

Vector Approximate Message Passing

Sundeep Rangan¹, *Fellow, IEEE*, Philip Schniter², *Fellow, IEEE*, and Alyson K. Fletcher, *Member, IEEE*

Abstract—The standard linear regression (SLR) problem is to recover a vector \mathbf{x}^0 from noisy linear observations $\mathbf{y} = \mathbf{A}\mathbf{x}^0 + \mathbf{w}$. The approximate message passing (AMP) algorithm proposed by Donoho, Maleki, and Montanari is a computationally efficient iterative approach to SLR that has a remarkable property: for large i.i.d. sub-Gaussian matrices \mathbf{A} , its per-iteration behavior is rigorously characterized by a scalar state-evolution whose fixed points, when unique, are Bayes optimal. The AMP algorithm, however, is fragile in that even small deviations from the i.i.d. sub-Gaussian model can cause the algorithm to diverge. This paper considers a “vector AMP” (VAMP) algorithm and shows that VAMP has a rigorous scalar state-evolution that holds under a much broader class of large random matrices \mathbf{A} : those that are right-orthogonally invariant. After performing an initial singular value decomposition (SVD) of \mathbf{A} , the per-iteration complexity of VAMP is similar to that of AMP. In addition, the fixed points of VAMP’s state evolution are consistent with the replica prediction of the minimum mean-squared error derived by Tulino, Caire, Verdú, and Shamai. Numerical experiments are used to confirm the effectiveness of VAMP and its consistency with state-evolution predictions.

Index Terms—Belief propagation, message passing, inference algorithms, random matrices, compressive sensing.

I. INTRODUCTION

CONSIDER the problem of recovering a vector $\mathbf{x}^0 \in \mathbb{R}^N$ from noisy linear measurements of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x}^0 + \mathbf{w} \in \mathbb{R}^M, \quad (1)$$

where \mathbf{A} is a known matrix and \mathbf{w} is an unknown, unstructured noise vector. In the statistics literature, this problem is known as *standard linear regression*, and in the signal processing literature this is known as *solving a linear inverse problem*, or as *compressive sensing* when $M \ll N$ and \mathbf{x}^0 is sparse.

Manuscript received October 11, 2016; revised June 10, 2018; accepted May 1, 2019. Date of publication May 13, 2019; date of current version September 13, 2019. This work was supported in part by the National Science Foundation under Grant 1302336, Grant 1564142, and Grant 1547332, in part by the National Science Foundation under Grant 1527162 and Grant 1716388, in part by the National Science Foundation under Grant 1254204 and Grant 1564278, and in part by the Office of Naval Research under Grant N00014-15-1-2677. This work was presented in part at the 2017 IEEE Symposium on Information Theory.

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

P. Schniter is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: schniter.1@osu.edu).

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

Communicated by N. Kiyavash, Associate Editor for Statistical Learning.

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2019.2916359

A. Problem Formulations

One approach to recovering \mathbf{x}^0 is *regularized quadratic loss minimization*, where an estimate $\hat{\mathbf{x}}$ of \mathbf{x}^0 is computed by solving an optimization problem of the form

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + f(\mathbf{x}). \quad (2)$$

Here, the penalty function or “regularization” $f(\mathbf{x})$ is chosen to promote a desired structure in $\hat{\mathbf{x}}$. For example, the choice $f(\mathbf{x}) = \lambda \|\mathbf{x}\|_1$ with $\lambda > 0$ promotes sparsity in $\hat{\mathbf{x}}$.

Another approach is through the Bayesian methodology. Here, one presumes a prior density $p(\mathbf{x})$ and likelihood function $p(\mathbf{y}|\mathbf{x})$ and then aims to compute the posterior density

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{\int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) d\mathbf{x}} \quad (3)$$

or, in practice, a summary of it [1]. Example summaries include the maximum *a posteriori* (MAP) estimate

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}), \quad (4)$$

the minimum mean-squared error (MMSE) estimate

$$\hat{\mathbf{x}}_{\text{MMSE}} = \arg \min_{\tilde{\mathbf{x}}} \int \|\mathbf{x} - \tilde{\mathbf{x}}\|^2 p(\mathbf{x}|\mathbf{y}) d\mathbf{x} = \mathbb{E}[\mathbf{x}|\mathbf{y}], \quad (5)$$

or the posterior marginal densities $\{p(x_n|\mathbf{y})\}_{n=1}^N$.

Note that, if the noise \mathbf{w} is modeled as $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \gamma_w^{-1}\mathbf{I})$, i.e., additive white Gaussian noise (AWGN) with some precision $\gamma_w > 0$, then the regularized quadratic loss minimization problem (2) is equivalent to MAP estimation under the prior $p(\mathbf{x}) \propto \exp[-\gamma_w f(\mathbf{x})]$, where \propto denotes equality up to a scaling that is independent of \mathbf{x} . Thus we focus on MAP, MMSE, and marginal posterior inference in the sequel.

B. Approximate Message Passing

Recently, the so-called *approximate message passing* (AMP) algorithm [2], [3] was proposed as an iterative method to recover \mathbf{x}^0 from measurements of the form (1). The AMP iterations are specified in Algorithm 1. There,¹ $\mathbf{g}_1(\cdot, \gamma_k) : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is a *denoising* function parameterized by γ_k , and $\langle \mathbf{g}'_1(\mathbf{r}_k, \gamma_k) \rangle$ is its *divergence* at \mathbf{r}_k . In particular, $\mathbf{g}'_1(\mathbf{r}_k, \gamma_k) \in \mathbb{R}^N$ is the diagonal of the Jacobian,

$$\mathbf{g}'_1(\mathbf{r}_k, \gamma_k) = \text{diag} \left[\frac{\partial \mathbf{g}_1(\mathbf{r}_k, \gamma_k)}{\partial \mathbf{r}_k} \right], \quad (6)$$

¹The subscript “1” in \mathbf{g}_1 is used to promote notational consistency with Vector AMP algorithm presented in the sequel.

Algorithm 1 AMP

Require: Matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$, measurement vector \mathbf{y} , denoiser $\mathbf{g}_1(\cdot, \gamma_k)$, and number of iterations K_{it} .

- 1: Set $\mathbf{v}_{-1} = \mathbf{0}$ and select initial \mathbf{r}_0, γ_0 .
- 2: **for** $k = 0, 1, \dots, K_{\text{it}}$ **do**
- 3: $\widehat{\mathbf{x}}_k = \mathbf{g}_1(\mathbf{r}_k, \gamma_k)$
- 4: $\alpha_k = \langle \mathbf{g}'_1(\mathbf{r}_k, \gamma_k) \rangle$
- 5: $\mathbf{v}_k = \mathbf{y} - \mathbf{A}\widehat{\mathbf{x}}_k + \frac{N}{M}\alpha_k\mathbf{v}_{k-1}$
- 6: $\mathbf{r}_{k+1} = \widehat{\mathbf{x}}_k + \mathbf{A}^\top\mathbf{v}_k$
- 7: Select γ_{k+1}
- 8: **end for**
- 9: Return $\widehat{\mathbf{x}}_{K_{\text{it}}}$.

and $\langle \cdot \rangle$ is the empirical averaging operation

$$\langle \mathbf{u} \rangle := \frac{1}{N} \sum_{n=1}^N u_n. \quad (7)$$

When \mathbf{A} is a large i.i.d. sub-Gaussian matrix, $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \gamma_{w0}^{-1}\mathbf{I})$, and $\mathbf{g}_1(\cdot, \gamma_k)$ is *separable*, i.e.,

$$[\mathbf{g}_1(\mathbf{r}_k, \gamma_k)]_n = g_1(r_{kn}, \gamma_k) \quad \forall n, \quad (8)$$

with identical Lipschitz components $g_1(\cdot, \gamma_k) : \mathbb{R} \rightarrow \mathbb{R}$, AMP displays a remarkable behavior, which is that \mathbf{r}_k behaves like a white-Gaussian-noise corrupted version of the true signal \mathbf{x}^0 [2]. That is,

$$\mathbf{r}_k = \mathbf{x}^0 + \mathcal{N}(\mathbf{0}, \tau_k \mathbf{I}), \quad (9)$$

for some variance $\tau_k > 0$. Moreover, the variance τ_k can be predicted through the following *state evolution* (SE):

$$\mathcal{E}(\gamma_k, \tau_k) = \frac{1}{N} \mathbb{E} \left[\left\| \mathbf{g}_1(\mathbf{x}^0 + \mathcal{N}(\mathbf{0}, \tau_k \mathbf{I}), \gamma_k) - \mathbf{x}^0 \right\|^2 \right] \quad (10a)$$

$$\tau_{k+1} = \gamma_{w0}^{-1} + \frac{N}{M} \mathcal{E}(\gamma_k, \tau_k), \quad (10b)$$

where $\mathcal{E}(\gamma_k, \tau_k)$ is the MSE of the AMP estimate $\widehat{\mathbf{x}}_k$.

The AMP SE (10) was rigorously established for i.i.d. Gaussian \mathbf{A} in [4] and for i.i.d. sub-Gaussian \mathbf{A} in [5] in the *large-system limit* (i.e., $N, M \rightarrow \infty$ and $N/M \rightarrow \delta \in (0, 1)$) under some mild regularity conditions. Because the SE (10) holds for generic $g_1(\cdot, \gamma_k)$ and generic γ_k -update rules, it can be used to characterize the application of AMP to many problems, as further discussed in Section II-A.

C. Limitations, Modifications, and Alternatives to AMP

An important limitation of AMP's SE is that it holds only under large i.i.d. sub-Gaussian \mathbf{A} . Although recent analysis [6] has rigorously analyzed AMP's performance under finite-sized i.i.d. Gaussian \mathbf{A} , there remains the important question of how AMP behaves with general \mathbf{A} .

Unfortunately, it turns out that the AMP Algorithm 1 is somewhat fragile with regard to the construction of \mathbf{A} . For example, AMP diverges with even mildly ill-conditioned or non-zero-mean \mathbf{A} [7]–[9]. Although damping [7], [9], mean-removal [9], sequential updating [10], and direct free-energy minimization [11] all help to prevent AMP from diverging, such strategies are limited in effectiveness.

Many other algorithms for standard linear regression (1) have been designed using approximations of belief propagation (BP) and/or free-energy minimization. Among these are the Adaptive Thouless-Anderson-Palmer (ADATAP) [12], Expectation Propagation (EP) [13], [14], Expectation Consistent Approximation (EC) [15]–[17], (S-transform AMP) S-AMP [18], [19], and (Orthogonal AMP) OAMP [20] approaches. Although numerical experiments suggest that some of these algorithms are more robust than AMP Algorithm 1 to the choice of \mathbf{A} , their convergence has not been rigorously analyzed. In particular, there remains the question of whether there exists an AMP-like algorithm with a rigorous SE analysis that holds for a larger class of matrices than i.i.d. sub-Gaussian. In the sequel, we describe one such algorithm.

D. Contributions

In this paper, we propose a computationally efficient iterative algorithm for the estimation of the vector \mathbf{x}^0 from noisy linear measurements \mathbf{y} of the form in (1). (See Algorithm 2.) We call the algorithm “*vector AMP*” (VAMP) because i) its behavior can be rigorously characterized by a scalar SE under large random \mathbf{A} , and ii) it can be derived using an approximation of BP on a factor graph with vector-valued variable nodes. We outline VAMP's derivation in Section III with the aid of some background material that is reviewed in Section II.

In Section IV, we establish the VAMP SE in the case of large *right-orthogonally invariant* random \mathbf{A} and separable Lipschitz denoisers $\mathbf{g}_1(\cdot, \gamma_k)$, using techniques similar to those used by Bayati and Montanari in [4]. Importantly, these right-orthogonally invariant \mathbf{A} allow arbitrary singular values and arbitrary left singular vectors, making VAMP much more robust than AMP in regards to the construction of \mathbf{A} . In Section V, we establish that the asymptotic MSE predicted by VAMP's SE agrees with the MMSE predicted by the replica method [21] when VAMP's priors are matched to the true data. Finally, in Section VI, we present numerical experiments demonstrating that VAMP's empirical behavior matches its SE at moderate dimensions, even when \mathbf{A} is highly ill-conditioned or non-zero-mean.

E. Relation to Existing Work

The idea to construct algorithms from graphical models with vector-valued nodes is not new, and in fact underlies the EC- and EP-based algorithms described in [13]–[17]. The use of vector-valued nodes is also central to the derivation of S-AMP [18], [19]. In the sequel, we present a simple derivation of VAMP that uses the EP methodology from [13], [14], which passes approximate messages between the nodes of a factor graph. But we note that VAMP can also be derived using the EC methodology, which formulates a variational optimization problem using a constrained version of the Kullback-Leibler distance and then relaxes the density constraints to moment constraints. For more details on the latter approach, we refer the interested reader to the discussion of “diagonal restricted EC” in [15, App. D] and “uniform diagonalized EC” in [17].

It was recently shown [16] that, for large right-orthogonally invariant \mathbf{A} , the fixed points of diagonal-restricted EC are “good” in the sense that they are consistent with a certain replica prediction of the MMSE that is derived in [16]. Since the fixed points of ADATAP and S-AMP are known [18] to coincide with those of diagonal-restricted EC (and thus VAMP), all of these algorithms can be understood to have good fixed points. The trouble is that these algorithms do not necessarily converge to their fixed points. For example, S-AMP diverges with even mildly ill-conditioned or non-zero-mean \mathbf{A} , as demonstrated in Section VI. *Our main contribution is establishing that VAMP’s behavior can be exactly predicted by an SE analysis analogous to that for AMP. This SE analysis then provides precise convergence guarantees for large right-orthogonally invariant \mathbf{A} .* The numerical results presented in Section VI confirm that, in practice, VAMP’s convergence is remarkably robust, even with very ill-conditioned or mean-perturbed matrices \mathbf{A} of finite dimension.

The main insight that leads to both the VAMP algorithm and its SE analysis comes from a consideration of the singular value decomposition (SVD) of \mathbf{A} . Specifically, take the “economy” SVD,

$$\mathbf{A} = \bar{\mathbf{U}}\text{Diag}(\bar{\mathbf{s}})\bar{\mathbf{V}}^T, \quad (11)$$

where $\bar{\mathbf{s}} \in \mathbb{R}^R$ for $R := \text{rank}(\mathbf{A}) \leq \min(M, N)$. The VAMP iterations can be performed by matrix-vector multiplications with $\bar{\mathbf{V}} \in \mathbb{R}^{N \times R}$ and $\bar{\mathbf{V}}^T$, yielding a structure very similar to that of AMP. Computationally, the SVD form of VAMP (i.e., Algorithm 2) has the benefit that, once the SVD has been computed, VAMP’s per-iteration cost will be dominated by $O(RN)$ floating-point operations (flops), as opposed to $O(N^3)$ for the EC methods from [15, App. D] or [17]. Furthermore, if these matrix-vector multiplications have fast implementations (e.g., $O(N)$ when $\bar{\mathbf{V}}$ is a discrete wavelet transform), then the complexity of VAMP reduces accordingly. We emphasize that VAMP uses a single SVD, not a per-iteration SVD. In many applications, this SVD can be computed off-line. In the case that SVD complexity may be an issue, we note that it costs $O(MNR)$ flops by classical methods or $O(MN \log R)$ by modern approaches [22].

The SVD offers more than just a fast algorithmic implementation. More importantly, it connects VAMP to AMP in such a way that the Bayati and Montanari’s SE analysis of AMP [4] can be extended to obtain a rigorous SE for VAMP. In this way, the SVD can be viewed as a proof technique. Since it will be useful for derivation/interpretation in the sequel, we note that the VAMP iterations can also be written without an explicit SVD (see Algorithm 3), in which case they coincide with the uniform-diagonalization variant of the generalized EC method from [17]. In this latter implementation, the linear MMSE (LMMSE) estimate (24) must be computed at each iteration, as well as the trace of its covariance matrix (25), which both involve the inverse of an $N \times N$ matrix.

The OAMP-LMMSE algorithm from [20] is similar to VAMP and diagonal-restricted EC, but different in that it approximates certain variance terms. This difference can be seen by comparing equations (30)-(31) in [20]

to lines 8 and 10 in Algorithm 2 (or lines 14 and 7 in Algorithm 3). Furthermore, OAMP-LMMSE differs from VAMP in its reliance on matrix inversion (see, e.g., the comments in the Conclusion of [20]).

Shortly after the initial publication of this work, [23] proved a very similar result for the complex case using a fully probabilistic analysis.

F. Notation

We use capital boldface letters like \mathbf{A} for matrices, small boldface letters like \mathbf{a} for vectors, $(\cdot)^T$ for transposition, and $a_n = [\mathbf{a}]_n$ to denote the n th element of \mathbf{a} . Also, we use $\|\mathbf{a}\|_p = (\sum_n |a_n|^p)^{1/p}$ for the ℓ_p norm of \mathbf{a} , $\|\mathbf{A}\|_2$ for the spectral norm of \mathbf{A} , $\text{Diag}(\mathbf{a})$ for the diagonal matrix created from vector \mathbf{a} , and $\text{diag}(\mathbf{A})$ for the vector extracted from the diagonal of matrix \mathbf{A} . Likewise, we use \mathbf{I}_N for the $N \times N$ identity matrix, $\mathbf{0}$ for the matrix of all zeros, and $\mathbf{1}$ for the matrix of all ones. For a random vector \mathbf{x} , we denote its probability density function (pdf) by $p(\mathbf{x})$, its expectation by $\mathbb{E}[\mathbf{x}]$, and its covariance matrix by $\text{Cov}[\mathbf{x}]$. Similarly, we use $p(\mathbf{x}|\mathbf{y})$, $\mathbb{E}[\mathbf{x}|\mathbf{y}]$, and $\text{Cov}[\mathbf{x}|\mathbf{y}]$ for the conditional pdf, expectation, and covariance, respectively. Also, we use $\mathbb{E}[\mathbf{x}|b]$ and $\text{Cov}[\mathbf{x}|b]$ to denote the expectation and covariance of $\mathbf{x} \sim b(\mathbf{x})$, i.e., \mathbf{x} distributed according to the pdf $b(\mathbf{x})$. We refer to the Dirac delta pdf using $\delta(\mathbf{x})$ and to the pdf of a Gaussian random vector $\mathbf{x} \in \mathbb{R}^N$ with mean \mathbf{a} and covariance \mathbf{C} using $\mathcal{N}(\mathbf{x}; \mathbf{a}, \mathbf{C}) = \exp(-(\mathbf{x}-\mathbf{a})^T \mathbf{C}^{-1}(\mathbf{x}-\mathbf{a})/2)/\sqrt{(2\pi)^N |\mathbf{C}|}$. Finally, $p(\mathbf{x}) \propto f(\mathbf{x})$ says that functions $p(\cdot)$ and $f(\cdot)$ are equal up to a scaling that is invariant to \mathbf{x} .

II. BACKGROUND ON THE AMP ALGORITHM

In this section, we provide background on the AMP algorithm that will be useful in the sequel.

A. Applications to Bayesian Inference

We first detail the application of the AMP Algorithm 1 to the Bayesian inference problems from Section I-A. Suppose that the prior on \mathbf{x} is i.i.d., so that it takes the form

$$p(\mathbf{x}) = \prod_{n=1}^N p(x_n). \quad (12)$$

Then AMP can be applied to MAP problem (4) by choosing the scalar denoiser as

$$g_1(r_{kn}, \gamma_k) = \arg \min_{x_n \in \mathbb{R}} \left[\frac{\gamma_k}{2} |x_n - r_{kn}|^2 - \ln p(x_n) \right]. \quad (13)$$

Likewise, AMP can be applied to the MMSE problem (5) by choosing

$$g_1(r_{kn}, \gamma_k) = \mathbb{E}[x_n | r_{kn}, \gamma_k], \quad (14)$$

where the expectation in (14) is with respect to the conditional density

$$p(x_n | r_{kn}, \gamma_k) \propto \exp \left[-\frac{\gamma_k}{2} |r_{kn} - x_n|^2 + \ln p(x_n) \right]. \quad (15)$$

In addition, $p(x_n | r_{kn}, \gamma_k)$ in (15) acts as AMP’s iteration- k approximation of the marginal posterior $p(x_n | \mathbf{y})$. For later use,

we note that the derivative of the MMSE scalar denoiser (14) w.r.t. its first argument can be expressed as

$$g'_1(r_{kn}, \gamma_k) = \gamma_k \text{var} \left[x_n \middle| r_{kn}, \gamma_k \right], \quad (16)$$

where the variance is computed with respect to the density (15) (see, e.g., [24]).

