

Multi-frequency sparse Bayesian learning for robust matched field processing

Kay L. Gemba,^{a)} Santosh Nannuru, Peter Gerstoft, and William S. Hodgkiss
*Marine Physical Laboratory of the Scripps Institution of Oceanography, University of California
at San Diego, La Jolla, California 92093-0238, USA*

(Received 9 February 2017; revised 27 April 2017; accepted 29 April 2017; published online 19 May 2017)

The multi-snapshot, multi-frequency sparse Bayesian learning (SBL) processor is derived and its performance compared to the Bartlett, minimum variance distortionless response, and white noise constraint processors for the matched field processing application. The two-source model and data scenario of interest includes realistic mismatch implemented in the form of array tilt and data snapshots not exactly corresponding to the range-depth grid of the replica vectors. Results demonstrate that SBL behaves similar to an adaptive processor when localizing a weaker source in the presence of a stronger source, is robust to mismatch, and exhibits improved localization performance when compared to the other processors. Unlike the basis or matching pursuit methods, SBL automatically determines sparsity and its solution can be interpreted as an ambiguity surface. Because of its computational efficiency and performance, SBL is practical for applications requiring adaptive and robust processing. © 2017 Acoustical Society of America.

[<http://dx.doi.org/10.1121/1.4983467>]

[ZHM]

Pages: 3411–3420

I. INTRODUCTION

Matched field processing (MFP) is a generalized beamforming method which matches received array data to a dictionary of replica vectors to localize one or several sources.¹ Traditionally, the Bartlett processor² is used as a point of departure to estimate source location parameters. This localization problem can be reformulated as an underdetermined system of linear equations which might be solved using algorithms such as orthogonal matching pursuit³ or with an optimization known as basis pursuit⁴ (BP). While a particular algorithm might be selected or even tailored to meet the need of a problem,⁵ estimating a sparse number of source locations and amplitudes recorded with few sensors from many more candidate Green's functions is coined compressive sensing.^{6,7}

Compressive sensing (CS) implemented with the BP method outperforms the high-resolution minimum variance distortionless response⁸ (MVDR) or MUSIC (Ref. 9) methods, discriminating between multiple, coherent plane-wave arrivals for the beamforming application.¹⁰ CS possesses properties similar to an adaptive processor and offers modest localization improvement when compared to the white noise constraint¹¹ (WNC) processor for the MFP application in single- and multiple-source scenarios.¹² Note that even in single source scenarios, adaptive processing may be desirable to improve localization performance over multiple cycles (generated by the interference pattern of adjacent modes).

However, CS has a shortcoming: for S unknown sources present, the number of sparse solutions K required to localize all sources is $K \geq S$.¹² Additional ambiguous solutions are

due to the presence of (environmental) mismatch, defined as a misalignment between the actual source field observed at the array and the modeled replica vector. Furthermore, several sparse solutions might correspond to a single source if it does not remain in a single range-depth cell for all processed data snapshots. This problem might also be encountered in realistic scenarios when jointly processing multiple frequencies and multiple snapshots in order to localize one or more sources.

An alternative CS implementation, known as sparse Bayesian learning (SBL),¹³ offers relief to this shortcoming. SBL has the advantage that it determines sparsity automatically without any user input. Following the CS approach, SBL also reformulates the parameter estimation problem as an underdetermined linear problem. The variables in the problem are treated as Gaussian random vectors and evidence maximization is performed using Bayesian analysis to obtain a sparse solution,¹⁴ e.g., demonstrated for plane wave beamforming.^{15,16} It has been shown that SBL can be interpreted as an iterative, re-weighted BP algorithm¹⁷ and that SBL exhibits similar sparse signal recovery when compared to the BP method.^{15,18} BP performance in the presence of perturbations or mismatch has been investigated.^{12,19,20}

This paper compares processor localization performance using simulated and SWellEx-96 data subject to mismatch. In particular, we investigate processor performance in the presence of array tilt and when data snapshots are distributed uniformly between range-depth grid points. SBL performance is compared to the non-adaptive Bartlett and adaptive WNC and MVDR processors. Section II introduces the processors followed by a description of data and simulation processing in Sec. III. Results are presented in Sec. IV followed by discussion in Sec. V and conclusions in Sec. VI.

