

An Expectation-Maximization Approach to Tuning Generalized Vector Approximate Message Passing

Christopher A. Metzler¹, Philip Schniter², and Richard G. Baraniuk¹

¹ Rice University, Department of Electrical and Computer Engineering,
6100 Main St. Houston, TX 77005

² The Ohio State University, Department of Electrical and Computer Engineering,
2015 Neil Ave. Columbus, OH 43210

Abstract. Generalized Vector Approximate Message Passing (GVAMP) is an efficient iterative algorithm for approximately minimum-mean-squared-error estimation of a random vector $\mathbf{x} \sim p_{\mathbf{x}}(\mathbf{x})$ from generalized linear measurements, i.e., measurements of the form $\mathbf{y} = Q(\mathbf{z})$ where $\mathbf{z} = \mathbf{A}\mathbf{x}$ with known \mathbf{A} , and $Q(\cdot)$ is a noisy, potentially nonlinear, componentwise function. Problems of this form show up in numerous applications, including robust regression, binary classification, quantized compressive sensing, and phase retrieval. In some cases, the prior $p_{\mathbf{x}}$ and/or channel $Q(\cdot)$ depend on unknown deterministic parameters $\boldsymbol{\theta}$, which prevents a direct application of GVAMP. In this paper we propose a way to combine expectation maximization (EM) with GVAMP to jointly estimate \mathbf{x} and $\boldsymbol{\theta}$. We then demonstrate how EM-GVAMP can solve the phase retrieval problem with unknown measurement-noise variance.

Keywords: Expectation Maximization, Generalized Linear Model, Compressive Sensing, Phase Retrieval

1 Introduction

We consider the problem of estimating a random vector $\mathbf{x} \in \mathbb{R}^N$ from observations $\mathbf{y} \in \mathbb{R}^M$ generated as shown in Figure 1, which is known as the *generalized linear model* (GLM) [1]. Under this model, \mathbf{x} has a prior density $p_{\mathbf{x}}$ and \mathbf{y} obeys a likelihood function of the form $p(\mathbf{y}|\mathbf{x}) = p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{A}\mathbf{x})$, where $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a known linear transform and $\mathbf{z} \triangleq \mathbf{A}\mathbf{x}$ are hidden transform outputs. The conditional density $p_{\mathbf{y}|\mathbf{z}}$ can be interpreted as a probabilistic measurement channel that accepts a vector \mathbf{z} and outputs a random vector \mathbf{y} . Although we have assumed real-valued quantities for the sake of simplicity, it is straightforward to generalize the methods in this paper to complex-valued quantities.

The GLM has many applications in statistics, computer science, and engineering. For example, in *statistical regression*, \mathbf{A} and \mathbf{y} contain experimental features and outcomes, respectively, and \mathbf{x} are coefficients that best predict \mathbf{y} from \mathbf{A} . The relationship between \mathbf{y} and the optimal scores $\mathbf{z} = \mathbf{A}\mathbf{x}$ is then characterized by $p_{\mathbf{y}|\mathbf{z}}$. In *imaging*-related inverse problems, \mathbf{x} is an image to recover, \mathbf{A} is often Fourier-based, and $p_{\mathbf{y}|\mathbf{z}}$ models the sensor(s). In *communications* problems, \mathbf{x} may be a vector of discrete symbols to recover, in which case \mathbf{A} is a function of

2 An EM Approach to Tuning GVAMP




Fig. 1. Generalized Linear Model (GLM): An unknown random vector \mathbf{x} is observed through a linear transform \mathbf{A} followed by a probabilistic measurement channel $p_{\mathbf{y}|\mathbf{z}}$, yielding the measured vector \mathbf{y} .

the modulation/demodulation scheme and the propagation physics. Or, \mathbf{x} may contain propagation-channel parameters to recover, in which case \mathbf{A} is a function of the modulation/demodulation scheme and the pilot symbols. In both cases, $p_{\mathbf{y}|\mathbf{z}}$ models receiver hardware and interference.

Below we give some examples of the measurement channels $p_{\mathbf{y}|\mathbf{z}}$ that are encountered in these applications.

- *Regression* often models $\mathbf{y} = \mathbf{z} + \mathbf{w}$ with additive noise \mathbf{w} , and so $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z}) = p_{\mathbf{w}}(\mathbf{y} - \mathbf{z})$, where $p_{\mathbf{w}}$ is the density of \mathbf{w} . The “standard linear model” treats \mathbf{w} as additive white Gaussian noise (AWGN) but is not robust to outliers. Robust methods typically use heavy-tailed models for \mathbf{w} .
- *Binary linear classification* can be modeled using $y_m = \text{sgn}(z_m + w_m)$, where $\text{sgn}(v) = 1$ for $v \geq 0$ and $\text{sgn}(v) = -1$ for $v < 0$, and w_m are errors. Gaussian w_m yields the “probit” model and logistic w_m yields the “logistic” model.
- *Quantized compressive sensing* models $y_m = Q(z_m + w_m)$, where $Q(\cdot)$ is a scalar quantizer and w_m is additive, often AWGN.
- *Phase retrieval* models $y_m = |z_m|$ in the noiseless case, where $z_n \in \mathbb{C}$. When noise is present, one approach is to model $y_m = |z_m + w_m|$ with $w_m \in \mathbb{C}$ and another is to model $y_m = |z_m|^2 + w_m$ with real-valued w_m .