In (13)-(15), γ_k can be interpreted as an estimate of τ_k^{-1} , the iteration- k precision of \mathbf{r}_k from (9). In the case that τ_k is known, the ‘‘matched’’ assignment

$$\gamma_k = \tau_k^{-1} \quad (17)$$

leads to the interpretation of (13) and (14) as the scalar MAP and MMSE denoisers of r_{kn} , respectively. Since, in practice, τ_k is usually not known, it has been suggested to use

$$\gamma_{k+1} = \frac{M}{\|\mathbf{v}_k\|^2}, \quad (18)$$

although other choices are possible [25].

B. Relation of AMP to IST

The AMP Algorithm 1 is closely related to the well-known *iterative soft thresholding* (IST) algorithm [26], [27] that can be used² to solve (2) with convex $f(\cdot)$. In particular, if the term

$$\frac{N}{M} \alpha_k \mathbf{v}_{k-1} \quad (19)$$

is removed from line 5 of Algorithm 1, then what remains is the IST algorithm.

The term (19) is known as the *Onsager* term in the statistical physics literature [28]. Under large i.i.d. sub-Gaussian \mathbf{A} , the Onsager correction ensures the behavior in (9). When (9) holds, the denoiser $g_1(\cdot, \gamma_k)$ can be optimized accordingly, in which case each iteration of AMP becomes very productive. As a result, AMP converges much faster than ISTA for i.i.d. Gaussian \mathbf{A} (see, e.g., [25] for a comparison).

C. Derivations of AMP

The AMP algorithm can be derived in several ways. One way is through approximations of loopy belief propagation (BP) [29], [30] on a bipartite factor graph constructed from the factorization

$$p(\mathbf{y}, \mathbf{x}) = \left[\prod_{m=1}^M \mathcal{N}(y_m; \mathbf{a}_m^T \mathbf{x}, \gamma_w^{-1}) \right] \left[\prod_{n=1}^N p(x_n) \right], \quad (20)$$

where \mathbf{a}_m^T denotes the m th row of \mathbf{A} . We refer the reader to [3], [24] for details on the message-passing derivation of AMP, noting connections to the general framework of *expectation propagation* (EP) [13], [14]. AMP can also be derived through a ‘‘free-energy’’ approach, where one i) proposes a cost function, ii) derives conditions on its stationary points, and iii) constructs an algorithm whose fixed points coincide with those stationary points. We refer the reader to [18], [31], [32] for details, and note connections to the general framework of *expectation consistent approximation* (EC) [15], [17].

²The IST algorithm is guaranteed to converge [27] when $\|\mathbf{A}\|_2 < 1$.

Algorithm 2 Vector AMP (SVD Form)

Require: Matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$; measurements $\mathbf{y} \in \mathbb{R}^M$; denoiser $\mathbf{g}_1(\cdot, \gamma_k)$; assumed noise precision $\gamma_w \geq 0$; and number of iterations K_{it} .

- 1: Compute economy SVD $\bar{\mathbf{U}} \text{Diag}(\bar{\mathbf{s}}) \bar{\mathbf{V}}^T = \mathbf{A}$ with $\bar{\mathbf{U}}^T \bar{\mathbf{U}} = \mathbf{I}_R$, $\bar{\mathbf{V}}^T \bar{\mathbf{V}} = \mathbf{I}_R$, $\bar{\mathbf{s}} \in \mathbb{R}_+^R$, $R = \text{rank}(\mathbf{A})$.
 - 2: Compute preconditioned $\tilde{\mathbf{y}} := \text{Diag}(\bar{\mathbf{s}})^{-1} \bar{\mathbf{U}}^T \mathbf{y}$
 - 3: Select initial \mathbf{r}_0 and $\gamma_0 \geq 0$.
 - 4: **for** $k = 0, 1, \dots, K_{\text{it}}$ **do**
 - 5: $\hat{\mathbf{x}}_k = \mathbf{g}_1(\mathbf{r}_k, \gamma_k)$
 - 6: $\alpha_k = \langle \mathbf{g}'_1(\mathbf{r}_k, \gamma_k) \rangle$
 - 7: $\tilde{\mathbf{r}}_k = (\hat{\mathbf{x}}_k - \alpha_k \mathbf{r}_k) / (1 - \alpha_k)$
 - 8: $\tilde{\gamma}_k = \gamma_k (1 - \alpha_k) / \alpha_k$
 - 9: $\mathbf{d}_k = \gamma_w \text{Diag}(\gamma_w \bar{\mathbf{s}}^2 + \tilde{\gamma}_k \mathbf{1})^{-1} \bar{\mathbf{s}}^2$
 - 10: $\gamma_{k+1} = \tilde{\gamma}_k \langle \mathbf{d}_k \rangle / (\frac{N}{R} - \langle \mathbf{d}_k \rangle)$
 - 11: $\mathbf{r}_{k+1} = \tilde{\mathbf{r}}_k + \frac{N}{R} \bar{\mathbf{V}} \text{Diag}(\mathbf{d}_k / \langle \mathbf{d}_k \rangle) (\tilde{\mathbf{y}} - \bar{\mathbf{V}}^T \tilde{\mathbf{r}}_k)$
 - 12: **end for**
 - 13: **Return** $\hat{\mathbf{x}}_{K_{\text{it}}}$.
-

III. THE VECTOR AMP ALGORITHM

The *Vector AMP* (VAMP) algorithm is stated in Algorithm 2. In line 9, ‘‘ $\bar{\mathbf{s}}^2$ ’’ refers to the componentwise square of vector $\bar{\mathbf{s}}$. Also, $\text{Diag}(\mathbf{a})$ denotes the diagonal matrix whose diagonal components are given by the vector \mathbf{a} .

A. Relation of VAMP to AMP

A visual examination of VAMP Algorithm 2 shows many similarities with AMP Algorithm 1. In particular, the denoising and divergence steps in lines 5-6 of Algorithm 2 are identical to those in lines 3-4 of Algorithm 1. Likewise, an Onsager term $\alpha_k \mathbf{r}_k$ is visible in line 7 of Algorithm 2, analogous to the one in line 5 of Algorithm 1. Finally, the per-iteration computational complexity of each algorithm is dominated by two matrix-vector multiplications: those involving \mathbf{A} and \mathbf{A}^T in Algorithm 1 and those involving $\bar{\mathbf{V}}$ and $\bar{\mathbf{V}}^T$ in Algorithm 2.

The most important similarity between the AMP and VAMP algorithms is not obvious from visual inspection and will be established rigorously in the sequel. It is the following: for certain large random \mathbf{A} , the VAMP quantity \mathbf{r}_k behaves like a white-Gaussian-noise corrupted version of the true signal \mathbf{x}^0 , i.e.,

$$\mathbf{r}_k = \mathbf{x}^0 + \mathcal{N}(\mathbf{0}, \tau_k \mathbf{I}), \quad (21)$$

for some variance $\tau_k > 0$. Moreover, the noise variance τ_k can be tracked through a scalar SE formalism whose details will be provided in the sequel. Furthermore, the VAMP quantity γ_k can be interpreted as an estimate of τ_k^{-1} in (21), analogous to the AMP quantity γ_k discussed around (17).

It should be emphasized that the class of matrices \mathbf{A} under which the VAMP SE holds is much bigger than the class under which the AMP SE holds. In particular, VAMP’s SE




Fig. 1. The factor graph used for the derivation of VAMP. The circles represent variable nodes and the squares represent factor nodes from (23).

holds for large random matrices \mathbf{A} whose right singular³ vector matrix $\mathbf{V} \in \mathbb{R}^{N \times N}$ is uniformly distributed on the group of orthogonal matrices. Notably, VAMP's SE holds for *arbitrary* (i.e., deterministic) left singular vector matrices \mathbf{U} and singular values, apart from some mild regularity conditions that will be detailed in the sequel. In contrast, AMP's SE is known to hold [4], [5] only for large i.i.d. sub-Gaussian matrices \mathbf{A} , which implies i) random orthogonal \mathbf{U} and \mathbf{V} and ii) a particular distribution on the singular values of \mathbf{A} .

B. EP Derivation of VAMP

As with AMP (i.e., Algorithm 1), VAMP (i.e., Algorithm 2) can be derived in many ways. Here we present a very simple derivation based on an EP-like approximation of the sum-product (SP) belief-propagation algorithm. Unlike the AMP algorithm, whose message-passing derivation uses a loopy factor graph with *scalar*-valued nodes, the VAMP algorithm uses a non-loopy graph with *vector*-valued nodes, hence the name “vector AMP.” We note that VAMP can also be derived using the “diagonal restricted” or “uniform diagonalization” EC approach [15], [17], but that derivation is much more complicated.

To derive VAMP, we start with the factorization

$$p(\mathbf{y}, \mathbf{x}) = p(\mathbf{x})\mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}, \gamma_w^{-1}\mathbf{I}), \quad (22)$$

and split \mathbf{x} into two identical variables $\mathbf{x}_1 = \mathbf{x}_2$, giving an equivalent factorization

$$p(\mathbf{y}, \mathbf{x}_1, \mathbf{x}_2) = p(\mathbf{x}_1)\delta(\mathbf{x}_1 - \mathbf{x}_2)\mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}_2, \gamma_w^{-1}\mathbf{I}), \quad (23)$$

where $\delta(\cdot)$ is the Dirac delta distribution. The factor graph corresponding to (23) is shown in Figure 1. We then pass messages on this factor graph according to the following rules.

- 1) Approximate beliefs: The approximate belief $b_{\text{app}}(\mathbf{x})$ on variable node \mathbf{x} is $\mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \eta^{-1}\mathbf{I})$, where $\hat{\mathbf{x}} = \mathbb{E}[\mathbf{x}|b_{\text{SP}}]$ and $\eta^{-1} = \langle \text{diag}(\text{Cov}[\mathbf{x}|b_{\text{SP}}]) \rangle$ are the mean and average variance of the corresponding SP belief $b_{\text{SP}}(\mathbf{x}) \propto \prod_i \mu_{f_i \rightarrow \mathbf{x}}(\mathbf{x})$, i.e., the normalized product of all messages impinging on the node. See Figure 2(a) for an illustration.
- 2) Variable-to-factor messages: The message from a variable node \mathbf{x} to a connected factor node f_i is $\mu_{\mathbf{x} \rightarrow f_i}(\mathbf{x}) \propto b_{\text{app}}(\mathbf{x}) / \mu_{f_i \rightarrow \mathbf{x}}(\mathbf{x})$, i.e., the ratio of the most recent approximate belief $b_{\text{app}}(\mathbf{x})$ to the most recent message from f_i to \mathbf{x} . See Figure 2(b) for an illustration.
- 3) Factor-to-variable messages: The message from a factor node f to a connected variable node \mathbf{x}_i is $\mu_{f \rightarrow \mathbf{x}_i}(\mathbf{x}_i) \propto$

³We use several forms of SVD in this paper. Algorithm 2 uses the “economy” SVD $\mathbf{A} = \bar{\mathbf{U}}\text{Diag}(\bar{\mathbf{s}})\bar{\mathbf{V}}^T \in \mathbb{R}^{M \times N}$, where $\bar{\mathbf{s}} \in \mathbb{R}_+^R$ with $R = \text{rank}(\mathbf{A})$, so that $\bar{\mathbf{U}}$ and/or $\bar{\mathbf{V}}$ may be tall. The discussion in Section III-A uses the “standard” SVD $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$, where $\mathbf{S} \in \mathbb{R}^{M \times N}$ and both \mathbf{U} and \mathbf{V} are orthogonal. Finally, the state-evolution proof in Section IV uses the standard SVD on square $\mathbf{A} \in \mathbb{R}^{N \times N}$.




Fig. 2. Factor graphs to illustrate (a) messaging through a factor node and (b) messaging through a variable node.

$\int f(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \neq i}) \prod_{j \neq i} \mu_{\mathbf{x}_j \rightarrow f}(\mathbf{x}_j) d\mathbf{x}_j$. See Figure 2(c) for an illustration.

By applying the above message-passing rules to the factor graph in Figure 1, one obtains Algorithm 3. (See Appendix A for a detailed derivation.) Lines 11–12 of Algorithm 3 use

$$\mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k}) := \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I} \right)^{-1} \left(\gamma_w \mathbf{A}^T \mathbf{y} + \gamma_{2k} \mathbf{r}_{2k} \right), \quad (24)$$

which can be recognized as the MMSE estimate of a random vector \mathbf{x}_2 under likelihood $\mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}_2, \gamma_w^{-1}\mathbf{I})$ and prior $\mathbf{x}_2 \sim \mathcal{N}(\mathbf{r}_{2k}, \gamma_{2k}^{-1}\mathbf{I})$. Since this estimate is linear in \mathbf{r}_{2k} , we will refer to it as the “LMMSE” estimator. From (6)–(7) and (24), it follows that line 12 of Algorithm 3 uses

$$\langle \mathbf{g}'_2(\mathbf{r}_{2k}, \gamma_{2k}) \rangle = \frac{\gamma_{2k}}{N} \text{Tr} \left[\left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I} \right)^{-1} \right]. \quad (25)$$

Algorithm 3 is merely a restatement of VAMP Algorithm 2. Their equivalence can then be seen by substituting the “economy” SVD $\mathbf{A} = \bar{\mathbf{U}}\text{Diag}(\bar{\mathbf{s}})\bar{\mathbf{V}}^T$ into Algorithm 3, simplifying, and equating $\hat{\mathbf{x}}_k \equiv \hat{\mathbf{x}}_{1k}$, $\mathbf{r}_k \equiv \mathbf{r}_{1k}$, $\gamma_k \equiv \gamma_{1k}$, $\tilde{\gamma}_k \equiv \gamma_{2k}$, and $\alpha_k \equiv \alpha_{1k}$.

As presented in Algorithm 3, the steps of VAMP exhibit an elegant symmetry. The first half of the steps perform denoising on \mathbf{r}_{1k} and then Onsager correction in \mathbf{r}_{2k} , while the second half of the steps perform LMMSE estimation \mathbf{r}_{2k} and Onsager correction in $\mathbf{r}_{1,k+1}$.

C. Implementation Details

For practical implementation with finite-dimensional \mathbf{A} , we find that it helps to make some small enhancements to VAMP. In the discussion below we will refer to Algorithm 2, but the same approaches apply to Algorithm 3.

First, we suggest to clip the precisions γ_k and $\tilde{\gamma}_k$ to a positive interval $[\gamma_{\min}, \gamma_{\max}]$. It is possible, though uncommon, for line 6 of Algorithm 2 to return a negative α_k , which will lead to negative γ_k and $\tilde{\gamma}_k$ if not accounted for. For the numerical results in Section VI, we used $\gamma_{\min} = 1 \times 10^{-11}$ and $\gamma_{\max} = 1 \times 10^{11}$.

Second, we find that a small amount of damping can be helpful when \mathbf{A} is highly ill-conditioned. In particular, we suggest to replace lines 5 and 10 of Algorithm 2 with

Algorithm 3 Vector AMP (LMMSE Form)

Require: LMMSE estimator $\mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k})$ from (24), denoiser $\mathbf{g}_1(\cdot, \gamma_{1k})$, and number of iterations K_{it} .

- 1: Select initial \mathbf{r}_{10} and $\gamma_{10} \geq 0$.
- 2: **for** $k = 0, 1, \dots, K_{\text{it}}$ **do**
- 3: // Denoising
- 4: $\widehat{\mathbf{x}}_{1k} = \mathbf{g}_1(\mathbf{r}_{1k}, \gamma_{1k})$
- 5: $\alpha_{1k} = \langle \mathbf{g}'_1(\mathbf{r}_{1k}, \gamma_{1k}) \rangle$
- 6: $\eta_{1k} = \gamma_{1k} / \alpha_{1k}$
- 7: $\gamma_{2k} = \eta_{1k} - \gamma_{1k}$
- 8: $\mathbf{r}_{2k} = (\eta_{1k} \widehat{\mathbf{x}}_{1k} - \gamma_{1k} \mathbf{r}_{1k}) / \gamma_{2k}$
- 9:
- 10: // LMMSE estimation
- 11: $\widehat{\mathbf{x}}_{2k} = \mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k})$
- 12: $\alpha_{2k} = \langle \mathbf{g}'_2(\mathbf{r}_{2k}, \gamma_{2k}) \rangle$
- 13: $\eta_{2k} = \gamma_{2k} / \alpha_{2k}$
- 14: $\gamma_{1,k+1} = \eta_{2k} - \gamma_{2k}$
- 15: $\mathbf{r}_{1,k+1} = (\eta_{2k} \widehat{\mathbf{x}}_{2k} - \gamma_{2k} \mathbf{r}_{2k}) / \gamma_{1,k+1}$
- 16: **end for**
- 17: **Return** $\widehat{\mathbf{x}}_{1K_{\text{it}}}$.

the damped versions

$$\widehat{\mathbf{x}}_k = \rho \mathbf{g}_1(\mathbf{r}_k, \gamma_k) + (1 - \rho) \widehat{\mathbf{x}}_{k-1} \quad (26)$$

$$\gamma_{k+1} = \rho \widetilde{\gamma}_k(\mathbf{d}_k) R / (N - \langle \mathbf{d}_k \rangle R) + (1 - \rho) \gamma_k \quad (27)$$

for all iterations $k > 1$, where $\rho \in (0, 1]$ is a suitably chosen damping parameter. Note that, when $\rho = 1$, the damping has no effect. For the numerical results in Section VI, we used $\rho = 0.97$.

Third, rather than requiring VAMP to complete K_{it} iterations, we suggest that the iterations are stopped when the normalized difference $\|\mathbf{r}_{1k} - \mathbf{r}_{1,k-1}\| / \|\mathbf{r}_{1k}\|$ falls below a tolerance τ . For the numerical results in Section VI, we used $\tau = 1 \times 10^{-4}$.

We note that the three minor modifications described above are standard features of many AMP implementations, such as the one in the GAMPmatlab toolbox [33]. However, as discussed in Section I-C, they are not enough to stabilize AMP for in the case of ill-conditioned or non-zero-mean \mathbf{A} .

Finally, we note that the VAMP algorithm requires the user to choose the measurement-noise precision γ_w and the denoiser $\mathbf{g}_1(\cdot, \gamma_k)$. Ideally, the true noise precision γ_{w0} is known and the signal \mathbf{x}^0 is i.i.d. with known prior $p(x_j)$, in which case the MMSE denoiser can be straightforwardly designed. In practice, however, γ_{w0} and $p(x_j)$ are usually unknown. Fortunately, expectation-maximization (EM)-based methods to estimate these quantities on-line were proposed in [34] and shown to have a rigorous state-evolution. Furthermore, under certain identifiability conditions, the hyperparameters learned by these methods are asymptotically consistent [34]. For the numerical results in Section VI, however, we assume that γ_{w0} and $p(x_j)$ are known.

Matlab implementations of VAMP and EM-VAMP can be found in the public-domain GAMPmatlab toolbox [33].

IV. STATE EVOLUTION

A. Large-System Analysis

Our primary goal is to understand the behavior of the VAMP algorithm for a certain class of matrices in the high-dimensional regime. We begin with an overview of our analysis framework and follow with more details in later sections.

1) *Linear Measurement Model:* Our analysis considers a sequence of problems indexed by the signal dimension N . For each N , we assume that there is a “true” vector $\mathbf{x}^0 \in \mathbb{R}^N$ which is observed through measurements of the form,

$$\mathbf{y} = \mathbf{A} \mathbf{x}^0 + \mathbf{w} \in \mathbb{R}^N, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \gamma_{w0}^{-1} \mathbf{I}_N), \quad (28)$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is a known transform and \mathbf{w} is Gaussian noise with precision γ_{w0} . Note that we use γ_{w0} to denote the “true” noise precision to distinguish it from γ_w , which is the noise precision postulated by the estimator.

For the transform \mathbf{A} , our key assumption is that it can be modeled as a large, *right-orthogonally invariant* random matrix. Specifically, we assume that it has an SVD of the form

$$\mathbf{A} = \mathbf{U} \mathbf{S} \mathbf{V}^T, \quad \mathbf{S} = \text{Diag}(\mathbf{s}), \quad (29)$$

where \mathbf{U} and \mathbf{V} are $N \times N$ orthogonal matrices such that \mathbf{U} is deterministic and \mathbf{V} is Haar distributed (i.e. uniformly distributed on the set of orthogonal matrices). We refer to \mathbf{A} as “right-orthogonally invariant” because the distribution of \mathbf{A} is identical to that of $\mathbf{A} \mathbf{V}_0$ for any fixed orthogonal matrix \mathbf{V}_0 . We will discuss the distribution of the singular values $\mathbf{s} \in \mathbb{R}^N$ below.

