. Bilinear EM-VAMP

From the descriptions of VAMP and EM in Section III, we see that they both minimize the same Gibbs free energy cost $J(q_1, q_2, q_3; \Theta)$ from (38), but w.r.t. different variables; VAMP minimizes J w.r.t. the moment-constrained beliefs $\{q_1, q_2, q_3\}$ for a given Θ , while EM minimizes J w.r.t. Θ for a given $\{q_1, q_2, q_3\}$. As a result, the two approaches can be straightforwardly merged for joint estimation of $\{q_1, q_2, q_3\}$ and Θ . In doing so, the goal is to solve the optimization problem

$$\arg \min_{\Theta} \min_{q_1, q_2, q_3} \max J(q_1, q_2, q_3; \Theta) \quad (42a)$$

$$\text{s.t. } \mathbb{E}[x|q_1] = \mathbb{E}[x|q_2] = \mathbb{E}[x|q_3] \quad (42b)$$

$$\text{tr}\{\text{Cov}[x|q_1]\} = \text{tr}\{\text{Cov}[x|q_2]\} = \text{tr}\{\text{Cov}[x|q_3]\}, \quad (42c)$$

and the proposed methodology is to “interleave” the VAMP and EM algorithms, as specified in Algorithm 2. There, the estimation functions g_1 in lines 3–4 and $g_{2,l}$ in lines 9–10 are defined as

$$\begin{aligned} g_1(r_{1,l}, \gamma_{1,l}; \theta_x) \\ \triangleq \frac{\int x p_X(x; \theta_x) \mathcal{N}(x; r_{1,l}, I/\gamma_{1,l}) dx}{\int p_X(x; \theta_x) \mathcal{N}(x; r_{1,l}, I/\gamma_{1,l}) dx} \end{aligned} \quad (43)$$

$$\begin{aligned} g_{2,l}(r_{2,l}, \gamma_{2,l}; \theta_A, \gamma_w) \\ \triangleq \frac{\int x p_{Y|X}(y_l | x; \theta_A, \gamma_w) \mathcal{N}(x; r_{2,l}, I/\gamma_{2,l}) dx}{\int p_{Y|X}(y_l | x; \theta_A, \gamma_w) \mathcal{N}(x; r_{2,l}, I/\gamma_{2,l}) dx}, \end{aligned} \quad (44)$$

The other lines in Algorithm 2 will be detailed in Section IV-C.

B. Bilinear Adaptive VAMP

The VAMP state-evolution from [45, Eq. (34), (35)] shows that when i) $A(\theta_A^t)$ is infinitely large and right-rotationally invariant and ii) the estimation functions g_1 and $g_{2,l}$ are “matched”

Algorithm 2: Bilinear EM-VAMP.

```

1: initialize:
2:    $\forall l : r_{1,l}^0, \gamma_{1,l}^0, \theta_x^0, \theta_A^0, \gamma_w^0$ 
3:   for  $t = 0, \dots, T_{\max}$  do
4:      $\forall l : x_{1,l}^t = g_1(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t)$ 
5:      $\forall l : 1/\eta_{1,l}^t = \langle g_1'(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t) \rangle / \gamma_{1,l}^t$ 
6:      $q_1^t(\mathbf{X}) \propto \prod_{l=1}^L p_{\mathbf{X}}(x_l; \theta_x^t) e^{-\frac{1}{2} \gamma_{1,l}^t \|x_l - r_{1,l}^t\|^2}$ 
7:      $\theta_x^{t+1} = \arg \max_{\theta_x} \mathbb{E} [\ln p_{\mathbf{X}}(\mathbf{X}; \theta_x) | q_1^t]$ 
8:      $\forall l : \gamma_{2,l}^t = \eta_{1,l}^t - \gamma_{1,l}^t$ 
9:      $\forall l : r_{2,l}^t = (\eta_{1,l}^t x_{1,l}^t - \gamma_{1,l}^t r_{1,l}^t) / \gamma_{2,l}^t$ 
10:     $\forall l : x_{2,l}^t = g_{2,l}(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t)$ 
11:     $\forall l : 1/\eta_{2,l}^t = \langle g_{2,l}'(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t) \rangle / \gamma_{2,l}^t$ 
12:     $q_2^t(\mathbf{X}) \propto \prod_{l=1}^L p_{\mathbf{Y}|\mathbf{X}}(y_l | x_l; \theta_A^t, \gamma_w^t) e^{-\frac{1}{2} \gamma_{2,l}^t \|x_l - r_{2,l}^t\|^2}$ 
13:     $\theta_A^{t+1} = \arg \max_{\theta_A} \mathbb{E} [\ln p_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y} | \mathbf{X}; \theta_A, \gamma_w^t) | Y, q_2^t]$ 
14:     $\gamma_w^{t+1} = \arg \max_{\gamma_w} \mathbb{E} [\ln p_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y} | \mathbf{X}; \theta_A^{t+1}, \gamma_w) | Y, q_2^t]$ 
15:     $\forall l : \gamma_{1,l}^{t+1} = \eta_{2,l}^t - \gamma_{2,l}^t$ 
16:     $\forall l : r_{1,l}^{t+1} = (\eta_{2,l}^t x_{2,l}^t - \gamma_{2,l}^t r_{2,l}^t) / \gamma_{1,l}^{t+1}$ 
17:   end for

```

(i.e., MMSE) for the statistical model generating (\mathbf{X}, \mathbf{Y}) , the VAMP quantities $\{(r_{i,l}^t, \gamma_{i,l}^t)\}_{l=1}^L$ for $i = 1, 2$ obey

$$r_{1,l}^t = x_l + \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_{1,l}^t) \quad \forall l \quad (45a)$$

$$x_l = r_{2,l}^t + \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_{2,l}^t) \quad \forall l, \quad (45b)$$

where x_l is the l th column of the true signal realization \mathbf{X} that we aim to recover. That is, $r_{1,l}^t$ is an AWGN-corrupted version of the true signal x_l with *known* AWGN precision $\gamma_{1,l}^t$, and the true signal x_l is an AWGN-corrupted version of $r_{2,l}^t$ with *known* AWGN precision $\gamma_{2,l}^t$. In the context of EM-VAMP under (9), this “matched” condition requires that θ_A^t , θ_x^t , and γ_w^t are all perfect estimates. When θ_A^t , θ_x^t , or γ_w^t are *not* perfect, so that g_1 and $g_{2,l}$ are mismatched, the VAMP state-evolution shows that $r_{1,l}^t$ is still an AWGN corrupted version of x_l , but with an AWGN precision *different* than $\gamma_{1,l}^t$. The impact on EM-VAMP is the following. While the algorithm is trying to learn $\Theta = \{\theta_A, \theta_x, \gamma_w\}$, the value of $\gamma_{1,l}^t$ does *not* correctly characterize the noise precision in $r_{1,l}^t$. As a result, the beliefs q_1^t and q_2^t in lines 5 and 11 of Algorithm 2 become mismatched, which compromises the EM updates of Θ^t .

To remedy this situation, it was proposed in [62] (in the context of EM-GAMP [63]) to explicitly estimate the precision of the AWGN corruption on $r_{1,l}^t$ and $r_{2,l}^t$ and use it in place of the AMP-supplied estimates $\gamma_{1,l}^t$ and $\gamma_{2,l}^t$. This approach was coined “Adaptive” GAMP in [62] and later extended to (linear) Adaptive VAMP in [51].

For Bilinear Adaptive VAMP, the first goal is to replace the estimation of θ_x^t in line 6 of Algorithm 2 with the joint ML estimation

$$(\theta_x^t, \gamma_1^t) = \arg \max_{\theta_x, \gamma_1} p(R_1^t; \gamma_1, \theta_x) \quad (46)$$

under the statistical model

$$r_{1,l}^t = x_l + \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_{1,l}) \quad \forall l, \quad x_l \sim p_{\mathbf{X}}(\cdot; \theta_x) \quad \forall l, \quad (47)$$

with independence across $l = 1, \dots, L$. For this subproblem, we propose to use (inner) EM iterations indexed by τ , i.e.,

$$(\theta_x^{\tau+1}, \gamma_1^{\tau+1}) = \arg \max_{\theta_x, \gamma_1} \mathbb{E} [\ln p(\mathbf{X}, R_1^t; \gamma_1, \theta_x) | R_1^t; \gamma_1^{\tau}, \theta_x^{\tau}] \quad (48)$$

$$= \arg \max_{\theta_x, \gamma_1} \left\{ \mathbb{E} [\ln p(\mathbf{X}; \theta_x) | R_1^t; \gamma_1^{\tau}, \theta_x^{\tau}] + \mathbb{E} [\ln p(R_1^t | \mathbf{X}; \gamma_1) | R_1^t; \gamma_1^{\tau}, \theta_x^{\tau}] \right\}. \quad (49)$$

