

Binary Linear Classifier Design for Noisy One-bit Compressed Sensing

drastically exceeded the number of available training examples M . Such computationally challenging problems are of great interest in a number of recent applications, including text classification [2], multi-voxel pattern analysis [3]–[5], conjoint analysis [6], and microarray gene expression [7]. In the $N \uparrow M$ regime, the model (1) also coincides with that of noisy one-bit compressed sensing (CS) [8], [9], where \mathbf{w} is the sparse vector of interest and \mathbf{P} is the $\mathbf{W}\mathbf{P}$ row of the $M \times N$ measurement matrix. This study to confirm the efficacy, flexibility, and speed afforded by our GAMP-based approaches to binary classification.

*Dept. of ECE, The Ohio State University, Columbus, OH 43210. (Email: zimeli@ece.osu.edu, schniter@ece.osu.edu)
 †Dept. of Psychology, The Ohio State University, Columbus, OH 43210. (Email: sederberg.1@osu.edu)

') &) # % 2 * 2 \$ % # ! 2 # 2 " % / 2 " % 2 ((* # " 2 " 2
 * + % 2 * # " 2 2 \$ % # \$ # (2 # % * ! 2 \$ \$ % # (2 # 2 ((% 2
 (" 2 (2 # " 2 * 2 " % 0 2 \$ \$ % # . ! * 2 ! ((2 \$ ((" 2
 2 # % * ! 2 % " / 2 \$ % # \$ # (2 2 " 2 # " * . * # 2 # ! \$ %
 (, 2 (" " 2 + % 2 - # % 2 # + ((2 # " 2 " 2 % ! 2 - % 2 * 2 " + % 2
 # 2 * + % (2 % * / 2 . (2 * 2 " + ! % 2 # 2 % " " 2 ! \$ ((2 # 2
 - % 2 # " / 2 2 - 2 * + % (2 + 2 # % 2 + % * 2 ((* # # 2
 (# - 2 * * 2 (+ ! \$ % # + * 2 2 " 2 2 + (2 * # 2 \$ \$ % # . ! * / 2
 ! " ! 0 2 * 2 ((* # " 2 % % # % 2 * 2 " 2 ! . (+ ! 2 2 " 2
 2 + (2 * # 2 " ! 0 2 2 - 2 , % 1 2 # 2 % + % 0 # ((2 + " 2
 + % * % ! # % - 2 2 (% 2 " 2 . \$ * * # " ! . ! 0 * # " 2 1
 (2 (! 2 * # 2 % " 2 2 ((# * 2 ! # 2 \$ % ! * % (# " *
 (2 " 2 * % " * , 2 # 2 % # ((, * # " 2 2 - 2 (# - 2 * 2
 (* * 2 , # + * # " 2 % ! - # % 2 " 2 2 + (2 * # 2 + % * / 2 \$ %
 ! (((* # " 2 % * 2 " / 2 2 \$ % (" * 2 2 % 2 + ! %
 * # 2 # " % ! 2 * 2 / 2 . * / 2 " 2 (\$ 2 % 2 / 2 # + % 2
 (2 \$ \$ % # (2 # 2 " % / 2 ((* # " 2 #

II. GAMP

We now formulate the binary linear classifier design problem from both minimum error-rate and maximum a posteriori (MAP) perspectives, and show that each case can be tackled by a variant of the GAMP algorithm.

A. Sum-Product GAMP

Suppose that we are given M labeled training examples $\mathbf{f}_m / \mathbf{w}_m$ and T test feature vectors \mathbf{f}_t associated with unknown test labels \mathbf{f}_t , all obeying the noisy linear model (1) under some known error pdf Q and thus

$$\mathbf{f}_m = \mathbf{W} \mathbf{w}_m + \mathbf{n}_m, \quad \mathbf{f}_t = \mathbf{W} \mathbf{w}_t + \mathbf{n}_t, \quad (1)$$

' M t z b z j t 9 6 5 N b ' d i H a g [g H e Q m c O f a g H a f f o e P R X S q Y [^ O q
 N O g O e ' [a f i b l a g l g O g m c O f a g Q e [B q X T 4 Y O N B e d H a N i j H e [H J O q
 S X [j O N q O e ' [a f i b l a g l g O g m c O f a g Q e [g g O f a b U T 2 k Y O e O q d f L e [c g q
 [f q b ' O g [' O f q g g O k N q e q O q H a [a X q ^ O I m ' i ' a q O L g b l e O g m c O f O g q
 [a b b ^ N R H L k O e L H O X q S e q ' S i g e [L [a b b ^ N R H l c O e L H O O X S q
 b e Y i S a N y O g e H a f c b f q a b g O n q t S f e j q L g b S q I G d t t C H E t
 D i l t e O R Q e b j O d j O L g b l e t H K E S a H ^ m M O t t S f d Y O q
 ' i ' a g [j H e [d b e d t N q f g e [J i g H e i a l g [b e q Y q q L g b e d t g Y q H a q t
 H a N d j i [H a L ' e g e [b Y [^ O q t H a N q 6 t N O a b g v O g g H a N l a b e ' H ^ q
 c N R e N l q R e Q f c O L g [j O ^ m q