In this work, we focus on the case that the prior $p_{\mathbf{x}}$ and the likelihood $p_{\mathbf{y}|\mathbf{z}}$ depend on parameters $\boldsymbol{\theta}$ that are a priori unknown. For example, the prior $p_{\mathbf{x}}$ might be Bernoulli-Gaussian with unknown sparsity rate and variance, and the likelihood might involve an additive noise of an unknown variance. We are interested in jointly estimating \mathbf{x} and $\boldsymbol{\theta}$ from \mathbf{y} , where $\boldsymbol{\theta}$ are treated as deterministic. In particular, we aim to compute the ML estimate of $\boldsymbol{\theta}$ and the MMSE estimate of \mathbf{x} under $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_{\text{ML}}$:

$$\hat{\boldsymbol{\theta}}_{\text{ML}} = \arg \max_{\boldsymbol{\theta}} p(\mathbf{y}; \boldsymbol{\theta}) \quad (1a)$$

$$\hat{\mathbf{x}}_{\text{MMSE}} = \mathbb{E}\{\mathbf{x}|\mathbf{y}; \hat{\boldsymbol{\theta}}_{\text{ML}}\}, \quad (1b)$$

sometimes referred to as the “empirical Bayesian” approach.

For most priors and likelihoods of interest, exact computation of the conditional mean in (1b) is intractable. Thus we might settle for an approximation of the MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}}$. In the case that \mathbf{A} is well modeled as a realization of a large rotationally invariant random matrix, the generalized vector approximate message passing (GVAMP) algorithm [2] is a computationally efficient approach to approximate-MMSE inference under the GLM in Figure 1. In the large system limit (i.e., $M, N \rightarrow \infty$ with $M/N \rightarrow \delta \in (0, 1)$), it is rigorously characterized by state-evolution whose fixed points, when unique, are Bayes optimal [3].

For the special case of an AWGN likelihood, i.e., $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{y}; \mathbf{z}, \nu_w \mathbf{I})$ for some $\nu_w > 0$, GVAMP reduces to the simpler VAMP algorithm [4]. By merging VAMP with expectation maximization (EM) [5], one obtains the ‘‘EM-VAMP’’ approach [6] to the empirical-Bayesian estimation problem (1). In fact, with large right-rotationally invariant \mathbf{A} , EM-VAMP is rigorously characterized by state-evolution [7]. Furthermore, under some identifiability conditions, it is possible to show that EM-VAMP yields an asymptotically efficient estimate of $\boldsymbol{\theta}$.

In this paper, we propose a way to merge EM and GVAMP to tackle GLMs of the form discussed above. This yields, for example, a way to handle phase retrieval with unknown measurement-noise variance. The proposed ‘‘EM-GVAMP’’ approach is described in the next section.

2 EM-GVAMP

In the sequel we assume a GLM of the form

$$p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z) = \prod_{i=1}^M p(y_i|z_i; \boldsymbol{\theta}_z), \quad \mathbf{z} = \mathbf{A}\mathbf{x}, \quad p(\mathbf{x}; \boldsymbol{\theta}_x) = \prod_{j=1}^N p(x_j; \boldsymbol{\theta}_x), \quad (2)$$

where $\boldsymbol{\theta} \triangleq [\boldsymbol{\theta}_x, \boldsymbol{\theta}_z]$ are unknown deterministic parameters, and where $\mathbf{z} \in \mathbb{R}^M$ and $\mathbf{x} \in \mathbb{R}^N$.

2.1 The EM Algorithm

Recalling the empirical-Bayesian methodology (1), the maximum-likelihood estimate of $\boldsymbol{\theta}$ given \mathbf{y} can be written as

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \{ -\ln p(\mathbf{y}; \boldsymbol{\theta}) \}, \quad (3)$$

where

$$\begin{aligned} p(\mathbf{y}; \boldsymbol{\theta}) &= \int p(\mathbf{y}, \mathbf{z}, \mathbf{x}; \boldsymbol{\theta}) \, d\mathbf{z} \, d\mathbf{x} = \int p(\mathbf{y}|\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}) p(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}) \, d\mathbf{z} \, d\mathbf{x} \\ &= \int p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z) \delta(\mathbf{z} - \mathbf{A}\mathbf{x}) p(\mathbf{x}; \boldsymbol{\theta}_x) \, d\mathbf{z} \, d\mathbf{x}. \end{aligned} \quad (4)$$

Although $p(\mathbf{y}; \boldsymbol{\theta})$ is difficult to work with directly, the expectation-maximization (EM) algorithm [8] offers an alternative. There, the idea is to write

$$-\ln p(\mathbf{y}; \boldsymbol{\theta}) = J(b; \boldsymbol{\theta}) - D(b \| p(\mathbf{x}|\mathbf{y}; \boldsymbol{\theta})) \quad (5)$$

for an arbitrary belief $b(\mathbf{x})$, where $D(\cdot \| \cdot)$ is KL divergence,

$$J(b; \boldsymbol{\theta}) \triangleq D(b \| p(\mathbf{x}; \boldsymbol{\theta}_x)) + D(b \| p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_z)) + H(b) \quad (6)$$

is known as the Gibbs free energy, and $H(b)$ is the entropy of b . Because $D(b \| p(\mathbf{x}|\mathbf{y}; \boldsymbol{\theta})) \geq 0$ for any b , we have that $J(b; \boldsymbol{\theta})$ is an upper bound on $-\ln p(\mathbf{y}; \boldsymbol{\theta})$, the quantity that ML seeks to minimize. Thus, if it is tractable to