Although we have assumed that \mathbf{A} is square to streamline the analysis, we make this assumption without loss of generality. For example, by setting

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_0 \\ \mathbf{0} \end{bmatrix},$$

our formulation can model a wide rectangular matrix whose SVD is $\mathbf{U}_0 \mathbf{S}_0 \mathbf{V}^T$ with $\text{diag}(\mathbf{S}_0) = \mathbf{s}_0$. A similar manipulation allows us to model a tall rectangular matrix.

2) *Denoiser:* Our analysis applies to a fairly general class of denoising functions $\mathbf{g}_1(\cdot, \gamma_{1k})$ indexed by the parameter $\gamma_{1k} \geq 0$. Our main assumption is that the denoiser is separable, meaning that it is of the form (8) for some scalar denoiser $g_1(\cdot, \gamma_{1k})$. As discussed above, this separability assumption will occur for the MAP and MMSE denoisers under the assumption of an i.i.d. prior. However, we do not require the denoiser to be MAP or MMSE for any particular prior. We will impose certain Lipschitz continuity conditions on $g_1(\cdot, \gamma_{1k})$ in the sequel.

3) *Asymptotic Distributions:* It remains to describe the distributions of the true vector \mathbf{x}^0 and the singular-value vector \mathbf{s} . A simple model would be to assume that they are random i.i.d. sequences that grow with N . However, following the Bayati-Montanari analysis [4], we will consider a more general framework where each of these vectors is modeled as deterministic sequence for which the empirical distribution of the components converges in distribution. When the vectors \mathbf{x}^0 and \mathbf{s} are i.i.d. random sequences, they will satisfy this

condition almost surely. Details of this analysis framework are reviewed in Appendix B.

Using the definitions in Appendix B, we assume that the components of the singular-value vector $\mathbf{s} \in \mathbb{R}^N$ in (29) converge empirically with second-order moments as

$$\lim_{N \rightarrow \infty} \{s_n\}_{n=1}^N \stackrel{PL(2)}{=} S, \quad (30)$$

for some positive random variable S . We assume that $\mathbb{E}[S] > 0$ and $S \in [0, S_{max}]$ for some finite maximum value S_{max} . Additionally, we assume that the components of the true vector, \mathbf{x}^0 , and the initial input to the denoiser, \mathbf{r}_{10} , converge empirically as

$$\lim_{N \rightarrow \infty} \{(r_{10,n}, x_n^0)\}_{n=1}^N \stackrel{PL(2)}{=} (R_{10}, X^0), \quad (31)$$

for some random variables (R_{10}, X^0) . Note that the convergence with second-order moments requires that $\mathbb{E}[(X^0)^2] < \infty$ and $\mathbb{E}[R_{10}^2] < \infty$, so they have bounded second moments. We also assume that the initial second-order term, if dependent on N , converges as

$$\lim_{N \rightarrow \infty} \gamma_{10}(N) = \bar{\gamma}_{10}, \quad (32)$$

for some $\bar{\gamma}_{10} > 0$.

As stated above, most of our analysis will apply to general separable denoisers $g_1(\cdot, \gamma_{1k})$. However, some results will apply specifically to MMSE denoisers. Under the assumption that the components of the true vector \mathbf{x}^0 are asymptotically distributed like the random variable X^0 , as in (31), the MMSE denoiser (14) and its derivative (16) reduce to

$$\begin{aligned} g_1(r_1, \gamma_1) &= \mathbb{E}\left[X^0 \middle| R_1 = r_1\right], \\ g'_1(r_1, \gamma_1) &= \gamma_1 \text{var}\left[X^0 \middle| R_1 = r_1\right], \end{aligned} \quad (33)$$

where R_1 is the random variable representing X^0 corrupted by AWGN noise, i.e.,

$$R_1 = X^0 + P, \quad P \sim \mathcal{N}(0, \gamma_1^{-1}),$$

with P being independent of X^0 . Thus, the MMSE denoiser and its derivative can be computed from the posterior mean and variance of X^0 under an AWGN measurement.

B. Error Functions

Before describing the state evolution (SE) equations and the analysis in the LSL, we need to introduce two key functions: *error functions* and *sensitivity functions*. We begin by describing the error functions.

The error functions, in essence, describe the mean squared error (MSE) of the denoiser and LMMSE estimators under AWGN measurements. Recall from Section IV-A, that we have assumed that the denoiser $\mathbf{g}_1(\cdot, \gamma_1)$ is separable with some componentwise function $g_1(\cdot, \gamma_1)$. For this function $g_1(\cdot, \gamma_1)$, define the error function as

$$\begin{aligned} \mathcal{E}_1(\gamma_1, \tau_1) &:= \mathbb{E}\left[(g_1(R_1, \gamma_1) - X^0)^2\right], \\ R_1 &= X^0 + P, \quad P \sim \mathcal{N}(0, \tau_1). \end{aligned} \quad (34)$$

The function $\mathcal{E}_1(\gamma_1, \tau_1)$ thus represents the MSE of the estimate $\hat{X} = g_1(R_1, \gamma_1)$ from a measurement R_1 corrupted by

Gaussian noise of variance τ_1 . For the LMMSE estimator, we define the error function as

$$\begin{aligned} \mathcal{E}_2(\gamma_2, \tau_2) &:= \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}\left[\|\mathbf{g}_2(\mathbf{r}_2, \gamma_2) - \mathbf{x}^0\|^2\right], \\ \mathbf{r}_2 &= \mathbf{x}^0 + \mathbf{q}, \quad \mathbf{q} \sim \mathcal{N}(0, \tau_2 \mathbf{I}), \\ \mathbf{y} &= \mathbf{A}\mathbf{x}^0 + \mathbf{w}, \quad \mathbf{w} \sim \mathcal{N}(0, \gamma_{w0}^{-1} \mathbf{I}), \end{aligned} \quad (35)$$

which is the average per component error of the vector estimate under Gaussian noise. Note that $\mathcal{E}_2(\gamma_2, \tau_2)$ is implicitly a function of the noise precision levels γ_{w0} and γ_w (through \mathbf{g}_2 from (24)), but this dependence is omitted to simplify the notation.

We will say that both estimators are “matched” when

$$\tau_1 = \gamma_1^{-1}, \quad \tau_2 = \gamma_2^{-1}, \quad \gamma_w = \gamma_{w0},$$

so that the noise levels used by the estimators both match the true noise levels. Under the matched condition, we will use the simplified notation

$$\mathcal{E}_1(\gamma_1) := \mathcal{E}_1(\gamma_1, \gamma_1^{-1}), \quad \mathcal{E}_2(\gamma_2) := \mathcal{E}_2(\gamma_2, \gamma_2^{-1}).$$

The following lemma establishes some basic properties of the error functions.

Lemma 1. Recall the error functions $\mathcal{E}_1, \mathcal{E}_2$ defined above.

(a) For the MMSE denoiser (33) under the matched condition $\tau_1 = \gamma_1^{-1}$, the error function is the conditional variance

$$\mathcal{E}_1(\gamma_1) = \text{var}\left[X^0 \middle| R_1 = X^0 + P\right], \quad P \sim \mathcal{N}(0, \gamma_1^{-1}). \quad (36)$$

(b) The LMMSE error function is given by

$$\mathcal{E}_2(\gamma_2, \tau_2) = \lim_{N \rightarrow \infty} \frac{1}{N} \text{Tr}\left[\mathbf{Q}^{-2} \tilde{\mathbf{Q}}\right], \quad (37)$$

where \mathbf{Q} and $\tilde{\mathbf{Q}}$ are the matrices

$$\mathbf{Q} := \gamma_w \mathbf{A}^T \mathbf{A} + \gamma_2 \mathbf{I}, \quad \tilde{\mathbf{Q}} := \frac{\gamma_w^2}{\gamma_{w0}} \mathbf{A}^T \mathbf{A} + \tau_2 \gamma_2^2 \mathbf{I}. \quad (38)$$

Under the matched condition $\tau_2 = \gamma_2^{-1}$ and $\gamma_w = \gamma_{w0}$,

$$\mathcal{E}_2(\gamma_2) = \lim_{N \rightarrow \infty} \frac{1}{N} \text{Tr}\left[\mathbf{Q}^{-1}\right]. \quad (39)$$

(c) The LMMSE error function is also given by

$$\mathcal{E}_2(\gamma_2, \tau_2) = \mathbb{E}\left[\frac{\gamma_w^2 S^2 / \gamma_{w0} + \tau_2 \gamma_2^2}{(\gamma_w S^2 + \gamma_2)^2}\right], \quad (40)$$

where S is the random variable (30) representing the distribution of the singular values of \mathbf{A} . For the matched condition $\tau_2 = \gamma_2^{-1}$ and $\gamma_w = \gamma_{w0}$,

$$\mathcal{E}_2(\gamma_2) = \mathbb{E}\left[\frac{1}{\gamma_w S^2 + \gamma_2}\right]. \quad (41)$$

Proof. See Appendix C. \square

C. Sensitivity Functions

The sensitivity functions describe the expected divergence of the estimator. For the denoiser, the sensitivity function is defined as

$$\begin{aligned} A_1(\gamma_1, \tau_1) &:= \mathbb{E} \left[g'_1(R_1, \gamma_1) \right], \\ R_1 &= X^0 + P, \quad P \sim \mathcal{N}(0, \tau_1), \end{aligned} \quad (42)$$

which is the average derivative under a Gaussian noise input. For the LMMSE estimator, the sensitivity is defined as

$$A_2(\gamma_2) := \lim_{N \rightarrow \infty} \frac{1}{N} \text{Tr} \left[\frac{\partial \mathbf{g}_2(\mathbf{r}_2, \gamma_2)}{\partial \mathbf{r}_2} \right]. \quad (43)$$

Lemma 2. *For the sensitivity functions above:*

(a) *For the MMSE denoiser (33) under the matched condition $\tau_1 = \gamma_1^{-1}$, the sensitivity function is given by*

$$A_1(\gamma_1, \gamma_1^{-1}) = \gamma_1 \text{var} \left[X^0 \middle| R_1 = X^0 + \mathcal{N}(0, \gamma_1^{-1}) \right], \quad (44)$$

which is the ratio of the conditional variance to the measurement variance γ_1^{-1} .

(b) *The LMMSE estimator's sensitivity function is given by*

$$A_2(\gamma_2) = \lim_{N \rightarrow \infty} \frac{1}{N} \gamma_2 \text{Tr} \left[(\gamma_w \mathbf{A}^\top \mathbf{A} + \gamma_2 \mathbf{I})^{-1} \right].$$

(c) *The LMMSE estimator's sensitivity function can also be written as*

$$A_2(\gamma_2) = \mathbb{E} \left[\frac{\gamma_2}{\gamma_w S^2 + \gamma_2} \right].$$

Proof. See Appendix C. \square

D. State Evolution Equations

We can now describe our main result, which is the SE equations for VAMP. For a given iteration $k \geq 1$, consider the set of components,

$$\{(\hat{x}_{1k,n}, r_{1k,n}, x_n^0), \quad n = 1, \dots, N\}.$$

This set represents the components of the true vector \mathbf{x}^0 , its corresponding estimate $\hat{\mathbf{x}}_{1k}$ and the denoiser input \mathbf{r}_{1k} . Theorem 1 below will show that, under certain assumptions, these components converge empirically as

$$\lim_{N \rightarrow \infty} \{(\hat{x}_{1k,n}, r_{1k,n}, x_n^0)\} \stackrel{PL(2)}{=} (\hat{X}_{1k}, R_{1k}, X^0), \quad (45)$$

where the random variables $(\hat{X}_{1k}, R_{1k}, X^0)$ are given by

$$R_{1k} = X^0 + P_k, \quad P_k \sim \mathcal{N}(0, \tau_{1k}), \quad (46a)$$

$$\hat{X}_{1k} = g_1(R_{1k}, \bar{\gamma}_{1k}), \quad (46b)$$

for constants $\bar{\gamma}_{1k}$ and τ_{1k} that will be defined below. Thus, each component $r_{1k,n}$ appears as the true component x_n^0 plus Gaussian noise. The corresponding estimate $\hat{x}_{1k,n}$ then appears as the denoiser output with $r_{1k,n}$ as the input. Hence, the asymptotic behavior of any component x_n^0 and its corresponding $\hat{x}_{1k,n}$ is identical to a simple scalar system. We will refer to (45)-(46) as the denoiser's *scalar equivalent model*.

For the LMMSE estimation function, we define the transformed error and transformed noise,

$$\mathbf{q}_k := \mathbf{V}^\top (\mathbf{r}_{2k} - \mathbf{x}^0), \quad \boldsymbol{\xi} := \mathbf{U}^\top \mathbf{w}, \quad (47)$$

where \mathbf{U} and \mathbf{V} are the matrices in the SVD decomposition (29). Theorem 1 will also show that these transformed errors and singular values s_n converge as,

$$\lim_{N \rightarrow \infty} \{(q_{k,n}, \xi_n, s_n)\} \stackrel{PL(2)}{=} (Q_k, \Xi, S), \quad (48)$$

to a set of random variables (Q_k, Ξ, S) . These random variables are independent, with S defined in the limit (30) and

$$Q_k \sim \mathcal{N}(0, \tau_{2k}), \quad \Xi \sim \mathcal{N}(0, \gamma_{w0}^{-1}), \quad (49)$$

where τ_{2k} is a variance that will be defined below and γ_{w0} is the noise precision in the measurement model (28). Thus (48)-(49) is a scalar equivalent model for the LMMSE estimator.

The variance terms are defined recursively through what are called *state evolution* equations,

$$\bar{\alpha}_{1k} = A_1(\bar{\gamma}_{1k}, \tau_{1k}) \quad (50a)$$

$$\bar{\eta}_{1k} = \frac{\bar{\gamma}_{1k}}{\bar{\alpha}_{1k}}, \quad \bar{\gamma}_{2k} = \bar{\eta}_{1k} - \bar{\gamma}_{1k} \quad (50b)$$

$$\tau_{2k} = \frac{1}{(1 - \bar{\alpha}_{1k})^2} \left[\mathcal{E}_1(\bar{\gamma}_{1k}, \tau_{1k}) - \bar{\alpha}_{1k}^2 \tau_{1k} \right], \quad (50c)$$

$$\bar{\alpha}_{2k} = A_2(\bar{\gamma}_{2k}, \tau_{2k}) \quad (50d)$$

$$\bar{\eta}_{2k} = \frac{\bar{\gamma}_{2k}}{\bar{\alpha}_{2k}}, \quad \bar{\gamma}_{1,k+1} = \bar{\eta}_{2k} - \bar{\gamma}_{2k} \quad (50e)$$

$$\tau_{1,k+1} = \frac{1}{(1 - \bar{\alpha}_{2k})^2} \left[\mathcal{E}_2(\bar{\gamma}_{2k}, \tau_{2k}) - \bar{\alpha}_{2k}^2 \tau_{2k} \right], \quad (50f)$$

which are initialized with

$$\tau_{10} = \mathbb{E}[(R_{10} - X^0)^2], \quad (51)$$

and $\bar{\gamma}_{10}$ defined from the limit (32).

Theorem 1. *Under the above assumptions and definitions, assume additionally that for all iterations k :*

(i) *The solution $\bar{\alpha}_{1k}$ from the SE equations (50) satisfies*

$$\bar{\alpha}_{1k} \in (0, 1). \quad (52)$$

(ii) *The functions $A_i(\gamma_i, \tau_i)$ and $\mathcal{E}_i(\gamma_i, \tau_i)$ are continuous at $(\gamma_i, \tau_i) = (\bar{\gamma}_{ik}, \tau_{ik})$.*

(iii) *The denoiser function $g_1(r_1, \gamma_1)$ and its derivative $g'_1(r_1, \gamma_1)$ are uniformly Lipschitz in r_1 at $\gamma_1 = \bar{\gamma}_{1k}$. (See Appendix B for a precise definition of uniform Lipschitz continuity.)*

Then, for any fixed iteration $k \geq 0$,

$$\lim_{N \rightarrow \infty} (\alpha_{ik}, \eta_{ik}, \gamma_{ik}) = (\bar{\alpha}_{ik}, \bar{\eta}_{ik}, \bar{\gamma}_{ik}) \quad (53)$$

almost surely. In addition, the empirical limit (45) holds almost surely for all $k > 0$, and (48) holds almost surely for all $k \geq 0$.

Proof. See Appendix G.

E. Mean Squared Error

One important use of the scalar equivalent model is to predict the asymptotic performance of the VAMP algorithm in the LSL. For example, define the asymptotic mean squared error (MSE) of the iteration- k estimate $\widehat{\mathbf{x}}_{ik}$ as

$$\text{MSE}_{ik} := \lim_{N \rightarrow \infty} \frac{1}{N} \|\widehat{\mathbf{x}}_{ik} - \mathbf{x}^0\|^2. \quad (54)$$

For this MSE, we claim that

$$\text{MSE}_{ik} = \mathcal{E}_i(\overline{\gamma}_{ik}, \tau_{ik}). \quad (55)$$

To prove (55) for $i = 1$, we write

$$\begin{aligned} \text{MSE}_{1k} &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N (\widehat{x}_{1k,n} - x_n^0)^2 \\ &\stackrel{(a)}{=} \mathbb{E}[(\widehat{X}_{1k} - X^0)^2] \\ &\stackrel{(b)}{=} \mathbb{E}[(g_1(R_1, \overline{\gamma}_{1k}) - X^0)^2] \stackrel{(c)}{=} \mathcal{E}_1(\overline{\gamma}_{1k}, \tau_{1k}) \end{aligned}$$

where (a) and (b) follow from the convergence in (45) and the scalar equivalent model (45), and where (c) follows from (34). Using the scalar equivalent model (48), the definition of $\mathcal{E}_2(\cdot)$ in (35), and calculations similar to the proof of Lemma 1, one can also show that (55) holds for $i = 2$.

Interestingly, this type of calculation can be used to compute any other componentwise distortion metric. Specifically, given any distortion function $d(x, \widehat{x})$ that is pseudo-Lipschitz of order two, its average value is given by

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N d(x_n^0, \widehat{x}_{1k,n}) = \mathbb{E}[d(X^0, \widehat{X}_{1k})],$$

where the expectation is from the scalar equivalent model (45).

F. Contractiveness of the Denoiser

An essential requirement of Theorem 1 is the condition (52) that $\overline{\alpha}_{1k} \in (0, 1)$. This assumption requires that, in a certain average, the denoiser function $g_1(\cdot, \gamma_1)$ is increasing (i.e., $g_1'(r_{1n}, \gamma_1) > 0$) and is a contraction (i.e., $g_1'(r_{1n}, \gamma_1) < 1$). If these conditions are not met, then $\overline{\alpha}_{1k} \leq 0$ or $\overline{\alpha}_{1k} \geq 1$, and either the estimated precision $\overline{\eta}_{1k}$ or $\overline{\gamma}_{2k}$ in (50b) may be negative, causing subsequent updates to be invalid. Thus, $\overline{\alpha}_{1k}$ must be in the range $(0, 1)$. There are two important conditions under which this increasing contraction property are provably guaranteed:

1) *Strongly Convex Penalties*: Suppose that $g_1(r_{1n}, \gamma_1)$ is either the MAP denoiser (13) or the MMSE denoiser (14) for a density $p(x_n)$ that is strongly log-concave. That is, there exists constants $c_1, c_2 > 0$ such that

$$c_1 \leq -\frac{\partial^2}{\partial x_n^2} \ln p(x_n) \leq c_2.$$

Then, using results from log-concave functions [35], it is shown in [11] that

$$g_1'(r_{1n}, \gamma_1) \in \left[\frac{\gamma_1}{c_2 + \gamma_1}, \frac{\gamma_1}{c_1 + \gamma_1} \right] \subset (0, 1),$$

for all r_{1n} and $\gamma_1 > 0$. Hence, from the definition of the sensitivity function (42), the sensitivity $\overline{\alpha}_{1k}$ in (50a) will be in the range $(0, 1)$.

