ADAPTIVE MATCHING PURSUIT FOR SPARSE SIGNAL RECOVERY

Tiep H. Vu, Hojjat S. Mousavi, Vishal Monga The Pennsylvania State University, University Park, PA

ABSTRACT

Spike and Slab priors have been of much recent interest in signal processing as a means of inducing sparsity in Bayesian inference. Applications domains that benefit from the use of these priors include sparse recovery, regression and classification. It is well-known that solving for the sparse coefficient vector to maximize these priors results in a hard non-convex and mixed integer programming problem. Most existing solutions to this optimization problem either involve simplifying assumptions/relaxations or are computationally expensive. We propose a new greedy and adaptive matching pursuit (AMP) algorithm to directly solve this hard problem. Essentially, in each step of the algorithm, the set of active elements would be updated by either adding or removing one index, whichever results in better improvement. In addition, the intermediate steps of the algorithm are calculated via an inexpensive Cholesky decomposition which makes the algorithm much faster. Results on simulated data sets as well as real-world image recovery challenges confirm the benefits of the proposed AMP, particularly in providing a superior cost-quality trade-off over existing alternatives.

1. INTRODUCTION

Over the past decade, sparsity has become one of the most prevalent themes in signal processing applications. In general, parsimony in signals describes the phenomenon where high dimensional data can be expressed by only a few measurements. Sparse models assume that a signal can be efficiently represented as sparse linear combination of atoms in a given or learned dictionary [1,2]. The presence of sparsity in signals often enables us to provide efficient algorithms for extracting relevant information from the underlying data and is often a natural assumption in inverse problems with variety of applications in image/signal classification [1,3], dictionary learning [4–7], signal recovery [8–10], image denoising and inpainting [11], etc.

A sparse reconstruction algorithm aims to recover a sparse signal $\mathbf{x} \in \mathbb{R}^p$ from a set of fewer linear measurements $\mathbf{y} \in \mathbb{R}^q$ ($q \ll p$) according to: $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, where $\mathbf{A} \in \mathbb{R}^{q \times p}$ is the measurement matrix and $\mathbf{n} \in \mathbb{R}^q$ represents Gaussian noise. Many solutions have been proposed for this problem and they include sparsity promoting optimization problems involving different regularizers such as ℓ_1 or ℓ_0 norms, greedy (e.g. matching pursuit) algorithms [9, 12], Bayesianbased methods [13, 14] or general sparse approximation algorithms – SpaRSA [8], ADMM [15], etc. Many of these sparse recovery methods have shown that adding structural constraints and prior information to the frameworks have value in terms of representation purposes [2, 16] and often leads to performance improvement. Introducing priors for capturing sparsity as an example of Bayesian inference has shown to be effective for signal recovery [17, 18]. Examples of such priors in statistics and signal processing are Laplacian [19], generalized Pareto [20], Spike and Slab [21], etc. Amongst these priors, a well-suited sparsity promoting prior is Spike and Slab prior which is widely used in Bayesian inference [22–24]. In fact, it is acknowledged that Spike and Slab prior is indeed the *gold standard* for inducing sparsity in Bayesian inference [24].

In this paper in particular, we focus on Spike and Slab priors, introduced by *Yen et al.* [25], where every coefficient x_i is modeled as a mixture of two densities as follows:

$$x_i \sim (1 - w_i) \mathbb{I}(x_i = 0) + w_i P_i(x_i).$$

 $\mathbb{I}(\cdot)$ is the indicator function at zero (spike) and P_i (slab) is a suitable prior distribution, e.g., Gaussian, for nonzero values of x_i . $w_i \in \{0, 1\}$ controls the activeness of x_i .

Optimization Problem (Hierarchical Bayesian Framework): As suggested by [20, 26], sparsity can be induced via a prior maximization procedure in a Bayesian setup. In this work, we employ Spike and Slab prior for inducing sparsity on x and formulate a hierarchical Bayesian framework as in [10].

More precisely, the Bayesian formulation is as follows:

$$\begin{aligned} \mathbf{y}|\mathbf{A}, \mathbf{x}, \boldsymbol{\gamma}, \sigma^2 &\sim & \mathcal{N}(\mathbf{A}\mathbf{x}, \sigma^2 \mathbf{I}), \\ \mathbf{x}|\boldsymbol{\gamma}, \lambda, \sigma^2 &\sim & \prod_{i=1}^p \gamma_i \mathcal{N}(0, \sigma^2 \lambda^{-1}) + (1 - \gamma_i) \mathbb{I}(x_i = 0), \\ \boldsymbol{\gamma} &\sim & \prod_{i=1}^p \text{Bernoulli}(\kappa_i), \end{aligned}$$

where $\mathcal{N}(.)$ represents the Gaussian distribution and γ is the indicator variable for vector \mathbf{x} , i.e., $\gamma_i = 0$ if x_i is zero, otherwise $\gamma_i = 1$. The parameter κ_i affects the sparsity level of the \mathbf{x} by separately controlling whether each indicator variable γ_i is active or not.