4 An EM Approach to Tuning GVAMP

construct and minimize $J(b; \boldsymbol{\theta})$, it makes sense to iterate the following two steps (over $k = 1, 2, \dots$):

$$\text{E step: } b^k(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}; \hat{\boldsymbol{\theta}}^k) \quad (7)$$

$$\text{M step: } \hat{\boldsymbol{\theta}}^{k+1} = \arg \min_{\boldsymbol{\theta}} J(b^k; \boldsymbol{\theta}) = \arg \min_{\boldsymbol{\theta}} D(b^k \| p(\mathbf{x}; \boldsymbol{\theta}_x)) + D(b^k \| p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_z)), \quad (8)$$

which together constitute the EM algorithm. The ‘‘E’’ step creates an upper bound on $-\ln p(\mathbf{y}; \boldsymbol{\theta})$ that is tight at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}^k$, and the ‘‘M’’ step finds the estimate of $\boldsymbol{\theta}$ that minimizes this bound.

Unfortunately, however, the posterior density required by the E-step (7),

$$p(\mathbf{x}|\mathbf{y}; \boldsymbol{\theta}) = \frac{p(\mathbf{x}; \boldsymbol{\theta}_x)p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_z)}{p(\mathbf{y}; \boldsymbol{\theta})} = \frac{p(\mathbf{x}; \boldsymbol{\theta}_x) \int p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)\delta(\mathbf{z} - \mathbf{Ax}) \, d\mathbf{z}}{\int p(\mathbf{x}; \boldsymbol{\theta}_x)p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)\delta(\mathbf{z} - \mathbf{Ax}) \, d\mathbf{z} \, d\mathbf{x}}, \quad (9)$$

is difficult to compute due to the high-dimensional integration. Thus we consider an approximation afforded by the GVAMP algorithm [2]. For this, we first reparameterize the GLM (2) as a standard linear model (SLM).

2.2 An SLM Equivalent

The GLM (2) can be written as an SLM using the following formulation:

$$\bar{\mathbf{y}} = \bar{\mathbf{A}}\bar{\mathbf{x}} + \bar{\mathbf{w}} \quad \text{with} \quad \bar{\mathbf{y}} \triangleq \mathbf{0}, \quad \bar{\mathbf{A}} \triangleq [\mathbf{A} \quad -\mathbf{I}_M], \quad \bar{\mathbf{x}} \triangleq \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \end{bmatrix}, \quad \bar{\mathbf{w}} \sim \mathcal{N}(\mathbf{0}, \epsilon \mathbf{I}_M) \text{ s.t. } \epsilon \rightarrow 0. \quad (10)$$

Here, \mathbf{x} is apriori independent of \mathbf{z} ; the dependence between \mathbf{x} and \mathbf{z} manifests only aposteriori, i.e., after the measurement $\bar{\mathbf{y}}$ is observed. For \mathbf{x} , we assign the prior $p(\mathbf{x}; \boldsymbol{\theta}_x)$, and for \mathbf{z} we assign the *improper* (i.e., unnormalized) prior $p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)$. The lack of normalization will not be an issue in GVAMP, because the ‘‘prior’’ $p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)$ is used only to compute posteriors of the form

$$p(\mathbf{z}|\mathbf{y}; \hat{\mathbf{p}}, \tau, \boldsymbol{\theta}_z) \propto p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)\mathcal{N}(\mathbf{z}; \hat{\mathbf{p}}, \mathbf{I}/\tau), \quad (11)$$

which are well defined because the right side is always integrable over \mathbf{z} .

Let us first consider direct ML estimation of $\boldsymbol{\theta}$ in the above SLM. The $\boldsymbol{\theta}$ -likelihood function is

$$\begin{aligned} p(\bar{\mathbf{y}}; \boldsymbol{\theta}) &= \int p(\bar{\mathbf{y}}, \bar{\mathbf{x}}; \boldsymbol{\theta}) \, d\bar{\mathbf{x}} = \int p(\bar{\mathbf{y}}|\bar{\mathbf{x}})p(\bar{\mathbf{x}}; \boldsymbol{\theta}) \, d\bar{\mathbf{x}} = \int \mathcal{N}(\bar{\mathbf{y}}; \bar{\mathbf{A}}\bar{\mathbf{x}}, \epsilon \mathbf{I})p(\bar{\mathbf{x}}; \boldsymbol{\theta}) \, d\bar{\mathbf{x}} \\ &= \int \underbrace{\mathcal{N}(\mathbf{z}; \mathbf{Ax}, \epsilon \mathbf{I})}_{\rightarrow \delta(\mathbf{z} - \mathbf{Ax})} p(\mathbf{x}; \boldsymbol{\theta}_x)p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z) \, d\mathbf{x} \, d\mathbf{z}, \end{aligned} \quad (12)$$

