

Sparse Reconstruction for RADAR

Lee C. Potter, Philip Schniter, and Justin Ziniel

Dept. of Electrical & Computer Engineering, The Ohio State University, Columbus, OH, US

ABSTRACT

Imaging is not itself a system goal, but is rather a means to support inference tasks. For data processing with linearized signal models, we seek to report all high-probability interpretations of the data and to report confidence labels in the form of posterior probabilities. A low-complexity recursive procedure is presented for Bayesian estimation in linear regression models. A Gaussian mixture is chosen as the prior on the unknown parameter vector. The algorithm returns both a set of high posterior probability mixing parameters and an approximate minimum mean squared error (MMSE) estimate of the parameter vector. Emphasis is given to the case of a sparse parameter vector. Numerical simulations demonstrate estimation performance and illustrate the distinctions between MMSE estimation and maximum *a posteriori* probability (MAP) model selection. The proposed tree-search algorithm provides exact ratios of posterior probabilities for a set of high probability solutions to the sparse reconstruction problem. These relative probabilities serve to reveal potential ambiguity among multiple candidate solutions that are ambiguous due to low signal-to-noise ratio and/or significant correlation among columns in the super-resolving regressor matrix.

Keywords: sparse reconstruction, radar imaging, Gaussian mixture models

1. INTRODUCTION

1.1 High frequency scattering

At high frequencies, the scattering response of an object can be well approximated as a sum of responses from individual scattering centers.¹ These scatters provide a concise, yet physically relevant, description of the object. Additionally, a parametric representation of the sum of reflectors offers a low dimensional manifold for inversion and compression of phase history data. For example, an approximate physical optics model for a set of canonical reflectors^{2,3} parameterizes responses as a functions of frequency, aspect, elevation and polarization. Fitting the models to phase history data is a nonlinear regression task, with the accompanying challenges of computational cost, model selection (order selection and reflector type detection), and local minima in the non-convex optimization. Yet, the nonlinear regression can provide a compressed, informative, super-resolved interpretation of the scattering data. For example, Fig. 1 shows a graphical representation of parameters estimated for vehicle in the 640 MHz bandwidth Gotcha public release data.⁴ The HH and VV polarization images from eight circular passes are jointly and coherently processed by automated model fitting. Estimated parameters for 3D position, orientation, length and polarization (even versus odd bounce) are displayed as icons in the figure. For visualization, a CAD model of the actual vehicle is superimposed in the figure.⁵ The eight passes provide only a sparse sampling across less than 2 degrees elevation, and hence present high side lobes and aliasing in backprojection or Fourier 3D imaging.

1.2 Sparse linear regression

An alternative to the nonlinear regression is to recast the inverse scattering problem as a linear regression. The unknown coefficients are taken as the reflections coefficients at discretely sampled 3D locations, and a regressor matrix, \mathbf{A} , is typically formed by sampling some parametric model (such as an ideal point scattering model). Thus, the linear model is given by

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\nu}, \quad (1)$$

Send correspondence to LCP, potter.36@osu.edu.




Figure 1. Nonlinear regression yields a parsimonious interpretation of 8-pass circular SAR data from the Gotcha Public Release. Icons depict the estimated parameters describing 3D position, orientation, length and polarization (even versus odd bounce) for individual scattering centers. A CAD model is superimposed for visualization.

with unknown parameter vector \mathbf{x} , unit norm columns in the regressor matrix \mathbf{A} , and additive noise $\boldsymbol{\nu}$. The high-frequency approximation motivates a preference towards a sparse solution in which few coefficients in \mathbf{x} are non-zero.

Sparse linear regression is a topic of long standing interest in statistics, geophysics, and signal processing. Algorithmic approaches have been proposed over several decades. One class of algorithms adopts a greedy search heuristic. Examples include CLEAN,⁶ projection pursuit,⁷ and orthogonal matched pursuit (OMP).⁸ Tropp and Gilbert,⁹ for example, provide sufficient conditions on the sparseness of \mathbf{x} and correlation among columns of \mathbf{A} such that the greedy OMP provides correct model selection with high probability in the noiseless measurement case. A second class of recursive algorithms solves a sequence of iteratively re-weighted linear least-squares (IRLS) problems;^{10, 11} recent results have likewise established sufficient conditions such that the sequence converges to the sparsest solution for a given residual energy.^{12, 13}

The method of sparse Bayesian learning^{14, 15} explicitly adopts a Bayesian framework with x_i independent, zero-mean, Gaussian with unknown variance σ_i^2 . The unknown variances are given the Gamma conjugate prior, and an expectation-maximization (EM) iteration computes a MAP estimate of \mathbf{x} .

In addition to greedy, IRLS and Bayesian learning approaches, penalized least-squares solutions for \mathbf{x} have likewise been presented in the past four decades.¹⁶ In this class of approaches, parameters are found via the optimization

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{x}\|_p^p, \quad (2)$$

or, equivalently for some $\epsilon > 0$

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_p \text{ s.t. } \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 < \epsilon. \quad (3)$$

Ridge regression¹⁷ (Tikhonov regularization) adopts $p = 2$, while basis pursuit¹⁸ and Lasso¹⁹ use $p = 1$. Equation (2) has been widely adopted, for example in radar imaging,²⁰ image reconstruction,^{21, 22} and elsewhere.^{23, 24} The large class of methods adopting (2) may be interpreted as implicitly seeking the Bayesian MAP estimate of \mathbf{x} under a sparsity inducing prior

$$p(\mathbf{x}) \propto \exp \left\{ -\frac{\lambda}{2} \|\mathbf{x}\|_p^p \right\}. \quad (4)$$

Elegant recent results by several authors²⁵⁻²⁷ have demonstrated sufficient conditions on \mathbf{A} , $\boldsymbol{\nu}$, and sparsity of \mathbf{x} such that the convex problem in (3) for $p = 1$ provides the unique solution to the non-convex task

$$\min \|\mathbf{x}\|_0 \text{ s.t. } \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 < \epsilon. \quad (5)$$

These proofs have validated the widespread use of (2)-(3), providing a deeper understanding, spurring a resurgent interest, and promoting the interpretation as “compressive sensing.” The sufficient conditions restrict the sparseness of \mathbf{x} and the correlation among the columns in the regressor matrix, \mathbf{A} . Example conditions are the restricted isometry property²⁷ (RIP) and the mutual coherence.²⁶ The low correlation required among columns of \mathbf{A} is incompatible with super-resolution at any SNR.

The constructive design of experiments consistent with RIP (design of \mathbf{A}) remains open.²⁸ But the compressive sensing hypotheses are met with high probability by draws from classes of random matrices. In this sense, compressive sensing trades the NP-hard ℓ_0 sparsest solution task for an intractable experiment design, then uses randomization for experiment design. Randomization has been used for over 40 years in array processing for low side-lobe responses.^{29,30} Thus, compressive sensing theorems offer an invitation to randomized sampling.

For radar imaging, traditional phase history sampling and point scattering models result in an \mathbf{A} matrix that has high correlation among columns; for example, the maximum correlation between columns of \mathbf{A} is 0.38 for MSTAR image chips. And, the aperture randomization used in magnetic resonance k -space sampling³⁰ would require random transport of an antenna across the full viewing hemisphere. For temporal sampling, phase jitter and sample-hold bandwidth limit SNR for randomized sampling. Thus, application of the randomization concept to radar imaging remains an open opportunity and challenge.