The MAP estimation based on the above mentioned Bayesian framework leads to the optimization problem below [10]:

$$(\mathbf{x}^*, \boldsymbol{\gamma}^*) = \arg\min_{\mathbf{x}, \boldsymbol{\gamma}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_2^2 + \sum_{i=1}^p \rho_i \gamma_i, \quad (1)$$

and $\rho_i \triangleq \sigma^2 \log(\frac{2\pi\sigma^2(1-\kappa_i)^2}{\lambda\kappa_i^2})$. Note that from this definition, ρ_i may be negative (if κ_i is large enough).

This is a more general sparsity inducing optimization problem than the ones containing only l_0 or l_1 regularizers

This work has been supported partially by the Office of Naval Research (ONR) under Grant 0401531 UP719Z0 and NSF CAREER award to (V.M.).

and has broad applicability in recovery and regression problems and is known to be a hard non-convex mixed integer programming. In [3, 23, 25] simplifications are pursued by assuming a common κ (ρ) for each coefficient, which reduces the last term in (1) to $\rho \| \mathbf{x} \|_0$. Further, relaxation of the $\| \mathbf{x} \|_0$ to $\|\mathbf{x}\|_1$ leads to the well-known Elastic Net [18]. More recently an iterative refinement solution is also proposed in [10] which refines the solution at each step by considering a history of solutions at previous iterations. The Main contributions of this paper are as follows: 1) To directly solve the hard non-convex problem in (1), we propose an adaptive matching pursuit (AMP) procedure which is supported with theoretical analysis that formally argues the effectiveness as well as computational benefits of the algorithm. 2) The procedure can be slightly modified to solve problems with non-negativity constraints, which is often required in real-world applications. 3) We perform experimental validation on both simulated data and a practical image recovery problem, that reveals the merits of the proposed AMP; the practical findings also support the aforementioned analytical results.

2. ADAPTIVE MATCHING PURSUIT

In this section, we propose a greedy solution for (1) by adding/removing elements to/from the support of **x**. First, let $\mathbf{D} = \begin{bmatrix} \mathbf{A} \\ \sqrt{\lambda}\mathbf{I} \end{bmatrix}$ and $\mathbf{z} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$ with $\mathbf{I} \in \mathbb{R}^{p \times p}$ and $\mathbf{0} \in \mathbb{R}^{p \times 1}$ being the identity and zero matrices, we can rewrite (1) as:

$$(\mathbf{x}^*, \boldsymbol{\gamma}^*) = \arg\min_{\mathbf{x}, \boldsymbol{\gamma}} \|\mathbf{z} - \mathbf{D}\mathbf{x}\|_2^2 + \sum_{i=1}^{P} \rho_i \gamma_i.$$
(2)

Note that we assume that each column of **A** has norm 1, i.e. $\|\mathbf{a}_i\|_2^2 = 1$ and subsequently, $\|\mathbf{d}_i\|_2^2 = \|\mathbf{a}_i\|_2^2 + \lambda = 1 + \lambda$, $\forall i = 1, ..., p$. It is also crucial to note that, if we know the true support of the signal, i.e. $S = \{i : \gamma_i \neq 0\}$, we can easily find the solution of (2) by calculating:

$$\mathbf{x}^{\mathcal{S}} = \arg\min_{\mathbf{x}^{\mathcal{S}}} \|\mathbf{z} - \mathbf{D}_{\mathcal{S}} \mathbf{x}^{\mathcal{S}}\|_{2}^{2} \Rightarrow \mathbf{D}_{\mathcal{S}}^{T} \mathbf{D}_{\mathcal{S}} \mathbf{x}^{\mathcal{S}} = \mathbf{w}^{\mathcal{S}}$$
(3)
and $\mathbf{r}_{\mathcal{S}} = \mathbf{z} - \mathbf{D}_{\mathcal{S}} \mathbf{x}^{\mathcal{S}}$ (residual generated by \mathcal{S}),

where \mathbf{w}^{S} and \mathbf{D}_{S} are a sub-vector of $\mathbf{w} = \mathbf{D}^{T}\mathbf{y}$ indexed by S and a sub-matrix of \mathbf{D} formed by collecting its columns indexed by S, respectively. \mathbf{x}^{S} is the vector containing only active coefficients of \mathbf{x} indexed by S. The bijection $S \leftrightarrow \mathbf{x}^{S}$ implies that solving (2) is equivalent to finding the active set S. This motivates us to utilize a greedy approach to find the support set S and then infer the solution \mathbf{x} . In particular, an Adaptive Matching Pursuit (AMP) is proposed to update the active set S at each step by either absorbing one of the *unselected* indices into S or removing one of active elements in S, whichever results in a larger decrease in the cost function. Concretely, for a given S, we define:

$$\sigma_{\mathcal{S}} = \sum_{i \in \mathcal{S}} \rho_i \text{ and } g(\mathcal{S}) = \min_{\mathbf{x}^{\mathcal{S}}} \|\mathbf{z} - \mathbf{D}_{\mathcal{S}} \mathbf{x}^{\mathcal{S}}\|_2^2 + \sigma_{\mathcal{S}}.$$
(4)

At each iteration, we calculate two "improvement" values:

$$U_{\mathcal{S}} = \min_{i \notin \mathcal{S}} g(\mathcal{S} \cup \{i\}) - g(\mathcal{S}), \tag{5}$$

$$V_{\mathcal{S}} = \min_{j \in \mathcal{S}} g(\mathcal{S} \setminus \{j\}) - g(\mathcal{S}).$$
(6)

(5) is the decrease in cost function if selecting one of unselected indices in the support and (6) is the decrease if removing one already-selected index. If both U_S and V_S greater than or equal to 0, we can stop the algorithm since no improvement is obtained. Otherwise, we compare U_S and V_S to update S by absorbing *i* (if $U_S < V_S$) or removing *j* (if $U_S > V_S$). This procedure guarantees that the cost function decreases after each iteration, and then, the algorithm will eventually stop after finite iterations.

Nevertheless, the cost of calculating $g(S \cup \{i\})$ and $g(S \setminus \{j\})$ is extremely high and this idea becomes hardly practical. In order to significantly reduce the computation cost while keeping our algorithm close to the idea above, instead of exactly calculating U_S and V_S , we aim to calculate their competent upper bounds \overline{U}_S and \overline{V}_S . The decision is then made based on these approximated values to obtain S^{new} . After that, $\mathbf{x}^{S^{\text{new}}}$ and $\mathbf{r}_{S^{\text{new}}}$ are calculated *precisely* before moving to the new iteration. The following lemmas support AMP in initializing S and choosing \overline{U}_S and \overline{V}_S .

Lemma 1: If $\rho_i < 0$, then $i \in \hat{S}$ – the optimal active set. *Proof:* Suppose that $i \notin \hat{S}$. Let $\mathbf{r}_{\hat{S}} = \mathbf{z} - \mathbf{D}_{\hat{S}} \mathbf{x}^{\hat{S}}$, we have:

$$g(\hat{\mathcal{S}} \cup \{i\}) \stackrel{\text{by}(4)}{\leq} \|\mathbf{r}_{\hat{\mathcal{S}}} - x_i \mathbf{d}_i\|_2^2 + \sigma_{\hat{\mathcal{S}}} + \rho_i$$

$$= g(\hat{\mathcal{S}}) + \mathbf{d}_i^T \mathbf{d}_i x_i^2 - 2\mathbf{r}_{\hat{\mathcal{S}}}^T \mathbf{d}_i x_i + \rho_i$$

$$= g(\hat{\mathcal{S}}) + (1 + \lambda)x_i^2 - 2\mathbf{r}_{\hat{\mathcal{S}}}^T \mathbf{d}_i x_i + \rho_i.$$
(7)

Let $h(x) = (1+\lambda)x^2 - 2\mathbf{r}_{\hat{\mathcal{S}}}^T \mathbf{d}_i x + \rho_i$, we observe that h(x) is continuous and $\lim_{x\to\infty} h(x) = +\infty$ and $h(0) = \rho_i < 0$. Then, there exists $\bar{x}_i \neq 0$ such that $\rho_i < h(\bar{x}_i) < 0$. Combining this with (7), we have:

$$g(\hat{\mathcal{S}} \cup \{i\}) \le g(\hat{\mathcal{S}}) + h(\bar{x}_i) < g(\hat{\mathcal{S}}).$$

In other words, $\hat{S} \cup \{i\}$ generates a lower cost than \hat{S} does, which contradicts the assumption that \hat{S} is the optimal solution. The contradiction suggests that *i* must be in \hat{S} . By using Lemma 1, we can initialize $S_0 = \{i : \rho_i < 0\}$.