The previous optimization decouples into

$$\theta_x^{\tau+1} = \arg \max_{\theta_x} \mathbb{E} [\ln p(\mathbf{X}; \theta_x) | R_1^t; \gamma_1^{\tau}, \theta_x^{\tau}] \quad (50)$$

and

$$\gamma_1^{\tau+1} = \arg \max_{\gamma_1} \mathbb{E} [\ln p(R_1^t | \mathbf{X}; \gamma_1) | R_1^t; \gamma_1^{\tau}, \theta_x^{\tau}] \quad (51)$$

$$= \arg \max_{\gamma_1} \sum_{l=1}^L \mathbb{E} [\ln p(r_{1,l}^t | x_l; \gamma_1) | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}], \quad (52)$$

where the latter optimization decouples further into

$$\gamma_1^{\tau+1} = \arg \max_{\gamma_1} \left\{ \frac{N}{2} \ln \gamma_1 + -\frac{\gamma_1}{2} \mathbb{E} [\|x_l - r_{1,l}^t\|_2^2 | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}] \right\} \quad (53)$$

$$= N \{ \mathbb{E} [\|x_l - r_{1,l}^t\|_2^2 | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}] \}^{-1} \quad (54)$$

$$= \left\{ \frac{1}{N} \sum_{n=1}^N \mathbb{E} [(x_{nl} - r_{1,nl}^t)^2 | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}] \right\}^{-1} \quad (55)$$

$$= \left\{ \frac{1}{N} \|x_{1,l}^t - r_{1,l}^t\|^2 + \frac{1}{\eta_{1,l}^{\tau}} \right\}^{-1}, \quad (56)$$

for $l = 1, \dots, L$ and

$$x_{1,l}^{\tau} \triangleq \mathbb{E} [x_l | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}] \quad (57)$$

$$= g_1(r_{1,l}^t, \gamma_1^{\tau}, \theta_x^{\tau}) \quad (58)$$

$$1/\eta_{1,l}^{\tau} \triangleq \text{tr} \{ \text{Cov} [x_l | r_{1,l}^t; \gamma_1^{\tau}, \theta_x^{\tau}] \} / N \quad (59)$$

$$= \langle g_1'(r_{1,l}^t, \gamma_1^{\tau}, \theta_x^{\tau}) \rangle / \gamma_1^{\tau}. \quad (60)$$

Above, we detailed the re-estimation of γ_1^{τ} . A similar procedure can be used for re-estimation of γ_2^{τ} . The resulting Bilinear Adaptive VAMP (BAd-VAMP) is summarized in Algorithm 3 using $\tau_{1,\max}$ EM iterations for the first inner loop and $\tau_{2,\max}$ EM iterations for the second inner loop. To avoid the complications of a dual-index notation (i.e., t and τ), we use only the single index t in Algorithm 3 and over-write the quantities in each inner loop. Note that, when $\tau_{1,\max} = \tau_{2,\max} = 0$, BAd-VAMP (i.e., Algorithm 3) reduces to bilinear EM-VAMP (i.e., Algorithm 2).

C. Algorithm Details

We now provide additional details on the steps in Algorithm 3.

Algorithm 3: Bilinear Adaptive VAMP.

```

1: initialize:
2:    $\forall l : r_{1,l}^0, \gamma_{1,l}^0, \theta_x^0, \theta_A^0, \gamma_w^0$ 
3:   for  $t = 0, \dots, T_{\max}$  do
4:     for  $\tau = 0, \dots, \tau_{1,\max}$  do
5:        $\forall l : x_{1,l}^t \leftarrow g_1(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t)$ 
6:        $\forall l : 1/\eta_{1,l}^t \leftarrow \langle g_1'(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t) \rangle / \gamma_{1,l}^t$ 
7:        $\forall l : 1/\gamma_{1,l}^t \leftarrow \frac{1}{N} \|x_{1,l}^t - r_{1,l}^t\|^2 + 1/\eta_{1,l}^t$ 
8:        $q_1^t(X) \propto \prod_{l=1}^L p_{\mathbf{X}}(x_l; \theta_x^t) e^{-\frac{1}{2} \gamma_{1,l}^t \|x_l - r_{1,l}^t\|^2}$ 
9:        $\theta_x^t \leftarrow \arg \max_{\theta_x} \mathbb{E}[\ln p_{\mathbf{X}}(X; \theta_x) | q_1^t]$ 
10:      end for
11:       $\theta_x^{t+1} = \theta_x^t$ 
12:       $\forall l : \gamma_{2,l}^t = \eta_{1,l}^t - \gamma_{1,l}^t$ 
13:       $\forall l : r_{2,l}^t = (\eta_{1,l}^t x_{1,l}^t - \gamma_{1,l}^t r_{1,l}^t) / \gamma_{2,l}^t$ 
14:      for  $\tau = 0, \dots, \tau_{2,\max}$  do
15:         $\forall l : x_{2,l}^t \leftarrow g_{2,l}(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t)$ 
16:         $\forall l : 1/\eta_{2,l}^t \leftarrow \langle g_{2,l}'(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t) \rangle / \gamma_{2,l}^t$ 
17:         $\forall l : 1/\gamma_{2,l}^t \leftarrow \frac{1}{N} \|x_{2,l}^t - r_{2,l}^t\|^2 + 1/\eta_{2,l}^t$ 
18:         $q_2^t(X) \propto \prod_l p_{\mathbf{Y}|\mathbf{X}}(y_l | x_l; \theta_A^t, \gamma_w^t) e^{-\frac{1}{2} \gamma_{2,l}^t \|x_l - r_{2,l}^t\|^2}$ 
19:         $\theta_A^t \leftarrow \arg \max_{\theta_A} \mathbb{E}[\ln p_{\mathbf{Y}|\mathbf{X}}(Y | X; \theta_A, \gamma_w) | Y, q_2^t]$ 
20:         $\gamma_w^t \leftarrow \arg \max_{\gamma_w} \mathbb{E}[\ln p_{\mathbf{Y}|\mathbf{X}}(Y | X; \theta_A^t, \gamma_w) | Y, q_2^t]$ 
21:      end for
22:       $\theta_A^{t+1} = \theta_A^t$ 
23:       $\gamma_w^{t+1} = \gamma_w^t$ 
24:       $\forall l : \gamma_{1,l}^{t+1} = \eta_{2,l}^t - \gamma_{2,l}^t$ 
25:       $\forall l : r_{1,l}^{t+1} = (\eta_{2,l}^t x_{2,l}^t - \gamma_{2,l}^t r_{2,l}^t) / \gamma_{1,l}^{t+1}$ 
end for

```

1) *Estimating X :* Recalling the definition of $g_{2,l}(\cdot)$ in (44), the form of $p_{\mathbf{Y}|\mathbf{X}}$ in (15) implies that

$$g_{2,l}(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t) = C_l^t (\gamma_{2,l}^t r_{2,l}^t + \gamma_w^t A(\theta_A^t)^T y_l) \quad (61)$$

$$\langle g_{2,l}'(r_{2,l}^t, \gamma_{2,l}^t; \theta_A^t, \gamma_w^t) \rangle = \gamma_{2,l}^t \text{tr}\{C_l^t\} / N \quad (62)$$

for

$$C_l^t \triangleq (\gamma_{2,l}^t I_N + \gamma_w^t A(\theta_A^t)^T A(\theta_A^t))^{-1}. \quad (63)$$

To avoid computing a separate matrix inverse (63) for each $l = 1, \dots, L$, one could instead compute the eigenvalue decomposition

$$A(\theta_A^t)^T A(\theta_A^t) = U^t \text{Diag}(s^t) U^{t\top}, \quad (64)$$

and then leverage the fact that

$$C_l^t = U^t \text{Diag}(\gamma_{2,l}^t \mathbf{1} + \gamma_w^t s^t)^{-1} U^{t\top}, \quad (65)$$

which reduces to the inversion of a diagonal matrix for each $l = 1, \dots, L$.