In the sparse reconstruction and compressive sensing literature, primary focus is placed on the detection of the few significant entries of the sparse \mathbf{x} —a task alternatively known as model selection or basis selection. In addition, an estimate of \mathbf{x} is also sought. Further, a single solution is reported, and typically without a report of uncertainty.

1.3 A third way

For data processing with linearized signal models, we adopt an alternative approach that embraces ambiguity among competing sparse interpretations of the noisy data. Linearization offers the potential to avoid the model order and computational difficulties of nonlinear regression. Our approach is motivated by the reality that reconstruction of \mathbf{x} (i.e., imaging) is not itself a system goal, but is rather a means to support inference tasks. Therefore, we seek to report all high-probability interpretations of the data, and to report confidence labels in the form of posterior probabilities. An analogy is found in digital communications, where capacity-achieving codes require soft-decisions on bits; the individual bits are not the goals, but are rather probabilistic evidence for decoding the message. Thus, we present a detection procedure that provides exact ratios of posterior probabilities for a set of high probability solutions to the sparse reconstruction problem. These relative probabilities serve to reveal uncertainty among multiple candidate solutions that are ambiguous due to low signal-to-noise ratio and/or significant correlation among columns in the (super-resolved) regressor matrix, \mathbf{A} . An additional output from the approach is a minimum mean-squared error (MMSE) estimator for \mathbf{x} from the noisy observations, \mathbf{y} . The MMSE estimation approach was likewise adopted in a crisp exposition by Larsson and Selén;³¹ as they noted, the Bayesian formulation requires *a priori* assumptions that are explicitly stated and admit precise interpretation. (We specifically identify similarities to³¹ in Section 5.)

The remainder of the paper is organized as follows. In Section 2, we state the signal model and explicitly identify the assumed priors. In Section 3, we describe our proposed technique, and in Section 4 we investigate its performance numerically. In Section 5, we discuss our findings, and in Section 6 we conclude.

2. SIGNAL MODEL

Consider observing $\mathbf{y} \in \mathbb{C}^M$, a noisy linear combination of the parameters in $\mathbf{x} \in \mathbb{C}^N$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\nu}, \tag{6}$$

where the noise $\boldsymbol{\nu}$ is assumed to be white circular Gaussian with variance σ^2 , i.e., $\boldsymbol{\nu} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_M)$, and the columns of \mathbf{A} are taken to be unit-norm. Our focus is on the under-determined case (i.e., $N \gg M$) with a suitably *sparse* parameter vector \mathbf{x} (i.e., $\|\mathbf{x}\|_0 \ll N$).

To model sparseness, we assume that the parameters are generated from a circular Gaussian mixture density:

$$\mathbf{x}|\mathbf{s} \sim \mathcal{CN}(\mathbf{0}, \mathbf{R}(\mathbf{s})), \quad (7)$$

where the covariance $\mathbf{R}(\mathbf{s})$ is determined by a discrete random vector $\mathbf{s} = [s_0, \dots, s_{N-1}]^T$ of mixture parameters. For simplicity, we take $\mathbf{R}(\mathbf{s})$ to be diagonal with $[\mathbf{R}(\mathbf{s})]_{n,n} = \sigma_{s_n}^2$, implying that $\{x_n|s_n\}_{n=0}^{N-1}$ are independent with $x_n|s_n \sim \mathcal{CN}(0, \sigma_{s_n}^2)$. Also for simplicity, we assume that the mixture parameters $\{s_n\}_{n=0}^{N-1}$ are* Bernoulli(p_1). To model sparseness in \mathbf{x} , we choose $\sigma_0^2 = 0$ and $p_1 \ll 1$.

From the model assumptions it can be seen that

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{x} \end{bmatrix} | \mathbf{s} \sim \mathcal{CN} \left(\mathbf{0}, \begin{bmatrix} \Phi(\mathbf{s}) & \mathbf{A}\mathbf{R}(\mathbf{s}) \\ \mathbf{R}(\mathbf{s})\mathbf{A}^H & \mathbf{R}(\mathbf{s}) \end{bmatrix} \right), \quad (8)$$

where

$$\Phi(\mathbf{s}) := \mathbf{A}\mathbf{R}(\mathbf{s})\mathbf{A}^H + \sigma^2 \mathbf{I}_M. \quad (9)$$

3. ESTIMATION OF BASIS AND PARAMETERS

In this section, we propose an efficient search procedure to find the most probable basis configurations along with their respective posterior probabilities. These posteriors can then be used to compute an MMSE estimate of sparse parameters \mathbf{x} .

3.1 Basis Selection Metric

As a consequence of the model described in Section 2, the locations of the nonzero elements in \mathbf{s} specify the set of “active” basis elements, i.e., columns of \mathbf{A} . Thus, basis selection reduces to estimation of \mathbf{s} . Because we have adopted a probabilistic model for $\{\mathbf{s}, \mathbf{y}\}$, we can not only compute *which* of the basis configurations are most probable, but also *how probable* these bases are. The latter is accomplished through the estimation of dominant posteriors $p(\mathbf{s}|\mathbf{y})$.

The posterior can be written, via Bayes rule, as

$$p(\mathbf{s}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{s})p(\mathbf{s})}{\sum_{\mathbf{s}' \in \mathcal{S}} p(\mathbf{y}|\mathbf{s}')p(\mathbf{s}')}, \quad (10)$$

where $\mathcal{S} = \{0, 1\}^N$, which shows that the estimation of $\{p(\mathbf{s}|\mathbf{y})\}_{\mathbf{s} \in \mathcal{S}}$ reduces to the estimation of $\{p(\mathbf{y}|\mathbf{s})p(\mathbf{s})\}_{\mathbf{s} \in \mathcal{S}}$. Although the size of \mathcal{S} makes it impractical to compute $\{p(\mathbf{s}|\mathbf{y})\}$ or $\{p(\mathbf{y}|\mathbf{s})p(\mathbf{s})\}$ for all $\mathbf{s} \in \mathcal{S}$, the subset $\mathcal{S}_* \subset \mathcal{S}$ that includes the *dominant* posteriors can be quite small and therefore practical to compute. Working in the log domain, we find

$$\mu(\mathbf{s}) := \ln p(\mathbf{y}|\mathbf{s})p(\mathbf{s}) \quad (11)$$

$$= \ln p(\mathbf{y}|\mathbf{s}) + \sum_{n=0}^{N-1} \ln p(s_n) \quad (12)$$

$$= \ln p(\mathbf{y}|\mathbf{s}) + \|\mathbf{s}\|_0 \ln p_1 + (N - \|\mathbf{s}\|_0) \ln(1 - p_1) \quad (13)$$

$$= -N \ln \pi - \ln \det(\Phi(\mathbf{s})) - \mathbf{y}^T \Phi(\mathbf{s})^{-1} \mathbf{y} \\ + \|\mathbf{s}\|_0 \ln \frac{p_1}{1-p_1} + N \ln(1 - p_1). \quad (14)$$

We will refer to $\mu(\mathbf{s})$ as the *basis selection metric*.

*In other words, s_n is binary with $\Pr\{s_n = 1\} = p_1$ and $\Pr\{s_n = 0\} = 1 - p_1$.

3.2 MMSE Parameter Estimation

For applications in which the identification of the most probable basis is the primary objective, the sparse coefficients \mathbf{x} can be regarded as nuisance parameters. In other applications, however, estimation of \mathbf{x} is the primary goal.