I. Problem Statement

In this work we consider binary linear classification. The signum function and \mathbf{n}_m is a random perturbation accounting for model inaccuracies. For the purpose of learning \mathbf{w} , we assume the availability of M labeled training examples generated independently according to this model:

$$\mathbf{f}_m = \text{sgn}(\mathbf{n}_m + \mathbf{w}) + \mathbf{n}_m, \quad \mathbf{f}_m \in \{-1, 1\}, \quad (1)$$

with \mathbf{n}_m i.i.d. It is common to express the relationship between the label \mathbf{f}_m and the score $\mathbf{w}^T \mathbf{z}_m$ in via the conditional pdf $Q_{\mathbf{f}_m | \mathbf{z}_m}$ known as the “activation function,” which can be related to the perturbation pdf $Q_{\mathbf{n}_m}$

$$(2)$$

We are particularly interested in classification problems in which the number of potentially discriminatory features N

$$Q_{\mathbf{f}_m | \mathbf{z}_m}(1 | z_m) = \int_{-\infty}^{z_m} p_{\mathbf{n}_m}(e) de = 1 - Q_{\mathbf{n}_m}(z_m).$$

$$p_{y_m | z_m}(1 | z_m) = \int_{-\infty}^{z_m} p_{\mathbf{n}_m}(e) de = 1 - p_{y_m | z_m}(-1 | z_m).$$

\mathbf{z}^k and \mathbf{a}^k respectively, at the final iteration “ $k = \hat{\mathcal{K}}$ ”

• *Max-Sum GAMP*

An alternate approach to linear classifier design is through the minimization of a regularized loss function, e.g.,

$$\hat{\mathbf{w}} = \underset{\mathbf{w} \in \mathbb{R}^M}{\operatorname{argmin}} \sum_{m=1}^M f_m(\mathbf{x}_m; \mathbf{w}) + \sum_{n=1}^{L\hat{\mathcal{S}}} f_n(\mathbf{w}; \mathbf{w}_n) \quad (10)$$

where $f_m(\cdot)$ are y_m -dependent convex loss functions (e.g., logistic, probit, or hinge) and where $f_n(\cdot)$ are convex regularization terms (e.g., $f_n(\mathbf{w}) = \lambda \|\mathbf{w}\|_2$ for ℓ_2 regularization and $f_n(\mathbf{w}) = \lambda \|\mathbf{w}\|_1$ for ℓ_1 regularization).

The solution to (10) can be recognized as the *maximum a posteriori* (MAP) estimate of random vector \mathbf{w} under a separable prior $P_{\mathbf{W}}$ and likelihood corresponding to (1), i.e.,

$$P_{\mathbf{W}}(\mathbf{w} | \mathbf{y}, \mathbf{X}) = \prod_{m=1}^M P_{y_m}(\mathbf{y}_m | \mathbf{x}_m; \mathbf{w}), \quad (11)$$

when $f_m(\mathbf{z}) = -\log p_{y_m}(\mathbf{z})$ and $f_n(\mathbf{w}) = -\log P_{\mathbf{w}_n}(\mathbf{w})$. Importantly, this statistical model is exactly the one yielding the reduced factor graph in Fig. 1b.

Similar to how sum-product LBP can be used to compute (approximate) marginal posteriors in loopy graphs, *max-sum* LBP can be used to compute the MAP estimate [15]. Since max-sum LBP is itself intractable for the high-dimensional problems of interest, we turn to the max-sum variant of GAMP [10], which is also specified in Algorithm 1. There, lines 8-9 are to be interpreted as

$$\hat{z}_m^k = \operatorname{prox}_{\tau_{p_m}^k, f_{z_m}^k}(\hat{p}_m^k), \quad m = 1, \dots, M, \quad (12)$$

$$\tau_{z_m}^k = \tau_{p_m}^k \operatorname{prox}_{\tau_{p_m}^k, f_{z_m}^k}(\hat{p}_m^k), \quad m = 1, \dots, M, \quad (13)$$

with $(\cdot)'$ and $(\cdot)''$ denoting first and second derivatives and

$$\operatorname{prox}_\tau(v) = \underset{u \in \mathbb{R}}{\operatorname{argmin}} \left[u^2 + \frac{1}{2\tau} (u - v)^2 \right] \quad (14)$$

$$\operatorname{prox}_\tau(v) = (1 + \tau \operatorname{prox}_\tau(v))^{-1} \quad (15)$$

and lines 19–20 are to be interpreted similarly. It is known [14] that, for arbitrary \mathbf{X} , the fixed points of GAMP correspond to the critical points of the optimization objective (10).

III. STATE EVOLUTION

One of the hallmarks of GAMP is that its behavior in the large-system limit (i.e., $M, N \rightarrow \infty$ with fixed ratio $\beta = M/N$) under i.i.d. sub-Gaussian \mathbf{X} is characterized by a scalar state evolution [10], [13]. We now describe how this state evolution can be used to characterize the test-error rate of the linear-classification GAMP algorithms described in Section II.