2) *Matched MMSE Denoising*: Suppose that $g_1(r_{1n}, \gamma_1)$ is the MMSE denoiser in the matched condition where $\overline{\gamma}_{1k} = \tau_{1k}^{-1}$ for some iteration k . From (44),

$$A_1(\gamma_1, \gamma_1^{-1}) = \gamma_1 \text{var} \left[X^0 \middle| R_1 = X^0 + \mathcal{N}(0, \gamma_1^{-1}) \right].$$

Since the conditional variance is positive, $A_1(\gamma_1, \gamma_1^{-1}) > 0$. Also, since the variance is bounded above by the MSE of a linear estimator,

$$\begin{aligned} \gamma_1 \text{var} \left[X^0 \middle| R_1 = X^0 + \mathcal{N}(0, \gamma_1^{-1}) \right] \\ \leq \gamma_1 \frac{\gamma_1^{-1} \tau_{x_0}}{\tau_{x_0} + \gamma_1^{-1}} = \frac{\gamma_1 \tau_{x_0}}{1 + \gamma_1 \tau_{x_0}} < 1, \end{aligned}$$

where $\tau_{x_0} = \text{var}(X^0)$. Thus, we have $A_1(\gamma_1, \gamma_1^{-1}) \in (0, 1)$ and $\overline{\alpha}_{1k} \in (0, 1)$.

In the case when the prior is not log-concave and the estimator uses an denoiser that is not perfectly matched, $\overline{\alpha}_{1k}$ may not be in the valid range $(0, 1)$. In these cases, VAMP may obtain invalid (i.e. negative) variance estimates.

V. MMSE DENOISING, OPTIMALITY, AND CONNECTIONS TO THE REPLICA METHOD

An important special case of the VAMP algorithm is when we apply the MMSE optimal denoiser under matched γ_w . In this case, the SE equations simplify considerably.

Theorem 2. *Consider the SE equations (50) with the MMSE optimal denoiser (33), matched $\gamma_w = \gamma_{w0}$, and matched initial condition $\overline{\gamma}_{10} = \tau_{10}^{-1}$. Then, for all iterations $k \geq 0$,*

$$\overline{\eta}_{1k} = \frac{1}{\mathcal{E}_1(\overline{\gamma}_{1k})}, \quad \overline{\gamma}_{2k} = \tau_{2k}^{-1} = \overline{\eta}_{1k} - \overline{\gamma}_{1k}, \quad (56a)$$

$$\overline{\eta}_{2k} = \frac{1}{\mathcal{E}_2(\overline{\gamma}_{2k})}, \quad \overline{\gamma}_{1,k+1} = \tau_{1,k+1}^{-1} = \overline{\eta}_{2k} - \overline{\gamma}_{2k}. \quad (56b)$$

In addition, for estimators $i = 1, 2$, $\overline{\eta}_{ik}$ is the inverse MSE:

$$\overline{\eta}_{ik}^{-1} = \lim_{N \rightarrow \infty} \frac{1}{N} \|\widehat{\mathbf{x}}_{ik} - \mathbf{x}^0\|^2. \quad (56c)$$

Proof. See Appendix H. \square

It is useful to compare this result with the work [21], which uses the *replica method* from statistical physics to predict the asymptotic MMSE in the LSL. To state the result, given a positive semidefinite matrix \mathbf{C} , we define its Stieltjes transform as

$$S_{\mathbf{C}}(\omega) = \frac{1}{N} \text{Tr} \left[(\mathbf{C} - \omega \mathbf{I}_N)^{-1} \right] = \frac{1}{N} \sum_{n=1}^N \frac{1}{\lambda_n - \omega}, \quad (57)$$

where λ_n are the eigenvalues of \mathbf{C} . Also, let $R_{\mathbf{C}}(\omega)$ denote the so-called R -transform of \mathbf{C} , given by

$$R_{\mathbf{C}}(\omega) = S_{\mathbf{C}}^{-1}(-\omega) - \frac{1}{\omega}, \quad (58)$$

where the inverse $S_{\mathbf{C}}^{-1}(\cdot)$ is in terms of composition of functions. The Stieltjes and R -transforms are discussed in detail in [36]. The Stieltjes and R -transforms can be extended

to random matrix sequences by taking limits as $N \rightarrow \infty$ (for matrix sequences where these limits converge almost surely).

Now suppose that $\hat{\mathbf{x}} = \mathbb{E}[\mathbf{x}^0|\mathbf{y}]$ is the MMSE estimate of \mathbf{x}^0 given \mathbf{y} . Let $\bar{\eta}^{-1}$ be the asymptotic inverse MSE

$$\bar{\eta}^{-1} := \lim_{N \rightarrow \infty} \frac{1}{N} \|\hat{\mathbf{x}} - \mathbf{x}^0\|^2.$$

Using a so-called replica symmetric analysis, it is argued in [21] that this MSE should satisfy the fixed point equations

$$\bar{\gamma}_1 = R_C(-\bar{\eta}^{-1}), \quad \bar{\eta}^{-1} = \mathcal{E}_1(\bar{\gamma}_1), \quad (59)$$

where $\mathbf{C} = \gamma_{w0} \mathbf{A}^\top \mathbf{A}$. A similar result is given in [16].

Theorem 3. *Let $\bar{\gamma}_i, \bar{\eta}_i$ be any fixed point solutions to the SE equations (56) of VAMP under MMSE denoising and matched $\gamma_w = \gamma_{w0}$. Then $\bar{\eta}_1 = \bar{\eta}_2$. If we define $\bar{\eta} := \bar{\eta}_i$ as the common value, then $\bar{\gamma}_1$ and $\bar{\eta}$ satisfy the replica fixed point equation (59).*

Proof. Note that we have dropped the iteration index k since we are discussing a fixed point. First, (56) shows that, at any fixed point,

$$\bar{\gamma}_1 + \bar{\gamma}_2 = \bar{\eta}_1 = \bar{\eta}_2,$$

so that $\bar{\eta}_1 = \bar{\eta}_2$. Also, in the matched case, (41) shows that

$$\mathcal{E}_2(\bar{\gamma}_2) = S_C(-\bar{\gamma}_2).$$

Since $\bar{\eta}^{-1} = \mathcal{E}_2(\bar{\gamma}_2)$, we have that

$$\bar{\gamma}_1 = \bar{\eta} - \bar{\gamma}_2 = \bar{\eta} + S_C^{-1}(\bar{\eta}^{-1}) = R_C(-\bar{\eta}^{-1}).$$

Also, $\bar{\eta}^{-1} = \bar{\eta}_1^{-1} = \mathcal{E}(\bar{\gamma}_1)$. \square

The consequence of Theorem 3 is that, if the replica equations (59) have a unique fixed point, then the MSE achieved by the VAMP algorithm exactly matches the Bayes optimal MSE as predicted by the replica method. Hence, if this replica prediction is correct, then the VAMP method provides a computationally efficient method for finding MSE optimal estimates under very general priors—including priors for which the associated penalty functions are not convex.

The replica method, however, is generally heuristic. But in the case of i.i.d. Gaussian matrices, it has recently been proven that the replica prediction is correct [37], [38].

VI. NUMERICAL EXPERIMENTS

In this section, we present numerical experiments that compare the VAMP⁴ Algorithm 2 to the VAMP state evolution from Section IV, the replica prediction from [21], the AMP Algorithm 1 from [3], the S-AMP algorithm from [18, Sec. IV], the adaptively damped (AD) GAMP algorithm from [9], and the support-oracle MMSE estimator, whose MSE lower bounds that achievable by any practical method. In all cases, we consider the recovery of vectors $\mathbf{x}^0 \in \mathbb{R}^N$ from AWGN-corrupted measurements $\mathbf{y} \in \mathbb{R}^M$ constructed from (1), where \mathbf{x}^0 was drawn i.i.d. zero-mean Bernoulli-Gaussian with $\Pr\{x_j^0 \neq 0\} = 0.1$, where $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}/\gamma_{w0})$, and where $M = 512$ and $N = 1024$. All

⁴A Matlab implementation of VAMP can be found in the public-domain GAMPmatlab toolbox [33].




Fig. 3. NMSE versus condition number $\kappa(\mathbf{A})$ at final algorithm iteration. The reported NMSE is the median over 500 realizations, with error bars shown on the VAMP trace.

methods under test were matched to the true signal and noise statistics. When computing the support-oracle MMSE estimate, the support of \mathbf{x}^0 is assumed to be known, in which case the problem reduces to estimating the non-zero coefficients of \mathbf{x}^0 . Since these non-zero coefficients are Gaussian, their MMSE estimate can be computed in closed form. For VAMP we used the implementation enhancements described in Section III-C. For line 7 of AMP Algorithm 1, we used $1/\gamma_{k+1} = 1/\gamma_{w0} + \frac{N}{M} \alpha_k / \gamma_k$, as specified in [3, Eq. (25)]. For the AMP, S-AMP, and AD-GAMP algorithms, we allowed a maximum of 1000 iterations, and for the VAMP algorithm we allowed a maximum of 100 iterations.

A. Ill-Conditioned \mathbf{A}

First we investigate algorithm robustness to the condition number of \mathbf{A} . For this study, realizations of \mathbf{A} were constructed from the SVD $\mathbf{A} = \bar{\mathbf{U}} \text{Diag}(\bar{\mathbf{s}}) \bar{\mathbf{V}}^\top \in \mathbb{R}^{M \times N}$ with geometric singular values $\bar{\mathbf{s}} \in \mathbb{R}^M$. That is, $\bar{s}_i / \bar{s}_{i-1} = \rho \forall i$, with ρ chosen to achieve a desired condition number $\kappa(\mathbf{A}) := \bar{s}_1 / \bar{s}_M$ and with \bar{s}_1 chosen so that $\|\mathbf{A}\|_F^2 = N$. The singular vector matrices $\bar{\mathbf{U}}, \bar{\mathbf{V}}$ were drawn uniformly at random from the group of orthogonal matrices, i.e., from the Haar distribution. Finally, the signal and noise variances were set to achieve a signal-to-noise ratio (SNR) $\mathbb{E}[\|\mathbf{A}\mathbf{x}\|^2] / \mathbb{E}[\|\mathbf{w}\|^2]$ of 40 dB.

Figure 3 plots the median normalized MSE (NMSE) achieved by each algorithm over 500 independent realizations of $\{\mathbf{A}, \mathbf{x}, \mathbf{w}\}$, where $\text{NMSE}(\hat{\mathbf{x}}) := \|\hat{\mathbf{x}} - \mathbf{x}^0\|^2 / \|\mathbf{x}^0\|^2$. To enhance visual clarity, NMSEs were clipped to a maximum value of 1. Also, error bars are shown that (separately) quantify the positive and negative standard deviations of VAMP's NMSE from the median value. The NMSE was evaluated for condition numbers $\kappa(\mathbf{A})$ ranging from 1 (i.e., row-orthogonal \mathbf{A}) to 1×10^6 (i.e., highly ill-conditioned \mathbf{A}).

In Figure 3, we see that AMP and S-AMP diverged for even mildly ill-conditioned \mathbf{A} . We also see that, while adaptive damping helped to extend the operating range of AMP, it had a limited effect. In contrast, Figure 3 shows that VAMP's

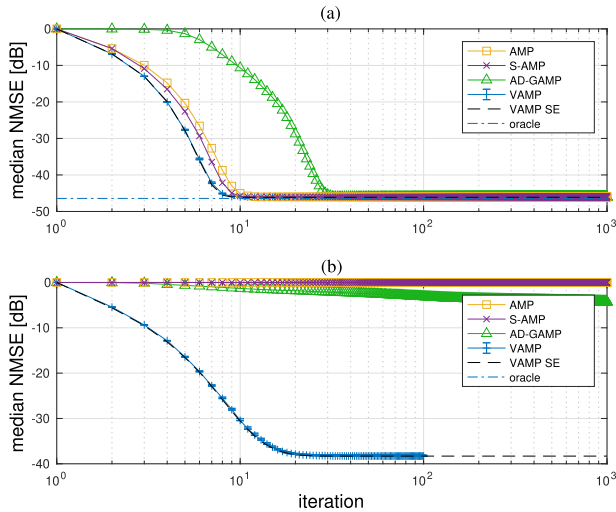


Fig. 4. NMSE versus algorithm iteration for condition number $\kappa(\mathbf{A}) = 1$ in (a) and $\kappa(\mathbf{A}) = 1000$ in (b). The reported NMSE is the median over 500 realizations, with error bars shown on the VAMP traces.

NMSE stayed relatively close to the replica prediction for all condition numbers $\kappa(\mathbf{A})$. The small gap between VAMP and the replica prediction is due to finite-dimensional effects; the SE analysis from Section IV establishes that this gap closes in the large-system limit. Finally, Figure 3 shows that the oracle bound is close to the replica prediction at small $\kappa(\mathbf{A})$ but not at large $\kappa(\mathbf{A})$.

Figure 4(a) plots NMSE versus algorithm iteration for condition number $\kappa(\mathbf{A}) = 1$ and Figure 4(b) plots the same for $\kappa(\mathbf{A}) = 1000$, again with error bars on the VAMP traces. Both figures show that the VAMP trajectory stayed very close to the VAMP-SE trajectory at every iteration. The figures also show that VAMP converges a bit quicker than AMP, S-AMP, and AD-GAMP when $\kappa(\mathbf{A}) = 1$, and that VAMP's convergence rate is relatively insensitive to the condition number $\kappa(\mathbf{A})$.

B. Non-Zero-Mean \mathbf{A}

In this section, we investigate algorithm robustness to the componentwise mean of \mathbf{A} . For this study, realizations of \mathbf{A} were constructed by first drawing an i.i.d. $\mathcal{N}(\mu, 1/M)$ matrix and then scaling it so that $\|\mathbf{A}\|_F^2 = N$ (noting that essentially no scaling is needed when $\mu \approx 0$). As before, the signal and noise variances were set to achieve an SNR of 40 dB. For AD-GAMP, we used the mean-removal trick proposed in [9].

Figure 5 plots the NMSE achieved by each algorithm over 200 independent realizations of $\{\mathbf{A}, \mathbf{x}, \mathbf{w}\}$. The NMSE was evaluated for mean parameters μ between 0.001 and 10. Note that, when $\mu > 0.044$, the mean is larger than the standard deviation. Thus, the values of μ that we consider are quite extreme relative to past studies like [8].

Figure 5 shows that AMP and S-AMP diverged for even mildly mean-perturbed \mathbf{A} . In contrast, the figure shows that VAMP and mean-removed AD-GAMP (MAD-GAMP) closely matched the replica prediction for all mean parameters μ . It also shows a relatively small gap between the replica prediction and the oracle bound, especially for small μ .




Fig. 5. NMSE versus mean μ at final algorithm iteration. The reported NMSE is the median over 200 realizations, with error bars shown on the VAMP trace.

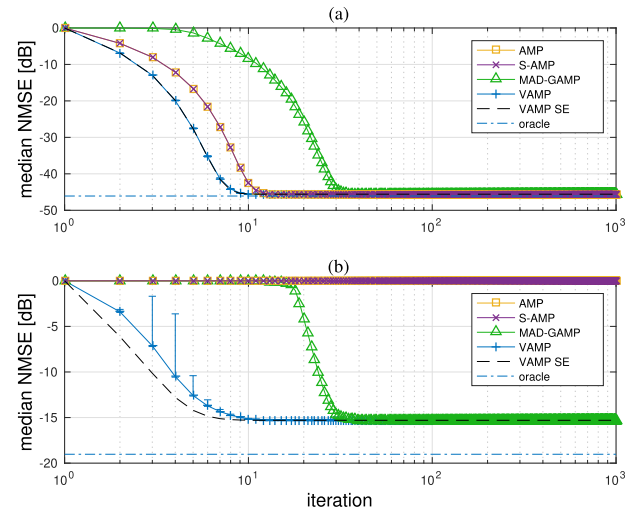


Fig. 6. NMSE versus algorithm iteration when \mathbf{A} has mean $\mu = 0.001$ in (a) and $\mu = 1$ in (b). The reported NMSE is the median over 200 realizations, with error bars shown on the VAMP traces.

Figure 6(a) plots NMSE versus algorithm iteration for matrix mean $\mu = 0.001$ and Figure 6(b) plots the same for $\mu = 1$. When $\mu = 0.001$, VAMP closely matched its SE at all iterations and converged noticeably quicker than AMP, S-AMP, and MAD-VAMP. When $\mu = 1$, there was a small but noticeable gap between VAMP and its SE for the first few iterations, although the gap closed after about 10 iterations. This gap may be due to the fact that the random matrix \mathbf{A} used for this experiment was not right-orthogonally invariant, since the dominant singular vectors are close to (scaled versions of) the $\mathbf{1}_s$ vector for sufficiently large μ .

C. Row-Orthogonal \mathbf{A}

In this section we investigate algorithm NMSE versus SNR for row-orthogonal \mathbf{A} , i.e., \mathbf{A} constructed as in Section VI-A but with $\kappa(\mathbf{A}) = 1$. Previous studies [16], [19] have demonstrated that, when \mathbf{A} is orthogonally invariant but not i.i.d.

TABLE I

AVERAGE NMSE VERSUS SNR FOR ROW-ORTHOGONAL \mathbf{A} , WHERE THE AVERAGE WAS COMPUTED FROM 1000 REALIZATIONS. STANDARD ERROR DEVIATIONS ARE ALSO REPORTED

SNR	replica	VAMP(stderr)	S-AMP(stderr)	AMP(stderr)
10 dB	5.09e-02	5.27e-02(4.3e-04)	5.27e-02(4.3e-04)	5.42e-02(4.2e-04)
20 dB	3.50e-03	3.57e-03(2.7e-05)	3.58e-03(2.7e-05)	3.62e-03(2.6e-05)
30 dB	2.75e-04	2.84e-04(2.2e-06)	2.85e-04(2.2e-06)	2.85e-04(2.1e-06)

Gaussian (e.g., row-orthogonal), the fixed points of S-AMP and diagonal-restricted EC are better than those of AMP because the former approaches exploit the singular-value spectrum of \mathbf{A} , whereas AMP does not.

Table I reports the NMSE achieved by VAMP, S-AMP, and AMP at three levels of SNR: 10 dB, 20 dB, and 30 dB. The NMSEs reported in the table were computed from an average of 1000 independent realizations of $\{\mathbf{A}, \mathbf{x}, \mathbf{w}\}$. Since the NMSE differences between the algorithms are quite small, the table also reports the standard error on each NMSE estimate to confirm its accuracy.

Table I shows that VAMP and S-AMP gave nearly identical NMSE at all tested SNRs, which is expected because these two algorithms share the same fixed points. The table also shows that VAMP's NMSE was strictly better than AMP's NMSE at low SNR (as expected), but that the NMSE difference narrows as the SNR increases. Finally, the table reports the replica prediction of the NMSE, which is about 3% lower (i.e., -0.15 dB) than VAMP's empirical NMSE at each SNR. We attribute this difference to finite-dimensional effects.

D. Discussion

Our numerical results confirm what is already known about the *fixed points* of diagonally restricted EC (via VAMP) and S-AMP. That is, when \mathbf{A} is large and right-orthogonally invariant, they agree with each other and with the replica prediction; and when \mathbf{A} is large i.i.d. Gaussian (which is a special case of right-orthogonally invariant [36]), they furthermore agree with the fixed points of AMP [16], [19].

But our numerical results also clarify that it is not enough for an algorithm to have good fixed points, because it may not converge to its fixed points. For example, although the fixed points of S-AMP are good (i.e., replica matching) for *any* large right-orthogonally invariant \mathbf{A} , our numerical results indicate that S-AMP converges only for a small subset of large right-orthogonally invariant \mathbf{A} : those with singular-value spectra similar (or flatter than) i.i.d. Gaussian \mathbf{A} .

The SE analysis from Section IV establishes that, in the large-system limit and under matched priors, VAMP is guaranteed to converge to a fixed point that is also a fixed point of the replica equation (59). Our numerical results suggest that, even with large but finite-dimensional right orthogonally invariant \mathbf{A} (i.e., 512×1024 in our simulations), VAMP attains NMSEs that are very close to the replica prediction.

VII. CONCLUSIONS

In this paper, we considered the standard linear regression (SLR) problem (1), where the goal is to recover the

vector \mathbf{x}^0 from noisy linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x}^0 + \mathbf{w}$. Our work is inspired by Donoho, Maleki, and Montanari's AMP algorithm [2], which offers a computationally efficient approach to SLR. AMP has the desirable property that its behavior is rigorously characterized under large i.i.d. sub-Gaussian \mathbf{A} by a scalar state evolution whose fixed points, when unique, are Bayes optimal [4]. A major shortcoming of AMP, however, is its fragility with respect to the i.i.d. sub-Gaussian model on \mathbf{A} : even small perturbations from this model can cause AMP to diverge.