The MMSE estimate of \mathbf{x} from \mathbf{y} is

$$\hat{\mathbf{x}}_{\text{mmse}} := \mathbb{E}\{\mathbf{x}|\mathbf{y}\} = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}|\mathbf{y}) \mathbb{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\} \quad (15)$$

where from (8) it is straightforward to obtain

$$\mathbb{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\} = \mathbf{R}(\mathbf{s}) \mathbf{A}^H \Phi(\mathbf{s})^{-1} \mathbf{y}. \quad (16)$$

Approximating (15) with a summation over only the dominant basis configurations \mathcal{S}_* yields

$$\hat{\mathbf{x}}_{\text{ammse}} := \sum_{\mathbf{s} \in \mathcal{S}_*} p(\mathbf{s}|\mathbf{y}) \mathbb{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\}. \quad (17)$$

The posterior probability is a Gaussian mixture, and hence possibly multi-modal. The second moment is the conditional covariance $\text{Cov}\{\mathbf{x}|\mathbf{y}\}$ and can be closely approximated as

$$\text{Cov}\{\mathbf{x}|\mathbf{y}\} \approx \sum_{\mathbf{s} \in \mathcal{S}_*} p(\mathbf{s}|\mathbf{y}) \{ \text{Cov}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\} + (\hat{\mathbf{x}}_{\text{ammse}} - \hat{\mathbf{x}}_{\text{sy}})(\hat{\mathbf{x}}_{\text{ammse}} - \hat{\mathbf{x}}_{\text{sy}})^H \} \quad (18)$$

$$\text{Cov}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\} = \mathbf{R}(\mathbf{s}) - \mathbf{R}(\mathbf{s}) \mathbf{A}^H \Phi(\mathbf{s})^{-1} \mathbf{A} \mathbf{R}(\mathbf{s}). \quad (19)$$

Here, $\hat{\mathbf{x}}_{\text{sy}} = \mathbb{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\}$ from Eqn. 16. Note that, in evaluating (15) and (18), the primary challenge becomes that of obtaining $p(\mathbf{s}|\mathbf{y})$ and $\Phi(\mathbf{s})^{-1}$ for each $\mathbf{s} \in \mathcal{S}_*$. In the sequel, we propose a fast algorithm to search for the dominant basis configurations \mathcal{S}_* that, as a byproduct, also generates $\mathbb{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\}$ and $\text{Cov}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\}$ for each of the configurations \mathbf{s} identified by the search.

3.3 Bayesian Matching Pursuit

We next consider the estimation of the set of dominant basis configurations \mathcal{S}_* . Here we present a prosaic description of the proposed search heuristic. Later, in Section 3.5, we detail a fast algorithm.

The proposed search starts with $\mathbf{s} = \mathbf{0}$ and first “turns on” one mixture parameter at a time, yielding a set of N binary vectors \mathbf{s} which we refer to as $\mathcal{S}^{(1)}$. The metrics $\mu(\mathbf{s})$ for the N vectors in $\mathcal{S}^{(1)}$ are then computed, and the elements of $\mathcal{S}^{(1)}$ with the D largest metrics are collected in $\mathcal{S}_*^{(1)}$. For each candidate in $\mathcal{S}_*^{(1)}$, all locations of a second active mixture parameter are then considered, yielding $(N-1) + (N-2) + \dots + (N-D) = ND - \frac{D(D+1)}{2}$ unique binary vectors to store in $\mathcal{S}^{(2)}$. The metrics $\mu(\mathbf{s})$ for all vectors in $\mathcal{S}^{(2)}$ are then computed, and the elements of $\mathcal{S}^{(2)}$ with the D largest metrics are collected in $\mathcal{S}_*^{(2)}$. Then, for each candidate vector in $\mathcal{S}_*^{(2)}$, all unique possibilities of a third active mixture parameter are considered, and those with the D largest metrics are stored in $\mathcal{S}_*^{(3)}$. The process continues until $\mathcal{S}_*^{(P)}$ is computed, where P can be chosen so that[†] $\Pr(\|\mathbf{s}\|_0 > P)$ is sufficiently small.[‡] The algorithm’s final estimate of \mathcal{S}_* is then $\hat{\mathcal{S}}_* := \bigcup_{i=0}^P \mathcal{S}_*^{(i)}$.

The proposed algorithm, which we refer to as Bayesian matching pursuit (BMP), can be interpreted as an application of Cotter and Rao’s “MP:K” algorithm³² to maximization of the Bayesian metric $\mu(\mathbf{s})$.

[†]One could also determine the stopping parameter P adaptively.

[‡]Notice that $\|\mathbf{s}\|_0$ is Binomial(N, p_1) distribution. When $Np_1 > 5$, it is common to use the approximation $\|\mathbf{s}\|_0 \sim \mathcal{N}(Np_1, Np_1(1-p_1))$, in which case $\Pr(\|\mathbf{s}\|_0 > P) = \frac{1}{2} \text{erfc}\left(\frac{P - Np_1}{\sqrt{2Np_1(1-p_1)}}\right)$.

3.4 Fast Metric Update

To accelerate the aforementioned BMP algorithm, we propose a fast metric update which computes the change in $\mu(\cdot)$ that results from the activation of a single basis element. More precisely, if we denote by \mathbf{s}_n the vector identical to \mathbf{s} except for the n^{th} element, which is active in \mathbf{s}_n but inactive in \mathbf{s} (i.e., $[\mathbf{s}_n]_n = 1$ and $[\mathbf{s}]_n = 0$), then we seek an efficient method of computing $\Delta_n(\mathbf{s}) := \mu(\mathbf{s}_n) - \mu(\mathbf{s})$. Note that the metric at the root node (i.e., $\mathbf{s} = \mathbf{0}$) is

$$\mu(\mathbf{0}) = -N \ln \pi - M \ln \sigma^2 - \frac{1}{\sigma^2} \|\mathbf{y}\|_2^2 + N \ln(1 - p_1) \quad (20)$$

via (14) and the fact that $\Phi(\mathbf{0}) = \sigma^2 \mathbf{I}_M$.

To derive the fast metric update, we start with the property that, for any n and \mathbf{s} ,

$$\Phi(\mathbf{s}_n) = \Phi(\mathbf{s}) + \sigma_1^2 \mathbf{a}_n \mathbf{a}_n^H, \quad (21)$$

from which the matrix inversion lemma implies

$$\Phi(\mathbf{s}_n)^{-1} = \Phi(\mathbf{s})^{-1} - \sigma_1^2 \beta_n \mathbf{b}_n \mathbf{b}_n^H \quad (22)$$

$$\mathbf{b}_n := \Phi(\mathbf{s})^{-1} \mathbf{a}_n \quad (23)$$

$$\beta_n := (1 + \sigma_1^2 \mathbf{a}_n^H \mathbf{b}_n)^{-1}. \quad (24)$$

Equations (21)-(24) then imply

$$\mathbf{y}^H \Phi(\mathbf{s}_n)^{-1} \mathbf{y} = \mathbf{y}^H \Phi(\mathbf{s})^{-1} \mathbf{y} - \sigma_1^2 \beta_n |\mathbf{y}^H \mathbf{b}_n|^2 \quad (25)$$

$$\ln \det(\Phi(\mathbf{s}_n)) = \ln \det(\Phi(\mathbf{s}) + \sigma_1^2 \mathbf{a}_n \mathbf{a}_n^H) \quad (26)$$

$$\begin{aligned} &= \ln \left[(1 + \sigma_1^2 \mathbf{a}_n^H \Phi(\mathbf{s})^{-1} \mathbf{a}_n) \det(\Phi(\mathbf{s})) \right] \\ &= \ln \det(\Phi(\mathbf{s})) - \ln \beta_n \end{aligned} \quad (27)$$

$$\|\mathbf{s}_n\|_0 \ln \frac{p_1}{1-p_1} = (\|\mathbf{s}\|_0 + 1) \ln \frac{p_1}{1-p_1} \quad (28)$$

$$= \|\mathbf{s}\|_0 \ln \frac{p_1}{1-p_1} + \ln \frac{p_1}{1-p_1}, \quad (29)$$

which, in conjunction with (14), yield

$$\mu(\mathbf{s}_n) = \underbrace{\mu(\mathbf{s}) + \ln \beta_n + \sigma_1^2 \beta_n |\mathbf{y}^H \mathbf{b}_n|^2}_{\Delta_n(\mathbf{s})} + \ln \frac{p_1}{1-p_1}. \quad (30)$$

In summary, $\Delta_n(\mathbf{s})$ in (30) quantifies the change in our basis selection metric $\mu(\cdot)$ due to the activation of the n^{th} tap of \mathbf{s} .