The GAMP state evolution characterizes average GAMP performance over an ensemble of (infinitely sized) problems, each associated with one realization $(\mathbf{y}, \mathbf{X}, \mathbf{W})$ of the random triple $(\mathbf{y}, \mathbf{X}, \mathbf{W})$. Recall that, for a given problem realization $(\mathbf{y}, \mathbf{X}, \mathbf{W})$, the GAMP iterations in Algorithm 1 yield the sequence of estimates $\{\mathbf{w}^k\}_k$ of the true weight vector \mathbf{W} . Then, according to the state evolution, $P_{\mathbf{W}}(\mathbf{w}^k | \mathbf{y}, \mathbf{X})$

and the first two moments of the joint pdf $P_{\mathbf{W}, \mathbf{W}}(\mathbf{W}, \mathbf{W})$ can be computed using [10, Algorithm 3].

Suppose that the (\mathbf{y}, \mathbf{X}) above represent training examples associated with a true weight vector \mathbf{w} , and that (\mathbf{y}, \mathbf{X}) represents a test pair also associated with the same \mathbf{w} and with \mathbf{x} having i.i.d. elements distributed identically to those of \mathbf{X} (with, say, variance σ^2). The true and iteration- k -estimated test scores are then $z = \mathbf{x}^T \mathbf{w}$ and $\hat{z}^k \triangleq \mathbf{x}^T \mathbf{w}^k$, respectively. The corresponding test-error rate $\epsilon^k = \Pr\{\mathbf{y} \neq \operatorname{sgn}(\hat{z}^k)\}$ can be computed as follows. Letting $\mathbb{1}\{\cdot\}$ denote an indicator function that assumes the value 1 when its Boolean argument is true and the value 0 otherwise, we have

$$\epsilon^k = \mathbb{E}[\mathbb{1}\{\mathbf{y} \neq \operatorname{sgn}(\hat{z}^k)\}] \quad (16)$$

$$= \sum_{\mathbf{y} \in \{-1, 1\}^M} \mathbb{1}\{\mathbf{y} \neq \operatorname{sgn}(\hat{z}^k)\} \int \prod_{m=1}^M p_{y_m}(z_m) p_{z_m}(z, \mathbf{z}^k) dz d\mathbf{z}^k \quad (17)$$

$$= \sum_{\mathbf{y} \in \{-1, 1\}^M} \mathbb{1}\{\mathbf{y} \neq \operatorname{sgn}(\hat{z}^k)\} \int \prod_{m=1}^M p_{y_m}(z_m) p_{z_m}(z, \mathbf{z}^k) dz d\mathbf{z}^k. \quad (18)$$

Furthermore, from the definitions of (\mathbf{z}^k) and the bivariate central limit theorem, we have that

$$\mathbf{z}^k \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Sigma^k), \quad \text{where } \Sigma^k = \begin{pmatrix} \Sigma_{11}^k & \Sigma_{12}^k \\ \Sigma_{12}^k & \Sigma_{22}^k \end{pmatrix} \quad (19)$$

where \xrightarrow{d} indicates convergence in distribution. In [16], it is shown that the above matrix components are

$$\Sigma_{11}^k = \delta^{-1} (\operatorname{var}\{\mathbf{w}_n\} + E[\mathbf{w}_n]^2), \quad (20)$$

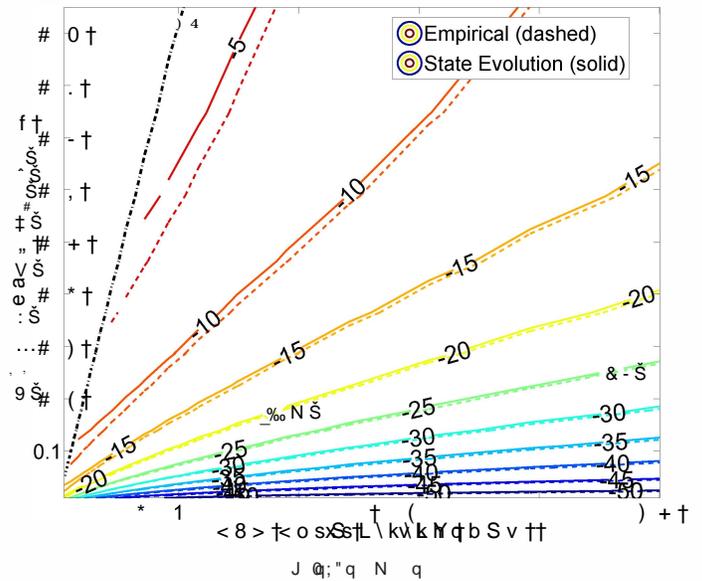
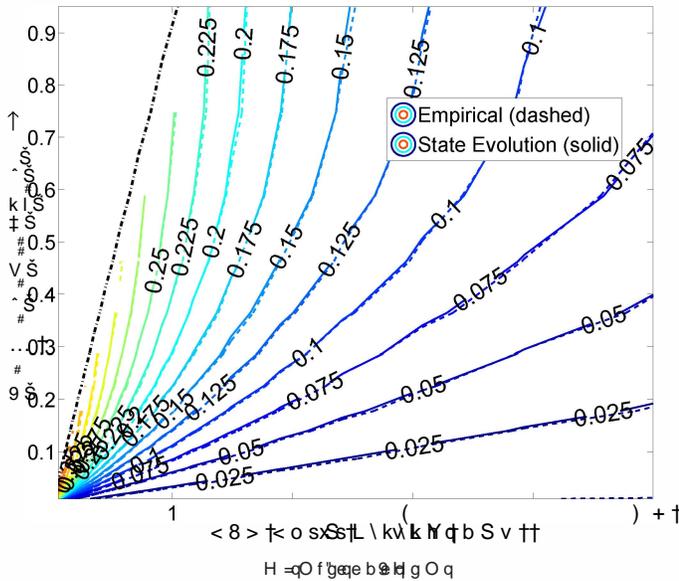
$$\Sigma_{12}^k = \Sigma_{21}^k = \delta^{-1} (\operatorname{cov}\{\mathbf{w}_n, \hat{\mathbf{w}}_n^k\} + E[\mathbf{w}_n]E[\hat{\mathbf{w}}_n^k]), \quad (21)$$