2) *Learning θ_A :* We now provide details on the update of θ_A and γ_w in lines 18-19 of Algorithm 3. Given the form of $p_{\mathbf{Y}|\mathbf{X}}$ in (14)-(15), we have that

$$\begin{aligned} \ln p_{\mathbf{Y}|\mathbf{X}}(Y | X; \theta_A, \gamma_w) \\ = \frac{ML}{2} \ln \gamma_w - \frac{\gamma_w}{2} \|Y - A(\theta_A)X\|_F^2 + \text{const} \end{aligned} \quad (66)$$

$$\begin{aligned} &= \frac{ML}{2} \ln \gamma_w - \frac{\gamma_w}{2} (\text{tr}\{YY^T\} - 2\text{tr}\{A(\theta_A)XY^T\} \\ &\quad + \text{tr}\{A(\theta_A)XX^T A(\theta_A)^T\}) + \text{const}. \end{aligned} \quad (67)$$

Since

$$\mathbb{E}[X | q_2^t] = X_2^t \quad (68)$$

$$\begin{aligned} \mathbb{E}[XX^T | q_2^t] &= \sum_{l=1}^L \mathbb{E}[x_l x_l^T | q_{2,l}^t] = X_2^t X_2^{t\top} + \sum_{l=1}^L C_l^t, \\ &\triangleq C^t \end{aligned} \quad (69)$$

we have that

$$\begin{aligned} \mathbb{E}[\ln p_{\mathbf{Y}|\mathbf{X}}(Y | X; \theta_A, \gamma_w) | Y, q_2^t] \\ = \frac{ML}{2} \ln \gamma_w - \frac{\gamma_w}{2} (\text{tr}\{YY^T\} - 2\text{tr}\{A(\theta_A)X_2^t Y^T\} \\ + \text{tr}\{A(\theta_A)X_2^t X_2^{t\top} A(\theta_A)^T\} + \text{tr}\{A(\theta_A)C^t A(\theta_A)^T\}) \\ = \frac{ML}{2} \ln \gamma_w - \frac{\gamma_w}{2} (\|Y - A(\theta_A)X_2^t\|_F^2 \\ + \text{tr}\{A(\theta_A)C^t A(\theta_A)^T\}) + \text{const}. \end{aligned} \quad (70)$$

To maximize (71) over $\theta_A = [\theta_{A,1}, \dots, \theta_{A,Q}]$ with fixed γ_w , we consider the affine-linear model

$$A(\theta_A) = A_0 + \sum_{i=1}^Q \theta_{A,i} A_i, \quad (72)$$

noting that non-linear models could be handled using similar techniques. Plugging (72) into (70), we get

$$\begin{aligned} \mathbb{E}[\ln p_{\mathbf{Y}|\mathbf{X}}(Y | X; \theta_A, \gamma_w) | Y, q_2^t] \\ = \text{const} - \frac{\gamma_w}{2} \sum_{i=1}^Q \sum_{j=1}^Q \theta_{A,i} \text{tr}\{A_i(C^t + X_2^t X_2^{t\top}) A_j^T\} \theta_{A,j} \\ - \gamma_w \sum_{i=1}^Q \theta_{A,i} (\text{tr}\{A_i(C^t + X_2^t X_2^{t\top}) A_0^T\} - \text{tr}\{A_i X_2^t Y^T\}) \\ = -\frac{\gamma_w}{2} (\theta_A^T H^t \theta_A - 2\theta_A^T \beta^t) + \text{const} \end{aligned} \quad (73)$$

for

$$[H^t]_{ij} = \text{tr}\{A_i(C^t + X_2^t X_2^{t\top}) A_j^T\} \quad (75)$$

$$= \text{tr}\{A_j^T A_i(C^t + X_2^t X_2^{t\top})\} \quad (76)$$

and

$$[\beta^t]_i = \text{tr}\{A_i X_2^t Y^T\} - \text{tr}\{A_i(C^t + X_2^t X_2^{t\top}) A_0^T\} \quad (77)$$

$$= \text{tr}\{Y^T A_i X_2^t\} - \text{tr}\{A_0^T A_i(C^t + X_2^t X_2^{t\top})\}, \quad (78)$$

where $A_j^\top A_i$ and $Y^\top A_i$ can be pre-computed. Zeroing the gradient of (74) w.r.t. θ_A , we find that the maximizer is

$$\theta_A^{t+1} = (H^t)^{-1} \beta^t. \quad (79)$$

A special case of (72) is where $A(\cdot)$ has no structure, i.e.,

$$A(\theta_A) = \sum_{m=1}^M \sum_{n=1}^N \theta_{A,m,n} e_m e_n^\top. \quad (80)$$

where e_m denotes the m th standard basis vector. In this case, it can be shown that

$$A(\theta_A^{t+1}) = Y X_2^{t\top} (C^t + X_2^t X_2^{t\top})^{-1}. \quad (81)$$

3) *Learning γ_w :* To maximize (71) over γ_w with fixed $\theta_A = \theta_A^{t+1}$, we search for the values of γ_w that zero the derivative of (71). The unique solution is straightforwardly shown to be

$$\begin{aligned} 1/\gamma_w^{t+1} = \frac{1}{ML} & \left(\|Y - A(\theta_A^{t+1}) X_2^t\|_F^2 \right. \\ & \left. + \text{tr}\{A(\theta_A^{t+1}) C^t A(\theta_A^{t+1})^\top\} \right). \end{aligned} \quad (82)$$

4) *Summary:* In Algorithm 4, BAd-VAMP is rewritten with detailed expressions for the updates of $x_{2,l}^t$, $\eta_{2,l}^t$, θ_A^t , and γ_w^t .

D. Algorithm Enhancements

We now propose several enhancements to the BAd-VAMP algorithm presented in Algorithm 3 and detailed in Algorithm 4.

1) *Damping:* For fixed Θ^t and infinitely large right-rotationally invariant $A(\theta_A^t)$, the state-evolution of VAMP guarantees its convergence. But when $A(\theta_A^t)$ deviates from this assumption, damping the VAMP iterations can help maintain convergence [45]. With damping, lines 25-26 of Algorithm 4 (or lines 23-24 of Algorithm 3) would be replaced by

$$\gamma_{1,l}^{t+1} = (1 - \zeta) \gamma_{1,l}^t + \zeta (\eta_{2,l}^t - \gamma_{2,l}^t) \quad (83)$$

$$\gamma_{2,l}^{t+1} = (1 - \zeta) r_{1,l}^t + \zeta (\eta_{2,l}^t x_{2,l}^t - \gamma_{2,l}^t r_{2,l}^t) / (\eta_{2,l}^t - \gamma_{2,l}^t) \quad (84)$$

for some $\zeta \in (0, 1)$. The case $\zeta = 1$ corresponds to no damping.

2) *Negative Precisions:* Sometimes the precisions $\{\gamma_{1,l}, \gamma_{2,l}\}_l$ can be negative. We suggest to restrict the precisions to the interval $[\gamma_{\min}, \infty)$, for very small $\gamma_{\min} > 0$, in lines 11 and 25 of Algorithm 4 (or lines 11 and 23 of Algorithm 3).

3) *Restarts:* Due to the non-convex nature of the bilinear inference problem, the algorithm may get stuck at local minima or slowed by saddle points. To mitigate these issues, it sometimes helps to restart the algorithm. For each restart, we suggest to initialize θ_A^0 at the final estimate of θ_A returned by the previous run.

E. Relation to Previous Work

The proposed Bilinear Adaptive VAMP algorithm extends the (linear) Adaptive VAMP algorithm of [51] from the case where $A(\theta_A)$ is known to the case where $A(\theta_A)$ is unknown. In the known- $A(\theta_A)$ setting, where $A(\theta_A)$ is infinitely large and right-rotationally invariant, it was rigorously established in [51] that Adaptive VAMP obeys a state-evolution similar to that

Algorithm 4: Bilinear Adaptive VAMP (Detailed).

```

1: initialize:
2:    $\forall l : r_{1,l}^0, \gamma_{1,l}^0, \theta_x^0, \theta_A^0, \gamma_w^0$ 
3:   for  $t = 0, \dots, T_{\max}$  do
4:     for  $\tau = 0, \dots, \tau_{1,\max}$  do
5:        $\forall l : x_{1,l}^t \leftarrow g_1(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t)$ 
6:        $\forall l : 1/\eta_{1,l}^t \leftarrow \langle g_1(r_{1,l}^t, \gamma_{1,l}^t; \theta_x^t) \rangle / \gamma_{1,l}^t$ 
7:        $\forall l : 1/\gamma_{1,l}^t \leftarrow \frac{1}{N} \|x_{1,l}^t - r_{1,l}^t\|^2 + 1/\eta_{1,l}^t$ 
8:        $q_1^t(X) \propto \prod_{l=1}^L p_X(x_l; \theta_x^t) e^{-\frac{1}{2} \gamma_{1,l}^t \|x_l - r_{1,l}^t\|^2}$ 
9:        $\theta_A^t \leftarrow \arg \max_{\theta_A} \mathbb{E} [\ln p_X(X; \theta_x) | q_1^t]$ 
10:      end for
11:       $\theta_x^{t+1} = \theta_x^t$ 
12:       $\forall l : \gamma_{2,l}^t = \eta_{1,l}^t - \gamma_{1,l}^t$ 
13:       $\forall l : r_{2,l}^t = (\eta_{1,l}^t x_{1,l}^t - \gamma_{1,l}^t r_{1,l}^t) / \gamma_{2,l}^t$ 
14:      for  $\tau = 0, \dots, \tau_{2,\max}$  do
15:         $\forall l : C_l^t \leftarrow (\gamma_{2,l}^t I_N + \gamma_w^t A(\theta_A^t)^\top A(\theta_A^t))^{-1}$ 
16:         $\forall l : x_{2,l}^t \leftarrow C_l^t (\gamma_{2,l}^t r_{2,l}^t + \gamma_w^t A(\theta_A^t)^\top y_l)$ 
17:         $\forall l : 1/\eta_{2,l}^t \leftarrow \text{tr}\{C_l^t\} / N$ 
18:         $C^t \leftarrow \sum_{l=1}^L C_l^t$ 
19:         $\forall i, j : [H^t]_{ij} \leftarrow \text{tr}\{A_j^\top A_i (C^t + X_2^t X_2^{t\top})\}$ 
20:         $\forall i : [\beta^t]_i \leftarrow \text{tr}\{Y^\top A_i X_2^t\}$ 
21:         $\quad - \text{tr}\{A_0^\top A_i (C^t + X_2^t X_2^{t\top})\}$ 
22:         $\theta_A^t \leftarrow (H^t)^{-1} \beta^t$ 
23:         $1/\gamma_w^t \leftarrow \frac{1}{ML} (\|Y - A(\theta^t) X_2^t\|_F^2$ 
24:         $\quad + \text{tr}\{A(\theta_A^t) C^t A(\theta_A^t)^\top\})$ 
25:      end for
26:       $\theta_A^{t+1} = \theta_A^t$ 
27:       $\gamma_w^{t+1} = \gamma_w^t$ 
28:       $\forall l : \gamma_{1,l}^{t+1} = \eta_{2,l}^t - \gamma_{2,l}^t$ 
29:       $\forall l : r_{1,l}^{t+1} = (\eta_{2,l}^t x_{2,l}^t - \gamma_{2,l}^t r_{2,l}^t) / \gamma_{1,l}^{t+1}$ 
30:    end for