which is consistent with (4) as $\epsilon \rightarrow 0$. Likewise, for any belief $b(\bar{\mathbf{x}})$, we can upper bound the negative log-likelihood by a Gibbs free energy $\bar{J}(b; \boldsymbol{\theta})$ of the form

$$\bar{J}(b; \boldsymbol{\theta}) \triangleq D(b \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})) + D(b \| p(\bar{\mathbf{y}}|\bar{\mathbf{x}})) + H(b), \quad (13)$$

since $-\ln p(\bar{\mathbf{y}}; \boldsymbol{\theta}) = \bar{J}(b; \boldsymbol{\theta}) - D(b \| p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta}))$ with $D(b \| p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta})) \geq 0$. The corresponding EM algorithm is

$$\text{E step: } b^k(\bar{\mathbf{x}}) = p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \hat{\boldsymbol{\theta}}^k) \quad (14)$$

$$\text{M step: } \hat{\boldsymbol{\theta}}^{k+1} = \arg \min_{\boldsymbol{\theta}} \bar{J}(b^k; \boldsymbol{\theta}) = \arg \min_{\boldsymbol{\theta}} D(b^k \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})). \quad (15)$$

As before, the posterior density required by the E-step (14)

$$p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta}) = \frac{p(\bar{\mathbf{x}}; \boldsymbol{\theta})p(\bar{\mathbf{y}}|\bar{\mathbf{x}})}{p(\bar{\mathbf{y}}; \boldsymbol{\theta})} = \frac{p(\mathbf{x}; \boldsymbol{\theta}_x)p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)\delta(\mathbf{z} - \mathbf{A}\mathbf{x})}{\int p(\mathbf{x}; \boldsymbol{\theta}_x)p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)\delta(\mathbf{z} - \mathbf{A}\mathbf{x}) d\mathbf{z} d\mathbf{x}}, \quad (16)$$

is difficult to compute due to the high-dimensional integral. Thus we consider an approximation afforded by the GVAMP algorithm [2], as described in the next section.

2.3 GVAMP

Recall that the exact posterior can (in principle) be found by solving the variational optimization problem

$$p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta}) = \arg \min_b D(b \| p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta})) \quad (17)$$

$$= \arg \min_b \bar{J}(b; \boldsymbol{\theta}) \quad (18)$$

$$= \arg \min_b D(b \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})) + D(b \| p(\bar{\mathbf{y}}|\bar{\mathbf{x}})) + H(b), \quad (19)$$

where (18) follows from $\bar{J}(b; \boldsymbol{\theta}) = D(b \| p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta})) - \ln p(\bar{\mathbf{y}}; \boldsymbol{\theta})$ and (19) follows from (13). But since the posterior computation problem is NP hard in general, (19) is no more tractable than any other approach. The GVAMP algorithm computes a posterior approximation using the expectation-consistent (EC) method [9, 10]. In this application of EC, we first split $b(\bar{\mathbf{x}})$ into three copies, i.e.,

$$p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta}) = \arg \min_{b_1=b_2=q} D(b_1 \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})) + D(b_2 \| p(\bar{\mathbf{y}}|\bar{\mathbf{x}})) + H(q), \quad (20)$$

and then relax the density-matching constraint $b_1 = b_2 = q$ to a moment-matching constraint:

$$p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \boldsymbol{\theta}) \approx \arg \min_{b_1, b_2, q} D(b_1 \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})) + D(b_2 \| p(\bar{\mathbf{y}}|\bar{\mathbf{x}})) + H(q) \quad (21)$$

$$\begin{aligned} \text{s.t. } \mathbb{E}[\bar{\mathbf{x}}|b_1] &= \mathbb{E}[\bar{\mathbf{x}}|b_2] = \mathbb{E}[\bar{\mathbf{x}}|q] \quad \text{and } \text{tr}_2\{\text{Cov}[\bar{\mathbf{x}}|b_1]\} \\ &= \text{tr}_2\{\text{Cov}[\bar{\mathbf{x}}|b_2]\} = \text{tr}_2\{\text{Cov}[\bar{\mathbf{x}}|q]\}, \end{aligned} \quad (22)$$

where $\mathbb{E}[\bar{\mathbf{x}}|b_i]$ and $\text{Cov}[\bar{\mathbf{x}}|b_i]$ denote the expectation and covariance of $\bar{\mathbf{x}}$ under $\bar{\mathbf{x}} \sim b_i(\bar{\mathbf{x}})$, and where

$$\text{tr}_2 \left(\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{bmatrix} \right) \triangleq \begin{bmatrix} \text{tr}(\mathbf{A}) \\ \text{tr}(\mathbf{C}) \end{bmatrix} \quad \text{for } \mathbf{A} \in \mathbb{R}^{N \times N} \text{ and } \mathbf{C} \in \mathbb{R}^{M \times M}. \quad (23)$$