3.5 Fast Bayesian Matching Pursuit

Notice that the parameters $\{\beta_n\}_{n=0}^{N-1}$, which are essential for the metric exploration step (30), require $\mathcal{O}(NM^2)$ operations to compute using (23)-(24) with standard matrix multiplication. As described next, the structure of $\Phi(\mathbf{s})^{-1}$ can be exploited to make this complexity in M .

Consider the $(i+1)^{\text{th}}$ stage of the sequential search described in Section 3.3, where the task is to evaluate the metrics corresponding to all valid single-tap extensions of a surviving hypothesis $\mathbf{s}^{(i)}$, where $\|\mathbf{s}^{(i)}\| = i$. Let us denote the (complete) set of these extensions by $\{\mathbf{s}_n^{(i+1)}\}_{n=0}^{N-1}$, remembering that particular extensions (corresponding, e.g., to re-activation of already-active taps) may be invalid. Say that, at the i^{th} stage, the quantities $\{\mathbf{b}_n^{(i)}\}_{n=0}^{N-1}$ and $\{\beta_n^{(i)}\}_{n=0}^{N-1}$ were computed in order to estimate the metrics $\{\mu(\mathbf{s}_n^{(i)})\}_{n=0}^{N-1}$ via (30), from which the surviving configuration $\mathbf{s}^{(i)}$ was chosen as

$$\mathbf{s}^{(i)} = \mathbf{s}_{n_*}^{(i)} \quad \text{for } n_* = \underset{n}{\operatorname{argmax}} \mu(\mathbf{s}_n^{(i)}).$$

Then, to generate $\{\mathbf{b}_n^{(i+1)}\}_{n=0}^{N-1}$ and $\{\beta_n^{(i+1)}\}_{n=0}^{N-1}$, from which the metrics $\{\mu(\mathbf{s}_n^{(i+1)})\}_{n=0}^{N-1}$ can be evaluated using (30), equations (21) and (23) suggest

$$\mathbf{b}_n^{(i+1)} = \Phi(\mathbf{s}^{(i)})^{-1} \mathbf{a}_n \quad (31)$$

$$= \left[\Phi(\mathbf{s}^{(i-1)})^{-1} - \sigma_1^2 \beta_{n_*}^{(i)} \mathbf{b}_{n_*}^{(i)} \mathbf{b}_{n_*}^{(i)H} \right] \mathbf{a}_n \quad (32)$$

$$= \mathbf{b}_n^{(i)} - \sigma_1^2 \beta_{n_*}^{(i)} \mathbf{b}_{n_*}^{(i)} \mathbf{b}_{n_*}^{(i)H} \mathbf{a}_n \quad (33)$$

$$\beta_n^{(i+1)} = \left(1 + \sigma_1^2 \mathbf{a}_n^H \mathbf{b}_n^{(i+1)} \right)^{-1}, \quad (34)$$

which requires $\mathcal{O}(NM)$ operations in total.

Going further, if we define $\mathbf{B}^{(i+1)} := [\mathbf{b}_0^{(i+1)}, \dots, \mathbf{b}_{N-1}^{(i+1)}]$ and notice that $\mathbf{B}^{(i+1)} = \Phi(\mathbf{s}^{(i)})^{-1} \mathbf{A}$, then we can compute the $\mathbf{s}^{(i)}$ -conditional MMSE estimate, and its covariance, as

$$\mathbf{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}^{(i)}\} = \mathbf{R}(\mathbf{s}^{(i)}) \mathbf{B}^{(i+1)H} \mathbf{y} \quad (35)$$

$$\text{Cov}\{\mathbf{x}|\mathbf{y}, \mathbf{s}^{(i)}\} = [\mathbf{I}_N - \mathbf{R}(\mathbf{s}^{(i)}) \mathbf{B}^{(i+1)H} \mathbf{A}] \mathbf{R}(\mathbf{s}^{(i)}), \quad (36)$$

using (15), (18), and the fact that $\Phi(\mathbf{s}^{(i)})$ is Hermitian. Because $\mathbf{R}(\mathbf{s}^{(i)}) \mathbf{B}^{(i+1)H}$ has only i nonzero rows and $\mathbf{A} \mathbf{R}(\mathbf{s}^{(i)})$ has only i nonzero columns, (35) and (36) can be computed using only $\mathcal{O}(iN)$ and $\mathcal{O}(i^2N)$ multiplications, respectively.

The recursive metric update (30), via recursive computation of $\{\mathbf{b}_n^{(i+1)}\}_{n=0}^{N-1}$ from (33), forms the foundation of the Fast Bayesian Matching Pursuit (FBMP) algorithm detailed in Table 1. From the table, it is straightforward to verify that the number of multiplications FBMP required to compute the set of dominant basis configurations $\hat{\mathcal{S}}_{\mathbf{x}}$, as well as the corresponding conditional means and covariances, is $\mathcal{O}(NMPD)$.

4. NUMERICAL EXPERIMENTS

4.1 FBMP Behavior

Numerical experiments were conducted to investigate the performance of FBMP from Table 1 for various values of model and algorithmic parameters, and the results are reported in Figs. 2-4. Unless otherwise noted, the experiments used $N = 256$, $M = 64$, $\text{SNR} = 15$ dB, $p_1 = 0.04$, and $P = \lceil \text{erfc}^{-1}(2P_0) \sqrt{2Np_1(1-p_1)} + Np_1 \rceil$ where $P_0 = 0.00005$ is the target value of $\Pr\{\|\mathbf{s}\|_0 > P\}$ as suggested in Section 3.3. Here we use $\text{SNR} := \frac{\sigma_1^2 p_1 N}{\sigma^2 M}$, as motivated by the unit-norm assumption on the columns of \mathbf{A} . The plots represent an average of 200 independent model realizations. For each realization of \mathbf{A} , an i.i.d. normal matrix was drawn and then scaled to make each of its columns unit-norm. Note that the average number of active coefficients $\mathbf{E}\{\|\mathbf{x}\|_0\} = p_1 N$ is approximately equal to 10 when $p_1 = 0.04$ and $N = 256$. When referring to the “normalized mean-squared error” (NMSE) of an estimate $\hat{\mathbf{x}}$, we mean $\mathbf{E}\{\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2 / \|\mathbf{x}\|_2^2\}$.