Lemma 2:
$$U_{\mathcal{S}} \leq \min_{i \notin \mathcal{S}} \left\{ \rho_i - \frac{(\mathbf{r}_{\mathcal{S}}^T \mathbf{d}_i)^2}{1+\lambda} \right\} \triangleq \overline{U}_{\mathcal{S}}.$$
 (8)

Proof: Since (7) holds for every x_i , we have:

$$g(\mathcal{S} \cup \{i\}) \leq g(\mathcal{S}) + \min_{x_i} \left\{ \rho_i + (1+\lambda)x_i^2 - 2\mathbf{r}_{\mathcal{S}}^T \mathbf{d}_i x_i \right\}$$
$$= g(\mathcal{S}) + \rho_i - \frac{(\mathbf{r}_{\mathcal{S}}^T \mathbf{d}_i)^2}{1+\lambda}. \tag{9}$$

Since (9) holds for every $i \notin S$, inequality (8) is true.

Lemma 3:

$$V_{\mathcal{S}} \le \min_{j \in \mathcal{S}} \left\{ (1+\lambda) (x_j^{\mathcal{S}})^2 + 2\mathbf{d}_j^T \mathbf{r}_{\mathcal{S}} x_j^{\mathcal{S}} - \rho_j \right\} \triangleq \overline{V}_{\mathcal{S}}.$$
(10)

where $x_i^{\mathcal{S}}$ is the element of $\mathbf{x}^{\mathcal{S}}$ corresponding to index $j \in \mathcal{S}$. *Proof:* For any $j \in S$,

$$g(\mathcal{S}) = \|\mathbf{r}_{\mathcal{S}} + \mathbf{d}_{j}x_{j}^{\mathcal{S}} - \mathbf{d}_{j}x_{j}^{\mathcal{S}}\|_{2}^{2} + \sigma_{\mathcal{S}\setminus\{j\}} + \rho_{j}$$

$$= \|\mathbf{r}_{\mathcal{S}} + \mathbf{d}_{j}x_{j}^{\mathcal{S}}\|_{2}^{2} + \sigma_{\mathcal{S}\setminus\{j\}} + (1+\lambda)(x_{j}^{\mathcal{S}})^{2} + \rho_{j} - 2\mathbf{d}_{j}^{T}(\mathbf{r}_{\mathcal{S}} + \mathbf{d}_{j}x_{j}^{\mathcal{S}})x_{j}^{\mathcal{S}}$$

$$\geq g(\mathcal{S}\setminus\{j\}) - (1+\lambda)(x_{j}^{\mathcal{S}})^{2} - 2\mathbf{d}_{j}^{T}\mathbf{r}_{\mathcal{S}}x_{j}^{\mathcal{S}} + \rho_{j}. \quad (11)$$

Combining (11) and (6) we can conclude (10).

From Lemma 2 and Lemma 3, we can consider \overline{U}_{S} and \overline{V}_{S} as good approximations of U_{S} and V_{S} and update the support set S based on those. It is worth to note here that calculating $\overline{U}_{\mathcal{S}}$ and $\overline{V}_{\mathcal{S}}$ requires low computations.

After obtaining updated S, we have to find \mathbf{x}^{S} using (3) where we present a computationally cheap solution for that. Let $\mathbf{D}_{S}^{T}\mathbf{D}_{S} = \mathbf{L}_{S}\mathbf{L}_{S}^{T}$ be the Cholesky decomposition of the symmetric positive definite matrix $\mathbf{D}_{\mathcal{S}}^T \mathbf{D}_{\mathcal{S}}$ where L_{S} is a low triangular matrix. Then, we can rewrite (3) as $\mathbf{L}_{\mathcal{S}}\mathbf{L}_{\mathcal{S}}^{T}\mathbf{x}^{\mathcal{S}} = \mathbf{w}^{\mathcal{S}}$. Subsequently, $\mathbf{x}^{\mathcal{S}}$ can be updated by solving two equations $\mathbf{L}_{\mathcal{S}}\mathbf{u} = \mathbf{w}^{\mathcal{S}}$ and $\mathbf{L}_{\mathcal{S}}^{T}\mathbf{x}^{\mathcal{S}} = \mathbf{u}$, consecutively. These two equations are simply solved by forward substitution and backward substitution, respectively, thanks to the triangularity of L_S . The remaining task now is that given $L_{\mathcal{S}}$, how do we effectively obtain $L_{\mathcal{S}\cup\{i\}}$ and $L_{S \setminus \{j\}}$? While the $L_{S \cup \{i\}}$ update procedure is quite classical, as far as we know, the $L_{S \setminus \{j\}}$ update procedure has not been widely addressed. Details of these two procedures are given below:

1. Given L_S , calculate $L_{S \cup \{i\}}$. This problem has been tackled in [27]. Specifically, by writing $\mathbf{D}_{S \cup \{i\}}$ = $|\mathbf{D}_{\mathcal{S}} \mathbf{d}_i|$, we have:

$$\begin{aligned} \mathbf{L}_{\mathcal{S}\cup\{i\}}\mathbf{L}_{\mathcal{S}\cup\{i\}}^T &= \mathbf{D}_{\mathcal{S}\cup\{i\}}^T\mathbf{D}_{\mathcal{S}\cup\{i\}} \\ &= \begin{bmatrix} \mathbf{D}_{\mathcal{S}}^T\mathbf{D}_{\mathcal{S}} & \mathbf{D}_{\mathcal{S}}^T\mathbf{d}_i \\ \mathbf{d}_i^T\mathbf{D}_{\mathcal{S}} & \mathbf{d}_i^T\mathbf{d}_i \end{bmatrix} = \begin{bmatrix} \mathbf{L}_{\mathcal{S}}\mathbf{L}_{\mathcal{S}}^T & \mathbf{D}_{\mathcal{S}}^T\mathbf{d}_i \\ \mathbf{d}_i^T\mathbf{D}_{\mathcal{S}} & \mathbf{1}+\lambda \end{bmatrix} \end{aligned}$$

If v is the solution of $\mathbf{L}_{\mathcal{S}} \mathbf{v} = \mathbf{D}_{\mathcal{S}}^T \mathbf{d}_i$ via forward substitution, then we can easily infer:

$$\mathbf{L}_{\mathcal{S}\cup\{i\}} = \begin{bmatrix} \mathbf{L}_{\mathcal{S}} & \mathbf{0} \\ \mathbf{v}^T & \sqrt{1+\lambda-\mathbf{v}^T\mathbf{v}} \end{bmatrix}.$$
 (12)

2. Given $L_{\mathcal{S}}$, calculate $L_{\mathcal{S}\setminus\{i\}}$. Suppose that: $\mathbf{D}_{\mathcal{S}}^T \mathbf{D}_{\mathcal{S}} =$

$$=\underbrace{\begin{bmatrix} \mathbf{L}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{l}_{21}^T & l_{22} & \mathbf{0} \\ \mathbf{L}_{31} & \mathbf{l}_{32} & \mathbf{L}_{33} \end{bmatrix}}_{\mathbf{L}_{\mathcal{S}}} \underbrace{\begin{bmatrix} \mathbf{L}_{11}^T & l_{21} & \mathbf{L}_{31}^T \\ \mathbf{0} & l_{22} & l_{32}^T \\ \mathbf{0} & \mathbf{0} & \mathbf{L}_{33}^T \end{bmatrix}}_{\mathbf{L}_{\mathcal{S}}^T} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{c}_{12} & \mathbf{C}_{13} \\ \mathbf{c}_{12}^T & \mathbf{c}_{22} & \mathbf{c}_{32}^T \\ \mathbf{C}_{31} & \mathbf{c}_{32} & \mathbf{C}_{33} \end{bmatrix},$$

where $[\mathbf{c}_{12}^T, c_{22}, \mathbf{c}_{32}^T]$ and $[\mathbf{c}_{12}^T, c_{22}, \mathbf{c}_{32}^T]^T$ are the row and column corresponding to the removed index j. By writing:

Algorithm 1 AMP algorithm for solving problem (1)

function
$$(\mathbf{x}^*, \boldsymbol{\gamma}^*) = AMP(\mathbf{y}, \mathbf{A}, \lambda, \boldsymbol{\rho}).$$

1. $\mathbf{D} = \begin{bmatrix} \mathbf{A} \\ \sqrt{\lambda}\mathbf{I} \end{bmatrix}$ and $\mathbf{z} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$
2. Initialize $S = \{i : \rho_i < 0\}$ and $\mathbf{L}_S : \mathbf{L}_S \mathbf{L}_S^T = \mathbf{D}_S^T \mathbf{D}_S.$
% if $S = \emptyset$, then $\mathbf{L}_S = []$.
while true do
3. Solve $\mathbf{u} : \mathbf{L}_S \mathbf{u} = \mathbf{w}^S$ then solve $\mathbf{x}^S : \mathbf{L}_S^T \mathbf{x}^S = \mathbf{u}.$
4. Update residual: $\mathbf{r}_S = \mathbf{z} - \mathbf{D}_S \mathbf{x}^S.$
5. Calculate: $[\overline{U}_S, i] = \min_{i \notin S} \left\{ \rho_i - \frac{(\mathbf{r}_S^T \mathbf{d}_i)^2}{1 + \lambda} \right\}.$
6. Calculate $\overline{V}_S:$
 $[\overline{V}_S, j] = \min_{j \in S} \left\{ (1 + \lambda) (x_j^S)^2 + 2\mathbf{d}_j^T \mathbf{r}_S x_j - \rho_j \right\}.$
7. Decide
if $\min(\overline{U}_S, \overline{V}_S) \ge 0$ then break the while loop.
else if $\overline{U}_S < \overline{V}_S$ then
Insert index: $S = S \cup \{i\}$ and update \mathbf{L}_S by (12).
else Remove index: $S = S \setminus \{j\}$ and update \mathbf{L}_S by (14).
end if
end while
8. OUTPUT: $S \Rightarrow \gamma^*$; $\mathbf{x}^S \Rightarrow \mathbf{x}^*$