```

of VAMP, and that its estimates of $\{\theta_x, \gamma_w\}$ are asymptotically consistent under certain identifiability conditions, i.e., they converge to the true values as $t \rightarrow \infty$. As future work, it would be interesting to understand whether Bilinear Adaptive VAMP also obeys a state evolution for certain classes of $A(\cdot)$.

The proposed BAd-VAMP algorithm targets the same class⁸ of bilinear recovery problems as the EM-PBiGAMP algorithm from [33], and both leverage EM for automated hyperparameter tuning. However, the “AMP” aspects of these algorithms are fundamentally different. PBiGAMP treats the vectors $\{a_{i,j}\}$ in (2) as i.i.d. Gaussian for its derivation, whereas BAd-VAMP treats the matrix $A(b) = \sum_{i=1}^Q b_i A_i$ as right rotationally-invariant for its derivation. The latter allows more freedom in the singular values of $A(b)$, which leads to increased robustness in practice, as demonstrated by the numerical experiments in Section V.

BAd-VAMP and lifted VAMP both leverage the VAMP approach from [45] to solve bilinear inference problems. However,

⁸Note that the BiGAMP algorithm [19]–[20] is a special case of the PBiGAMP algorithm [33]. BiGAMP applies to the recovery of A and X from measurements of the form $Y = AX + W$, whereas PBiGAMP applies to the recovery of b and X from $Y = A(b)X + W$ under known affine linear $A(\cdot)$. Both can be combined with EM for hyperparameter learning.

they do so in very different ways. As discussed in Section I, lifted VAMP “lifts” the bilinear problem into a higher-dimensional linear problem, and then uses non-separable denoising to jointly estimate b and c in (5). An unfortunate consequence of lifting is a possibly significant increase in computational complexity and memory. In contrast, BAd-VAMP avoids lifting, and it employs EM to estimate b and VAMP to estimate c . Interestingly, Section V shows BAd-VAMP performing equal or better to lifted VAMP in all experiments.

V. NUMERICAL SIMULATIONS

In this section, we present the results of numerical simulations that study the behavior of the BAd-VAMP algorithm from Algorithm 4, in comparison to other state-of-the-art algorithms, on several bilinear recovery problems. In all cases, we ran BAd-VAMP with $\tau_{1,\max} = 1$ and $\tau_{2,\max} = 0$ inner EM iterations and we assumed that the signal prior p_x is fully known (i.e., θ_x is known). We nominally used a damping coefficient of $\zeta = 0.8$ and minimum precision of $\gamma_{\min} = 10^{-6}$. We initialized BAd-VAMP by $\gamma_w^0 = 0.1$, $\gamma_{1,l}^0 = 10^{-3} \forall l$, and we set $r_{1,l}^0$ and θ_A^0 to random vectors drawn i.i.d. from $\mathcal{N}(0, 10)$ and $\mathcal{N}(0, 1)$ respectively. Unless otherwise noted, no restarts were used.

A. CS With Matrix Uncertainty

In compressive sensing (CS) with matrix uncertainty [6], the goal is to recover the K -sparse signal $c \in \mathbb{R}^N$ and the uncertainty parameters b from measurements $y = A(b)c + w$ with $w \sim \mathcal{N}(0, I/\gamma_w)$. Here, $A(b) = A_0 + \sum_{i=1}^Q b_i A_i$, where $\{A_i\}_{i=0}^Q$ are known. For our experiments, we used $Q = 10$, $N = 256$, and $K = 10$, and we selected γ_w so that SNR $\triangleq \mathbb{E}[\|Ac\|^2]/\mathbb{E}\|w\|^2 = 40$ dB. Also, the uncertainty parameters b were drawn $\mathcal{N}(0, I)$, and c was drawn with uniformly random support and with K non-zero elements from $\mathcal{N}(0, I)$. We measured performance using $\text{NMSE}(\hat{b}) \triangleq \|\hat{b} - b\|^2/\|b\|^2$ and $\text{NMSE}(\hat{c}) \triangleq \|\hat{c} - c\|^2/\|c\|^2$. As a reference, we considered two oracle estimators: the MMSE estimator for b assuming c is known, and the MMSE estimator for c assuming b and the support of c are known.

For our first experiment, the elements of $\{A_i\}_{i=1}^Q$ were drawn i.i.d. $\mathcal{N}(0, 1)$ and the elements of A_0 were drawn $\mathcal{N}(0, 20)$. BAd-VAMP was run for a maximum of 200 iterations with a maximum of 2 restarts and damping $\zeta = 0.86$.

Figure 1 shows that the AMP-based algorithms gave near-oracle performance for the tested range of M/N , although lifted VAMP performed slightly worse than the others when $M/N = 0.2$. In contrast, the performance of WSS-TLS from the award-winning paper [6] was significantly worse than the AMP approaches. WSS-TLS aims to solve the non-convex optimization problem

$$(\hat{b}, \hat{c}) = \arg \min_{b, c} \left\| \left(A_0 + \sum_{i=1}^Q b_i A_i \right) c - y \right\|^2 + \|b\|^2/\gamma_w + \lambda \|c\|_1 \quad (85)$$

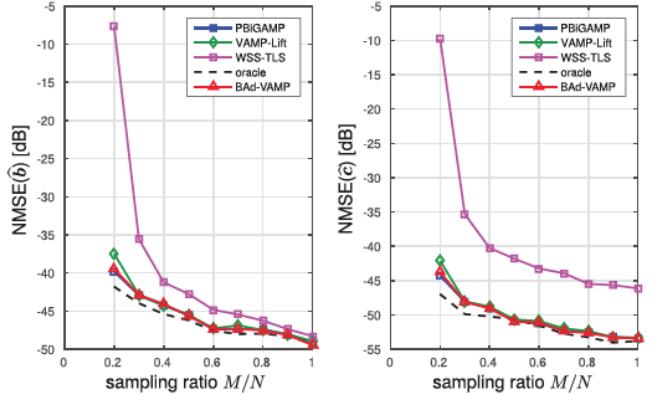


Fig. 1. CS with matrix uncertainty: Median NMSE (over 50 trials) on signal c and uncertainty parameters b versus sampling ratio M/N .

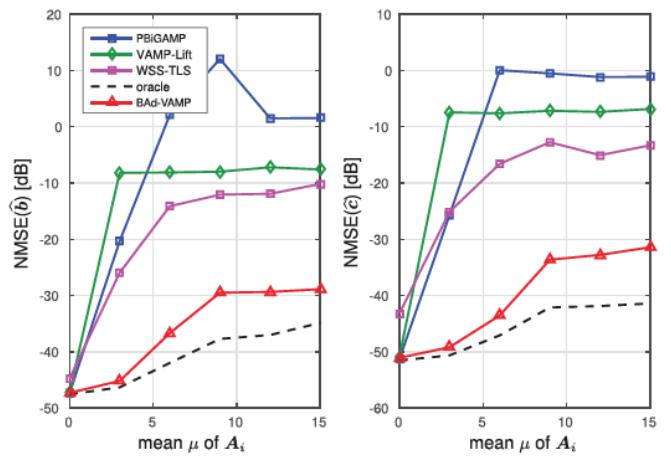


Fig. 2. CS with matrix uncertainty: Median NMSE (over 50 trials) on signal c and uncertainty parameters b versus mean of matrices A_i at $M/N = 0.6$.

using alternating minimization. For WSS-TLS, we used oracle knowledge of γ_w , oracle tuning of the regularization parameter λ , and code from the authors’ website.