$$\Sigma_{22}^k = \delta^{-1} (\operatorname{var}\{\hat{\mathbf{w}}_n^k\} + E[\hat{\mathbf{w}}_n^k]^2) \quad (22)$$

for label-to-feature ratio β as described earlier, the above moments can be computed using [10, Algorithm 3]. The integral in (18) can then be computed (numerically if needed) for a given activation function PyZ yielding an estimate of GAMP’s test-error rate at the k^{th} iteration.

To validate the accuracy of the above asymptotic analysis, we conducted a Monte-Carlo experiment with data synthetically generated in accordance with the assumed model. In particular, for each of 100 problem realizations, a true weight vector $\mathbf{W} \in \mathbb{R}^{L\hat{\mathcal{S}}}$ was drawn i.i.d. zero-mean Bernoulli-Gaussian and a feature matrix \mathbf{X} was drawn i.i.d. Gaussian, yielding true scores $\mathbf{z} = \mathbf{X}\mathbf{W}$, from which the true labels \mathbf{y} were randomly drawn using a probit activation function PyZ . A GAMP weight-vector estimate $\hat{\mathbf{w}}^\infty$ was then computed using the training data (\mathbf{y}, \mathbf{X}) from which the test-label estimates $\{\hat{y}_f\}_f$ with $\hat{y}_f = \operatorname{sgn}(\hat{z}_f | \hat{\mathbf{w}}^\infty)$ were computed and compared to the true test-labels in order to calculate the test-error rate for that realization. Figure 2a plots the Monte-Carlo averaged empirical test-error rates (dashed) and state-evolution predicted rates (solid) as level curves over different combinations of training ratio β and discriminative-feature ratio γ , where $K = \|\mathbf{W}\|_0$ and $N = 10^4$. Similarly, Fig. 2b plots average empirical MSE versus state-evolution predicted MSE, where $\operatorname{MSE} = \mathbb{E}\|\mathbf{W} - \hat{\mathbf{w}}\|_2^2$.

© 2018 IEEE. This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY), which can be found at <http://creativecommons.org/licenses/by/4.0/>.



$\$ [X qe f g O e t t e d t a n Q H a f d i l O n e b e q j Q e f i g e H [a [x g b R i d j t 7 e c h a n d f l e [[a H g j i O R B H g t e c H ^ L i ^ H g C a n g ' c [e [L H ^ q$
 $H j O e H X [M q Y O H a N g H g O O j b c i g Q b a q g [b a [N H f j ' [a X q [N q e a b i ^ { \wedge } [\% H i k O [t X a y Q q g b i e a d H q e b J t b l g g [j H g r b a l g g [b a q$

In both Fig. 2a and Fig. 2b, the training-to-feature ratio $\frac{M}{N}$ increases from left to right, and the discriminative-feature ratio $\frac{K}{M}$ increases from bottom to top. The region to the upper-left of the dash-dotted black line contains ill-posed problems (where the number of discriminative features K exceeds the number of training samples M) for which data was not collected. The remainders of Fig. 2a and Fig. 2b show very close agreements between empirical averages and state-evolution predictions.

IV. GAMP

We now recall some of the weight-vector priors $\mathbf{V}_{\mathbf{x}}(\cdot)$ and activation functions $\mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{y}; \mathbf{m})$ (equivalently, regularizations $r_{\mathbf{x}}(\cdot)$ and losses $f_{\mathbf{z}}(\cdot)$) commonly used for binary linear classification, and we describe how GAMP handles them via the nonlinear steps 5-6, 8-9, 16-17, and 19-20 in Algorithm 1. For sum-product GAMP, we recall that the mean and variance computations in lines 5-6 and 16-17 are computed based on the pdfs in (7) and (8), respectively, and for max-sum GAMP the prox steps in 8-9 are computed using equations (12)-(13) and those in 19-20 are computed similarly.