^{a)}Electronic mail: gemba@ucsd.edu

II. PROCESSORS

A. Bartlett, WNC, and MVDR processors

Bartlett is a spatial matched-filter processor which matches replica vectors $\mathbf{a}(\boldsymbol{\theta})$ (corresponding to the complex wavefield of a source at frequency f and position $\boldsymbol{\theta}$ received at an array of N elements) to the data \mathbf{y} ,

$$P_B(\boldsymbol{\theta}) = \mathbf{a}^H(\boldsymbol{\theta})\mathbf{K}\mathbf{a}(\boldsymbol{\theta}). \quad (1)$$

The superscript H denotes the Hermitian operator and $P_B(\boldsymbol{\theta})$ denotes the Bartlett power at position $\boldsymbol{\theta}$ (i.e., range and depth) using normalized replicas vectors [i.e., $\|\mathbf{a}(\boldsymbol{\theta})\|_2 = 1$]. The sample covariance matrix (SCM) $\mathbf{K} \in \mathbb{C}^{N \times N}$ is obtained using L snapshots,

$$\mathbf{K} = \frac{1}{L} \sum_{l=1}^L \mathbf{y}_l \mathbf{y}_l^H. \quad (2)$$

The snapshot $\mathbf{y}_l \in \mathbb{C}^N$ consists of a vector of Fourier coefficients at a single frequency f obtained via a fast Fourier transform (FFT) of the l th data segment from each of the array elements. While Bartlett does not invert \mathbf{K} and thus does not have a minimum number of required snapshots, it suffers from high sidelobes. Sidelobe suppression is important if several sources (or a combination of sources and interferers) are present.

The WNC processor P_{wnc} discriminates against other sources/interferers while offering a degree of robustness in frequently encountered replica-data mismatch scenarios. The WNC is versatile because of its ability to adjust its behavior (thus resolution and sidelobe suppression) from Bartlett to MVDR (Ref. 21) at the expense of inverting \mathbf{K} . To have \mathbf{K} invertible, we require $L \geq N$ (diagonal loading of \mathbf{K} can be used to mitigate this requirement). The WNC processor is given by

$$P_{wnc}(\boldsymbol{\theta}) = \mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{K}\mathbf{a}_w(\boldsymbol{\theta}), \quad (3)$$

where

$$\mathbf{a}_w(\boldsymbol{\theta}) = \frac{(\mathbf{K} + \epsilon\mathbf{I})^{-1}\mathbf{a}(\boldsymbol{\theta})}{\mathbf{a}^H(\boldsymbol{\theta})(\mathbf{K} + \epsilon\mathbf{I})^{-1}\mathbf{a}(\boldsymbol{\theta})}.$$

The adaptive weights $\mathbf{a}_w(\boldsymbol{\theta})$ correspond to diagonally loaded MVDR weights and are obtained by solving

$$\begin{aligned} & \min_{\mathbf{a}_w} \mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{K}\mathbf{a}_w(\boldsymbol{\theta}) \\ & \text{subject to } \mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{a}(\boldsymbol{\theta}) = 1, \\ & |\mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{a}_w(\boldsymbol{\theta})|^{-1} \geq \delta^2, \end{aligned} \quad (4)$$

for each replica vector at position $\boldsymbol{\theta}$. The constraining value δ^2 imposes a gain constraint on the adaptive weights and the white noise gain constraint G_{wng} such that

$$\delta^2 \leq G_{wng} = |\mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{a}_w(\boldsymbol{\theta})|^{-1} < N, \quad (5)$$

which in practice is normalized and expressed as $10 \log_{10}(\delta^2/N) \leq 0$ dB.

The data used to construct the SCM is not normalized, which requires robust selection of an initialization value ϵ_0 for this barrier optimization problem such that $|\mathbf{a}_w^H(\boldsymbol{\theta})\mathbf{a}_w(\boldsymbol{\theta})|_{\epsilon_0}^{-1} < \delta^2$. Since $\epsilon(\boldsymbol{\theta})$ can span many orders of magnitude, ϵ is parameterized using the decibel scale. The search routine is implemented in terms of a singular value decomposition of \mathbf{K} with $\epsilon_0 = 10 \log_{10}(\text{Tr}(\mathbf{K})/N) - 30$. Tr denotes the trace of a matrix. The iterative algorithm converges when a selected constraint is satisfied within ± 0.1 dB. Thus, $P_{wnc}(\boldsymbol{\theta})$ denotes the WNC power at position $\boldsymbol{\theta}$ for a selected (white noise gain) constraint. The constraint frequently falls within $[-6, -2]$ dB²⁸ and we use a constraint of -3 dB for processing.

The WNC solution approaches the MVDR as $\epsilon \rightarrow 0$. In practice, it is common to load the diagonal of the SCM when using the MVDR. We approximate the MVDR using a white noise gain constraint of -20 dB and denote the loaded MVDR processor by $P_{MV}(\boldsymbol{\theta})$. Selecting such a low constraint essentially bypasses the optimization in Eq. (4) for most candidate positions $\boldsymbol{\theta}$.