For our second experiment, we tested algorithm robustness to non-zero mean in $A(b)$,⁹ since this is a known issue with many AMP algorithms [42]–[44]. For this, we fixed the sampling ratio at $M/N = 0.6$, drew the elements of $\{A_i\}_{i=1}^Q$ from i.i.d. $\mathcal{N}(\mu, 1)$, and drew the elements of A_0 from i.i.d. $\mathcal{N}(\mu, 20)$. Figure 2 reports the median NMSE versus mean μ , and shows that BAd-VAMP is much more robust to $\mu > 0$ than the other tested AMP algorithms as well as WSS-TLS.

Figure 3 shows the runtime of the algorithms in Figure 1. Our implementation used MATLAB (R2015b) on an RHEL workstation with an 8-core Intel i7 processor. Although, for WSS-TLS, we used a grid-search to optimize λ in (85), Figure 3 only shows the runtime of WSS-TLS after λ was chosen. Figure 3 shows BAd-VAMP running much faster than lifted VAMP and WSS-TLS, and slightly slower than PBiGAMP.

⁹For the simpler case where b is known and the objective is to recover c from $y = Ac + w$, modifications of AMP that temporarily remove the mean from A have been proposed [44]. However, it is not clear how to extend this approach to the bilinear problem of recovering b and c from $y = A(b)c + w$.

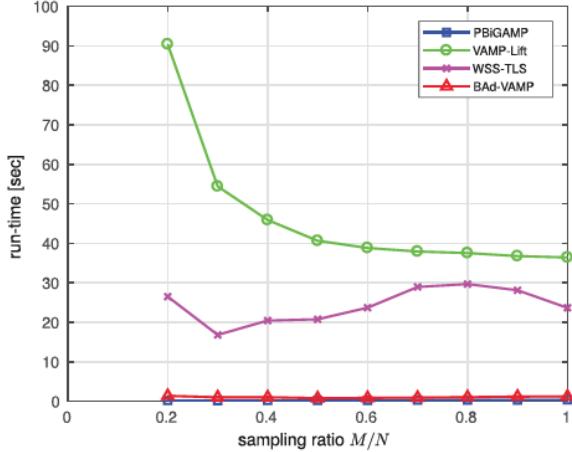


Fig. 3. CS with matrix uncertainty: Median run-time in seconds (over 10 trials) versus sampling ratio M/N .

B. Self-Calibration

In self-calibration [5], the goal is to recover the K -sparse signal vector c and the calibration parameters b from measurements of the form $y = \text{Diag}(Hb)\Psi c$ with known $H \in \mathbb{R}^{M \times Q}$ and $\Psi \in \mathbb{R}^{M \times N}$. Here, Hb represents an unknown vector of gains on the measurements, where the gain vector is believed to lie in the Q -dimensional subspace spanned by the columns of H . For our experiment, $M = 128$, $N = 256$, Ψ and b were drawn i.i.d. $\mathcal{N}(0, 1)$, H was constructed using Q randomly selected columns of the Hadamard matrix, and c was drawn with uniformly random support and with K non-zero elements from $\mathcal{N}(0, I)$.

Figure 4 shows the rate of successful recovery versus subspace dimension Q and sparsity K for several algorithms. A recovery (\hat{b}, \hat{c}) was considered “successful” when $\|\hat{b}\hat{c}^\top - bc^\top\|_F^2/\|bc^\top\|_F^2 \leq -50$ dB. From the figure, we see that the performance of BAd-VAMP is similar to that of EM-PBiGAMP, and even slightly better when Q is small and K is large. Meanwhile, BAd-VAMP appears significantly better than both lifted VAMP and SparseLift from [5]. SparseLift is a convex relaxation with provable guarantees [5]. For computational reasons (recall the discussion in Section I-B), it was difficult to simulate lifted VAMP for $Q \geq 10$.

C. Calibration in Tomography

We consider the problem of reconstructing an image from a sequence of tomographic projections, where the projections along each direction are scaled by an unknown calibration gain. In particular, let Ψ_ω be the tomographic projection matrix¹⁰ (i.e., Radon transform) corresponding to angle $\omega \in [0, \pi]$. Our goal is to reconstruct the image x from measurements

$$y = \begin{bmatrix} b_1 \Psi_{\omega_1} \\ \vdots \\ b_K \Psi_{\omega_K} \end{bmatrix} x + w \text{ for } w \sim \mathcal{N}(0, \gamma_w^{-1} I), \quad (86)$$

¹⁰We used the matrix form of the Radon transform instead of the operator form to avoid numerical error when implementing the adjoint.

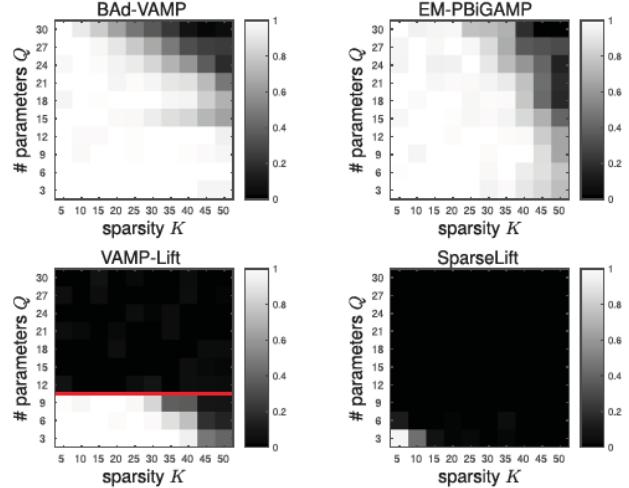


Fig. 4. Self-calibration: Empirical success rate (over 50 trials) for several algorithms versus number of calibration parameters Q and sparsity K . For computational reasons, VAMP-Lift was simulated only for $Q < 10$.

where $b_k \sim \mathcal{N}(1, \sigma_b^2)$ are unknown. Note that, by defining $A_k = [0, \dots, \Psi_{\omega_k}^\top, 0]^\top$, we can write $y = A(b)x + \mathcal{N}(0, \gamma_w^{-1} I)$ for $A(b) = \sum_{k=1}^K b_k A_k$, which matches (8).

In an attempt to solve the above problem, we used BAd-VAMP to recover the image x while simultaneously learning the calibration gains b . For this, we used BAd-VAMP in “plug-and-play” mode, where the BM3D image denoiser [64] was used to implement the $g_1(\cdot)$ function in Algorithm 1. Due to the inherent scaling ambiguity of the problem (i.e., if (\hat{x}, \hat{b}) is a solution then so is $(\alpha\hat{x}, \alpha^{-1}\hat{b})$ for any $\alpha > 0$), we scaled the image estimate \hat{x} by the α that minimized $\|x - \alpha\hat{x}\|$ before computing the PSNR.

As baselines, we also tested the VAMP [45], total variation (TV) [65], [66] and regularization-by-denoising (RED) [67], [68] approaches (see descriptions below). These approaches all assume a noisy *linear* data model of the form $y = \hat{A}x + \mathcal{N}(0, \gamma_w^{-1} I)$ with known \hat{A} . To apply them to (86), we considered two cases: the *genie-calibrated* (GC) case, where a genie supplies the true gains b and the algorithm uses $\hat{A} = A(b)$, and the *un-calibrated* (UC) case, where b is unknown and the algorithms assume $\hat{A} = A(1)$. In the latter case, the \hat{A} -based model is mismatched to the data-generation model (86). For fair comparison, we scaled the GC and UC image estimates \hat{x} by the α that minimized $\|x - \alpha\hat{x}\|_2$ before computing the PSNR.

We now provide additional details on the experimental setup. For x , we used the modified Shepp-Logan phantom of size 64×64 , shown in the top-left panel of Fig. 5. For $A(\cdot)$, we used $K = 25$ projections spaced uniformly in $\omega \in [0, \pi]$. The calibration gains b were generated using $\sigma_b = 0.06$, and the noise precision γ_w was set to achieve an SNR of $\mathbb{E}[\|A(b)x\|^2]/\mathbb{E}[\|w\|^2] = 40$ dB. The TV method [65] computes

$$\hat{x} = \arg \min_x \left\{ \frac{1}{2} \|y - \hat{A}x\|_2^2 + \lambda_t \|\nabla x\|_{2,1} \right\} \quad (87)$$

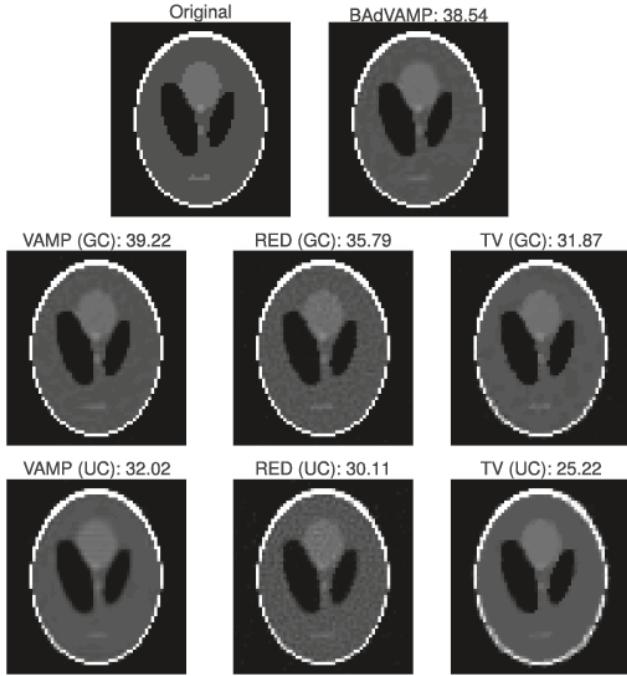


Fig. 5. Calibration in tomography: Reconstruction PSNR (dB) of 64×64 Shepp-Logan phantom from 25 equally spaced tomographic projections. In the genie-calibrated (GC) case, $\hat{\mathbf{A}} = \mathbf{A}(b)$, while in the un-calibrated (UC) case, $\hat{\mathbf{A}} = \mathbf{A}(1)$.