Sparse reconstruction with Gaussian ensembles asymptotically requires measurements numbering³³

$$M > 2K \log(N - K) + K + 1 \quad (37)$$

Here, we have $K = p_1 N$ is the (average) number of active coefficients. We note that the approximate inequality $M > 2K \log(N/K)$ is satisfied for $2p_1 N/M < -1/\log(p_1)$. Our nominal parameters yield $2p_1 N/M = 0.32$ and $-1/\log(p_1) = 0.31$, and below we see threshold behavior consistent with the asymptotic results.

In Fig. 2(a), we plot NMSE versus observation length M for FBMP under several values of the search parameter D . Recall that D effects a tradeoff between search accuracy and search complexity (the latter of which is expected to grow linearly in D). There we see that NMSE performance improves as M gets larger, i.e., as the average number of unknown parameters per measurement $\frac{p_1 N}{M}$ decreases. For $D = 1$ (i.e., the simplest possible search), Fig. 2(a) shows a “knee” in the curve at $M = 64$ (i.e., $\frac{p_1 N}{M} = 0.16$) below which NMSE degrades quickly. By increasing search complexity D , the knee shifts so that the FBMP is robust to a wider range of M (e.g., $M = 48$ or $\frac{p_1 N}{M} = 0.21$ when $D = 5$). The benefits of increased D diminish quickly for $D > 5$, however.

```

 $\mu^{(0,1)} = -N \ln 2\pi - M \ln \sigma^2 - \frac{1}{\sigma^2} \|\mathbf{y}\|_2^2 + N \ln(1 - p_1);$ 
 $\mathbf{n}^{(0,0)} = [];$ 
for  $n = 0 : N - 1,$ 
   $\mathbf{b}_n^{(1,1)} = \frac{1}{\sigma^2} \mathbf{a}_n;$ 
   $\beta_n = (1 + \sigma_1^2 \mathbf{a}_n^H \mathbf{b}_n^{(1,1)})^{-1};$ 
   $\mu_n = \mu^{(0,1)} + \log \beta_n + \sigma_1^2 \beta_n |\mathbf{y}^H \mathbf{b}_n^{(1,1)}|^2 + \log \frac{p_1}{1-p_1};$ 
end
for  $q = 1 : D,$ 
   $n_* = n$  indexing  $q^{\text{th}}$  largest element in  $\{\mu_n\};$ 
   $\mu^{(1,q)} = \mu_{n_*};$ 
   $\mathbf{n}^{(1,q)} = n_*;$ 
  for  $n = 0 : N - 1,$ 
     $\mathbf{b}_n^{(2,q)} = \mathbf{b}_n^{(1,1)} + \sigma_1^2 \beta_{n_*} \mathbf{b}_{n_*}^{(1,1)} \mathbf{b}_{n_*}^{(1,1)H} \mathbf{a}_n;$ 
  end
end
for  $i = 2 : P,$ 
  for  $d = 1 : D,$ 
    for  $n = 0 : N - 1,$ 
       $\beta_{d,n} = (1 + \sigma_1^2 \mathbf{a}_n^H \mathbf{b}_n^{(i,d)})^{-1};$ 
       $\mu_{d,n} = \mu^{(i-1,d)} + \log \beta_{d,n} + \sigma_1^2 \beta_{d,n} |\mathbf{y}^H \mathbf{b}_n^{(i,d)}|^2 + \log \frac{p_1}{1-p_1};$ 
      if  $n \in \mathbf{n}^{(i-1,d)}$  then  $\mu_{d,n} = -\infty;$ 
    end
  end
  for  $q = 1 : D,$ 
     $(d_*, n_*) = (d, n)$  indexing  $q^{\text{th}}$  largest unique element in  $\{\mu_{d,n}\};$ 
     $\mu^{(i,q)} = \mu_{d_*, n_*};$ 
     $\mathbf{n}^{(i,q)} = [\mathbf{n}^{(i,d_*)}, n_*];$ 
    for  $n = 0 : N - 1,$ 
       $\mathbf{b}_n^{(i+1,q)} = \mathbf{b}_n^{(i,d_*)} + \sigma_1^2 \beta_{d_*, n_*} \mathbf{b}_{n_*}^{(i,d_*)} \mathbf{b}_{n_*}^{(i,d_*)H} \mathbf{a}_n;$ 
    end
     $\hat{\mathbf{x}}^{(i,q)} = \sigma_1^2 \sum_{j=1}^i \delta_{[\mathbf{n}^{(i,q)}]_j} \mathbf{b}_{[\mathbf{n}^{(i,q)}]_j}^{(i+1,q)T} \mathbf{y};$ 
     $\hat{\Sigma}^{(i,q)} = \sigma_1^2 \sum_{j=1}^i \sum_{k=1}^i \delta_{[\mathbf{n}^{(i,q)}]_j} [\delta_{[\mathbf{n}^{(i,q)}]_j} - \delta_{[\mathbf{n}^{(i,q)}]_k}]$ 
       $-\sigma_1^2 \mathbf{b}_{[\mathbf{n}^{(i,q)}]_j}^{(i+1,q)T} \mathbf{a}_{[\mathbf{n}^{(i,q)}]_k} \delta_{[\mathbf{n}^{(i,q)}]_k}^T;$ 
  end
end
end

```

Table 1. Fast Bayesian Matching Pursuit

In Fig. 2(b), we plot the number of active basis elements missing from FBMP’s estimate of the MAP basis configuration:

$$\hat{\mathbf{s}}_{\text{map}} := \underset{\mathbf{s} \in \hat{\mathcal{S}}_*}{\operatorname{argmax}} p(\mathbf{s}|\mathbf{y}). \quad (38)$$

In particular, the traces in Fig. 2(b) show number-of-misses versus observation length M for FBMP under several values of search parameter D . Because the number-of-misses closely parallel the NMSE values, we conjecture that the sub-optimality of FBMP’s greedy search is to blame for the relatively large NMSE values that occur when $M < 64$ (i.e., when $\frac{p_1 N}{M} > 0.16$). In Fig. 3(a), we plot NMSE versus $p_1 N$, the expected number of active coefficients, for FBMP under several values of search parameter D . There we see that NMSE performance quickly degrades as $p_1 N$ increases above $p_1 N = 10$ (i.e., above $\frac{p_1 N}{M} = 0.16$), mirroring the results in Fig. 2. As in Fig. 2, when $\frac{p_1 N}{M} > 0.16$, increasing D from 1 to 10 can yield an NMSE improvement of 3 dB. When $\frac{p_1 N}{M} \leq 0.16$, however, $D = 1$ appears to suffice. In Fig. 3(b), we plot NMSE versus SNR for FBMP under several values of search parameter D (where $\{M, p_1\}$ correspond to the aforementioned breakpoints in the NMSE-vs- M and NMSE-vs- $p_1 N$ curves). Fig. 3(b) shows a satisfying linear relationship between NMSE and SNR (in dB). As expected, the effect of increasing D from 1 to 10 is negligible because $\frac{p_1 N}{M} = 0.16$; a more significant effect would be expected if $\frac{p_1 N}{M}$ had been larger.