In response, we proposed a vector AMP (VAMP) algorithm that (after performing an initial SVD) has similar complexity to AMP but is much more robust with respect to the matrix \mathbf{A} . Our main contribution is establishing that VAMP's behavior can be rigorously characterized by a scalar state-evolution that holds for large, right-orthogonally invariant \mathbf{A} . The fixed points of VAMP's state evolution are, in fact, consistent with the replica prediction of the minimum mean-squared error recently derived in [21]. We also showed how VAMP can be derived as an approximation of belief propagation on a factor graph with vector-valued nodes, hence the name "vector AMP." Finally, we presented numerical experiments to demonstrate VAMP's robust convergence for ill-conditioned and mean-perturbed matrices \mathbf{A} that cause earlier AMP algorithms to diverge.

As future work, it would be interesting to extend VAMP to the generalized linear model, where the outputs $\mathbf{A}\mathbf{x}^0$ are non-linearly mapped to \mathbf{y} . Also, it would be interesting to design and analyze extensions of VAMP that are robust to more general models for \mathbf{A} , such as the case where \mathbf{A} is statistically coupled to \mathbf{x}^0 .

APPENDIX A

MESSAGE-PASSING DERIVATION OF VAMP

In this appendix, we detail the message-passing derivation of Algorithm 3. Below, we will use k to denote the VAMP iteration and n to index the elements of N -dimensional vectors like $\mathbf{x}_1, \mathbf{r}_{1k}$ and $\widehat{\mathbf{x}}_{1k}$. We start by initializing the message-passing with $\mu_{\delta \rightarrow \mathbf{x}_1}(\mathbf{x}_1) = \mathcal{N}(\mathbf{x}_1; \mathbf{r}_{10}, \gamma_{10}^{-1}\mathbf{I})$. The following steps are then repeated for $k = 0, 1, 2, \dots$.

From Rule 1, we first set the approximate belief on \mathbf{x}_1 as $\mathcal{N}(\mathbf{x}_1; \widehat{\mathbf{x}}_{1k}, \eta_{1k}^{-1}\mathbf{I})$, where $\widehat{\mathbf{x}}_{1k} = \mathbb{E}[\mathbf{x}_1 | b_{\text{sp}}(\mathbf{x}_1)]$ and $\eta_{1k}^{-1} = \langle \text{diag}(\text{Cov}[\mathbf{x}_1 | b_{\text{sp}}(\mathbf{x}_1)]) \rangle$ for the SP belief $b_{\text{sp}}(\mathbf{x}_1) \propto p(\mathbf{x}_1)\mathcal{N}(\mathbf{x}_1; \mathbf{r}_{1k}, \gamma_{1k}^{-1}\mathbf{I})$. With an i.i.d. prior $p(\mathbf{x}_1)$ as in (12), we have that $[\widehat{\mathbf{x}}_{1k}]_n = g_1(r_{1k,n}, \gamma_{1k})$ for the conditional-mean estimator $g_1(\cdot, \gamma_{1k})$ given in (14), yielding line 4 of Algorithm 3. Furthermore, from (16) we see that the corresponding conditional covariance is $\gamma_{1k}^{-1} g_1'(r_{1k,n}, \gamma_{1k})$, yielding lines 5-6 of Algorithm 3.

Next, Rule 2 says to set the message $\mu_{\mathbf{x}_1 \rightarrow \delta}(\mathbf{x}_1)$ proportional to $\mathcal{N}(\mathbf{x}_1; \widehat{\mathbf{x}}_{1k}, \eta_{1k}^{-1}\mathbf{I}) / \mathcal{N}(\mathbf{x}_1; \mathbf{r}_{1k}, \gamma_{1k}^{-1}\mathbf{I})$. Since

$$\begin{aligned} & \mathcal{N}(\mathbf{x}; \widehat{\mathbf{x}}, \eta^{-1}\mathbf{I}) / \mathcal{N}(\mathbf{x}; \mathbf{r}, \gamma^{-1}\mathbf{I}) \\ & \propto \mathcal{N}(\mathbf{x}; (\widehat{\mathbf{x}}\eta - \mathbf{r}\gamma) / (\eta - \gamma), (\eta - \gamma)^{-1}\mathbf{I}), \end{aligned} \quad (60)$$

we have $\mu_{\mathbf{x}_1 \rightarrow \delta}(\mathbf{x}_1) = \mathcal{N}(\mathbf{x}_1; \mathbf{r}_{2k}, \gamma_{2k}^{-1}\mathbf{I})$ for $\mathbf{r}_{2k} = (\widehat{\mathbf{x}}_{1k}\eta_{1k} - \mathbf{r}_{1k}\gamma_{1k}) / (\eta_{1k} - \gamma_{1k})$ and $\gamma_{2k} = \eta_{1k} - \gamma_{1k}$, yielding lines 7-8 of Algorithm 3. Rule 3 then implies that the message $\mu_{\mathbf{x}_1 \rightarrow \delta}(\mathbf{x}_1)$

will flow rightward through the δ node unchanged, manifesting as $\mu_{\delta \rightarrow \mathbf{x}_2}(\mathbf{x}_2) = \mathcal{N}(\mathbf{x}_2; \mathbf{r}_{2k}, \gamma_{2k}^{-1} \mathbf{I})$ on the other side.

Rule 1 then says to set the approximate belief on \mathbf{x}_2 at $\mathcal{N}(\mathbf{x}_2; \widehat{\mathbf{x}}_{2k}, \eta_{2k}^{-1} \mathbf{I})$, where $\widehat{\mathbf{x}}_{2k} = \mathbb{E}[\mathbf{x}_2 | b_{\text{sp}}(\mathbf{x}_2)]$ and $\eta_{2k}^{-1} = \langle \text{diag}(\text{Cov}[\mathbf{x}_2 | b_{\text{sp}}(\mathbf{x}_2)]) \rangle$ for the SP belief $b_{\text{sp}}(\mathbf{x}_2) \propto \mathcal{N}(\mathbf{x}_2; \mathbf{r}_{2k}, \gamma_{2k}^{-1} \mathbf{I}) \mathcal{N}(\mathbf{y}; \mathbf{A} \mathbf{x}_2, \gamma_w^{-1} \mathbf{I})$. Using standard manipulations, it can be shown that this belief is Gaussian with mean

$$\widehat{\mathbf{x}}_{2k} = \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I} \right)^{-1} \left(\gamma_w \mathbf{A}^T \mathbf{y} + \gamma_{2k} \mathbf{r}_{2k} \right) \quad (61)$$

and covariance $(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I})^{-1}$. The equivalence between (61) and (24) explains line 11 of Algorithm 3. Furthermore, it can be seen by inspection that the average of the diagonal of this covariance matrix coincides with $\gamma_{2k}^{-1} \langle \mathbf{g}'_2(\mathbf{r}_{2k}, \gamma_{2k}) \rangle$ for $\langle \mathbf{g}'_2(\mathbf{r}_{2k}, \gamma_{2k}) \rangle$ from (25), thus explaining lines 12-13 of Algorithm 3.

Rule 2 then says to set the message $\mu_{\mathbf{x}_2 \rightarrow \delta}(\mathbf{x}_2)$ at $\mathcal{N}(\mathbf{x}_2; \widehat{\mathbf{x}}_{2k}, \eta_{2k}^{-1} \mathbf{I}) / \mathcal{N}(\mathbf{x}_2; \mathbf{r}_{2k}, \gamma_{2k}^{-1} \mathbf{I})$, which (60) simplifies to $\mathcal{N}(\mathbf{x}_2; \mathbf{r}_{1,k+1}, \gamma_{1,k+1}^{-1} \mathbf{I})$ for $\mathbf{r}_{1,k+1} = (\widehat{\mathbf{x}}_{2k} \eta_{2k} - \mathbf{r}_{2k} \gamma_{2k}) / (\eta_{2k} - \gamma_{2k})$ and $\gamma_{1,k+1} = \eta_{2k} - \gamma_{2k}$, yielding lines 14-15 of Algorithm 3. Finally, Rule 3 implies that the message $\mu_{\mathbf{x}_2 \rightarrow \delta}(\mathbf{x}_2)$ flows left through the δ node unchanged, manifesting as $\mu_{\delta \rightarrow \mathbf{x}_1}(\mathbf{x}_1) = \mathcal{N}(\mathbf{x}_1; \mathbf{r}_{1,k+1}, \gamma_{1,k+1}^{-1} \mathbf{I})$ on the other side. The above messaging sequence is then repeated with $k \leftarrow k + 1$.

APPENDIX B

CONVERGENCE OF VECTOR SEQUENCES

We review some definitions from the Bayati-Montanari paper [4], since we will use the same analysis framework in this paper. Fix a dimension $r > 0$, and suppose that, for each N , $\mathbf{x}(N)$ is a vector of the form

$$\mathbf{x}(N) = (\mathbf{x}_1(N), \dots, \mathbf{x}_N(N)),$$

with vector sub-components $\mathbf{x}_n(N) \in \mathbb{R}^r$. Thus, the total dimension of $\mathbf{x}(N)$ is rN . In this case, we will say that $\mathbf{x}(N)$ is a *block vector sequence that scales with N under blocks* $\mathbf{x}_n(N) \in \mathbb{R}^r$. When $r = 1$, so that the blocks are scalar, we will simply say that $\mathbf{x}(N)$ is a *vector sequence that scales with N* . Such vector sequences can be deterministic or random. In most cases, we will omit the notational dependence on N and simply write \mathbf{x} .

Now, given $p \geq 1$, a function $\mathbf{f} : \mathbb{R}^s \rightarrow \mathbb{R}^r$ is called *pseudo-Lipschitz of order p* , if there exists a constant $C > 0$ such that for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^s$,

$$\|\mathbf{f}(\mathbf{x}_1) - \mathbf{f}(\mathbf{x}_2)\| \leq C \|\mathbf{x}_1 - \mathbf{x}_2\| \left[1 + \|\mathbf{x}_1\|^{p-1} + \|\mathbf{x}_2\|^{p-1} \right].$$

Observe that in the case $p = 1$, pseudo-Lipschitz continuity reduces to the standard Lipschitz continuity.

Now suppose that $\mathbf{x} = \mathbf{x}(N)$ is a block vector sequence, which may be deterministic or random. Given $p \geq 1$, we will say that $\mathbf{x} = \mathbf{x}(N)$ converges *empirically with p -th order moments* if there exists a random variable $X \in \mathbb{R}^r$ such that

- (i) $\mathbb{E}|X|^p < \infty$; and
- (ii) for any scalar-valued pseudo-Lipschitz continuous function $f(\cdot)$ of order p ,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(x_n(N)) = \mathbb{E}[f(X)] \quad \text{a.s.} \quad (62)$$

Thus, the empirical mean of the components $f(x_n(N))$ converges to the expectation $\mathbb{E}[f(X)]$. When \mathbf{x} converges empirically with p -th order moments, we will write, with some abuse of notation,

$$\lim_{N \rightarrow \infty} \{x_n\}_{n=1}^N \stackrel{PL(p)}{=} X, \quad (63)$$

where, as usual, we have omitted the dependence $x_n = x_n(N)$. Note that the almost sure convergence in condition (ii) applies to the case where $\mathbf{x}(N)$ is a random vector sequence. Importantly, this condition holds pointwise over each function $f(\cdot)$. It is shown in [4, Lemma 4] that, if condition (i) is true and condition (ii) is true for any bounded continuous function $f(x)$ as well as $f(x) = x^p$, then condition (ii) holds for all pseudo-Lipschitz functions of order p .

We conclude with one final definition. Let $\phi(\mathbf{r}, \gamma)$ be a function on $\mathbf{r} \in \mathbb{R}^s$ and $\gamma \in \mathbb{R}$. We say that $\phi(\mathbf{r}, \gamma)$ is *uniformly Lipschitz continuous* in \mathbf{r} at $\gamma = \bar{\gamma}$ if there exists constants L_1 and $L_2 \geq 0$ and an open neighborhood U of $\bar{\gamma}$, such that

$$\|\phi(\mathbf{r}_1, \gamma) - \phi(\mathbf{r}_2, \gamma)\| \leq L_1 \|\mathbf{r}_1 - \mathbf{r}_2\|, \quad (64)$$

for all $\mathbf{r}_1, \mathbf{r}_2 \in \mathbb{R}^s$ and $\gamma \in U$; and

$$\|\phi(\mathbf{r}, \gamma_1) - \phi(\mathbf{r}, \gamma_2)\| \leq L_2 (1 + \|\mathbf{r}\|) |\gamma_1 - \gamma_2|, \quad (65)$$

for all $\mathbf{r} \in \mathbb{R}^s$ and $\gamma_1, \gamma_2 \in U$.

APPENDIX C

PROOF OF LEMMAS 1 AND 2

For Lemma 1, part (a) follows immediately from (33) and (34). To prove part (b), suppose

$$\mathbf{y} = \mathbf{A} \mathbf{x}^0 + \mathbf{w}, \quad \mathbf{r}_2 = \mathbf{x}^0 + \mathbf{q}.$$

Then, the error is given by

$$\begin{aligned} \mathbf{g}_2(\mathbf{r}_2, \gamma_2) - \mathbf{x}^0 &\stackrel{(a)}{=} \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_2 \mathbf{I} \right)^{-1} \\ &\quad \times \left(\gamma_w \mathbf{A}^T \mathbf{A} \mathbf{x}^0 + \gamma_w \mathbf{A}^T \mathbf{w} + \gamma_2 \mathbf{r}_2 \right) - \mathbf{x}^0 \\ &\stackrel{(b)}{=} \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_2 \mathbf{I} \right)^{-1} \left(\gamma_2 \mathbf{q} + \gamma_w \mathbf{A}^T \mathbf{w} \right), \\ &\stackrel{(c)}{=} \mathbf{Q}^{-1} \left(\gamma_2 \mathbf{q} + \gamma_w \mathbf{A}^T \mathbf{w} \right), \end{aligned}$$

where (a) follows by substituting $\mathbf{y} = \mathbf{A} \mathbf{x}^0 + \mathbf{w}$ into (24); part (b) follows from the substitution $\mathbf{r}_2 = \mathbf{x}^0 + \mathbf{q}$ and collecting the terms with \mathbf{x}^0 ; and (c) follows from the definition of \mathbf{Q} in (38). Hence, the error covariance matrix is given

$$\begin{aligned} &\mathbb{E} \left[(\mathbf{g}_2(\mathbf{r}_2, \gamma_2) - \mathbf{x}^0) (\mathbf{g}_2(\mathbf{r}_2, \gamma_2) - \mathbf{x}^0)^T \right] \\ &= \mathbf{Q}^{-1} \left[\gamma_2^2 \mathbb{E}[\mathbf{q} \mathbf{q}^T] + \gamma_w^2 \mathbf{A} \mathbb{E}[\mathbf{w} \mathbf{w}^T] \mathbf{A}^T \right] \mathbf{Q}^{-1} \\ &= \mathbf{Q}^{-1} \widetilde{\mathbf{Q}} \mathbf{Q}^{-1}, \end{aligned}$$

where we have used the the fact that \mathbf{q} and \mathbf{w} are independent Gaussians with variances τ_2 and γ_w^{-1} . This proves (37). Then, under the matched condition, we have that $\mathbf{Q} = \widetilde{\mathbf{Q}}$, which proves (39). Part (c) of Lemma 1 follows from part (b) by using the SVD (29).

For Lemma 2, part (a) follows from averaging (33) over r_1 . Part (b) follows by taking the derivative in (24) and part (c) follows from using the SVD (29).

APPENDIX D

ORTHOGONAL MATRICES UNDER LINEAR CONSTRAINTS

In preparation for proving Theorem 1, we derive various results on orthogonal matrices subject to linear constraints. To this end, suppose $\mathbf{V} \in \mathbb{R}^{N \times N}$ is an orthogonal matrix satisfying linear constraints

$$\mathbf{A} = \mathbf{V}\mathbf{B}, \quad (66)$$

for some matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{N \times s}$ for some s . Assume \mathbf{A} and \mathbf{B} are full column rank (hence $s \leq N$). Let

$$\mathbf{U}_A = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1/2}, \quad \mathbf{U}_B = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2}. \quad (67)$$

Also, let \mathbf{U}_{A^\perp} and \mathbf{U}_{B^\perp} be any $N \times (N - s)$ matrices whose columns are an orthonormal bases for $\text{Range}(\mathbf{A})^\perp$ and $\text{Range}(\mathbf{B})^\perp$, respectively. Define

$$\tilde{\mathbf{V}} := \mathbf{U}_{A^\perp}^\top \mathbf{V} \mathbf{U}_{B^\perp}, \quad (68)$$

which has dimension $(N - s) \times (N - s)$.

Lemma 3. *Under the above definitions $\tilde{\mathbf{V}}$ satisfies*

$$\mathbf{V} = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{B}^\top + \mathbf{U}_{A^\perp} \tilde{\mathbf{V}} \mathbf{U}_{B^\perp}^\top. \quad (69)$$

Proof. Let $\mathbf{P}_A := \mathbf{U}_A \mathbf{U}_A^\top$ and $\mathbf{P}_A^\perp := \mathbf{U}_{A^\perp} \mathbf{U}_{A^\perp}^\top$ are the orthogonal projections onto $\text{Range}(\mathbf{A})$ and $\text{Range}(\mathbf{A})^\perp$ respectively. Define \mathbf{P}_B and \mathbf{P}_B^\perp similarly. Since, $\mathbf{A} = \mathbf{V}\mathbf{B}$, we have $\mathbf{V}^\top \mathbf{A} = \mathbf{B}$ and therefore,

$$\mathbf{P}_A^\perp \mathbf{V} \mathbf{P}_B = \mathbf{0}, \quad \mathbf{P}_A \mathbf{V} \mathbf{P}_B^\perp = \mathbf{0}. \quad (70)$$

Therefore,

$$\begin{aligned} \mathbf{V} &= (\mathbf{P}_A + \mathbf{P}_A^\perp) \mathbf{V} (\mathbf{P}_B + \mathbf{P}_B^\perp) \\ &= (\mathbf{P}_A \mathbf{V} \mathbf{P}_B + \mathbf{P}_A^\perp \mathbf{V} \mathbf{P}_B^\perp). \end{aligned} \quad (71)$$

Now,

$$\begin{aligned} \mathbf{P}_A \mathbf{V} \mathbf{P}_B &= \mathbf{P}_A \mathbf{V} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \\ &= \mathbf{P}_A \mathbf{A} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \\ &= \mathbf{A} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top = \mathbf{A} (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{B}^\top, \end{aligned} \quad (72)$$

where, in the last step we used the fact that

$$\mathbf{A}^\top \mathbf{A} = \mathbf{B}^\top \mathbf{V}^\top \mathbf{V} \mathbf{B} = \mathbf{B}^\top \mathbf{B}.$$

Also, using the definition of $\tilde{\mathbf{V}}$ in (68),

$$\mathbf{P}_A^\perp \mathbf{V} \mathbf{P}_B^\perp = \mathbf{U}_{A^\perp} \tilde{\mathbf{V}} \mathbf{U}_{B^\perp}^\top. \quad (73)$$

Substituting (72) and (73) into (71) obtains (69). To prove that $\tilde{\mathbf{V}}$ is orthogonal,

$$\begin{aligned} \tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} &\stackrel{(a)}{=} \mathbf{U}_{B^\perp}^\top \mathbf{V} \mathbf{P}_A^\perp \mathbf{V} \mathbf{U}_{B^\perp} \\ &\stackrel{(b)}{=} \mathbf{U}_{B^\perp}^\top \mathbf{V}^\top \mathbf{V} \mathbf{U}_{B^\perp} \stackrel{(c)}{=} \mathbf{I}, \end{aligned}$$

where (a) uses (68); (b) follows from (70) and (c) follows from the fact that \mathbf{V} and \mathbf{U}_{B^\perp} have orthonormal columns. \square

Lemma 4. *Let $\mathbf{V} \in \mathbb{R}^{N \times N}$ be a random matrix that is Haar distributed. Suppose that \mathbf{A} and \mathbf{B} are deterministic and G is the event that \mathbf{V} satisfies linear constraints (66). Then,*

the conditional distribution given G , $\tilde{\mathbf{V}}$ is Haar distributed matrix independent of G . Thus,

$$\mathbf{V}|_G \stackrel{d}{=} \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{B}^\top + \mathbf{U}_{A^\perp} \tilde{\mathbf{V}} \mathbf{U}_{B^\perp}^\top,$$

where $\tilde{\mathbf{V}}$ is Haar distributed and independent of G .