TABLE I
PSNR (dB) IN THE TOMOGRAPHY EXPERIMENT

measurements	BAd-VAMP	VAMP	RED	TV
genie calibrated (GC)	—	39.57	36.56	33.27
un-calibrated (UC)	38.27	31.62	31.24	26.79

for the isotropic TV operator

$$\|\nabla x\|_{2,1} = \sum_{i,j} \sqrt{(x_{i,j} - x_{i,j-1})^2 + (x_{i,j} - x_{i-1,j})^2}. \quad (88)$$

We solved (87) using FASTA [69], and tuned λ_t to maximize the PSNR. RED [67], [68] solves the fixed-point equation

$$\hat{\mathbf{A}}^H(\hat{\mathbf{A}}\hat{x} - \mathbf{y}) + \lambda_r(\hat{x} - \rho(\hat{x}, \tau)) = \mathbf{0} \quad (89)$$

for \hat{x} , where $\rho(\cdot, \tau)$ is an image denoising algorithm with noise-variance τ . For our experiment, we used BM3D for $\rho(\cdot, \tau)$, solved (89) using the ADMM method from [67] with 200 iterations, and tuned both λ_r and τ to maximize PSNR. For BAd-VAMP, we initialized b to 1, used damping $\zeta = 0.1$, assumed known noise precision γ_w , and used at most 100 iterations. For VAMP, we initialized $\gamma_1^0 = 10^{-4}$ and used at most 100 iterations.

Table I reports the median PSNR achieved by each algorithm across 10 random draws of b and w . Fig. 5 shows example image recoveries, and Fig. 6 shows the corresponding error images. From Table I, we see that the PSNR performance of BAd-VAMP (which does not know b) is nearly as good as genie-calibrated VAMP, and 1.7 dB better than genie-calibrated RED. Furthermore, the PSNR performance of BAd-VAMP is

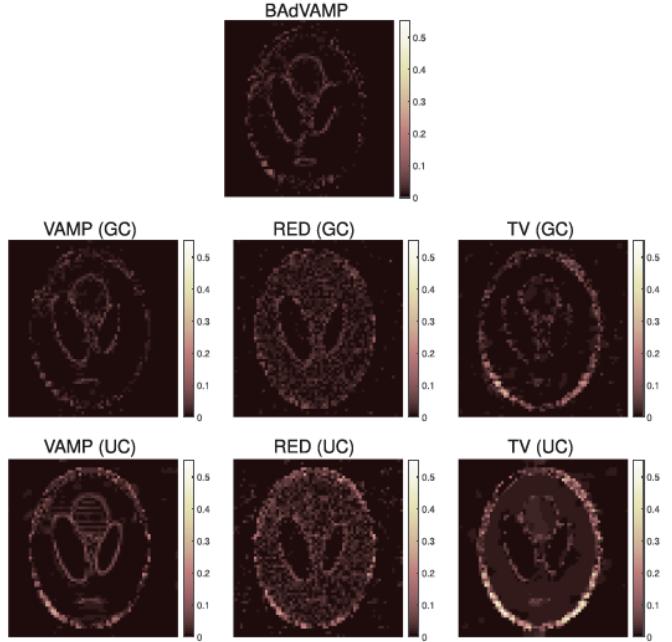


Fig. 6. Calibration in tomography: Error images for the reconstructions shown in Fig. 5.

more than 6.6 dB better than un-calibrated VAMP and 7 dB better than un-calibrated RED. The uncalibrated VAMP, RED, and TV recoveries in Fig. 5 are plagued by either streaking artifacts and/or loss of detail (e.g., note the disappearance of the small white dots in uncalibrated TV). But the BAd-VAMP image recovery in Fig. 5 shows no streaking artifacts and a high level of detail. Likewise, Fig. 6 shows that TV has trouble correctly recovering the white outer ellipse, RED has trouble in the interior region, but BAd-VAMP does well throughout.

D. Noiseless Dictionary Learning

In dictionary learning (DL) [3], the goal is to find a dictionary matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ and a sparse matrix $\mathbf{X} \in \mathbb{R}^{N \times L}$ such that a given matrix $\mathbf{Y} \in \mathbb{R}^{M \times L}$ can be approximately factored as $\mathbf{Y} \approx \mathbf{AX}$. In this section, we test the proposed BAd-VAMP algorithm for DL by generating $\mathbf{Y} = \mathbf{AX}$ such that \mathbf{X} has K -sparse columns, and measuring the NMSE on the resulting estimates of \mathbf{A} and \mathbf{X} .

We consider two cases: i) where the true \mathbf{A} is structured as $\mathbf{A} = \sum_{i=1}^Q b_i \mathbf{A}_i$ with known $\{\mathbf{A}_i\}_{i=1}^Q$ (recall (72)), and ii) where the true \mathbf{A} is unstructured (recall (80)). In either case, the pair (\mathbf{A}, \mathbf{X}) is recoverable only up to an ambiguity: a scalar ambiguity in the structured case and a generalized permutation ambiguity in the unstructured case. Thus, when measuring reconstruction quality, we consider

$$\text{NMSE}(\hat{\mathbf{A}}) \triangleq \min_{\lambda \in \mathbb{R}} \frac{\|\mathbf{A} - \lambda \hat{\mathbf{A}}\|_F^2}{\|\mathbf{A}\|_F^2} \quad (90)$$

in the structured case and

$$\text{NMSE}(\hat{\mathbf{A}}) \triangleq \min_{\mathbf{P} \in \mathcal{P}} \frac{\|\mathbf{A} - \hat{\mathbf{A}}\mathbf{P}\|_F^2}{\|\mathbf{A}\|_F^2} \quad (91)$$

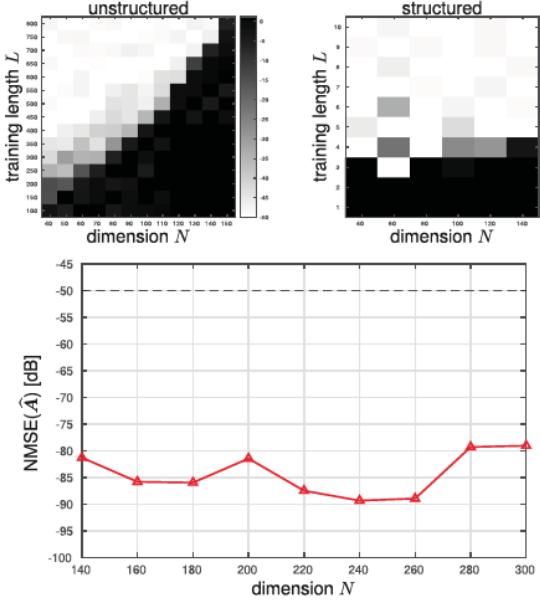


Fig. 7. BAd-VAMP dictionary learning with $N \times N$ dictionary A and $N \times L$ code matrix X with sparsity rate $K/N = 0.2$. Top left: success-rate for unstructured dictionary versus N and L . Top right: success-rate for structured dictionary with $Q = N$ free parameters. Bottom: Median $\text{NMSE}(\hat{A})$ over 20 trials versus N for an unstructured dictionary with $L = 6N \ln N$.

in the unstructured case, where \mathcal{P} denotes the set of generalized permutation matrices.¹¹ For our experiments, we drew the coefficients of $\{A_i\}_{i=1}^Q$ and b as i.i.d. $\mathcal{N}(0, 1)$ with $Q = N$ in the structured case, and we drew the coefficients of A as i.i.d. $\mathcal{N}(0, 1)$ in the unstructured case.

In our first experiment, we fixed the sparsity rate at $K/N = 0.2$ and we varied both the dictionary dimension N and the training length L . The top-right panel of Fig. 7 suggests that, as the dimension N grows, a *fixed* training length L is sufficient to successfully recover A in the structured case with $Q = N$. By “successfully recover,” we mean that $\text{NMSE}(\hat{A}) \leq -50$ dB. Note that this latter prescription for L is consistent with the theoretical analysis in [18]. In the unstructured case, the bottom panel of Figure 7 shows the median $\text{NMSE}(\hat{A})$ versus N when $L = 6N \ln N$. Together, the top-left and bottom panels of Fig. 7 suggest that a training length of $L = O(N \ln N)$ suffices to successfully recover A .