In Fig. 4(a), we plot NMSE versus SNR for two FBMP-derived estimates: the (approximate) MMSE estimate $\hat{\mathbf{x}}_{\text{ammse}}$ from (17) and the quasi-MAP estimate $\hat{\mathbf{x}}_{\text{amap}}$ from (39):

$$\hat{\mathbf{x}}_{\text{amap}} := \operatorname{E}\{\mathbf{x}|\mathbf{y}, \hat{\mathbf{s}}_{\text{map}}\}. \quad (39)$$

Whereas $\hat{\mathbf{x}}_{\text{ammse}}$ is the *average* of the conditional MMSE estimates $\operatorname{E}\{\mathbf{x}|\mathbf{y}, \mathbf{s}\}$ over $\mathbf{s} \in \hat{\mathcal{S}}_*$, the estimate $\hat{\mathbf{x}}_{\text{amap}}$ is MMSE conditioned on (FBMP’s estimate of) the MAP basis-configuration $\hat{\mathbf{s}}_{\text{map}}$. In terms of average NMSE, Fig. 4(a) demonstrates that $\hat{\mathbf{x}}_{\text{ammse}}$ are about 1 dB better than $\hat{\mathbf{x}}_{\text{amap}}$ at SNR ≤ 10 dB and about 0.5 dB better at SNR > 10 dB. The improvement reflects the advantage of allowing for *uncertainty* in the estimated basis.

Finally, in Fig. 4(b), we plot average FBMP runtime versus search parameter D . As expected from the algorithmic description in Table 1, the runtime scales linearly in D .

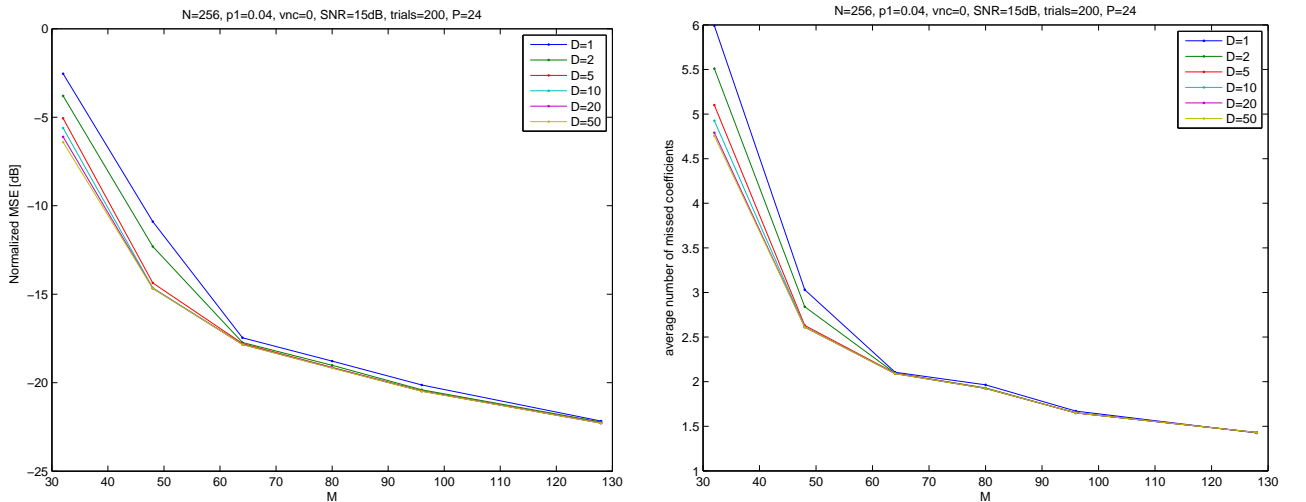


Figure 2. (a) Normalized MSE for FBMP versus observation length M . (b) Number of active coefficients in \mathbf{x} missing from FBMP’s estimate of the most probable basis configuration versus observation length M . In each subfigure, from top to bottom, the traces reflect increasing D . (See the graph title for configuration.)




Figure 3. (a) Normalized MSE versus average number of active coefficients (i.e., $p_1 N$) for FBMP. (b) Normalized MSE versus SNR for FBMP. In each subfigure, from top to bottom, the traces reflect increasing D . (See the graph title for configuration.)




Figure 4. (a) Normalized MSE versus SNR of FBMP-returned MMSE estimate $\hat{\mathbf{x}}_{\text{ammse}}$ and quasi-MAP estimate $\hat{\mathbf{x}}_{\text{amap}}$. (b) Average FBMP runtime versus search parameter D . (See the graph title for configuration.)

4.2 Comparison To Other Algorithms

In Figs. 5-7 we compare FBMP to several other popular sparse estimation algorithms, including SparseBayes,¹⁴ OMP,⁹ StOMP,³⁴ GPSR-Basic,³⁵ and BCS.³⁶ Unless otherwise noted, the model parameters were set at $N = 512$, $M = 128$, $p_1 = 0.04$, and $\sigma^2 = 0.001$ (which corresponds to SNR = 19 dB at the nominal values of N , M , and p_1). Our plots represent an average of 100 independent model realizations.

For FBMP, we used *non*-optimized MATLAB code (which we plan to optimize in the near future), and unless otherwise noted used $D = 5$ and the same P specified in Section 4.1. For the other algorithms, we used the publicly available implementations that were found at the web-sites listed in our bibliography. The algorithmic parameters were chosen largely in accordance with suggested values provided by the authors of the software, or in accordance with values used in examples that accompanied the algorithms. The SparseBayes algorithm was tested with the initial hyper-parameter set to $\alpha = 1$. StOMP was tested using the “False Alarm Control”

thresholding strategy, with the thresholding parameter set to $\frac{M}{NQ}(1 - \frac{1}{M}\|\mathbf{x}\|_0)$, where the default number of iterations, $Q = 10$, was used. The ℓ_1 -penalty in the GPSR algorithm was chosen as $\tau = 0.1\|\mathbf{A}^H\mathbf{y}\|_\infty$, and the MSE kept for comparison purposes was the smaller of the MSEs of the un-debiased and debiased estimates. The BCS algorithm was tested with the ‘‘Adaptive CS’’ option turned off.

In Fig. 5 we plot NMSE versus observation length M (at $\sigma^2 = 0.001$) for the various sparse estimation algorithms. There we see that FBMP achieved significantly lower NMSE than the other algorithms for $M > 64$. In particular, for $M > 85$ it outperformed BCS by approximately 4 dB and outperformed OMP by 6 dB. In Fig. 6 we plot NMSE versus SNR (at $M = 128$) for the various sparse estimation algorithms. Again, the NMSEs achieved by FBMP were significantly lower than those achieved by the other algorithms. At SNR = 22 dB, FBMP outperformed BCS by approximately 3 dB and the other algorithms by > 9 dB; at SNR = 15 dB, FBMP outperformed all other algorithms by > 5 dB; and, at SNR = 3 dB, FBMP outperformed GPSR by approximately 1 dB and the other algorithms by > 5 dB.

Finally, in Fig. 7, we plot average runtime versus observation length M for the various sparse estimation algorithms. For FBMP, we used $D = 1$. Fig. 7 shows that FBMP is about an order of magnitude faster than SparseBayes, on the same order of complexity as BCS, and about an order of magnitude slower than OMP, StOMP, and GPSR. We anticipate that optimized FBMP code will yield improved runtimes.




Figure 5. Normalized MSE versus observation length M for several algorithms. (See the graph title for configuration.)