Proof. Let O_N be the set of $N \times N$ orthogonal matrices and let \mathcal{L} be the set of matrices $\mathbf{V} \in O_N$ that satisfy the linear constraints (66). If $p_{\mathbf{V}}(\mathbf{V})$ is the uniform density on O_N (i.e. the Haar measure), the conditional density on \mathbf{V} given the event G ,

$$p_{\mathbf{V}|G}(\mathbf{V}|G) = \frac{1}{Z} p_{\mathbf{V}}(\mathbf{V}) \mathbb{1}_{\{\mathbf{V} \in \mathcal{L}\}},$$

where Z is the normalization constant. Now let $\phi : \tilde{\mathbf{V}} \mapsto \mathbf{V}$ be the mapping described by (69) which maps O_{N-s} to \mathcal{L} . This mapping is invertible. Since ϕ is affine, the conditional density on $\tilde{\mathbf{V}}$ is given by

$$\begin{aligned} p_{\tilde{\mathbf{V}}|G}(\tilde{\mathbf{V}}|G) &\propto p_{\mathbf{V}|G}(\phi(\tilde{\mathbf{V}})|G) \\ &\propto p_{\mathbf{V}}(\phi(\tilde{\mathbf{V}})) \mathbb{1}_{\{\phi(\tilde{\mathbf{V}}) \in \mathcal{L}\}} = p_{\mathbf{V}}(\phi(\tilde{\mathbf{V}})), \end{aligned} \quad (74)$$

where in the last step we used the fact that, for any matrix $\tilde{\mathbf{V}}$, $\phi(\tilde{\mathbf{V}}) \in \mathcal{L}$ (i.e. satisfies the linear constraints (66)). Now to show that $\tilde{\mathbf{V}}$ is conditionally Haar distributed, we need to show that for any orthogonal matrix $\mathbf{W}_0 \in O_{N-s}$,

$$p_{\tilde{\mathbf{V}}|G}(\mathbf{W}_0 \tilde{\mathbf{V}}|G) = p_{\tilde{\mathbf{V}}|G}(\tilde{\mathbf{V}}|G). \quad (75)$$

To prove this, given $\mathbf{W}_0 \in O_{N-s}$, define the matrix,

$$\mathbf{W} = \mathbf{U}_A \mathbf{U}_A^\top + \mathbf{U}_{A^\perp} \mathbf{W}_0 \mathbf{U}_{A^\perp}^\top.$$

One can verify that $\mathbf{W} \in O_N$ (i.e. it is orthogonal) and

$$\phi(\mathbf{W}_0 \tilde{\mathbf{V}}) = \mathbf{W} \phi(\tilde{\mathbf{V}}). \quad (76)$$

Hence,

$$\begin{aligned} p_{\tilde{\mathbf{V}}|G}(\mathbf{W}_0 \tilde{\mathbf{V}}|G) &\stackrel{(a)}{\propto} p_{\mathbf{V}}(\phi(\mathbf{W}_0 \tilde{\mathbf{V}})) \\ &\stackrel{(b)}{\propto} p_{\mathbf{V}}(\mathbf{W} \phi(\tilde{\mathbf{V}})) \stackrel{(c)}{\propto} p_{\mathbf{V}}(\phi(\tilde{\mathbf{V}})), \end{aligned}$$

where (a) follows from (74); (b) follows from (76); and (c) follows from the orthogonal invariance of \mathbf{V} . Hence, the conditional density of $\tilde{\mathbf{V}}$ is invariant under orthogonal transforms and is thus Haar distributed. \square

We will use Lemma 4 in conjunction with the following simple result.

Lemma 5. *Fix a dimension $s \geq 0$, and suppose that $\mathbf{x}(N)$ and $\mathbf{U}(N)$ are sequences such that for each N ,*

- (i) $\mathbf{U} = \mathbf{U}(N) \in \mathbb{R}^{N \times (N-s)}$ is a deterministic matrix with $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$;
- (ii) $\mathbf{x} = \mathbf{x}(N) \in \mathbb{R}^{N-s}$ a random vector that is isotropically distributed in that $\mathbf{V}\mathbf{x} \stackrel{d}{=} \mathbf{x}$ for any orthogonal $(N - s) \times (N - s)$ matrix \mathbf{V} .
- (iii) *The normalized squared Euclidean norm converges almost surely as*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{x}\|^2 = \tau,$$

for some $\tau > 0$.

Then, if we define $\mathbf{y} = \mathbf{U}\mathbf{x}$, we have that the components of \mathbf{y} converge empirically to a Gaussian random variable

$$\lim_{N \rightarrow \infty} \{y_n\} \stackrel{PL(2)}{=} Y \sim \mathcal{N}(0, \tau). \quad (77)$$

Proof. Since \mathbf{x} is isotropically distributed, it can be generated as a normalized Gaussian, i.e.

$$\mathbf{x} \stackrel{d}{=} \frac{\|\mathbf{x}\|}{\|\mathbf{w}_0\|} \mathbf{w}_0, \quad \mathbf{w}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N-s}).$$

For each N , let \mathbf{U}_\perp be an $N \times s$ matrix such that $\mathbf{S} := [\mathbf{U} \ \mathbf{U}_\perp]$ is orthogonal. That is, the s columns of \mathbf{U}_\perp are an orthonormal basis of the orthogonal complement of the Range(\mathbf{U}). If we let $\mathbf{w}_1 \sim \mathcal{N}(0, \mathbf{I}_s)$ independent of \mathbf{w}_0 and define

$$\mathbf{w} = \begin{bmatrix} \mathbf{w}_0 \\ \mathbf{w}_1 \end{bmatrix},$$

then $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_N)$. With this definition, we can write \mathbf{y} as

$$\mathbf{y} = \mathbf{U}\mathbf{x} \stackrel{d}{=} \frac{\|\mathbf{x}\|}{\|\mathbf{w}_0\|} [\mathbf{S}\mathbf{w} - \mathbf{U}_\perp \mathbf{w}_1]. \quad (78)$$

Now,

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{x}\|}{\|\mathbf{w}_0\|} = \sqrt{\tau},$$

almost surely. Also, since $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$ and \mathbf{S} is orthogonal, $\mathbf{S}\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$. Finally, since \mathbf{w}_1 is s -dimensional,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{U}_\perp \mathbf{w}_1\|^2 = \lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{w}_1\|^2 = 0,$$

almost surely. Substituting these properties into (78), we obtain (77). \square

APPENDIX E

A GENERAL CONVERGENCE RESULT

To analyze the VAMP method, we consider the following more general recursion. For each dimension N , we are given an orthogonal matrix $\mathbf{V} \in \mathbb{R}^{N \times N}$, and an initial vector $\mathbf{u}_0 \in \mathbb{R}^N$. Also, we are given disturbance vectors

$$\mathbf{w}^p = (w_1^p, \dots, w_{n_p}^p), \quad \mathbf{w}^q = (w_1^q, \dots, w_{n_q}^q),$$

where the components $w_n^p \in \mathbb{R}^{n_p}$ and $w_n^q \in \mathbb{R}^{n_q}$ for some finite dimensions n_p and n_q that do not grow with N . Then we generate a sequence of iterates by the following recursion:

$$\mathbf{p}_k = \mathbf{V}\mathbf{u}_k \quad (79a)$$

$$\alpha_{1k} = \langle \mathbf{f}'_p(\mathbf{p}_k, \mathbf{w}^p, \gamma_{1k}), \gamma_{2k} = \Gamma_1(\gamma_{1k}, \alpha_{1k}) \quad (79b)$$

$$\mathbf{v}_k = C_1(\alpha_{1k}) [\mathbf{f}_p(\mathbf{p}_k, \mathbf{w}^p, \gamma_{1k}) - \alpha_{1k} \mathbf{p}_k] \quad (79c)$$

$$\mathbf{q}_k = \mathbf{V}^T \mathbf{v}_k \quad (79d)$$

$$\alpha_{2k} = \langle \mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}), \gamma_{1,k+1} = \Gamma_2(\gamma_{2k}, \alpha_{2k}) \quad (79e)$$

$$\mathbf{u}_{k+1} = C_2(\alpha_{2k}) [\mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}) - \alpha_{2k} \mathbf{q}_k], \quad (79f)$$

which is initialized with some vector \mathbf{u}_0 and scalar γ_{10} . Here, $\mathbf{f}_p(\cdot)$ and $\mathbf{f}_q(\cdot)$ are separable functions, meaning

$$\begin{aligned} [\mathbf{f}_p(\mathbf{p}, \mathbf{w}^p, \gamma_1)]_n &= f_p(p_n, w_n^p, \gamma_1) \quad \forall n, \\ [\mathbf{f}_q(\mathbf{q}, \mathbf{w}^q, \gamma_2)]_n &= f_q(q_n, w_n^q, \gamma_2) \quad \forall n, \end{aligned} \quad (80)$$

for scalar-valued functions $f_p(\cdot)$ and $f_q(\cdot)$. The functions $\Gamma_i(\cdot)$ and $C_i(\cdot)$ are also scalar-valued. In the recursion (79), the variables γ_{1k} and γ_{2k} represent some parameter of the update functions $\mathbf{f}_p(\cdot)$ and $\mathbf{f}_q(\cdot)$, and the functions $\Gamma_i(\cdot)$ represent how these parameters are updated.

Similar to our analysis of the VAMP, we consider the following large-system limit (LSL) analysis. We consider a sequence of runs of the recursions indexed by N . We model the initial condition \mathbf{u}_0 and disturbance vectors \mathbf{w}^p and \mathbf{w}^q as deterministic sequences that scale with N and assume that their components converge empirically as

$$\lim_{N \rightarrow \infty} \{u_{0n}\} \stackrel{PL(2)}{=} U_0, \quad (81)$$

and

$$\lim_{N \rightarrow \infty} \{w_n^p\} \stackrel{PL(2)}{=} W^p, \quad \lim_{N \rightarrow \infty} \{w_n^q\} \stackrel{PL(2)}{=} W^q, \quad (82)$$

to random variables U_0 , W^p and W^q . The vectors W^p and W^q are random vectors in \mathbb{R}^{n_p} and \mathbb{R}^{n_q} , respectively. We assume that the initial constant converges as

$$\lim_{N \rightarrow \infty} \gamma_{10} = \bar{\gamma}_{10}, \quad (83)$$

for some $\bar{\gamma}_{10}$. The matrix $\mathbf{V} \in \mathbb{R}^{N \times N}$ is assumed to be uniformly distributed on the set of orthogonal matrices independent of \mathbf{r}_0 , \mathbf{w}^p and \mathbf{w}^q . Since \mathbf{r}_0 , \mathbf{w}^p and \mathbf{w}^q are deterministic, the only randomness is in the matrix \mathbf{V} .

Under the above assumptions, define the SE equations

$$\bar{\alpha}_{1k} = \mathbb{E} \left[f'_p(P_k, W^p, \bar{\gamma}_{1k}) \right], \quad (84a)$$

$$\tau_{2k} = C_1^2(\bar{\alpha}_{1k}) \left\{ \mathbb{E} \left[f_p^2(P_k, W^p, \bar{\gamma}_{1k}) \right] - \bar{\alpha}_{1k}^2 \tau_{1k} \right\} \quad (84b)$$

$$\bar{\gamma}_{2k} = \Gamma_1(\bar{\gamma}_{1k}, \bar{\alpha}_{1k}) \quad (84c)$$

$$\bar{\alpha}_{2k} = \mathbb{E} \left[f'_q(Q_k, W^q, \bar{\gamma}_{2k}) \right], \quad (84d)$$

$$\tau_{1,k+1} = C_2^2(\bar{\alpha}_{2k}) \left\{ \mathbb{E} \left[f_q^2(Q_k, W^q, \bar{\gamma}_{2k}) \right] - \bar{\alpha}_{2k}^2 \tau_{2k} \right\} \quad (84e)$$

$$\gamma_{1,k+1} = \Gamma_2(\bar{\gamma}_{2k}, \bar{\alpha}_{2k}), \quad (84f)$$

which are initialized with $\bar{\gamma}_{10}$ in (83) and

$$\tau_{10} = \mathbb{E}[U_0^2], \quad (85)$$

where U_0 is the random variable in (81). In the SE equations (84), the expectations are taken with respect to random variables

$$P_k \sim \mathcal{N}(0, \tau_{1k}), \quad Q_k \sim \mathcal{N}(0, \tau_{2k}),$$

where P_k is independent of W^p and Q_k is independent of W^q .

Theorem 4. Consider the recursions (79) and SE equations (84) under the above assumptions. Assume additionally that, for all k :

(i) For $i = 1, 2$, the functions

$$C_i(\alpha_i), \quad \Gamma_i(\gamma_i, \alpha_i),$$

are continuous at the points $(\gamma_i, \alpha_i) = (\bar{\gamma}_{ik}, \bar{\alpha}_{ik})$ from the SE equations; and

(ii) The function $f_p(p, w^p, \gamma_1)$ and its derivative $f'_p(p, w^p, \gamma_1)$ are uniformly Lipschitz continuous in (p, w^p) at $\gamma_1 = \bar{\gamma}_{1k}$.

(iii) The function $f_q(q, w^q, \gamma_2)$ and its derivative $f'_q(q, w^q, \gamma_2)$ are uniformly Lipschitz continuous in (q, w^q) at $\gamma_2 = \bar{\gamma}_{2k}$.

Then,

(a) For any fixed k , almost surely the components of $(\mathbf{w}^p, \mathbf{p}_0, \dots, \mathbf{p}_k)$ empirically converge as

$$\lim_{N \rightarrow \infty} \{(w_n^p, p_{0n}, \dots, p_{kn})\} \stackrel{PL(2)}{=} (W^p, P_0, \dots, P_k), \quad (86)$$

where W^p is the random variable in the limit (82) and (P_0, \dots, P_k) is a zero mean Gaussian random vector independent of W^p , with $\mathbb{E}[P_k^2] = \tau_{1k}$. In addition, we have that

$$\lim_{N \rightarrow \infty} (\alpha_{1k}, \gamma_{1k}) = (\bar{\alpha}_{1k}, \bar{\gamma}_{1k}), \quad (87)$$

almost surely.

(b) For any fixed k , almost surely the components of $(\mathbf{w}^q, \mathbf{q}_0, \dots, \mathbf{q}_k)$ empirically converge as

$$\lim_{N \rightarrow \infty} \{(w_n^q, q_{0n}, \dots, q_{kn})\} \stackrel{PL(2)}{=} (W^q, Q_0, \dots, Q_k), \quad (88)$$

where W^q is the random variable in the limit (82) and (Q_0, \dots, Q_k) is a zero mean Gaussian random vector independent of W^q , with $\mathbb{E}[P_k^2] = \tau_{2k}$. In addition, we have that

$$\lim_{N \rightarrow \infty} (\alpha_{2k}, \gamma_{2k}) = (\bar{\alpha}_{2k}, \bar{\gamma}_{2k}), \quad (89)$$

almost surely.

Proof. We will prove this in the next Appendix, Appendix F. \square

APPENDIX F

PROOF OF THEOREM 4

A. Induction Argument

We use an induction argument. Given iterations $k, \ell \geq 0$, define the hypothesis, $H_{k,\ell}$ as the statement:

- Part (a) of Theorem 4 is true up to k ; and
- Part (b) of Theorem 4 is true up to ℓ .

The induction argument will then follow by showing the following three facts:

- $H_{0,-1}$ is true;
- If $H_{k,k-1}$ is true, then so is $H_{k,k}$;
- If $H_{k,k}$ is true, then so is $H_{k+1,k}$.

B. Induction Initialization

We first show that the hypothesis $H_{0,-1}$ is true. That is, we must show (86) and (87) for $k = 0$. This is a special case of Lemma 5. Specifically, for each N , let $\mathbf{U} = \mathbf{I}_N$, the $N \times N$ identity matrix, which trivially satisfies property (i) of Lemma 5 with $s = 0$. Let $\mathbf{x} = \mathbf{p}_0$. Since $\mathbf{p}_0 = \mathbf{V}\mathbf{u}_0$ and \mathbf{V} is Haar distributed independent of \mathbf{u}_0 , we have that \mathbf{p}_0 is orthogonally invariant and satisfies property (ii) of Lemma 5. Also,

$$\lim_{N \rightarrow \infty} \|\mathbf{p}_0\|^2 \stackrel{(a)}{=} \lim_{N \rightarrow \infty} \|\mathbf{u}_0\|^2 \stackrel{(b)}{=} \mathbb{E}[U_0^2] \stackrel{(c)}{=} \tau_{10},$$

where (a) follows from the fact that $\mathbf{p}_0 = \mathbf{V}\mathbf{u}_0$ and \mathbf{V} is orthogonal; (b) follows from the assumption (81) and (c) follows from the definition (85). This proves property (iii) of Lemma 5. Hence, $\mathbf{p}_0 = \mathbf{U}\mathbf{p}_0$, we have that the components of \mathbf{p}_0 converge empirically as

$$\lim_{N \rightarrow \infty} \{p_{0n}\} \stackrel{PL(2)}{=} P_0 \sim \mathcal{N}(0, \tau_{10}),$$

for a Gaussian random variable P_0 . Moreover, since \mathbf{V} is independent of \mathbf{w}^p , and the components of \mathbf{w}^p converge empirically as (82), we have that the components of $\mathbf{p}_n, \mathbf{w}^p$ almost surely converge empirically as

$$\lim_{N \rightarrow \infty} \{w_n^p, p_{0n}\} \stackrel{PL(2)}{=} (W^p, P_0),$$

where W^p is independent of P_0 . This proves (86) for $k = 0$.

Now, we have assumed in (83) that $\gamma_{10} \rightarrow \bar{\gamma}_{10}$ as $N \rightarrow \infty$. Also, since $f'_p(p, w^p, \gamma_1)$ is uniformly Lipschitz continuous in (p, w^p) at $\gamma_1 = \bar{\gamma}_{10}$, we have that $\alpha_{10} = \langle \mathbf{f}'_p(\mathbf{p}_0, \mathbf{w}^p, \gamma_{10}) \rangle$ converges to $\bar{\alpha}_{10}$ in (84a) almost surely. This proves (87).

C. The Induction Recursion

We next show the implication $H_{k,k-1} \Rightarrow H_{k,k}$. The implication $H_{k,k} \Rightarrow H_{k+1,k}$ is proven similarly. Hence, fix k and assume that $H_{k,k-1}$ holds. Since $\Gamma_1(\gamma_i, \alpha_i)$ is continuous at $(\bar{\gamma}_{1k}, \bar{\alpha}_{1k})$, the limits (87) combined with (84c) show that

$$\lim_{N \rightarrow \infty} \gamma_{2k} = \lim_{N \rightarrow \infty} \Gamma_1(\gamma_{1k}, \alpha_{1k}) = \bar{\gamma}_{2k}.$$

In addition, the induction hypothesis shows that for $\ell = 0, \dots, k$, the components of $(\mathbf{w}^p, \mathbf{p}_\ell)$ almost surely converge empirically as

$$\lim_{N \rightarrow \infty} \{(w_n^p, p_{\ell n})\} \stackrel{PL(2)}{=} (W^p, P_\ell),$$

where $P_\ell \sim \mathcal{N}(0, \tau_{1\ell})$ for $\tau_{1\ell}$ given by the SE equations. Since $f_p(\cdot)$ is Lipschitz continuous and $C_1(\alpha_{1\ell})$ is continuous at $\alpha_{1\ell} = \bar{\alpha}_{1\ell}$, one may observe that the definition of \mathbf{v}_ℓ in (79c) and the limits (87) show that

$$\lim_{N \rightarrow \infty} \{(w_n^p, p_{\ell n}, v_{\ell n})\} \stackrel{PL(2)}{=} (W^p, P_\ell, V_\ell),$$

where V_ℓ is the random variable

$$V_\ell = g_p(P_\ell, W^p, \bar{\gamma}_{1\ell}, \bar{\alpha}_{1\ell}), \quad (90)$$

and $g_p(\cdot)$ is the function

$$g_p(p, w^p, \gamma_1, \alpha_1) := C_1(\alpha_1) [f_p(p, w^p, \gamma_1) - \alpha_1 p]. \quad (91)$$

Similarly, we have the limit

$$\lim_{N \rightarrow \infty} \{(w_n^q, q_{\ell n}, u_{\ell n})\} \stackrel{PL(2)}{=} (W^q, Q_\ell, U_\ell),$$

where U_ℓ is the random variable,

$$U_\ell = g_q(Q_\ell, W^q, \bar{\gamma}_{2\ell}, \bar{\alpha}_{2\ell}) \quad (92)$$

and $g_q(\cdot)$ is the function

$$g_q(q, w^q, \gamma_2, \alpha_2) := C_2(\alpha_2) [f_q(q, w^q, \gamma_2) - \alpha_2 q]. \quad (93)$$

We next introduce the notation

$$\mathbf{U}_k := [\mathbf{u}_0 \cdots \mathbf{u}_k] \in \mathbb{R}^{N \times (k+1)},$$

to represent the first $k+1$ values of the vector \mathbf{u}_ℓ . We define the matrices \mathbf{V}_k , \mathbf{Q}_k and \mathbf{P}_k similarly. Using this notation, let G_k be the tuple of random matrices,

$$G_k := \{\mathbf{U}_k, \mathbf{P}_k, \mathbf{V}_k, \mathbf{Q}_{k-1}\}. \quad (94)$$

With some abuse of notation, we will also use G_k to denote the sigma-algebra generated by these variables. The set (94) contains all the outputs of the algorithm (79) immediately before (79d) in iteration k .