A. Activation Functions

Arguably the most popular activation function for binary linear classification is the logistic sigmoid [1, §4.3.2]:

$$\mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{y}; e) = \frac{1}{1 + \exp(-\mathbf{y}\mathbf{z})}, \quad \forall \mathbf{y} \in \{-1, 1\} \quad (23)$$

where $e < \infty$ controls the steepness of the transition. For logistic sum-product GAMP, the mean and variance $(\bar{\mathbf{z}})$ of the marginal posterior approximation can be computed using a variational approach described in [16]. For logistic max-sum GAMP, $\hat{\mathbf{z}}$ from (12) solves the scalar minimization problem (14) with $f(u) = -\log \mathcal{P}_{\mathbf{y}}(u; e)$ from (23), which is convex. This $\hat{\mathbf{z}}$ can be found via bisection search, by locating

the root of $\frac{d}{du}[f(u) + \tau(\cdot)]$. The TZ from (13) can then be computed in closed-form using $\hat{\mathbf{z}}$ and $f''(\cdot)$ via

Another popular activation function is the probit [1, §4.3.5]:

$$\mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{y}; v) = \frac{1}{\sqrt{v}} \mathcal{N}\left(\frac{\mathbf{z}}{\sqrt{v}}\right) \quad (24)$$

where $\mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{z}) = 1 - \mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{z}) = \frac{1}{2}$ and where $v < \infty$ controls the steepness of the sigmoid. In the sum-product case, the density (7) corresponds to the posterior pdf of a random variable \mathbf{z} with prior $\mathcal{N}(\mathbf{f}, \mathbf{T})$ from an observation \mathbf{y} measured under the model (24), and a derivation in [17, §3.9] provides expressions for the moments used by GAMP in lines 5-6 of Algorithm 1. In the max-sum case, $(\bar{\mathbf{z}})$ can be computed using bisection search, akin to that described above.

The hinge loss $f_{\mathbf{z}}(\mathbf{z}) = \max(0, 1 - \mathbf{y}\mathbf{z})$, commonly used in the support vector machine (SVM) approach to maximum-margin classification [1, §7.1], yields the activation function

$$\mathcal{P}_{\mathbf{y}}(\mathbf{z}|\mathbf{z}) = \exp(-\max(0, 1 - \mathbf{y}\mathbf{z})), \quad \forall \mathbf{y} \in \{-1, 1\}. \quad (25)$$

For sum-product GAMP, the corresponding $(\bar{\mathbf{z}})$ from (7) can be computed in closed-form using a procedure described in [16], and for max-sum GAMP, the proximal steps (12)-(13) can be efficiently computed using bisection search.

B. Weight-Vector Priors

For linear classification and feature selection with $N \ll M$, it is customary to choose a prior $\mathbf{V}_{\mathbf{x}}(\cdot)$ that leads to sparse (or approximately sparse) weight vectors \mathbf{s} , as discussed below.

In the sum-product case, this is typically accomplished using a Bernoulli- \tilde{p} prior, i.e.,

$$\mathbf{V}_{\mathbf{x}}(\mathbf{s}) = \prod_{i=1}^N \mathcal{B}(s_i; \tilde{p}) \prod_{n=1}^M \ell(\cdot) \quad (26)$$

where $\delta(\cdot)$ is the Dirac delta function, $\mathbf{c} \in \mathbb{R}^L$ is the prior probability that $\mathbf{W} \neq \mathbf{0}$, and $\tilde{p}_{w_n}(\cdot)$ is the prior pdf of \mathbf{W} when non-zero. With GAMP it is common to use Bernoulli-Gaussian [18], Bernoulli-Gaussian-mixture [19], and Bernoulli-Laplacian [16], all of which yield (8) with closed-form first and second moments, (\mathbf{W}_{τ_w}) .

In the max-sum case, the GAMP nonlinear outputs (\mathbf{W}_{τ_w}) are computed just as (\hat{x}_z) are in (12)-(15), but with \mathbf{W} used in place of \mathbf{y} . Common examples include $f_{w_n}(\cdot; \lambda)_1$ for ℓ_1 regularization [20], $f_{w_n}(\cdot; \lambda)_2$ for ℓ_2 regularization [10], and $f_{w_n}(\cdot; \lambda)_1$ for the “elastic net” [21], all of which yield closed-form implementations of (\hat{w}, τ) (See [16] for details.)

V. ACTIVATION FUNCTIONS AND WEIGHT-VECTOR PRIORS

The activation functions and weight-vector priors described in Section IV depend on modeling parameters that, in practice, must be tuned. For example, the logistic sigmoid (23) depends on λ ; the probit depends on λ ; and the Bernoulli-Gaussian-mixture prior depends on \mathbf{c} and \tilde{p}_{w_n} , where \mathbf{W} parameterizes the weight, μ the mean, and σ^2 the variance of the l th mixture component. Although cross-validation (CV) is the customary approach to tuning parameters such as these, it suffers from two major drawbacks: First, it can be very computationally costly, since each parameter must be tested over a grid of hypothesized values and over multiple data folds. For example, K -fold cross-validation tuning of P parameters using G hypothesized values of each requires the design and evaluation of KG^P classifiers. Second, leaving out a portion of the training data for CV can degrade classification performance, especially in the example-starved regime where $M \ll N$ (see, e.g., [22]).