5. DISCUSSION

The Bayesian framework provides a report on the confidence of estimates of both the coefficients \mathbf{x} and the basis configuration \mathbf{s} . In particular, the basis selection metric $\mu(\mathbf{s})$ provides a posterior confidence label for a candidate basis configuration \mathbf{s} , in addition to providing the MMSE estimate $\hat{\mathbf{x}}_{\text{mmse}}$ through (15). Specifically, from (10), we can write the posterior probability of basis configuration \mathbf{s} as

$$p(\mathbf{s}|\mathbf{y}) = \frac{\exp\{\mu(\mathbf{s})\}}{\sum_{\mathbf{s}' \in \mathcal{S}} \exp\{\mu(\mathbf{s}')\}} \approx \frac{\exp\{\mu(\mathbf{s})\}}{\sum_{\mathbf{s}' \in \mathcal{S}_*} \exp\{\mu(\mathbf{s}')\}}, \quad (40)$$

where the approximation in (40) includes only the basis configurations $\mathcal{S}_* \subset \mathcal{S}$ that account for the dominant values of $\exp\{\mu(\mathbf{s})\}$. Likewise, $p(\mathbf{x}|\mathbf{y})$ provides an approximate density function describing the uncertainty in




Figure 6. Normalized MSE versus SNR for several algorithms. (See the graph title for configuration.)




Figure 7. Average runtime versus observation length M for several algorithms. (See the graph title for configuration.)

resolving \mathbf{x} with the noisy observations, \mathbf{y}

$$p(\mathbf{x}|\mathbf{y}) \approx \sum_{\mathbf{s}' \in \mathcal{S}_*} p(\mathbf{s}'|\mathbf{y})p(\mathbf{x}|\mathbf{y}, \mathbf{s}') \quad (41)$$

The posterior density function reflects the multi-modal ambiguity inherently present in the sparse inference problem—an ambiguity especially evident when the SNR is low and/or the correlation among the columns of \mathbf{A} is high.

In contrast, confidence labels for estimated $\hat{\mathbf{x}}$ are largely absent in the compressive sensing and sparse linear regression literature. Exceptions are found in^{19,36} which give the error covariance for the simple linear problem conditioned on *perfect knowledge of the active basis elements*. As noted by Tibshirani,¹⁹ such a measure of posterior uncertainty has dubious value, because “a difficulty with this formula is that it gives an estimated variance of 0 for predictors with” $s_i = 0$. The Gaussian confidence intervals in^{19,36} are constructed by assuming that $\hat{\mathbf{s}}_{\text{map}}$ is the true sequence of active coefficients. However, in our simulations the true \mathbf{s} typically does not rank first in $p(\mathbf{s}|\mathbf{y})$, and instead usually ranks between second and tenth (or lower). In this light, we expect certain advantages for algorithms that consider the active basis as implicitly uncertain.

A Gaussian mixture model similar to that in Section 2 was likewise adopted by Larsson and Selén,³¹ who also constructed the MMSE estimate in the manner of (17) but with an \mathcal{S}_* that contains exactly one sequence \mathbf{s} for each Hamming weight 0 to N . They proposed to find these \mathbf{s} via greedy deflation, i.e., starting with an all-active basis configuration and recursively deactivating one element at a time. Thus, the $D = 1$ version of the BMP heuristic from Section 3.3 recalls the heuristic of,³¹ but in reverse. Note, however, that the *fast* $D = 1$ BMP presented in Section 3.5 has a complexity of only $\mathcal{O}(NMP)$, in comparison to $\mathcal{O}(N^3M^2)$ for the technique in.³¹ Given the typically large values of N encountered in practice, the complexity of FBMP can be several orders of magnitude lower than that of.³¹

As a caveat, we should emphasize that our greedy FBMP search returns only $\hat{\mathcal{S}}_*$, an *estimate* of the dominant subset \mathcal{S}_* , along with the values of $\mu(\mathbf{s})$ for $\mathbf{s} \in \hat{\mathcal{S}}_*$. Thus, while the values $\mu(\mathbf{s})$ returned by FBMP can be used to compute exact *ratios* between the posterior probabilities of the configurations in $\hat{\mathcal{S}}_*$, the *absolute* posteriors of these configurations (as approximated by (40) with $\hat{\mathcal{S}}_*$ in place of \mathcal{S}_*) will only be accurate when $\hat{\mathcal{S}}_*$ indeed contains \mathcal{S}_* . For example, if FBMP somehow missed the MAP configuration $\hat{\mathbf{s}}_{\text{map}}$ (i.e., $\hat{\mathbf{s}}_{\text{map}} \notin \hat{\mathcal{S}}_*$), then we would expect a large discrepancy between $\sum_{\mathbf{s}' \in \mathcal{S}_*} \exp\{\mu(\mathbf{s}')\}$ and $\sum_{\mathbf{s}' \in \hat{\mathcal{S}}_*} \exp\{\mu(\mathbf{s}')\}$ which would in turn corrupt the FBMP estimates of $p(\mathbf{s}|\mathbf{y})$ and $\text{Cov}\{\mathbf{x}|\mathbf{y}\}$. Fortunately, the proposed greedy FBMP basis search seems to perform quite well, as least for $\frac{p_1 N}{M} \leq -1/\log(p_1)$.

Although the model in Section 2 assumed that each x_i is generated according to a binary mixture of zero-mean Gaussians, one can imagine extending the model to, e.g., a mixture of finitely many Gaussians with non-zero means. In this case, one would need to generalize the BMP search heuristic of Section 3.3 to handle several types of active coefficient (e.g., one for each allowed mean).

6. CONCLUSION

At high frequencies, radar imaging can be viewed as sparse reconstruction from noisy observations on a sparse aperture. On the one hand, nonlinear regression techniques can provide super-resolution and posterior confidence labels, but such techniques may suffer from model selection, model mismatch, and computational cost. On the other hand, sparse linear regression techniques are computationally inexpensive and without model selection difficulties; however, the linear techniques provide neither super-resolution of high SNR reflectors nor posterior confidence labels.

In this paper, we proposed an algorithm for joint basis selection and sparse parameter estimation which we call fast Bayesian matching pursuit (FBMP). In brief, FBMP models each unknown coefficient x_i as either inactive or active (with prior probability p_1), where an i.i.d. Gaussian distribution (with zero mean and variance σ_1^2) is assigned to the values of active coefficients. The observation \mathbf{y} is then modeled as an AWGN-corrupted version of the unknown coefficients that has been mixed by a known matrix \mathbf{A} . FBMP navigates through the tree of active/inactive configurations \mathcal{S} with the goal of finding the configurations with dominant posterior probability, \mathcal{S}_* . The search is controlled by a parameter D which effects a tradeoff between complexity and accuracy. Numerical experiments suggest that the estimates returned by FBMP outperform (in normalized MSE) those of other popular algorithms (e.g., SparseBayes, OMP, StOMP, GPSR-Basic, BCS) by several dB in typical situations.

Thus, we presented a detection procedure that provides exact ratios of posterior probabilities for a set of high probability solutions to the sparse reconstruction problem. These relative probabilities serve to reveal potential ambiguity among multiple candidate solutions that are ambiguous due to low signal-to-noise ratio and/or significant correlation among columns in the (super-resolved) regressor matrix. The explicit management

of ambiguity and uncertainty is essential for applications in which the estimated detections, $\hat{\mathbf{s}}$, and estimated coefficients, $\hat{\mathbf{x}}$, are not final products, but rather are statistics for use in making inference from the noisy observations, \mathbf{y} .

We plan, in the near future, to extend FBMP to the case where the active coefficients have non-zero means chosen from a finite set according to some prior probabilities. Additionally, we are developing a large-scale version of FBMP in which $M \times M$ matrices are not explicitly maintained and the operator \mathbf{A} is only given implicitly. An example of the non-zero mean sparse signal model can be found in electron paramagnetic resonance imaging,³⁷ where micro-liter particulate probes are inserted into a tumor and fill less than 0.25% volume in the field of view. The fabrication of the paramagnetic signal probes results in variable shape, size and electron spin density, giving rise to a non-zero-mean and nearly Gaussian distribution of signal strength in a very few active voxels.