6 An EM Approach to Tuning GVAMP

Essentially, $\text{tr}_2\{\text{Cov}[\bar{\mathbf{x}}]\}$ separately computes the trace of the covariance of \mathbf{x} and the trace of the covariance of \mathbf{z} . The right side of (21) yields three different approximations of the posterior:

$$b_1(\bar{\mathbf{x}}; \boldsymbol{\theta}) \propto p(\bar{\mathbf{x}}; \boldsymbol{\theta}) \mathcal{N}\left(\bar{\mathbf{x}}; \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{p}_1 \end{bmatrix}, \begin{bmatrix} \mathbf{I}_N/\gamma_1 & \\ & \mathbf{I}_M/\tau_1 \end{bmatrix}\right) \quad (24)$$

$$b_2(\bar{\mathbf{x}}; \boldsymbol{\theta}) \propto p(\bar{\mathbf{y}}|\bar{\mathbf{x}}) \mathcal{N}\left(\bar{\mathbf{x}}; \begin{bmatrix} \mathbf{r}_2 \\ \mathbf{p}_2 \end{bmatrix}, \begin{bmatrix} \mathbf{I}_N/\gamma_2 & \\ & \mathbf{I}_M/\tau_2 \end{bmatrix}\right) \quad (25)$$

$$q(\bar{\mathbf{x}}; \boldsymbol{\theta}) \propto \mathcal{N}\left(\bar{\mathbf{x}}; \begin{bmatrix} \hat{\mathbf{x}} \\ \hat{\mathbf{z}} \end{bmatrix}, \begin{bmatrix} \mathbf{I}_N/\eta & \\ & \mathbf{I}_M/\zeta \end{bmatrix}\right), \quad (26)$$

where the form of (24)-(26) can be deduced by analyzing the stationary points of the Lagrangian of (21), as shown in [9].

The GVAMP algorithm is an iterative approach to finding the values of $\mathbf{r}_1, \gamma_1, \mathbf{p}_1, \tau_1, \mathbf{r}_2, \gamma_2, \mathbf{p}_2, \tau_2, \hat{\mathbf{x}}, \eta, \hat{\mathbf{z}}, \zeta$ under which the three beliefs in (24)-(26) obey the moment constraints in (21). When \mathbf{A} is large and rotationally invariant, GVAMP is rigorously characterized by a state evolution [3]. Empirically, we find that the algorithm converges quickly in this scenario (e.g., on the order of 10 iterations).

Note that the values of $\mathbf{r}_1, \gamma_1, \mathbf{p}_1, \tau_1, \mathbf{r}_2, \gamma_2, \mathbf{p}_2, \tau_2, \hat{\mathbf{x}}, \eta, \hat{\mathbf{z}}, \zeta$ that satisfy the moment constraints are interdependent, and thus they all depend on the assumed value of $\boldsymbol{\theta}$ through (24).

2.4 EM-GVAMP

Recall that our current motivation for using GVAMP is to compute an approximation to the posterior $b^k(\bar{\mathbf{x}}) = p(\bar{\mathbf{x}}|\bar{\mathbf{y}}; \hat{\boldsymbol{\theta}}^k)$ in the EM algorithm (14)-(15). Of the three posterior approximations produced by GVAMP, the Gaussian approximation from (26) is the simplest to use for this purpose. Plugging the Gaussian approximation into (14)-(15) yields

$$\text{E step: } b^k(\bar{\mathbf{x}}) = \mathcal{N}\left(\bar{\mathbf{x}}; \begin{bmatrix} \hat{\mathbf{x}}^k \\ \hat{\mathbf{z}}^k \end{bmatrix}, \begin{bmatrix} \mathbf{I}_N/\eta^k & \\ & \mathbf{I}_M/\zeta^k \end{bmatrix}\right) \text{ found via GVAMP with } \boldsymbol{\theta} = \hat{\boldsymbol{\theta}}^k \quad (27)$$

$$\text{M step: } \hat{\boldsymbol{\theta}}^{k+1} = \arg \min_{\boldsymbol{\theta}} D(b^k \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})). \quad (28)$$

The difference between the EM algorithm (14)-(15) and the EM algorithm (27)-(28) is that, in the former case, the bound is tight at each EM iteration k , whereas in the latter case the bound is only approximately tight.

Due to the form of $b^k(\bar{\mathbf{x}})$ in (27), the M-step is relatively easy to compute:

$$\hat{\boldsymbol{\theta}}^{k+1} = \arg \min_{\boldsymbol{\theta}} D(b^k \| p(\bar{\mathbf{x}}; \boldsymbol{\theta})) \quad (29)$$

$$= \arg \min_{\boldsymbol{\theta}} D(\mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}^k, \mathbf{I}_N/\eta^k) \mathcal{N}(\mathbf{z}; \hat{\mathbf{z}}^k, \mathbf{I}_M/\zeta^k) \| p(\mathbf{x}; \boldsymbol{\theta}_x) p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z)) \quad (30)$$

$$= \arg \max_{\boldsymbol{\theta}} \int \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}^k, \mathbf{I}_N/\eta^k) \ln p(\mathbf{x}; \boldsymbol{\theta}_x) d\mathbf{x} \\ + \int \mathcal{N}(\mathbf{z}; \hat{\mathbf{z}}^k, \mathbf{I}_M/\zeta^k) \ln p(\mathbf{y}|\mathbf{z}; \boldsymbol{\theta}_z) d\mathbf{z} \quad (31)$$

$$= \arg \max_{\boldsymbol{\theta}} \sum_{j=1}^N \int \mathcal{N}(x_j; \hat{x}_j^k, 1/\eta^k) \ln p(x_j; \boldsymbol{\theta}_x) dx_j \\ + \sum_{i=1}^M \int \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \ln p(y_i|z_i; \boldsymbol{\theta}_z) dz_i. \quad (32)$$

The resulting $(\hat{\boldsymbol{\theta}}_x^{k+1}, \hat{\boldsymbol{\theta}}_z^{k+1})$ are necessarily values of $(\boldsymbol{\theta}_x, \boldsymbol{\theta}_z)$ that zero the gradient of the right side of (32) with respect to $\boldsymbol{\theta}_x$ and to $\boldsymbol{\theta}_z$.