Now, the actions of the matrix \mathbf{V} in the recursions (79) are through the matrix-vector multiplications (79a) and (79d). Hence, if we define the matrices

$$\mathbf{A}_k := [\mathbf{P}_k \ \mathbf{V}_{k-1}], \quad \mathbf{B}_k := [\mathbf{U}_k \ \mathbf{Q}_{k-1}], \quad (95)$$

the output of the recursions in the set G_k will be unchanged for all matrices \mathbf{V} satisfying the linear constraints

$$\mathbf{A}_k = \mathbf{V}\mathbf{B}_k. \quad (96)$$

Hence, the conditional distribution of \mathbf{V} given G_k is precisely the uniform distribution on the set of orthogonal matrices satisfying (96). The matrices \mathbf{A}_k and \mathbf{B}_k are of dimensions $N \times s$ where $s = 2k + 1$. From Lemma 4, this conditional distribution is given by

$$\mathbf{V}|_{G_k} \stackrel{d}{=} \mathbf{A}_k(\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{B}_k^\top + \mathbf{U}_{\mathbf{A}_k^\perp} \tilde{\mathbf{V}} \mathbf{U}_{\mathbf{B}_k^\perp}^\top, \quad (97)$$

where $\mathbf{U}_{\mathbf{A}_k^\perp}$ and $\mathbf{U}_{\mathbf{B}_k^\perp}$ are $N \times (N-s)$ matrices whose columns are an orthonormal basis for $\text{Range}(\mathbf{A}_k)^\perp$ and $\text{Range}(\mathbf{B}_k)^\perp$. The matrix $\tilde{\mathbf{V}}$ is Haar distributed on the set of $(N-s) \times (N-s)$ orthogonal matrices and independent of G_k .

Using (97) we can write \mathbf{q}_k in (79d) as a sum of two terms

$$\mathbf{q}_k = \mathbf{V}^\top \mathbf{v}_k = \mathbf{q}_k^{\text{det}} + \mathbf{q}_k^{\text{ran}}, \quad (98)$$

where $\mathbf{q}_k^{\text{det}}$ is what we will call the *deterministic* part:

$$\mathbf{q}_k^{\text{det}} = \mathbf{B}_k(\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{A}_k^\top \mathbf{v}_k, \quad (99)$$

and $\mathbf{q}_k^{\text{ran}}$ is what we will call the *random* part:

$$\mathbf{q}_k^{\text{ran}} = \mathbf{U}_{\mathbf{B}_k^\perp} \tilde{\mathbf{V}}^\top \mathbf{U}_{\mathbf{A}_k^\perp}^\top \mathbf{v}_k. \quad (100)$$

The next few lemmas will evaluate the asymptotic distributions of the two terms in (98).

Lemma 6. *Under the induction hypothesis $H_{k,k-1}$, there exist constants $\beta_{k,0}, \dots, \beta_{k,k-1}$ such that the components of $\mathbf{q}_k^{\text{det}}$ along with $(\mathbf{q}_0, \dots, \mathbf{q}_{k-1})$ converge empirically as*

$$\lim_{N \rightarrow \infty} \left\{ w_n^q, q_{0n}, \dots, q_{k-1,n}, q_{kn}^{\text{det}} \right\} \stackrel{PL(2)}{=} (W^q, Q_0, \dots, Q_{k-1}, Q_k^{\text{det}}), \quad (101)$$

where Q_ℓ , $\ell = 0, \dots, k-1$ are the Gaussian random variables in induction hypothesis (88) and Q_k^{det} is a linear combination,

$$Q_k^{\text{det}} = \beta_{k,0} Q_0 + \dots + \beta_{k,k-1} Q_{k-1}. \quad (102)$$

Proof. We evaluate the asymptotic values of various terms in (99). Using the definition of \mathbf{A}_k in (95),

$$\mathbf{A}_k^\top \mathbf{A}_k = \begin{bmatrix} \mathbf{P}_k^\top \mathbf{P}_k & \mathbf{P}_k^\top \mathbf{V}_{k-1} \\ \mathbf{V}_{k-1}^\top \mathbf{P}_k & \mathbf{V}_{k-1}^\top \mathbf{V}_{k-1} \end{bmatrix}$$

We can then easily evaluate the asymptotic value of these terms as follows. For example, the asymptotic value of the (i, j) component of the matrix $\mathbf{P}_k^\top \mathbf{P}_k$ is given by

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} [\mathbf{P}_k^\top \mathbf{P}_k]_{ij} &\stackrel{(a)}{=} \frac{1}{N} \mathbf{p}_i^\top \mathbf{p}_j \\ &= \frac{1}{N} \sum_{n=1}^N p_{in} p_{jn} \stackrel{(b)}{=} E(P_i P_j) \stackrel{(c)}{=} [\mathbf{Q}_k^p]_{ij}, \end{aligned}$$

where (a) follows since the i -th column of \mathbf{P}_k is precisely the vector \mathbf{p}_i ; (b) follows due to convergence assumption in (86); and in (c), we use \mathbf{Q}_k^p to denote the covariance matrix of (P_0, \dots, P_k) . Similarly

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{V}_{k-1}^\top \mathbf{V}_{k-1} = \mathbf{Q}_{k-1}^v,$$

where \mathbf{Q}_{k-1}^v has the components,

$$[\mathbf{Q}_{k-1}^v]_{ij} = \mathbb{E}[V_i V_j],$$

where V_i is the random variable in (90). Finally, the expectation for the cross-terms are given by

$$\begin{aligned} \mathbb{E}[V_i X_j] &\stackrel{(a)}{=} \mathbb{E}[g_p(P_i, W^p, \bar{\gamma}_{1i}, \bar{\alpha}_{1i}) X_j] \\ &\stackrel{(b)}{=} \mathbb{E}\left[g'_p(P_i, W^p, \bar{\gamma}_{1i}, \bar{\alpha}_{1i})\right] \mathbb{E}[X_i X_j] \\ &\stackrel{(c)}{=} \mathbb{E}[X_i X_j] C_1(\bar{\alpha}_{1i}) \left(\mathbb{E}\left[f'_p(P_i, W^p, \bar{\gamma}_{1i})\right] - \bar{\alpha}_{1i}\right) \\ &\stackrel{(d)}{=} 0, \end{aligned}$$

where (a) follows from (90); (b) follows from Stein's Lemma; and (c) follows from the definition of $g_p(\cdot)$ in (91); and (d) follows from (84a). The above calculations show that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{A}_k^\top \mathbf{A}_k \stackrel{a.s.}{=} \begin{bmatrix} \mathbf{Q}_k^p & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_{k-1}^v \end{bmatrix}, \quad (103)$$

A similar calculation shows that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{A}_k^\top \mathbf{s}_k = \begin{bmatrix} \mathbf{0} \\ \mathbf{b}_k^s \end{bmatrix}, \quad (104)$$

where \mathbf{b}_k^s is the vector of correlations

$$\mathbf{b}_k^s = [\mathbb{E}[V_0 V_k] \ \mathbb{E}[V_1 V_k] \ \dots \ \mathbb{E}[V_{k-1} V_k]]^\top. \quad (105)$$

Combining (103) and (104) shows that

$$\lim_{N \rightarrow \infty} (\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{A}_k^\top \mathbf{v}_k \stackrel{a.s.}{=} \begin{bmatrix} \mathbf{0} \\ \beta_k \end{bmatrix}, \quad (106)$$

where

$$\beta_k := [\mathbf{Q}_{k-1}^v]^{-1} \mathbf{b}_k^s.$$

Therefore,

$$\begin{aligned} \mathbf{q}_k^{\text{det}} &= \mathbf{B}_k(\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{A}_k^\top \mathbf{v}_k \\ &= [\mathbf{U}_k \ \mathbf{Q}_{k-1}] \begin{bmatrix} \mathbf{0} \\ \beta_k \end{bmatrix} + \xi \\ &= \sum_{\ell=0}^{k-1} \beta_{k\ell} \mathbf{q}_\ell + \xi, \end{aligned} \quad (107)$$

where $\boldsymbol{\xi} \in \mathbb{R}^N$ is the error,

$$\boldsymbol{\xi} = \mathbf{B}_k \mathbf{s}, \quad \mathbf{s} := (\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{A}_k^\top \mathbf{v}_k - \begin{bmatrix} \mathbf{0} \\ \beta_k \end{bmatrix}. \quad (108)$$

We next need to bound the norm of the error term $\boldsymbol{\xi}$. Since $\boldsymbol{\xi} = \mathbf{B}_k \mathbf{s}$, the definition of \mathbf{B}_k in (95) shows that

$$\boldsymbol{\xi} = \sum_{i=0}^k s_i \mathbf{u}_i + \sum_{j=0}^{k-1} s_{k+j+1} \mathbf{q}_j, \quad (109)$$

where we have indexed the components of \mathbf{s} in (108) as $\mathbf{s} = (s_0, \dots, s_{2k})$. From (106), the components $s_j \rightarrow 0$ almost surely, and therefore

$$\lim_{N \rightarrow \infty} \max_{j=0, \dots, 2k} |s_j| \stackrel{a.s.}{=} 0.$$

Also, by the induction hypothesis,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{u}_i\|^2 \stackrel{a.s.}{=} E(U_i^2), \quad \lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{q}_j\|^2 \stackrel{a.s.}{=} E(Q_j^2).$$

Therefore, from (109),

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \|\boldsymbol{\xi}\|^2 &\leq \lim_{N \rightarrow \infty} \left[\max_{j=0, \dots, 2k} |s_j|^2 \right] \\ &\times \frac{1}{N} \left[\sum_i \|\mathbf{u}_i\|^2 + \sum_j \|\mathbf{q}_j\|^2 \right] \stackrel{a.s.}{=} 0. \end{aligned} \quad (110)$$

Therefore, if $f(q_1, \dots, q_k)$ is pseudo-Lipschitz continuous of order 2,

$$\begin{aligned} &\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(q_{0n}, \dots, q_{k-1,n}, q_k^{\det}) \\ &\stackrel{(a)}{=} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f\left(q_{0n}, \dots, q_{k-1,n}, \sum_{\ell=0}^{k-1} \beta_{k\ell} q_{\ell n}\right) \\ &\stackrel{(b)}{=} \mathbb{E} \left[f\left(Q_0, \dots, Q_{k-1}, \sum_{\ell=0}^{k-1} \beta_{k\ell} Q_{\ell n}\right) \right], \end{aligned}$$

where (a) follows from the (107), the bound (110), and the pseudo-Lipschitz continuity of $f(\cdot)$; and (b) follows from the fact that $f(\cdot)$ is pseudo-Lipschitz continuous and the induction hypothesis that

$$\lim_{N \rightarrow \infty} \{q_{0n}, \dots, q_{k-1,n}\} \stackrel{PL(2)}{=} (Q_0, \dots, Q_{k-1}).$$

This proves (101). \square

Lemma 7. *Under the induction hypothesis $H_{k,k-1}$, the following limit holds almost surely*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{U}_{\mathbf{A}_k^\perp}^\top \mathbf{s}_k\|^2 = \rho_k, \quad (111)$$

for some constant $\rho_k \geq 0$.

Proof. From (95), the matrix \mathbf{A}_k has $s = 2k + 1$ columns. From Lemma 4, $\mathbf{U}_{\mathbf{A}_k^\perp}$ is an orthonormal basis of $N - s$ in the $\text{Range}(\mathbf{A}_k)^\perp$. Hence, the energy $\|\mathbf{U}_{\mathbf{A}_k^\perp} \mathbf{s}_k\|^2$ is precisely

$$\|\mathbf{U}_{\mathbf{A}_k^\perp} \mathbf{s}_k\|^2 = \mathbf{s}_k^\top \mathbf{s}_k - \mathbf{s}_k^\top \mathbf{A}_k (\mathbf{A}_k^\top \mathbf{A}_k)^{-1} \mathbf{A}_k^\top \mathbf{s}_k.$$

Using similar calculations as the previous lemma, we have

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{U}_{\mathbf{A}_k} \mathbf{s}_k\|^2 = \mathbb{E}[S_k^2] - (\mathbf{b}_k^s)^\top [\mathbf{Q}_k^s]^{-1} \mathbf{b}_k^s.$$

Hence, the lemma is proven if we define ρ_k as the right hand side of this equation. \square

Lemma 8. *Under the induction hypothesis $H_{k,k-1}$, the components of the “random” part $\mathbf{q}_k^{\text{ran}}$ along with the components of $(\mathbf{w}^q, \mathbf{q}_0, \dots, \mathbf{q}_{k-1})$ almost surely converge empirically as*

$$\begin{aligned} &\lim_{N \rightarrow \infty} \{(w_n^q, q_{0n}, \dots, q_{k-1,n}, q_{kn}^{\text{ran}})\} \\ &\stackrel{PL(2)}{=} (W^q, Q_0, \dots, Q_{k-1}, U_k), \end{aligned} \quad (112)$$

where $U_k \sim \mathcal{N}(0, \rho_k)$ is a Gaussian random variable independent of $(W^q, Q_0, \dots, Q_{k-1})$ and ρ_k is the constant in Lemma 7.

Proof. This is a direct application of Lemma 5. Let $\mathbf{x} = \tilde{\mathbf{V}}^\top \mathbf{U}_{\mathbf{A}_k^\perp}^\top \mathbf{s}_k$ so that

$$\mathbf{q}_k^{\det} = \mathbf{U}_{\mathbf{B}_k^\perp} \mathbf{x}_k.$$

For each N , $\mathbf{U}_{\mathbf{B}_k^\perp} \in \mathbb{R}^{N \times (N-s)}$ is a matrix with orthonormal columns spanning $\text{Range}(\mathbf{B}_k)^\perp$. Also, since $\tilde{\mathbf{V}}$ is uniformly distributed on the set of $(N-s) \times (N-s)$ orthogonal matrices, and independent of G_k , the conditional distribution \mathbf{x}_k given G_k is orthogonally invariant in that

$$\mathbf{U} \mathbf{x}_k |_{G_k} \stackrel{d}{=} \mathbf{x}_k |_{G_k},$$

for any orthogonal matrix \mathbf{U} . Lemma 7 also shows that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{x}_k\|^2 = \rho_k,$$

almost surely. The limit (112) now follows from Lemma 5. \square

Using the partition (98) and Lemmas 6 and 8, the components of $(\mathbf{w}^q, \mathbf{q}_0, \dots, \mathbf{q}_k)$ almost surely converge empirically as

$$\begin{aligned} &\lim_{N \rightarrow \infty} \{(w_n^q, q_{0n}, \dots, q_{kn})\} \\ &\stackrel{PL(2)}{=} \lim_{N \rightarrow \infty} \{(w_n^q, q_{0n}, \dots, q_{kn}^{\det} + q_{kn}^{\text{ran}})\} \\ &\stackrel{PL(2)}{=} (W^q, Q_0, \dots, Q_k), \end{aligned}$$

where Q_k is the random variable

$$Q_k = \beta_{k0} Q_0 + \dots + \beta_{k,k-1} Q_{k-1} + U_k.$$

Since (Q_0, \dots, Q_{k-1}) is jointly Gaussian and U_k is Gaussian independent of (Q_0, \dots, Q_{k-1}) we have that (Q_0, \dots, Q_k) is Gaussian. This proves (88).

Now the function $\Gamma_1(\gamma_1, \alpha_1)$ is assumed to be continuous at $(\bar{\gamma}_{1k}, \bar{\alpha}_{1k})$. Also, the induction hypothesis assumes that $\alpha_{1k} \rightarrow \bar{\alpha}_{1k}$ and $\gamma_{1k} \rightarrow \bar{\gamma}_{1k}$ almost surely. Hence,

$$\lim_{N \rightarrow \infty} \gamma_{2k} = \lim_{N \rightarrow \infty} \Gamma_1(\gamma_{1k}, \alpha_{1k}) = \bar{\gamma}_{2k}. \quad (113)$$

In addition, since we have assumed that $\mathbf{f}'_q(\mathbf{q}, \mathbf{w}^q, \gamma_1)$ is Lipschitz continuous in $(\mathbf{q}, \mathbf{w}^q)$ and continuous in γ_1 ,

$$\begin{aligned} \lim_{N \rightarrow \infty} \alpha_{2k} &= \lim_{N \rightarrow \infty} \langle \mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{1k}) \rangle \\ &= \mathbb{E} \left[f'_q(Q_k, W^q, \bar{\gamma}_{1k}) \right] = \bar{\alpha}_{1k}. \end{aligned} \quad (114)$$

The limits (113) and (114) prove (89).

Finally, we need to show that $\mathbb{E}[Q_k^2] = \tau_{2k}$ is the variance from the SE equations.

$$\begin{aligned}
\mathbb{E}[Q_k^2] &\stackrel{(a)}{=} \lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{q}_k\|^2 \\
&\stackrel{(b)}{=} \lim_{N \rightarrow \infty} \frac{1}{N} \|\mathbf{v}_k\|^2 \\
&\stackrel{(c)}{=} \mathbb{E} [g_p(P_k, W^p, \bar{\gamma}_{1k}, \bar{\alpha}_{1k})] \\
&\stackrel{(d)}{=} C_1^2(\bar{\alpha}_{1k}) \mathbb{E} \left[(f_p(P_k, W^p, \bar{\gamma}_{1k}) - \bar{\alpha}_{1k} P_k)^2 \right] \\
&= C_1^2(\bar{\alpha}_{1k}) \left\{ \mathbb{E} \left[f_p^2(P_k, W^p, \bar{\gamma}_{1k}) \right] \right. \\
&\quad \left. - 2\bar{\alpha}_{1k} \mathbb{E} [P_k f_p(P_k, W^p, \bar{\gamma}_{1k})] + \bar{\alpha}_{1k}^2 \mathbb{E} [P_k^2] \right\} \\
&\stackrel{(e)}{=} C_1^2(\bar{\alpha}_{1k}) \left\{ \mathbb{E} \left[f_p^2(P_k, W^p, \bar{\gamma}_{1k}) \right] \right. \\
&\quad \left. - 2\bar{\alpha}_{1k} \tau_{1k} \mathbb{E} \left[f_p'(P_k, W^p, \bar{\gamma}_{1k}) \right] + \bar{\alpha}_{1k}^2 \tau_{1k} \right\} \\
&\stackrel{(f)}{=} C_1^2(\bar{\alpha}_{1k}) \left\{ \mathbb{E} \left[f_p^2(P_k, W^p, \bar{\gamma}_{1k}) \right] - \bar{\alpha}_{1k}^2 \tau_{1k} \right\} \\
&\stackrel{(g)}{=} \tau_{2k}, \tag{115}
\end{aligned}$$

where (a) follows from the fact that the components of \mathbf{q}_k converge empirically to Q_k ; (b) follows from (79d) and the fact that \mathbf{V} is orthogonal; (c) follows from the limit (90); and (d) follows from (91); (e) follows from Stein's Lemma and the fact that $\mathbb{E}[P_k^2] = \tau_{1k}$; (f) follows from the definition of $\bar{\alpha}_{1k}$ in (84a); and (g) follows from (84b). Thus, $\mathbb{E}[Q_k^2] = \tau_{2k}$, and we have proven the implication $H_{k,k-1} \Rightarrow H_{k,k}$.