ACKNOWLEDGMENTS

LCP was supported by AFOSR award FA9550-06-1-0324; PS was supported by National Science Foundation Grant 0237037 and the Office of Naval Research grant N00014-07-1-0209. Cited work by S. Sharma was supported by the SET Corporation. A preliminary version of Fast Bayes Matched Pursuits was presented at the Workshop on Information Theory and Applications (ITA), La Jolla, CA, January 2008.

REFERENCES

- [1] Keller, J. B., “Geometrical theory of diffraction,” *J. Opt. Soc. Amer.*, 116–130 (1962).
- [2] Gerry, M. J., Potter, L. C., Gupta, I. J., and van der Merwe, A., “A parametric model for synthetic aperture radar measurements,” *IEEE Transactions on Antennas and Propagation* **47**, 1179–1188 (1999).
- [3] Jackson, J. A., Rigling, B. D., and Moses, R. L., “Canonical scattering feature models for 3D and bistatic SAR,” *IEEE Trans. Aerospace and Electronic Systems* (in review).
- [4] Casteel, Jr., C. H., Gorham, L. A., Minardi, M. J., Scarborough, S. M., Naidu, K. D., and Majumder, U. K., “A challenge problem for 2D/3D imaging of targets from a volumetric data set in an urban environment,” in *Algorithms for Synthetic Aperture Radar Imagery XIV*, Zelnio, E. G. and Garber, F. D., eds., *Proc. SPIE* **6568**, 65680D (2007).
- [5] Sharma, S., *Three Dimensional Parameter Estimation from Sparse, Multipass Synthetic Aperture Radar*, Master’s thesis, The Ohio State University (2007).
- [6] Hogbom, J., “Aperture synthesis with a non-regular distribution of interferometer baselines,” *Astrophys. J. Suppl. Ser* **15**, 417–426 (1974).
- [7] Huber, P. J., “Projection pursuit,” *The Annals of Statistics* **13**, 435–475 (1985).
- [8] Pati, Y. C., Rezaiifar, R., and Krishnaprasad, P. S., “Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition,” in *Proc. 27th Ann. Asilomar Conf. Signals, Systems, and Computers*, (1993).
- [9] Tropp, J. and Gilbert, A., “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Trans. Information Theory* **53**, 4655–4666 (2007). (software available at <http://sparselab.stanford.edu/>).
- [10] Lawson, C. L., “Contributions to the theory of linear least maximum approximations,” PhD dissertation, UCLA (1961).
- [11] Lee, H., Sullivan, D., and Huang, T., “Improvement of discrete band-limited signal extrapolation by iterative subspace modification,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 1569 – 1572 (1987).
- [12] Chartrand and Yin, ICASSP 2008, Las Vegas, NV, to appear.
- [13] R. DeVore, “Performance of iterated least squares,” IMA Seminar, June 2007.
- [14] Tipping, M. E., “Sparse Bayesian learning and the relevance vector machine,” *J. Machine Learning Res.* **1**, 211–244 (2001). (software available at <http://www.miketipping.com/index.php?page=rvm>).
- [15] Wipf, D. and Rao, B., “Sparse Bayesian learning for basis selection,” *IEEE Trans. Signal Process.* **52**, 2153 – 2164 (2004).

- [16] Taylor, H. L., Banks, S. C., and McCoy, J. F., “Deconvolution with the ℓ_1 norm,” *Geophysics* **44**, 39–52 (1979).
- [17] Hoerl, A. E. and Kennard, R. W., “Ridge regression: Biased estimation for nonorthogonal problems,” *Technometrics* **12**, 55–67 (1970).
- [18] Chen, S. S., Donoho, D. L., and Saunders, M. A., “Atomic decomposition by basis pursuit,” *SIAM Journal on Scientific Computing* **20**(1), 33–61 (1998).
- [19] Tibshirani, R., “Regression shrinkage and selection via the lasso,” *J. R. Statist. Soc. B* **58**(1), 267 – 288 (1996).
- [20] Çetin, M. and Karl, W. C., “Feature-enhanced synthetic aperture radar image formation based on non-quadratic regularization,” *IEEE Trans. Image Process.* **10**, 623–631 (2001).
- [21] Bouman, C. and Sauer, K., “A generalized Gaussian image model for edge-preserving MAP estimation,” *IEEE Trans. Image Process.* **2**, 296–310 (1993).
- [22] Delaney, A. H. and Bresler, Y., “A fast and accurate Fourier algorithm for iterative parallel-beam tomography,” *IEEE Trans. Image Process.* **5**, 740–753 (May 1996).
- [23] Levy, S. and Fullagar, P. K., “Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution,” *Geophysics* **46**(9), 1235–1243 (1981).
- [24] Rao, B. D. and Kreutz-Delgado, K., “An affine scaling methodology for best basis selection,” *IEEE Trans. Signal Processing* **47**, 187–200 (1999).
- [25] Donoho, D. L., Elad, M., and Temlyakov, V. N., “Stable recovery of sparse overcomplete representations in the presence of noise,” *IEEE Trans. Information Theory* **52**(1), 6–18 (2006).
- [26] Tropp, J. A., “Just relax: Convex programming methods for identifying sparse signal,” *IEEE Trans. Info. Theory* **51**(3), 1030–1051 (2006).
- [27] Candès, E., Romberg, J., and Tao, T., “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics* **59**(8), 1207–1223 (2006).
- [28] Tao, T., “Open question: deterministic UUP matrices,” <http://terrytao.wordpress.com/2007/07/02/open-question-deterministic-uup-matrices> (July 2007).
- [29] Lo, Y. T., “A mathematical theory of antenna arrays with randomly spaced elements,” *IEEE Transactions on Antennas and Propagation* **12**, 257–268 (1964).
- [30] Marseille, G. J., de Beer, R., Fuderer, M., Mehlkopf, A. F., and van Ormondt, D., “Nonuniform phase-encode distributions for MRI scan time reduction,” *J. Magn. Reson.* **111**, 70–75 (1996).
- [31] Larsson, E. and Selén, Y., “Linear regression with a sparse parameter vector,” *IEEE Trans. Signal Process.* **55**, 451 – 460 (2007).
- [32] Cotter, S. F. and Rao, B. D., “Application of tree-based searches to matching pursuit,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 3933–3936, Salt Lake City, UT (2001).
- [33] Wainwright, M. J., “Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming,” Tech. Rep. 709, Dept. of Statistics, UC-Berkeley (May 2006).
- [34] Donoho, D. L., Tsaig, Y., Drori, I., and Starck, J.-C., “Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit,” Tech. Rep. 2006-02, Dept. of Statistics, Stanford University, Stanford, CA (2006). (software available at <http://sparselab.stanford.edu/>).
- [35] Figueiredo, M. A. T., Nowak, R. D., and Wright, S. J., “Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems,” *IEEE Journal of Selected Topics in Signal Processing* **1**(4), 586–597 (2007). (software available at <http://www.lx.it.pt/~mtf/GPSR/>).
- [36] Ji, S. and Carin, L., “Bayesian compressive sensing and projection optimization,” in *Proc. 24th Intern. Conf. Machine Learning (ICML)*, (2007). (software available at <http://www.ece.duke.edu/~shji/BCS.html>).
- [37] Kuppusamy, P., “EPR spectroscopy in biology and medicine,” *Antioxid. Redox Signal* **6**, 583–585 (2004).