In our second experiment, we focused on the unstructured case, fixed the training length at $L = 5N \ln N$, and varied both the dictionary dimension N and the sparsity K in the columns of X . Figure 8 shows that BAd-VAMP performed similarly to EM-BiGAMP [20] for all but very small N , and much better than K-SVD [16] and SPAMS [17]. The advantage of BAd-VAMP over EM-BiGAMP for DL will be illustrated in the sequel.

E. Noisy, Ill-Conditioned Dictionary Learning

In this section, we show the robustness of BAd-VAMP over EM-BiGAMP [19] when learning ill-conditioned dictionaries

¹¹If P is a generalized permutation matrix then $P = \Pi D$, where Π is a permutation matrix and D is a diagonal matrix.

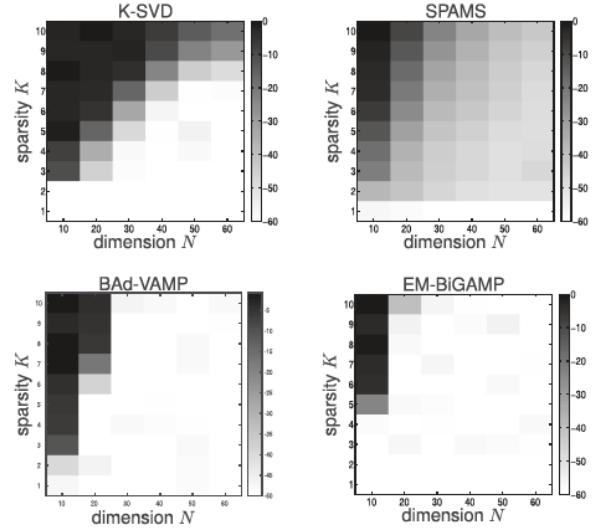


Fig. 8. Noiseless dictionary learning: Median $\text{NMSE}(\hat{A})$ in dB (over 20 trials) versus dictionary dimension N and sparsity K with unstructured dictionary $A \in \mathbb{R}^{N \times N}$ and training length $L = 5N \ln N$.

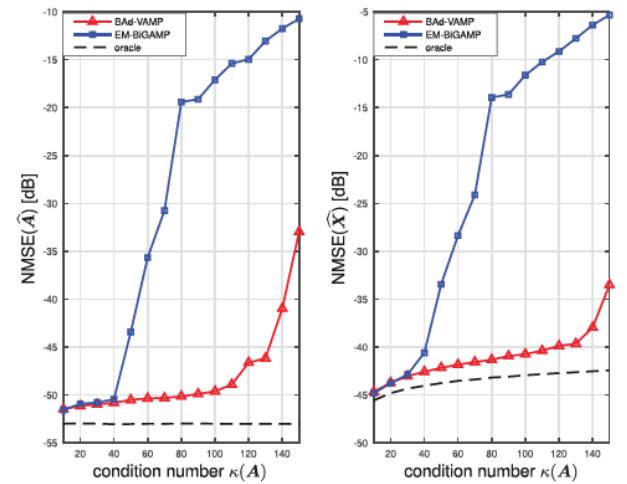


Fig. 9. Noisy dictionary learning: Median NMSE in dB (over 50 trials) versus condition number $\kappa(A)$ for unstructured $A \in \mathbb{R}^{N \times N}$ with $N = 64$, sparsity $K = 13$, training length $L = 5N \ln N$, and SNR = 40 dB.

from noisy measurements. To do so, we generated the measurements as $Y = AX + W$ and tested the algorithms in recovering A and X (up to appropriate ambiguities). The elements of W were drawn i.i.d. $\mathcal{N}(0, 1/\gamma_w)$ with γ_w chosen to achieve $\text{SNR} \triangleq \mathbb{E}[\|AX\|_F^2]/\mathbb{E}\|W\|_F^2 = 40$ dB. The true dictionary was generated as $A = U \text{Diag}(s)V^\top$, where U and V were drawn uniformly over the group of orthogonal matrices, and where the singular values in s were chosen so that $s_i/s_{i-1} = \rho \forall i$. The values of s_0 and ρ were selected to obtain a desired condition number $\kappa(A)$ while also ensuring $\|A\|_F^2 = N$.

Figure 9 reports median $\text{NMSE}(\hat{A})$ and $\text{NMSE}(\hat{X})$ versus condition number $\kappa(A)$ for the recovery of $A \in \mathbb{R}^{N \times N}$ and K -sparse $X \in \mathbb{R}^{N \times L}$ from noisy measurements Y . For this

figure, we used $K = 13$, $N = 64$, $L = 5N \ln N$, the unstructured definition of $\text{NMSE}(\hat{A})$ from (91), and a similar definition for $\text{NMSE}(\hat{X})$. In addition to showing the performance of BAd-VAMP and EM-PBiGAMP, the figure shows the performance of the known- X oracle for the estimation of A , as well as the known- A and known-support oracle for the estimation of X . Figure 9 shows that EM-BiGAMP gave near-oracle NMSE for $\kappa(A) \leq 40$, but its performance degraded significantly for larger $\kappa(A)$. In contrast, BAd-VAMP gave near-oracle NMSE for $\kappa(A) \leq 110$, which suggests increased robustness to ill-conditioned dictionaries A .

VI. CONCLUSION

In this paper, we considered the problem of jointly recovering the vector b and the matrix C from noisy measurements $Y = A(b)C + W$, where $A(\cdot)$ is a known affine linear function of b (i.e., $A(b) = A_0 + \sum_{i=1}^Q b_i A_i$ with known matrices A_i). To solve this problem, we proposed the BAd-VAMP algorithm, which combines the VAMP algorithm [45], the EM algorithm [61], and variance auto-tuning [62] in a manner appropriate for bilinear recovery. We demonstrated numerically that the proposed approach has robustness advantages over other state-of-the-art bilinear recovery algorithms, including lifted VAMP [50] and EM-PBiGAMP [33]. As future work, we plan to rigorously analyze BAd-VAMP through the state-evolution formalism.