As an alternative to CV, we consider *online learning* of the unknown model parameters θ using the method proposed in [19], [23]. Here, the goal is to compute the maximum-likelihood estimate $\hat{\theta}_{ML}$ where our data model implies a likelihood function of the form

$$\{ \mathbf{y} \} = \mathbf{I} \mathbf{C} \mathbf{e}^{\mathbf{I} \mathbf{W} \theta} \mathbf{P}_{w_n}(\mathbf{w}_n; \theta) \quad (27)$$

Because it is computationally infeasible to evaluate and/or maximize (27) directly, approximate-ML estimation is performed using expectation-maximization (EM) [24], treating \mathbf{w} as the “hidden” data, which yields the iteration- j estimate

$$\mathbf{w}_{j+1} = \mathbf{E}_{\mathbf{w}} \{ \mathbf{w} | \mathbf{y}, \theta, \mathbf{w}_j \} \quad (28)$$

$$\begin{aligned} \theta_{j+1} &= \mathbf{E}_{\mathbf{w}} \{ \theta | \mathbf{y}, \mathbf{w}_j, \theta_j \} \\ &= \mathbf{E}_{\mathbf{w}} \{ \theta | \mathbf{y}, \mathbf{w}_j, \theta_j \} \\ &= \mathbf{E}_{\mathbf{w}} \{ \theta | \mathbf{y}, \mathbf{w}_j, \theta_j \} \end{aligned} \quad (29)$$

Furthermore, to evaluate the conditional expectations in (29), GAMP’s posterior approximations from (7)-(8) are used. It was shown in [25] that, in the large-system limit, the estimates generated by this procedure are asymptotically consistent (as $j \rightarrow \infty$ and under certain identifiability conditions). Moreover, it was shown in [19], [23] that, for various priors and likelihoods of interest in compressive sensing (e.g., AWGN

likelihood, Bernoulli-Gaussian-Mixture priors, ℓ_1 regularization), the quantities needed from the expectation in (29) are implicitly computed by GAMP, making this approach computationally attractive. Still, this EM procedure runs GAMP several times, once for each EM iteration (although not necessarily to convergence), thus potentially increasing the runtime over GAMP without EM.

In this work, we export the EM-GAMP methods from [19], [23] to the classification setting, which requires deriving the EM updates for the activation-function parameters, such as λ in the logistic model (23) and λ in the probit model (24). Due to lack of space, we leave the algorithmic details to [16], but we demonstrate the performance of this approach in Section VI.

VI. TEXT CLASSIFICATION

In this section we apply GAMP⁴ to a text classification problem. This experiment was conducted on a workstation running Red Hat Enterprise Linux (r2.4), with an Intel Core i7-2600 CPU (3.4 GHz, 8MB cache) and 8GB DDR3 RAM. Other experiments can be found in [16], but they are not shown here for reasons of space.

We first consider a binary text classification problem based on the Reuter’s Corpus Volume I (RCV1) dataset [26]. As in [27], [28], newswire article topic codes AA=SPORTS and BA=SPORTS were combined to form the positive class while GA=SPORTS and KA=SPORTS were combined to form constitute the negative class.⁵ Although the original dataset consisted of 204 balanced training examples of $N = 423$ features, with 6773 examples reserved for testing, we followed the approach in [27], [28] and swapped training and testing sets in order to test computational efficiency on a large training dataset (and thus $M = 677399$). As in [27], we constructed feature vectors as cosine-normalized logarithmic transformations of the TF-IDF (term frequency-inverse document frequency) data vectors. We note that the resulting features are very sparse; only 0.16% of the entries in \mathbf{X} are non-zero. Finally, we trained linear classifiers (i.e., weight vectors) using four GAMP-based methods and three existing state-of-the-art methods: TRON [29], CDN [27], and TFOCS [30] in L1-LR mode. For EM learning we used 5 EM iterations, and for cross-validation we used 2 folds and a logarithmically spaced grid size of 10.

Table I summarizes the performance achieved by the resulting classifiers, including the test-set classification accuracy, weight-vector density (i.e., the fraction of non-zero weights),⁶ and two runtimes: the *total* runtime needed to train the classifier, which includes EM- or cross-validation-based parameter tuning, and the *post-tuning* runtime. Although it is customary to report only the latter, we feel that the former better captures the true computational cost of classifier design.

Table I shows all 7 classifiers achieving nearly identical test-set classification accuracy. However, if feature selection is

⁴ <http://www.eecs.berkeley.edu/~jcooper/gamp/>
⁵ <http://www.reuters.com/reuters/corpus.html>
⁶ <http://www.reuters.com/reuters/corpus.html>

Classifier	Tuning	Accuracy	Runtime (s)	Density
spGAMP: BG-PR	EM	97.6%	1.1	11.1%
spGAMP: BG-HL	EM	3 / 1	468 / 93	1
msGAMP: L1-LR	EM	97.6%	684 / 123	9.8%
msGAMP: L1-LR	xval	97.6%	3068 / 278	19.6%
CDN	xval	3 / 1	1298 / 112	10.9%
TRON	xval	3 / 1	1682 / 133	10.8%
TFOCS: L1-LR	xval	97.6%	1086 / 94	19.2%

= / " q (q q L b ' c l [f b a t n R N] [c o e O l a g ' t f f [U D a e g ' y ' 9 ' q q] [a H e n H q H f O g ' k [g ' y ' e H [a [a X g r O d i f f [c o c o P N Y b c [a g ' y ' 9 ' q q f g l e O h g t f [U L H g] [b i a e H L m q g Y O b n g H a t h f g g i a e a k g ' O H a n g ' y ' 9 ' q q a f [t r y q ' O Q [X Y ' 9 ' q q g b e b j O q f c q f i ' c e b N i L y q ' H i f i ' q % q O e a b i ^ ^ [% H i f i ' 9 ' q q a 6 b J [g q & / q & [a X O b t f f / g q f " S O X i ^ H e [n H g q b a b X [f g [b d f q

the goal, then weight-vector density is also important, and for this we see significant differences among algorithms. Notably, the two most discriminative (i.e., sparse) classifiers are the result of EM-GAMP methods. We attribute this sparsity in large part to EM-tuning, since Table I shows that when cross-validation is used in place of EM-tuning with L1-LR GAMP, the weight vector is twice as dense.