$$\mathbf{D}_{\mathcal{S}\backslash\{j\}}^{T}\mathbf{D}_{\mathcal{S}\backslash\{j\}} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{13} \\ \mathbf{C}_{31} & \mathbf{C}_{33} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{\bar{L}}_{11} & \mathbf{0} \\ \mathbf{\bar{L}}_{31} & \mathbf{\bar{L}}_{33} \end{bmatrix}}_{\mathbf{L}_{\mathcal{S}\backslash\{j\}}} \underbrace{\begin{bmatrix} \mathbf{\bar{L}}_{11}^{T} & \mathbf{\bar{L}}_{31}^{T} \\ \mathbf{0} & \mathbf{\bar{L}}_{33}^{T} \end{bmatrix}}_{\mathbf{L}_{\mathcal{S}\backslash\{j\}}}$$

we obtain:

fun

$$\begin{aligned}
 \mathbf{C}_{11} &= \mathbf{L}_{11} \mathbf{L}_{11}^T = \bar{\mathbf{L}}_{11} \bar{\mathbf{L}}_{11}^T \Rightarrow \bar{\mathbf{L}}_{11} = \mathbf{L}_{11}, \quad (13) \\
 \mathbf{C}_{13} &= \mathbf{L}_{11} \mathbf{L}_{31}^T = \bar{\mathbf{L}}_{11} \bar{\mathbf{L}}_{31}^T \Rightarrow \bar{\mathbf{L}}_{31} = \mathbf{L}_{31}, \\
 \mathbf{C}_{33} &= \mathbf{L}_{31} \mathbf{L}_{31}^T + \mathbf{l}_{32} \mathbf{l}_{32}^T + \mathbf{L}_{33} \mathbf{L}_{33}^T = \bar{\mathbf{L}}_{31} \bar{\mathbf{L}}_{31}^T + \bar{\mathbf{L}}_{33} \bar{\mathbf{L}}_{33}^T, \\
 \Rightarrow \bar{\mathbf{L}}_{33} \bar{\mathbf{L}}_{33}^T = \mathbf{L}_{33} \mathbf{L}_{33}^T + \mathbf{l}_{32} \mathbf{l}_{32}^T.
 \end{aligned}$$

Note that since Cholesky decomposition is unique we can obtain (13). Also since l_{32} is a vector, $l_{32}l_{32}^T$ is a rank-one matrix. Therefore, $\overline{\mathbf{L}}_{33}$ can be obtained quickly from \mathbf{L}_{33} and l_{32} using rank one update for the Cholesky decomposition [28] and we obtain:

$$\mathbf{L}_{\mathcal{S}\setminus\{j\}} = \begin{bmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{L}_{31} & \bar{\mathbf{L}}_{33} \end{bmatrix}.$$
 (14)

Altogether, the AMP algorithm is presented in Algorithm 1.

Remark: Problem (1) with non-negativity constraint can also be solved by slightly modifying Step 3 and Step 5 of Algorithm 1. In Step 3, \mathbf{x}^{S} is instead solved via: $\mathbf{x}^{S} = \arg \min_{\mathbf{x}^{S} \succeq 0} \|\mathbf{z} - \mathbf{D}_{S} \mathbf{x}^{S}\|_{2}^{2}$; while in Step 5, the alternative is: $[\overline{U}_{\mathcal{S}}, i] = \min_{i \notin \mathcal{S}} \left\{ \rho_i - (\max\{\mathbf{r}_{\mathcal{S}}^T \mathbf{d}_i, 0\})^2 / (1+\lambda) \right\}.$ The non-negative quadratic programming in Step 3 can be solved by an ADMM [15] procedure as proposed in [29] which can also benefit from our Cholesky decomposition to make it computationally cheap¹.

¹Details of this procedure will be discussed in our future work.



Fig. 1: Comparison of (from left to right): running time, mean squared error (MSE), cost function and Support Match (SM) obtained by each method versus sparsity level of \mathbf{x}_0 with non-negativity constraint (p = 512, q = 256).

Table 1: Comparison of methods for p = 512, q = 256.