APPENDIX G PROOF OF THEOREM 1

Theorem 1 is essentially a special case of Theorem 4. We need to simply rewrite the recursions in Algorithm 3 in the form (79). To this end, define the error terms

$$\mathbf{p}_k := \mathbf{r}_{1k} - \mathbf{x}^0, \quad \mathbf{v}_k := \mathbf{r}_{2k} - \mathbf{x}^0, \tag{116}$$

and their transforms,

$$\mathbf{u}_k := \mathbf{V}^T \mathbf{p}_k, \quad \mathbf{q}_k := \mathbf{V}^T \mathbf{v}_k. \tag{117}$$

Also, define the disturbance terms

$$\mathbf{w}^q := (\boldsymbol{\xi}, \mathbf{s}), \quad \mathbf{w}^p := \mathbf{x}^0, \quad \boldsymbol{\xi} := \mathbf{U}^T \mathbf{w}, \tag{118}$$

and the componentwise update functions

$$f_q(q, (\boldsymbol{\xi}, \mathbf{s}), \gamma_2) := \frac{\gamma_w s \boldsymbol{\xi} + \gamma_2 q}{\gamma_w s^2 + \gamma_2}, \tag{119a}$$

$$f_p(p, \mathbf{x}^0, \gamma_1) = g_1(p + \mathbf{x}^0, \gamma_1) - \mathbf{x}^0. \tag{119b}$$

With these definitions, we claim that the outputs satisfy the recursions:

$$\mathbf{p}_k = \mathbf{V} \mathbf{u}_k \tag{120a}$$

$$\alpha_{1k} = (\mathbf{f}'_p(\mathbf{p}_k, \mathbf{x}^0, \gamma_{1k})), \quad \gamma_{2k} = \frac{(1 - \alpha_{1k}) \gamma_{1k}}{\alpha_{1k}} \tag{120b}$$

$$\mathbf{v}_k = \frac{1}{1 - \alpha_{1k}} \left[\mathbf{f}_p(\mathbf{p}_k, \mathbf{x}^0, \gamma_{1k}) - \alpha_{1k} \mathbf{p}_k \right] \tag{120c}$$

$$\mathbf{q}_k = \mathbf{V}^T \mathbf{v}_k \tag{120d}$$

$$\alpha_{2k} = (\mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k})), \quad \gamma_{1,k+1} = \frac{(1 - \alpha_{2k}) \gamma_{2k}}{\alpha_{2k}} \tag{120e}$$

$$\mathbf{u}_{k+1} = \frac{1}{1 - \alpha_{2k}} \left[\mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}) - \alpha_{2k} \mathbf{q}_k \right] \tag{120f}$$

Before we prove (120), we can see that (120) is a special case of the general recursions in (79) if we define

$$C_i(\alpha_i) = \frac{1}{1 - \alpha_i}, \quad \Gamma_i(\gamma_i, \alpha_i) = \gamma_i \left[\frac{1}{\alpha_i} - 1 \right].$$

It is also straightforward to verify the continuity assumptions in Theorem 4. The assumption of Theorem 1 states that $\bar{\alpha}_{ik} \in (0, 1)$. Since $\bar{\gamma}_{10} > 0$, $\bar{\gamma}_{ik} > 0$ for all k and i . Therefore, $C_i(\alpha_i)$ and $\Gamma_i(\gamma_i, \alpha_i)$ are continuous at all points $(\gamma_i, \alpha_i) = (\bar{\gamma}_{ik}, \bar{\alpha}_{ik})$. Also, since $s \in [0, S_{max}]$ and $\gamma_{2k} > 0$ for all k , the function $f_q(q, (\boldsymbol{\xi}, \mathbf{s}), \gamma_2)$ in (119) is uniformly Lipschitz continuous in $(q, \boldsymbol{\xi}, \mathbf{s})$ at all $\gamma_2 = \bar{\gamma}_{2k}$. Similarly, since the denoiser function $g_1(r_1, \gamma_1)$ is assumed to be uniformly Lipschitz continuous in r_1 at all $\gamma_1 = \bar{\gamma}_{1k}$, so is the function $f_p(r_1, \mathbf{x}^0, \gamma_1)$ in (119b). Hence all the conditions of Theorem 4 are satisfied. The SE equations (50) immediately follow from the general SE equations (84). In addition, the limits (45) and (48) are special cases of the limits (86) and (88). This proves Theorem 1.

So, it remains only to show that the updates in (120) indeed hold. Equations (120a) and (120d) follow immediately from the definitions (116) and (117). Next, observe that we can rewrite the LMMSE estimation function (24) as

$$\begin{aligned}
&\mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k}) \\
&\stackrel{(a)}{=} \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I} \right)^{-1} \left(\gamma_w \mathbf{A}^T \mathbf{A} \mathbf{x}^0 + \gamma_w \mathbf{A}^T \mathbf{w} + \gamma_{2k} \mathbf{r}_{2k} \right) \\
&\stackrel{(b)}{=} \mathbf{x}^0 + \left(\gamma_w \mathbf{A}^T \mathbf{A} + \gamma_{2k} \mathbf{I} \right)^{-1} \left(\gamma_{2k} (\mathbf{r}_{2k} - \mathbf{x}^0) + \gamma_w \mathbf{A}^T \mathbf{w} \right) \\
&\stackrel{(c)}{=} \mathbf{x}^0 + \mathbf{V} \left(\gamma_w \mathbf{S}^2 + \gamma_{2k} \mathbf{I} \right)^{-1} \left(\gamma_{2k} \mathbf{q}_k + \mathbf{S} \boldsymbol{\xi} \right), \\
&\stackrel{(d)}{=} \mathbf{x}^0 + \mathbf{V} \mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}), \tag{121}
\end{aligned}$$

where (a) follows by substituting (28) into (24); (b) is a simple algebraic manipulation; (c) follows from the SVD definition (29) and the definitions $\boldsymbol{\xi}$ in (118) and \mathbf{q}_k in (117); and (d) follows from the definition of componentwise function $f_q(\cdot)$ in (119a). Therefore, the divergence α_{2k} satisfies

$$\begin{aligned}
\alpha_{2k} &\stackrel{(a)}{=} \frac{1}{N} \text{Tr} \left[\frac{\partial \mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k})}{\partial \mathbf{r}_{2k}} \right] \\
&\stackrel{(b)}{=} \frac{1}{N} \text{Tr} \left[\mathbf{V} \text{Diag}(\mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k})) \frac{\partial \mathbf{q}_k}{\partial \mathbf{r}_{2k}} \right] \\
&\stackrel{(c)}{=} \frac{1}{N} \text{Tr} \left[\mathbf{V} \text{Diag}(\mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k})) \mathbf{V}^T \right] \\
&\stackrel{(d)}{=} (\mathbf{f}'_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k})), \tag{122}
\end{aligned}$$

where (a) follows from line 12 of Algorithm 3 and (6)–(7); (b) follows from (121); (c) follows from (117); and (d) follows from $\mathbf{V}^\top \mathbf{V} = \mathbf{I}$ and (6)–(7). Also, from lines 13–14 of Algorithm 3,

$$\gamma_{1,k+1} = \eta_{2k} - \gamma_{2k} = \gamma_{2k} \left[\frac{1}{\alpha_{2k}} - 1 \right]. \quad (123)$$

Equations (122) and (123) prove (120e). In addition,

$$\begin{aligned} \mathbf{p}_{k+1} &\stackrel{(a)}{=} \mathbf{r}_{1,k+1} - \mathbf{x}^0 \\ &\stackrel{(b)}{=} \frac{1}{1 - \alpha_{2k}} [\mathbf{g}_2(\mathbf{r}_{2k}, \gamma_{2k}) - \alpha_{2k} \mathbf{r}_{2k}] - \mathbf{x}^0 \\ &\stackrel{(c)}{=} \frac{1}{1 - \alpha_{2k}} \left[\mathbf{x}^0 + \mathbf{V} \mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}) - \alpha_{2k} (\mathbf{x}^0 + \mathbf{v}_k) \right] - \mathbf{x}^0 \\ &\stackrel{(d)}{=} \frac{1}{1 - \alpha_{2k}} [\mathbf{V} \mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}) - \alpha_{2k} \mathbf{v}_k] \\ &\stackrel{(e)}{=} \mathbf{V} \left[\frac{1}{1 - \alpha_{2k}} [\mathbf{f}_q(\mathbf{q}_k, \mathbf{w}^q, \gamma_{2k}) - \alpha_{2k} \mathbf{q}_k] \right], \end{aligned} \quad (124)$$

where (a) follows from (116); (b) follows from lines 11–15 of Algorithm 3; (c) follows from (121) and the definition of \mathbf{v}_k in (116); (d) follows from collecting the terms with \mathbf{x}^0 ; and (e) follows from the definition $\mathbf{q}_k = \mathbf{V}^\top \mathbf{v}_k$ in (117). Combining (124) with $\mathbf{u}_{k+1} = \mathbf{V}^\top \mathbf{p}_{k+1}$ proves (120f).

The derivation for the updates for \mathbf{v}_k are similar. First,

$$\alpha_{1k} \stackrel{(a)}{=} \langle \mathbf{g}'_1(\mathbf{r}_{1k}, \gamma_{1k}) \rangle \stackrel{(b)}{=} \langle \mathbf{f}'_p(\mathbf{p}_k, \mathbf{x}^0) \rangle, \quad (125)$$

where (a) follows from line 5 of Algorithm 3 and (b) follows from the vectorization of $\mathbf{f}_p(\cdot)$ in (119b) and the fact that $\mathbf{p}_k = \mathbf{r}_{1k} + \mathbf{x}^0$. Also, from lines 6–7 of Algorithm 3,

$$\gamma_{2k} = \eta_{1k} - \gamma_{1k} = \gamma_{1k} \left[\frac{1}{\alpha_{1k}} - 1 \right]. \quad (126)$$

Equations (125) and (126) prove (120b). Also,

$$\begin{aligned} \mathbf{v}_k &\stackrel{(a)}{=} \mathbf{r}_{2k} - \mathbf{x}^0 \\ &\stackrel{(b)}{=} \frac{1}{1 - \alpha_{1k}} [\mathbf{g}_1(\mathbf{r}_{1k}, \gamma_{1k}) - \alpha_{1k} \mathbf{r}_{1k}] - \mathbf{x}^0 \\ &\stackrel{(c)}{=} \frac{1}{1 - \alpha_{1k}} \left[\mathbf{f}_p(\mathbf{p}_k, \mathbf{x}^0, \gamma_{1k}) + \mathbf{x}^0 - \alpha_{1k} (\mathbf{p}_k + \mathbf{x}^0) \right] - \mathbf{x}^0 \\ &\stackrel{(d)}{=} \frac{1}{1 - \alpha_{1k}} [\mathbf{f}_p(\mathbf{p}_k, \mathbf{x}^0, \gamma_{1k}) - \alpha_{1k} \mathbf{p}_k] \end{aligned} \quad (127)$$

where (a) is the definition of \mathbf{v}_k in (116); (b) follows from lines 4–8 of Algorithm 3; (c) follows from the vectorization of $f_p(\cdot)$ in (119b) and the definition of \mathbf{p}_k in (116); and (d) follows from collecting the terms with \mathbf{x}^0 . This proves (120c). All together, we have proven (120) and the proof is complete.

APPENDIX H PROOF OF THEOREM 2

We use induction. Suppose that, for some k , $\bar{\gamma}_{1k} = \tau_{1k}^{-1}$. From (50a), (44) and (36),

$$\bar{\alpha}_{1k} = \bar{\gamma}_{1k} \mathcal{E}_1(\bar{\gamma}_{1k}). \quad (128)$$

Hence, from (50b), $\bar{\eta}_{1k}^{-1} = \mathcal{E}_1(\bar{\gamma}_{1k})$ and $\bar{\gamma}_{2k} = \bar{\eta}_{1k} - \bar{\gamma}_{1k}$. Also,

$$\begin{aligned} \tau_{2k} &\stackrel{(a)}{=} \frac{1}{(1 - \bar{\alpha}_{1k})^2} \left[\mathcal{E}_1(\bar{\gamma}_{1k}, \tau_{1k}) - \bar{\alpha}_{1k}^2 \tau_{1k} \right] \\ &\stackrel{(b)}{=} \frac{1}{(1 - \bar{\gamma}_{1k} \mathcal{E}_1(\bar{\gamma}_{1k}))^2} \left[\mathcal{E}_1(\bar{\gamma}_{1k}, \tau_{1k}) - \bar{\gamma}_{1k} \mathcal{E}_1^2(\bar{\gamma}_{1k}) \right] \\ &\stackrel{(c)}{=} \frac{\mathcal{E}_1(\bar{\gamma}_{1k}, \tau_{1k})}{1 - \bar{\gamma}_{1k} \mathcal{E}_1(\bar{\gamma}_{1k})} \\ &\stackrel{(d)}{=} \frac{1}{\bar{\eta}_{1k} - \bar{\gamma}_{1k}}, \end{aligned}$$

where (a) follows from (50c); (b) follows from (128) and the matched condition $\bar{\gamma}_{1k} = \tau_{1k}^{-1}$; (c) follows from canceling terms in the fraction and (d) follows from the fact that $\bar{\eta}_{1k}^{-1} = \mathcal{E}_1(\bar{\gamma}_{1k})$ and $\bar{\gamma}_{1k} = \bar{\eta}_{1k} / \bar{\alpha}_{1k}$. This proves (56a). A similar argument shows that (56b) holds if $\bar{\gamma}_{2k} = \tau_{2k}^{-1}$. Finally, (56c) follows from (56) and (55).

REFERENCES

- [1] M. Pereyra *et al.*, “A survey of stochastic simulation and optimization methods in signal processing,” *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 2, pp. 224–241, Mar. 2016.
- [2] D. L. Donoho, A. Maleki, and A. Montanari, “Message-passing algorithms for compressed sensing,” *Proc. Nat. Acad. Sci. USA*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.
- [3] D. L. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing: I. Motivation and construction,” in *Proc. IEEE Inf. Theory Workshop Inf. Theory (ITW)*, Cairo, Egypt, Jan. 2010, pp. 1–5.
- [4] 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.
- [5] M. Bayati, M. Lelarge, and A. Montanari, “Universality in polytope phase transitions and message passing algorithms,” *Ann. Appl. Probab.*, vol. 25, no. 2, pp. 753–822, 2015.
- [6] C. Rush and R. Venkataramanan, “Finite-sample analysis of approximate message passing,” in *Proc. IEEE ISIT*, Jul. 2016, pp. 755–759.
- [7] S. Rangan, P. Schniter, and A. Fletcher, “On the convergence of approximate message passing with arbitrary matrices,” in *Proc. IEEE ISIT*, Jul. 2014, pp. 236–240.
- [8] F. Caltagirone, L. Zdeborová, and F. Krzakala, “On convergence of approximate message passing,” in *Proc. IEEE ISIT*, Jun./Jul. 2014, pp. 1812–1816.
- [9] 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 ICASSP*, Apr. 2015, pp. 2021–2025.
- [10] A. Manoel, F. Krzakala, E. W. Tramel, and L. Zdeborová, “Swept approximate message passing for sparse estimation,” in *Proc. ICML*, 2015, pp. 1123–1132.
- [11] S. Rangan, A. K. Fletcher, P. Schniter, and U. S. Kamilov, “Inference for generalized linear models via alternating directions and Bethe free energy minimization,” in *Proc. IEEE ISIT*, Jun. 2015, pp. 1640–1644.
- [12] M. Opper and O. Winther, “Adaptive and self-averaging Thouless-Anderson-Palmer mean-field theory for probabilistic modeling,” *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 64, no. 5, 2001, Art. no. 056131.
- [13] T. P. Minka, “A family of algorithms for approximate Bayesian inference,” Ph.D. dissertation, Dept. Comp. Sci. Eng., MIT, Cambridge, MA, USA, 2001.
- [14] M. Seeger, “Expectation propagation for exponential families,” EPFL, EPFL-Rep. 161464, Lausanne, Switzerland, 2005.
- [15] M. Opper and O. Winther, “Expectation consistent approximate inference,” *J. Mach. Learn. Res.*, vol. 1, pp. 2177–2204, Dec. 2005.
- [16] Y. Kabashima and M. Vehkaperä, “Signal recovery using expectation consistent approximation for linear observations,” in *Proc. IEEE ISIT*, Jun. 2014, pp. 226–230.
- [17] A. K. Fletcher, M. Sahræe-Ardakan, S. Rangan, and P. Schniter, “Expectation consistent approximate inference: Generalizations and convergence,” in *Proc. IEEE ISIT*, Jul. 2016, pp. 190–194.

- [18] B. Çakmak, O. Winther, and B. H. Fleury, "S-AMP: Approximate message passing for general matrix ensembles," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Nov. 2014, pp. 192–196.
- [19] B. Çakmak, O. Winther, and B. H. Fleury, "S-AMP for non-linear observation models," in *Proc. IEEE ISIT*, Jun. 2015, pp. 2807–2811.
- [20] J. Ma and L. Ping, "Orthogonal AMP," *IEEE Access*, vol. 5, pp. 2020–2033, 2017.
- [21] 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.
- [22] N. Halko, P. G. Martinsson, and J. A. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions," *SIAM Rev.*, vol. 53, no. 2, pp. 217–288, 2011.
- [23] K. Takeuchi, "Rigorous dynamics of expectation-propagation-based signal recovery from unitarily invariant measurements," in *Proc. IEEE ISIT*, Jun. 2017, pp. 501–505.
- [24] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE ISIT*, Jul./Aug. 2011, pp. 2174–2178.
- [25] A. Montanari, "Graphical model concepts in compressed sensing," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2012, pp. 394–438.
- [26] A. Chambolle, R. A. De Vore, N.-Y. Lee, and B. J. Lucier, "Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 319–335, Mar. 1998.
- [27] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, Nov. 2004.
- [28] D. J. Thouless, P. W. Anderson, and R. G. Palmer, "Solution of 'Solvable model of a spin glass,'" *Philos. Mag.*, vol. 35, no. 3, pp. 983–990, 1977.
- [29] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA, USA: Morgan Kaufmann, 1988.
- [30] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Understanding belief propagation and its generalizations," in *Exploring Artificial Intelligence in the New Millennium*. San Francisco, CA, USA: Morgan Kaufmann, 2003, pp. 239–269.
- [31] S. Rangan, P. Schniter, E. Riegler, A. K. Fletcher, and V. Cevher, "Fixed points of generalized approximate message passing with arbitrary matrices," in *Proc. IEEE ISIT*, Jul. 2013, pp. 664–668.
- [32] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová, "Variational free energies for compressed sensing," in *Proc. IEEE ISIT*, Jun./Jul. 2014, pp. 1499–1503.
- [33] *Generalized Approximate Message Passing, Source-Forge. Net Project GAMP/MATLAB*. Accessed: May 24, 2019. [Online]. Available: <http://gampmatlab.sourceforge.net/>
- [34] A. K. Fletcher, M. Sahræe-Ardakan, S. Rangan, and P. Schniter, "Rigorous dynamics and consistent estimation in arbitrarily conditioned linear systems," in *Proc. NIPS*, 2017, pp. 2542–2551.
- [35] H. J. Brascamp and E. H. Lieb, "On extensions of the brunn-minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation," in *Inequalities*. Berlin, Germany: Springer, 2002, pp. 441–464.
- [36] A. Tulino and S. Verdú, "Random matrix theory and wireless communications," *Found. Trends Commun. Inf. Theory*, vol. 1, no. 1, pp. 1–182, 2004.
- [37] G. Reeves and H. D. Pfister, "The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact," in *Proc. IEEE ISIT*, Jul. 2016, pp. 665–669.
- [38] J. Barbier, M. Dia, N. Macris, and F. Krzakala, "The mutual information in random linear estimation," in *Proc. 54th Annu. Allerton Conf. Commun., Control, Comput.*, Sep. 2016, pp. 625–632. [Online]. Available: <https://arxiv.org/abs/1607.02335>

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, 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.

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 at The Ohio State University, Columbus, OH, USA.

Alyson K. Fletcher (S'03–M'04) received the B.S. degree in mathematics from the University of Iowa, Iowa City, IA, USA, and the M.S. degree in mathematics and electrical engineering (EE) and the Ph.D. degree in EE, 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, EE, and Computer Science at the University of California, Los Angeles.