REFERENCES

- [1] E. J. Candès and Y. Plan, "Matrix completion with noise," *Proc. IEEE*, vol. 98, no. 6, pp. 925–936, Jun. 2010.
- [2] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, May 2011, Art. no. 11.
- [3] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionary for sparse representation modeling," *Proc. IEEE*, vol. 98, no. 6, pp. 1045–1057, Jun. 2010.
- [4] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Proc. 13th Int. Conf. Neural Inform. Process. Syst.*, 2001, pp. 556–562.
- [5] S. Ling and T. Strohmer, "Self-calibration and biconvex compressive sensing," *Inverse Problems*, vol. 31, no. 11, 2015, Art. no. 115002.
- [6] H. Zhu, G. Leus, and G. B. Giannakis, "Sparsity-cognizant total least-squares for perturbed compressive sampling," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2002–2016, May 2011.
- [7] G. K. Kaleh and R. Vallet, "Joint parameter estimation and symbol detection for linear or nonlinear unknown channels," *IEEE Trans. Commun.*, vol. 42, pp. 2406–2413, Jul. 1994.
- [8] E. J. Candès and Y. Plan, "Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2342–2359, Apr. 2011.
- [9] A. E. Waters, A. C. Sankaranarayanan, and R. G. Baraniuk, "SpaRCS: Recovering low-rank and sparse matrices from compressive measurements," in *Proc. 24th Int. Conf. Neural Inform. Process. Syst.*, 2011, pp. 1089–1097.
- [10] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices," 2010, arXiv:1009.5055.
- [11] Z. Wen, W. Yin, and Y. Zhang, "Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm," *Math. Program. Comput.*, vol. 4, pp. 333–361, 2012.
- [12] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Proc. Allerton Conf. Commun. Control Comput.*, Sep. 2010, pp. 704–711.
- [13] A. Kyriolidis and V. Cevher, "Matrix recipes for hard thresholding methods," *J. Math. Imag. Vis.*, vol. 48, pp. 235–265, 2014.
- [14] S. D. Babacan, M. Luessi, R. Molina, and A. K. Katsaggelos, "Sparse Bayesian methods for low-rank matrix estimation," *IEEE Trans. Signal Process.*, vol. 60, no. 8, pp. 3964–3977, Aug. 2012.
- [15] J. He, L. Balzano, and A. Szlam, "Incremental gradient on the Grassmannian for online foreground and background separation in subsampled video," in *Proc. IEEE Conf. Comp. Vis. Pattern Recognit.*, 2012, pp. 1568–1575.
- [16] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [17] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, Jan. 2010.
- [18] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proc. 25th Annu. Conf. Learn. Theory*, 2012, pp. 37.1–37.18.
- [19] J. T. Parker, P. Schniter, and V. Cevher, "Bilinear generalized approximate message passing—Part I: Derivation," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 5839–5853, Nov. 2014.
- [20] J. T. Parker, P. Schniter, and V. Cevher, "Bilinear generalized approximate message passing—Part II: Applications," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 5854–5867, Nov. 2014.
- [21] Y. Kabashima, F. Krzakala, M. Mézard, A. Sakata, and L. Zdeborová, "Phase transitions and sample complexity in Bayes-optimal matrix factorization," *IEEE Trans. Inf. Theory*, vol. 62, no. 7, pp. 4228–4265, Jul. 2016.
- [22] 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.
- [23] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inf. Theory*, Aug. 2011, pp. 2168–2172.
- [24] R. Matsushita and T. Tanaka, "Low-rank matrix reconstruction and clustering via approximate message passing," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 917–925.
- [25] T. Lesieur, F. Krzakala, and L. Zdeborová, "Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications," *J. Stat. Mech.*, vol. 2017, no. 7, 2017, Art. no. 073403.
- [26] L. Miolane, "Fundamental limits of low-rank matrix estimation: The non-symmetric case," 2017, arXiv:1702.00473.
- [27] C. Bilen, G. Puy, and R. Gribonval, "Convex optimization approaches for blind sensor calibration using sparsity," *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4847–4856, Sep. 2014.
- [28] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *IEEE Trans. Inf. Theory*, vol. 60, no. 3, pp. 1711–1732, Mar. 2014.
- [29] C. Hegde and R. G. Baraniuk, "Sampling and recovery of pulse streams," *IEEE Trans. Signal Process.*, vol. 59, no. 14, pp. 1505–1517, Apr. 2011.
- [30] A. Agarwal, S. Negahban, and M. J. Wainwright, "Matrix decomposition via convex relaxation: Optimal rates in high dimensions," *Ann. Statist.*, vol. 40, no. 2, pp. 1171–1197, 2012.
- [31] J. Wright, A. Ganesh, K. Min, and Y. Ma, "Compressive principal component pursuit," *Inf. Inference*, vol. 2, no. 1, pp. 32–68, 2013.
- [32] M. A. Davenport and J. Romberg, "An overview of low-rank matrix recovery from incomplete observations," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 608–622, Jun. 2016.
- [33] J. T. Parker and P. Schniter, "Parametric bilinear generalized approximate message passing," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 795–808, Jun. 2016.
- [34] C. Schüller, P. Schniter, and L. Zdeborová, "Phase diagram of matrix compressed sensing," *Phys. Rev. E*, vol. 94, no. 6, Dec. 2016, Art. no. 062136.
- [35] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [36] E. Romanov and M. Gavish, "Near-optimal matrix recovery from random linear measurements," *Proc. Nat. Acad. Sci.*, vol. 115, pp. 7200–7205, 2018.
- [37] C. A. Metzler, A. Maleki, and R. G. Baraniuk, "From denoising to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 62, no. 9, pp. 5117–5144, Sep. 2016.
- [38] R. Berthier, A. Montanari, and P.-M. Nguyen, "State evolution for approximate message passing with non-separable functions," *Inf. Inference*, 2019.
- [39] C. Rush and R. Venkataraman, "Finite-sample analysis of approximate message passing," in *Proc. IEEE Int. Symp. Inf. Theory*, 2016, pp. 755–759.

[40] 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.

[41] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and message passing algorithms," *Ann. App. Prob.*, vol. 25, no. 2, pp. 753–822, 2015.

[42] S. Rangan, P. Schniter, E. Riegler, A. Fletcher, and V. Cevher, "Fixed points of generalized approximate message passing with arbitrary matrices," *IEEE Trans. Inf. Theory*, vol. 62, no. 12, pp. 7464–7474, Dec. 2016.

[43] F. Caltagirone, F. Krzakala, and L. Zdeborová, "On convergence of approximate message passing," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2014, pp. 1812–1816.

[44] J. Vila, P. Schniter, S. Rangan, F. Krzakala, and L. Zdeborová, "Adaptive damping and mean removal for the generalized approximate message passing algorithm," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2015, pp. 2021–2025.

[45] S. Rangan, P. Schniter, and A. K. Fletcher, "Vector approximate message passing," in *Proc. IEEE Int. Symp. Inf. Theory*, 2017, pp. 1588–1592.

[46] M. Opper and O. Winther, "Expectation consistent free energies for approximate inference," in *Proc. 17th Int. Conf. Neural Inf. Process. Syst.*, 2005, pp. 1001–1008.

[47] A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Expectation consistent approximate inference: Generalizations and convergence," in *Proc. IEEE Int. Symp. Inf. Theory*, 2016, pp. 190–194.

[48] J. Ma and L. Ping, "Orthogonal AMP," *IEEE Access*, vol. 5, pp. 2020–2033, 2017.

[49] K. Takeuchi, "Rigorous dynamics of expectation-propagation-based signal recovery from unitarily invariant measurements," in *Proc. IEEE Int. Symp. Inf. Theory*, 2017, pp. 501–505.

[50] A. K. Fletcher, S. Rangan, S. Sarkar, and P. Schniter, "Plug-in estimation in high-dimensional linear inverse problems: A rigorous analysis," in *Proc. Neural Inf. Process. Syst. Conf.*, 2018, pp. 7440–7449.

[51] A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Rigorous dynamics and consistent estimation in arbitrarily conditioned linear systems," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 2542–2551.

[52] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Found. Trends Mach. Learn.*, vol. 1, pp. 1–305, Jan. 2008.

[53] T. Minka, "A family of approximate algorithms for Bayesian inference," Ph.D. dissertation, Dept. Comp. Sci. Eng., MIT, Cambridge, MA, USA, Jan. 2001.

[54] M. Seeger, "Expectation propagation for exponential families," Ecole polytechnique fédrale de Lausanne, Lausanne, Switzerland, Tech. Rep. 161464, 2005.

[55] T. Heskes and O. Zoeter, "Expectation propagation for approximate inference in dynamic Bayesian networks," in *Proc. Uncertainty Artif. Intell.*, 2002, pp. 313–320.

[56] A. M. Tulino, G. Caire, S. Verdú, and S. Shamai (Shitz), "Support recovery with sparsely sampled free random matrices," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4243–4271, Jul. 2013.

[57] G. Reeves, "Additivity of information in multilayer networks via additive Gaussian noise transforms," in *Proc. Allerton Conf. Commun. Control Comput.*, 2017, pp. 1064–1070.

[58] G. Reeves and H. D. Pfister, "The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact," in *Proc. IEEE Int. Symp. Inf. Theory*, 2016, pp. 665–669.

[59] J. Barbier, M. Dia, N. Macris, and F. Krzakala, "The mutual information in random linear estimation," in *Proc. Allerton Conf. Commun. Control Comput.*, 2016, pp. 625–632.

[60] 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.

[61] R. Neal and G. Hinton, "A view of the EM algorithm that justifies incremental, sparse, and other variants," in *Learning in Graphical Models*, M. I. Jordan, Ed. Cambridge, MA, USA: MIT Press, 1998, pp. 355–368.

[62] U. S. Kamilov, S. Rangan, A. K. Fletcher, and M. Unser, "Approximate message passing with consistent parameter estimation and applications to sparse learning," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2969–2985, May 2014.

[63] J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Trans. Signal Process.*, vol. 61, no. 19, pp. 4658–4672, Oct. 2013.

[64] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.

[65] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, pp. 259–268, 1992.

[66] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problem," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.

[67] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, 2017.

[68] E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Trans. Comp. Imag.*, vol. 5, no. 1, pp. 52–67, Mar. 2019.

[69] T. Goldstein, C. Studer, and R. Baraniuk, "Forward-backward splitting with a FASTA implementation," 2014, arXiv:1411.3406.

Subrata Sarkar (S'12) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Guwahati, India, in 2014, and the M.S. degree in electrical engineering, in 2016, from The Ohio State University, Columbus, OH, USA, where he is currently working toward the Ph.D. degree.

Alison K. Fletcher (S'03–M'04) received the B.S. degree in mathematics from the University of Iowa, Iowa City, IA, USA, the M.S. degree in mathematics and electrical engineering, and the Ph.D. degree in electrical engineering, both from the University of California at Berkeley, Berkeley, CA, USA. Since 2016, she has been on the Faculty of the Departments of Statistics, Mathematics, Electrical Engineering, and Computer Science, University of California, Los Angeles, Los Angeles, CA, USA.

Sundeep Rangan (S'94–M'98–SM'13–F'16) received the B.A.Sc. degree from the University of Waterloo, Waterloo, ON, Canada, and the M.Sc. and Ph.D. degrees from the University of California, Berkeley, Berkeley, CA, USA, all in electrical engineering. Since 2010, he has been on the Faculty of the Department of Electronics and Communication Engineering, New York University Polytechnic School of Engineering, Brooklyn, NY, USA.

Philip Schniter (S'92–M'93–SM'05–F'14) received the B.S. and M.S. degrees in electrical engineering from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, and the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, USA. Since 2000, he has been on the Faculty of the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, USA.