Average	CoSaMP	E-NET	SpaRSA	ICR	AMP
Time (s)	1.1E0	7.9E-2	2.2E-1	2.9E1	1.6E-1
MSE	3.0E-4	4.5E-3	5.4E-4	3.1E-3	6.1E-5
Cost	-	-	1.6E-1	2.1E-1	9.5E-2
SM (%)	93.52	85.55	93.87	83.79	97.32

3. EXPERIMENTAL RESULTS

To illustrate the effectiveness of our AMP algorithm, we apply it to sparse recovery problems in two different scenarios: i) simulated data and ii) a real-world image recovery problem. Comparisons are made against state-of-the-art alternatives: 1) CoSaMP [30] / NNOMP [31]²; 2) Elastic Net (E-NET) [18] 3) SpaRSA [8] and 4) ICR [10].

Simulated data: We set up a typical experiment for sparse recovery as in [25, 32] with a randomly generated Gaussian matrix $\mathbf{A} \in \mathbb{R}^{256 \times 512}$ and a sparse vector $\mathbf{x}_0 \in \mathbb{R}^{512}$ with 100 non-zeros. Based on A and x_0 , we form the observation vector $\mathbf{y} \in \mathbb{R}^{256}$ according to the additive noise model with $\sigma = 0.01$: y = Ax₀ + n. Table 1 reports the experimental results (averages of 100 trials with different \mathbf{A}, \mathbf{x}_0 and \mathbf{n}) for this problem. As can be seen from Table 1, AMP outperforms the competing methods in many different aspects. AMP runs almost two hundred times faster than ICR and seven times faster than CoSaMP. Results in the second row of Table 1 reveal that the AMP solution is the closest to the "ground truth" x_0 in terms of mean square error (MSE). In terms of the cost function, we compare AMP with two other methods, ICR and SpaRSA, which solve the same optimization problem (1). The third row confirms that AMP achieves the lowest cost function (by more than one order of magnitude). Finally, we use "support match" (SM) to measure how much the support of each solution matches to that of x_0 . It is evident that AMP provides significantly higher support match (SM = 97.32%) than other state-of-the-art methods. Next, in order to see how each method performs in the presence of non-negativity constraints, we perform one more experiment with sparse



Fig. 2: Examples of reconstructed images from MNIST dataset using different methods. In each pair of images: left is the original image x_0 (for the first row) or the reconstructed image (for other rows), right is its "support" image. Numbers on the left are MSEs and on the right are Support Matches.

data that is naturally non-negative and vary the sparsity level of \mathbf{x}_0 from 10 to 120 and compare the running time, MSE, optimal cost function and SM of different methods. Results are shown in Figure 1. Similar trends can be seen in this figure where AMP requires less running time than ICR does while it consistently outperforms others in the remaining aspects. It is worth to mention that AMP and ICR obtain almost the identical cost which is better than what SpaRSA achieves.

Image recovery: Next, we apply the different sparse recovery algorithms to real data for image reconstruction. We work with the well-known handwritten digit images MNIST [33]. The dataset contains 60,000 digit images (0 to 9) of size 28×28 . Since most of pixels in each image are inactive (0), each image is naturally sparse. The experiment is set up such that a vectorized sparse signal $\mathbf{x} \in \mathbb{R}^{784 \times 1}$ is to be reconstructed from a smaller set of random measurements y. We randomly generate a Gaussian matrix $\mathbf{A} \in \mathbb{R}^{350 \times 784}$. For each signal, we add a Gaussian noise \mathbf{n} with $\sigma = 0.03$. Recovered images and their supports are demonstrated in Figure 2 along with the corresponding MSE and Support Match. Clearly, AMP and ICR outperform other methods with slightly better but much faster results provided by AMP.

²CoSaMP for unconstrained OMP, NNOMP for non-negative OMP.