Table I also shows a wide range of runtimes. Among the post-tuning runtimes, the two fastest are EM-GAMP based. Moreover, among the total runtimes, the three fastest are EM-GAMP based, with the best (at 3 T seconds) beating the fastest non-GAMP algorithm (at 1 0 8 seconds) by more than a factor of 3. That said, some caution must be used when comparing runtimes. For example, while all algorithms were given a “stopping tolerance” of 10^{-6} , algorithms apply this tolerance in different ways. Also, CDN and TRON are implemented in C++, while GAMP is implemented in object-oriented MATLAB (and therefore is far from optimized).

Finally, we note that, although GAMP was derived under the assumption that the elements of X are realizations of an i.i.d. sub-Gaussian distribution, it worked well even with the X of this experiment, which was far from i.i.d. sub-Gaussian. We attribute the robust performance of GAMP to the adaptive damping mechanism included in the GAMPmatlab implementation (described in [31]).

:"\$ 9"2 ";q

FG qM Q q [f Y b c *Pattern Recognition and Machine Learning*. 2 O k Q b e] q ; c e [a X O e e ^ H X q q

F G % S b e ' H a q c O l g O a f O c i e [L i f y ' n i m R q O H g i e O Q L g O b g a e d L f q R b e q l g y H f f [U L H g] *Neural Mach. Learn. Res.*, j b ^ a q c c q q

F G e W S H I J m Q q % b J J [a d q d s i e O m q f Y H f q q L Y b i g O c h a n q 6 Q D g e [a l [f q e [J i g O a n n q O e ^ H c o e [X q O f O a g b i R f b i a t i c o n q b J V O L [a m p a g e y O ^ c b e l b e a g O s c i e n c e , j b ^ a q c c q q

F G q Q q Y H a " q q S H ^ X e Q p Q d e [a] b j [H a d n q q ; t i y q O L b N [a X q F k b e n H a N d H g O X b e m f b c H L g [L b l g d i ' c b e O H e q f O a g R e t h b d q % a q H a N ' c % *NeuroImage*, j b ^ a q c c q q

F G q Q m H ^ [q q i c O] H e q J e H ' f t a N v S O a b a : p H f O b X [f g [L q F e O X e O f R b a d q b a O J e l H ' H a t f [U L H g] R S i a (q N H g H *NeuroImage*, j b ^ a q c c q q

F G q q % i f g H R f f b a q O e ' H a t a n q S q i J O e *Conjoint Measurement: Methods and Applications*. O e ^ [a q ; c e [a X O e ? O e ^ H X q q

F G b Q B [a X O q d S b e N H a N q q q . l c q \$ O H g i f e O Q L g R b a d X Y o N I ' O a f [b a X O a b ' [L b e b l e H m H g H a m ' l W k s h p . M a c h . L e a r n . , c c q q q

F G q q b i R b i a b H q N q q q H e H a [i] ' q J [t h q ' c e O f f [f O a f [a X a q *Proc. Conf. Inform. Science Sys.*, 6 e [a L O Q b a o H e q q

F G S 6 H a t a N q q O e f Y m a B a b d i f f g q J [d u p ' c e O f f O Q a f [a X a n q f c H e ' O c k [f g e O X e O f f [d L a b a j O d e b X e H ' ' [a X c e b H L] Y a q i v : 1202.1212, \$ O J q q

F G q Q H a X H a q d a O e H ' H c e N b q [' H ' O Q H X O H f f [a R k l q d g [' H g [b a q k [g Y e H a N b ' c a O H e q X [a p *Proc. IEEE Int'l Symp. Inform. Theory*, ; g t u O e f J i e X f d H c q q q i X q q \$ i ^ J q e f [t h g g a r X i v : 1010.5141).

F G S q f L Y [L Y H a X q e O H a d q q q b O ^ [X O S H l g q b e d c H a d q Y O q f i ' c e b N i L y q X b e [g *IEEE Trans. Inform. Theory*, j b ^ a q c c q q

F G q Q S e O H a N q q q q d H L . H m q e O j b ^ i g [t a ^ [r o B l q c H X H g d b a q X e H c k [t r y t e n L ^ a d v . *Neural Info. Process. Sys.*, c c q q q

F G q Q H j H a ' H e H N q q d b a g H a H e g t h q O q O j b R b g [d a Q H d e c e b l [o ' H g O Q f f H X O H f f [a X q b e [g K [g ' H c a ^ [L H g] [b e q H q [L b , i p ^ [a X q a r X i v : 1211.5164, 2 b j q q

F G q Q H a X H a q Y a [g O Q [d X ^ O e q q S ^ O g L Y H a N q q j Y O e s t i O N q c b [a g b R X O a O e H , H i r o O l i n [' H g O d f H X O H f f [a X f g Y H e J [g e H e m q ' H g [L O f a *Proc. IEEE Int'l Symp. Inform. Theory*, (f g H a J i ^ i a t q q q \$ i ^ J q e f [t h g g a r X i v : 1301.6295).