3 Application to Noise-Variance Estimation in Phase Retrieval

In this section we will demonstrate how the EM procedure can be used to estimate noise variances in the context of phase retrieval. Noise variance estimation in this setting has also been performed in [11] and [12]. The below derivation is related to, but distinct from, these previous works.

Phase retrieval is a problem that can be formulated in the GLM setting [11], allowing application of the GVAMP algorithm [13]. We denote the special case of GVAMP applied to phase retrieval as prVAMP.

One way to model the i th measured intensity y_i is via

$$y_i = |z_i + w_i| \text{ for i.i.d. } w_i \sim \mathcal{N}(0, \nu_w), \quad (33)$$

where $z_i, w_i \in \mathbb{C}$ and $\mathcal{N}(w_i; \mu, \nu) = \frac{1}{\pi\nu} \exp(-|w_i - \mu|^2/\nu)$ represents a circular complex-Gaussian density with mean $\mu \in \mathbb{C}$ and variance $\nu > 0$. In this case, the measurement noise variance ν_w may be unknown in practice, and so we might try to estimate it using the methods described in this report. In that case, the unknown \mathbf{z} -likelihood parameters " $\boldsymbol{\theta}_z$ " reduce to ν_w . In the sequel, we will use the notation ν_w instead of $\boldsymbol{\theta}_z$.

It was shown [11] that, under (33), the z_i -likelihood function $p(y_i|z_i; \nu_w)$ takes the form

$$p(y_i|z_i; \nu_w) = \mathbf{1}_{y_i \geq 0} y_i \int_0^{2\pi} \mathcal{N}(y_i e^{j\theta_i}; z_i, \nu_w) d\theta_i \quad (34)$$

$$= \frac{2y_i}{\nu_w} \exp\left(-\frac{y_i^2 + |z_i|^2}{\nu_w}\right) I_0\left(\frac{2y_i|z_i|}{\nu_w}\right) \mathbf{1}_{y_i \geq 0}, \quad (35)$$

8 An EM Approach to Tuning GVAMP

where $I_0(\cdot)$ is the 0th-order modified Bessel function of the first kind. If we view $p(y_i|z_i; \nu_w)$ as a density on y_i , then y_i is Rician (conditional on z_i). Note that θ_i above denotes the (hidden) phase on $z_i + w_i$; it should not be confused with the statistical parameters θ described earlier in this paper.

From (32), we see that the EM estimate $\hat{\nu}_w^{k+1}$ of ν_w must obey

$$0 = \frac{\partial}{\partial \nu_w} \sum_{i=1}^M \int_{\mathbb{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \ln p(y_i|z_i; \hat{\nu}_w^{k+1}) dz_i \quad (36)$$

$$= \sum_{i=1}^M \int_{\mathbb{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \frac{\partial}{\partial \nu_w} \ln \int_0^{2\pi} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i dz_i \quad (37)$$

$$= \sum_{i=1}^M \int_{\mathbb{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \frac{\int_0^{2\pi} \frac{\partial}{\partial \nu_w} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i}{\int_0^{2\pi} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i} dz_i. \quad (38)$$

Plugging in the derivative expression (see [14])

$$\frac{\partial}{\partial \nu_w} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) = \frac{\mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1})}{2(\hat{\nu}_w^{k+1})^2} (|y_i e^{j\theta_i} - z_i|^2 - \hat{\nu}_w^{k+1}) \quad (39)$$

into (38) and multiplying both sides by $2(\hat{\nu}_w^{k+1})^2$, we find

$$\hat{\nu}_w^{k+1} = \frac{1}{M} \sum_{i=1}^M \int_{\mathbb{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \frac{\int_0^{2\pi} |y_i e^{j\theta_i} - z_i|^2 \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i}{\int_0^{2\pi} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i} dz_i \quad (40)$$

$$= \frac{1}{M} \sum_{i=1}^M \int_{\mathbb{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \int_0^{2\pi} |y_i e^{j\theta_i} - z_i|^2 p(\theta_i; z_i, \hat{\nu}_w^{k+1}) d\theta_i dz_i \quad (41)$$

with the newly defined pdf

$$p(\theta_i; z_i, \hat{\nu}_w^{k+1}) \triangleq \frac{\mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1})}{\int_0^{2\pi} \mathcal{N}(y_i e^{j\theta_i}; z_i, \hat{\nu}_w^{k+1}) d\theta_i} \propto \exp\left(-\frac{|z_i - y_i e^{j\theta_i}|^2}{\hat{\nu}_w^{k+1}}\right) \quad (42)$$

$$\propto \exp(\kappa_i \cos(\theta_i - \phi_i)) \text{ for } \kappa_i \triangleq \frac{2|z_i| |y_i|}{\hat{\nu}_w^{k+1}}, \quad (43)$$

where ϕ_i denotes the phase of z_i . The expression (43) identifies this pdf as a von Mises distribution [15], which can be stated in normalized form as

$$p(\theta_i; z_i, \hat{\nu}_w^{k+1}) = \frac{\exp(\kappa_i \cos(\theta_i - \phi_i))}{2\pi I_0(\kappa_i)}. \quad (44)$$