4. REFERENCES

- [1] J. Wright, Allen Y Yang, Arvind Ganesh, Shankar S Sastry, and Yi Ma, "Robust face recognition via sparse representation," *IEEE Trans. on Pattern Analysis and Machine Int.*, vol. 31, no. 2, pp. 210–227, 2009.
- [2] P. Sprechmann, I. Ramirez, G. Sapiro, and Y. C. Eldar, "C-HiLASSO: A collaborative hierarchical sparse modeling framework," *IEEE Trans. on Signal Processing*, vol. 59, no. 9, pp. 4183–4198, 2011.
- [3] U. Srinivas, Y. Suo, Minh Dao, V. Monga, and T. D Tran, "Structured sparse priors for image classification," *IEEE Transactions on Image Processing*, vol. 24, no. 6, pp. 1763–1776, 2015.
- [4] F. Pourkamali Anaraki and S. M. Hughes, "Compressive K-SVD," in Proc. IEEE Int. on Conf. Acoustics, Speech, and Signal Processing. IEEE, 2013, pp. 5469–5473.
- [5] T. H. Vu, H. S. Mousavi, V. Monga, UK Rao, and G. Rao, "DFDL: Discriminative feature-oriented dictionary learning for histopathological image classification," *Proc. IEEE Int. Symp. Biomed. Imag.*, pp. 990–994, 2015.
- [6] T. H. Vu, H. S. Mousavi, V. Monga, UK Rao, and G. Rao, "Histopathological image classification using discriminative feature-oriented dictionary learning," *IEEE Transactions on Medical Imaging*, vol. 35, no. 3, pp. 738–751, March, 2016.
- [7] T. H. Vu and V. Monga, "Learning a low-rank shared dictionary for object classification," *Proc. IEEE Conf. on Image Processing*, pp. 4428–4432, 2016.
- [8] S. J. Wright, R. D. Nowak, and M. AT Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. on Signal Processing*, vol. 57, no. 7, pp. 2479–2493, 2009.
- [9] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans.* on Info. Theory, vol. 53, no. 12, pp. 4655–4666, 2007.
- [10] H. S. Mousavi, V. Monga, and T. D. Tran, "Iterative convex refinement for sparse recovery," *IEEE Signal Processing Letters*, vol. 22, no. 11, pp. 1903–1907, 2015.
- [11] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. on Im. Processing*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [12] H. Mohimani, M. Babaie-Zadeh, and C. Jutten, "A fast approach for overcomplete sparse decomposition based on smoothed norm," *IEEE Trans. on Signal Processing*, vol. 57, no. 1, pp. 289–301, 2009.
- [13] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. on Signal Processing*, vol. 56, no. 6, pp. 2346– 2356, 2008.
- [14] X. Lu, Y. Wang, and Y. Yuan, "Sparse coding from a Bayesian perspective," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 24, no. 6, pp. 929–939, 2013.
- [15] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends* (R) in *Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

- [16] U. Srinivas, Y. Suo, Minh Dao, V. Monga, and T. D Tran, "Structured sparse priors for image classification.," in *Proc. IEEE Conf. on Image Processing*, 2013, pp. 3211–3215.
- [17] R. Jenatton, J-Y Audibert, and F. Bach, "Structured variable selection with sparsity-inducing norms," *The Journal of Machine Learning Research*, vol. 12, pp. 2777–2824, 2011.
- [18] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Stat. Society: Series B (Stat. Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [19] S.D. Babacan, R. Molina, and A.K. Katsaggelos, "Bayesian compressive sensing using laplace priors," *IEEE Trans. on Image Processing*, vol. 19, no. 1, pp. 53–63, 2010.
- [20] V. Cevher, P. Indyk, L. Carin, and R. G. Baraniuk, "Sparse signal recovery and acquisition with graphical models," *Signal Processing Magazine*, *IEEE*, vol. 27, no. 6, pp. 92–103, 2010.
- [21] T. J. Mitchell and J. J. Beauchamp, "Bayesian variable selection in linear regression," *Journal of the American Statistical Association*, vol. 83, no. 404, pp. 1023–1032, 1988.
- [22] H. Ishwaran and J. S. Rao, "Spike and Slab variable selection: frequentist and Bayesian strategies," *Annals of Statistics*.
- [23] M. R. Andersen, O. Winther, and L. K. Hansen, "Bayesian inference for structured Spike and Slab priors," in Advances in Neural Information Processing Systems, 2014, pp. 1745–1753.
- [24] M. Lázaro-gredilla and M. K. Titsias, "Spike and Slab variational inference for multi-task and multiple kernel learning," in Advances in neural information processing systems, 2011, pp. 2339–2347.
- [25] T. J. Yen et al., "A majorization-minimization approach to variable selection using Spike and Slab priors," *The Annals of Statistics*, vol. 39, no. 3, pp. 1748–1775, 2011.
- [26] V. Cevher, "Learning with compressible priors," in Advances in Neural Information Processing Systems, 2009, pp. 261–269.
- [27] R. Rubinstein, M. Zibulevsky, and M. Elad, "Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit," *CS Technion*, vol. 40, no. 8, pp. 1–15, 2008.
- [28] M. Seeger, "Low rank updates for the Cholesky decomposition," Tech. Rep., 2004.
- [29] D. Boley, "Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2183–2207, 2013.
- [30] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied* and Computational Harmonic Analysis, vol. 26, no. 3, pp. 301– 321, 2009.
- [31] A. M Bruckstein, M. Elad, and M. Zibulevsky, "On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations," *IEEE Transactions on Information Theory*, vol. 54, no. 11, pp. 4813–4820, 2008.
- [32] A. Beck and M. Teboulle, "A fast iterative shrinkagethresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [33] Y. LeCun, C. Cortes, and Burges C. J.C., "MNIST dataset," http://yann.lecun.com/exdb/mnist/, Accessed: 2016-08-30.