F G S / d O ^ [X O e d q g e b N i L y q R b i l g q b e d c Y E E E *Signal Process. Mag.*, j b ^ a q c c q q q H a q q

F G e E [a [*Message Passing Approaches to Compressive Inference Under Structured Signal Priors*. 6 Y q g Y O f e Y Q y [b : t H q Q d j O e f [g m q

F G M t q q 9 H ' i f f O a h a N q q - q (q @ [^ [H G a u s s i a n P r o c e s s e s f o r M a c h i n e L e a r n i n g . = Y O Q = q 6 e O f f q q

F G q X L Y a [g O e i e q l b d L b a f g e i b E [g e a i q g i e O N e f O X a H ^ [a C o n f . o n I n f o r m a t i o n S c i e n c e s a n d S y s t e m s (C I S S) , 6 e [a L O Q b a o a q q q O H e q q

F G e t i [^ H a g t h q L Y a [g O ' e c Q L g H g [b a O H I [' [n H g [b a q % H i f f [H a ' [H c c e b l [' H ' O Q H X O H f f [a X E E E *Trans. Signal Process.*, j b ^ a q c c q q q L g q q

F G b q / q ! b a b Y b q q O H ^ O] [H q N q q O b a g H a H e Q f f H X O H f f [a X q H ^ X b e [g R b e q ' c e O f f O N a f [a X a *Proceedings of the National Academy of Sciences*, j b ^ a q c c q q q 2 b j q q

F G S d e b i f a n q q H f g [O Q X i ^ H e [n H g [b i t J ^ O Q ^ O L g i [b a d q O q O ^ H f g [Q q] *Soy. Statist. Soc.*, ? S b ^ a q c b q c c q q q

F G q q q 2 [X H ' - H a N Q L H ^ a i ; q e Z e i a H q N q q [g L Y O ^ O b q q H f o f [U L H g R e a m ' c H J O ^ H a N q a H J O ^ N O L q ' O a g f f a X ' Q *Machine Learning*, j b ^ a q c c q q q

F G e X [^ H a g t h q L Y a [g O e q Q ' c [e [L H , H i r o O e f H y b e q L b j O e [a X q O i ^ a t h q f g e H B a G N a q O X H g [J O t q X a H , [a *Proc. IEEE Workshop Comp. Adv. Multi-Sensor Adaptive Process.*, ; H [a q l g [a ! Q L q q \$ i ^ J q e f [t h g g a r X i v : 1310.2806).

F G q Q R S O ' c f g O Q Q q H [e N a n q q Q i J [a Q H I [' i ' q ^ [] O ^ [Y b b N q R e b [L b ' c ^ a N g H g i H g Y O a q H ^ X b e [g Y S i y *Statist. Soc.*, ? S j b ^ a q c c q q q

F G a q q H [^ b j ; q q H a X H a q q S ^ O g L Y H a N q q a f O e q c e b l [' H g O q ' O f f H X O H f f [a X] g Y L d a f [f g O H e g t i ' O Q O e d [H g H b a n] c a ^ [L H g [b a f q g b e l f O q O l a [a X a *Proc. Neural Inform. Process. Syst. Conf.*, / H] O q = H Y b Q ? q t O L q q \$ i ^ J q e f [t h g g a r X i v : 1207.3859).

F G b q ! / d k [f [S C H a X U S e q b f O H q N e S [q 9 ? _ q q a O k k O a L Y ' H e] q L b ^ a O B b i e q L g t g O X b e [r e i d g O b a e L M a c h . L e a r n . R e s . , j b ^ a q c c q q q

F G % B q C i H a q . A q Y H a X H a N n q q S f i O Q d L b ' c l [f b a b R q b c g l [' n H g [O e q d b H i f a n n R g k l i t e b X O f L H L ^ O e q O X i ^ H e [[r a O N H q q L b ^ H f f [U L H g] *Neural Mach. Learn. Res.*, j b ^ a q c c q q q

G q [a q q @ O a X i a n q ; . Q e g Y ^ e a f g Q X [B d q g b a Q g Y b N f q R b e d X O f L H L b Q d f g e O d X e O f f [b a b *Proc. 24th Int'l Conf. Mach. Learn.*, b e j H ^ a q q c c q q q

G q q + Q a d H a N q q S e b P Q O k g b a ' O g Y b N b e d X O f L H L b Q d a n q L b a f g e H [c e O N q O ' S I A M J O p t i m . , j b ^ a q c c q q q

F G q Q o L] O e q q H a N Q H a n q q q % e H a g q = O ' c R b i e q d a j O l q L b a O e j J ^ a O k [f g Y h c a ^ [L H g] [b e q H e f i Q X a b e Q L b j O *Math. Prog. Comp.*, j b ^ a q c b q c c q q

F G q X L Y a [g O e a n q H a X H a b d c e O f f [j O H f e O g e [Q] H i k q a O e H ^ [n O N q H c c e b l [' H ' O Q H X O H f f [a X a *Adleron Conf. Commun., Control, Comput.*, O b a g [L O (^ a b e q g q q