Expanding the quadratic in (41) and plugging in (44), we get

$$\begin{aligned} \hat{v}_w^{k+1} &= \frac{1}{M} \sum_{i=1}^M \int_{\mathcal{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \left(y_i^2 + |z_i|^2 \right. \\ &\quad \left. - 2y_i |z_i| \int_0^{2\pi} \cos(\theta_i - \phi_i) \frac{\exp(\kappa_i \cos(\theta_i - \phi_i))}{2\pi I_0(\kappa_i)} d\theta_i \right) dz_i \end{aligned} \quad (45)$$

$$= \frac{1}{M} \sum_{i=1}^M \int_{\mathcal{C}} \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) \left(y_i^2 + |z_i|^2 - 2y_i |z_i| R_0 \left(\frac{2|z_i|y_i}{\hat{v}_w^{k+1}} \right) \right) dz_i, \quad (46)$$

where $R_0(\cdot)$ is the modified Bessel function ratio $R_0(\kappa_i) \triangleq I_1(\kappa_i)/I_0(\kappa_i)$ and (46) follows from [16, 9.6.19].

Simplifying approximations of (46) could be taken as needed. For example, in the high-SNR case, the expansion $R_0(\kappa) = 1 - \frac{1}{2\kappa} - \frac{1}{8\kappa^2} - \frac{1}{8\kappa^3} + o(\kappa^{-3})$ from [17, Lemma 5] could be used to justify

$$R_0(\kappa) \approx 1 - \frac{1}{2\kappa}, \quad (47)$$

which, when applied to (46), yields

$$\hat{v}_w^{k+1} \approx \frac{2}{M} \sum_{i=1}^M \int_{\mathcal{C}} (y_i - |z_i|)^2 \mathcal{N}(z_i; \hat{z}_i^k, 1/\zeta^k) dz_i. \quad (48)$$

Approximation (48) can be reduced to an expression that involves the mean of a Rician distribution. In particular, using $z_i = \rho_i e^{j\phi_i}$, the integral in (48) can be converted to polar coordinates as follows:

$$\begin{aligned} \int_0^\infty (y_i - \rho_i)^2 \underbrace{\int_0^{2\pi} \mathcal{N}(\rho_i e^{j\phi_i}; \hat{z}_i^k, 1/\zeta^k) d\phi_i}_{\text{Ricean density}} \rho_i d\rho_i &= y_i^2 - 2y_i \mathbb{E}[\rho_i] + \mathbb{E}[\rho_i^2], \quad (49) \\ \frac{2\rho_i}{1/\zeta^k} \exp\left(-\frac{\rho_i^2 + |\hat{z}_i^k|^2}{1/\zeta^k}\right) I_0\left(\frac{2\rho_i |\hat{z}_i^k|}{1/\zeta^k}\right) &1_{\rho_i \geq 0} \end{aligned}$$

where, for the expectations, ρ_i has the Rician density under the brace. For this density, it is known that

$$\mathbb{E}[\rho_i] = \sqrt{\frac{\pi}{4\zeta^k}} L_{1/2}(-\zeta^k |\hat{z}_i^k|^2) \quad (50)$$

$$\mathbb{E}[\rho_i^2] = 1/\zeta^k + |\hat{z}_i^k|^2, \quad (51)$$

where the Laguerre polynomial $L_{1/2}(x)$ can be computed as

$$L_{1/2}(x) = \exp\left(\frac{x}{2}\right) \left[(1-x)I_0\left(-\frac{x}{2}\right) - xI_1\left(-\frac{x}{2}\right) \right]. \quad (52)$$

Note that, for reasons of numerical precision, $\exp(x/2)I_d(-x/2)$ is computed using “`besseli(d, -x/2, 1)`” in Matlab, not “`exp(x/2) .* besseli(d, -x/2)`.”

10 An EM Approach to Tuning GVAMP

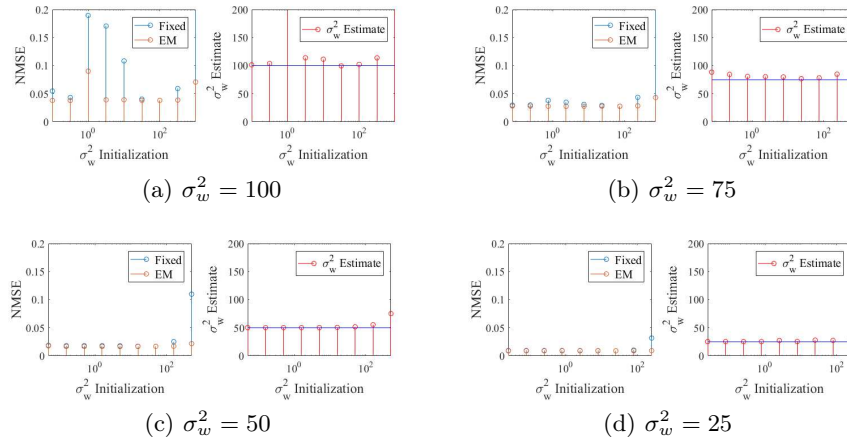


Fig. 2. Reconstruction errors (left subplots) and estimates of σ_w^2 (right subplots) with different initial estimates of σ_w^2 . The EM procedure is capable of estimating the true noise variance over a range of operating conditions. Using this estimate of the noise variance incrementally improves recovery accuracy.

4 Simulations

In this section, we demonstrate the effectiveness of the EM procedure in simulation. In particular, we show how EM can approximately recover the noise variance even when initialized by estimates far from the ground truth. This in turn enables improved signal reconstruction when the noise variance is a priori unknown.

We set up our simulations as follows. We aim to recover an i.i.d. circular Gaussian random vector $x \in \mathbb{C}^n$, with variance $\sqrt{2}$, from phaseless noisy measurements of the form $y = |\mathbf{A}x + w|$. Our measurement matrix \mathbf{A} is 8192×1024 and the elements of A are i.i.d. circular Gaussian with variance $\sqrt{2}$. The elements of the noise vector w also follow an i.i.d. circular Gaussian distribution, but with variance σ_w^2 . We test the cases of $\sigma_w^2 = 100$, $\sigma_w^2 = 75$, $\sigma_w^2 = 50$, and $\sigma_w^2 = 25$. prVAMP was provided with initial estimates of σ_w^2 ranging from 1% to $10\times$ the true variance. Using these initializations, we reconstructed the signal with and without the EM procedure.

Figure 2 presents our reconstructions. The results demonstrate that EM can be used to estimate σ_w^2 . Moreover, it shows that this estimate lets prVAMP accurately reconstruct the signal even when σ_w is not known a priori.

Code demonstrating the EM procedure will be made available at http://gampmatlab.wikia.com/wiki/Generalized_Approximate_Message_Passing.

5 Conclusion

This paper combines EM and GVAMP to estimate the unknown channel parameters associated with GLMs. This in turn enables GVAMP to estimate signals from their generalized linear measurements. In this paper we applied the proposed technique to phase retrieval and showed that it is effective at estimating unknown noise variances, thus enabling noise robust phase retrieval over a range of operating conditions.

References

1. P. McCullagh and J. A. Nelder, *Generalized Linear Models*. Chapman & Hall, 2nd ed., 1989.
2. P. Schniter, S. Rangan, and A. K. Fletcher, "Vector approximate message passing for the generalized linear model," in *Asilomar Conf. Sig., Sys., Comput.*, pp. 1525–1529, 2016.
3. A. K. Fletcher, S. Rangan, and P. Schniter, "Inference in deep networks in high dimensions," *arXiv preprint arXiv:1706.06549*, 2017.
4. S. Rangan, P. Schniter, and A. K. Fletcher, "Vector approximate message passing," in *Proc. IEEE ISIT*, pp. 1588–1592, 2017.
5. A. Dempster, N. M. Laird, and D. B. Rubin, "Maximum-likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc.*, vol. 39, pp. 1–17, 1977.
6. A. K. Fletcher and P. Schniter, "Learning and free energies for vector approximate message passing," in *IEEE Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4247–4251, 2017.
7. A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Rigorous dynamics and consistent estimation in arbitrarily conditioned linear systems," in *Proc. NIPS*, pp. 2542–2551, 2017.
8. R. Neal and G. Hinton, "A view of the EM algorithm that justifies incremental, sparse, and other variants," in *Learning in Graphical Models* (M. I. Jordan, ed.), pp. 355–368, MIT Press, 1998.
9. M. Opper and O. Winther, "Expectation consistent approximate inference," *J. Mach. Learning Res.*, vol. 1, pp. 2177–2204, 2005.
10. A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Expectation consistent approximate inference: Generalizations and convergence," in *Proc. IEEE ISIT*, pp. 190–194, 2016.
11. P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1043–1055, 2015.
12. A. Drémeau and F. Krzakala, "Phase recovery from a bayesian point of view: the variational approach," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 3661–3665, IEEE, 2015.
13. C. A. Metzler, M. K. Sharma, S. Nagesh, R. G. Baraniuk, O. Cossairt, and A. Veeraraghavan, "Coherent inverse scattering via transmission matrices: Efficient phase retrieval algorithms and a public dataset," in *Proc. Int. Conf. Computational Photography (ICCP)*, May 2017.
14. J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Trans. Signal Processing*, vol. 61, pp. 4658–4672, Oct. 2013.
15. K. V. Mardia and P. E. Jupp, *Directional Statistics*. New York: Wiley, 2000.
16. M. Abramowitz and I. A. Stegun, eds., *Handbook of Mathematical Functions*. New York: Dover, 1964.
17. C. Robert, "Modified Bessel functions and their applications in probability and statistics," *Stat. Prob. Lett.*, vol. 9, pp. 155–161, 1990.