To localize a source, Eqs. (1) and (3) are evaluated at M range-depth positions or cells $\boldsymbol{\theta}$. Computed ambiguity surfaces for F processed frequencies are averaged,

$$P^F(\boldsymbol{\theta}) = \sum_{f=1}^F P(\boldsymbol{\theta}, f). \quad (6)$$

Processing additional frequencies improves source localization performance for a weaker source in the presence of a stronger source and environmental or model mismatch.

B. Sparse Bayesian learning

In this section we summarize the multi-frequency SBL algorithm, discuss its estimate of the noise, and present a pseudocode for its implementation.

1. SBL formulation

The l th data snapshot \mathbf{y}_l can be expressed with an under-determined system of linear equations

$$\mathbf{y}_l = \mathbf{A}\mathbf{x}_l + \mathbf{n}_l. \quad (7)$$

In Eq. (7), $\mathbf{A} = [\mathbf{a}(\boldsymbol{\theta}_1), \dots, \mathbf{a}(\boldsymbol{\theta}_M)]$ is the dictionary containing M replica vectors and $\mathbf{n}_l \sim \mathcal{CN}(\mathbf{n}_l|0, \sigma^2\mathbf{I}_N)$ is complex Gaussian noise. \mathbf{x}_l is the vector of complex source amplitudes with entries corresponding to the same range-depth cells as in $\boldsymbol{\theta}$. SBL models the unknown source amplitudes as complex Gaussian random variables with prior density $p(\mathbf{x}_l) = \mathcal{CN}(\mathbf{x}_l|0, \boldsymbol{\Gamma})$, where $\boldsymbol{\Gamma}$ is a diagonal covariance matrix, i.e., $\boldsymbol{\Gamma} = \text{diag}(\gamma_1, \dots, \gamma_M) = \text{diag}(\boldsymbol{\gamma})$. The vector $\boldsymbol{\gamma}$ is the source power in each range-depth cell $\boldsymbol{\theta}$. Since the noise is Gaussian, the likelihood is expressed as $p(\mathbf{y}_l|\mathbf{x}_l; \mathbf{A}) = \mathcal{CN}(\mathbf{y}_l|\mathbf{A}\mathbf{x}_l, \sigma^2\mathbf{I}_N)$.

Next we extend the single-frequency SBL approach^{15,22} to include multiple frequencies.^{16,23,24} Denote the collection of L snapshots at the f th frequency as $\mathbf{Y}_f = [\mathbf{y}_1, \dots, \mathbf{y}_L]$. Let the corresponding collection of source amplitude vectors and dictionaries be denoted by \mathbf{X}_f and \mathbf{A}_f , respectively. Then

$$\mathbf{Y}_f = \mathbf{A}_f \mathbf{X}_f + \mathbf{N}_f, \quad f = 1, 2, \dots, F, \quad (8)$$

where \mathbf{N}_f are additive noise contributions. For approximately stationary sources, $\mathbf{x}_{f,l}$ are assumed independent over time. Hence, we have

$$p(\mathbf{X}_f) = \prod_{l=1}^L \mathcal{CN}(\mathbf{x}_{f,l} | 0, \mathbf{\Gamma}_f), \quad f = 1, 2, \dots, F, \quad (9)$$

where $\mathbf{\Gamma}_f = \text{diag}(\gamma_f)$ is the covariance of the source amplitudes at frequency f . We assume that \mathbf{X}_f are independent for F processed frequencies. There is no assumption of sparsity in the frequency domain, which makes this formulation attractive for localizing a few broadband sources from many candidate replica vectors. Because vectors contained in $\mathbf{X}_1, \dots, \mathbf{X}_F$ and $\mathbf{N}_1, \dots, \mathbf{N}_F$ are independent, the joint evidence $p(\mathbf{Y}_1, \dots, \mathbf{Y}_F)$ over all frequencies is

$$p(\mathbf{Y}_1, \dots, \mathbf{Y}_F) = \prod_{f=1}^F p(\mathbf{Y}_f) = \prod_{f=1}^F \prod_{l=1}^L \mathcal{CN}(\mathbf{y}_{f,l} | 0, \mathbf{\Sigma}_{y_f}), \quad (10)$$

where the model covariance $\mathbf{\Sigma}_{y_f} = \sigma_f^2 \mathbf{I} + \mathbf{A}_f \mathbf{\Gamma}_f \mathbf{A}_f^H$. To estimate γ_f (denoted by $\hat{\gamma}_f$), we maximize the joint evidence

$$\begin{aligned} \{\hat{\gamma}_1 \cdots \hat{\gamma}_F\} &= \arg \max_{\{\gamma_1 \cdots \gamma_F\}} p(\mathbf{Y}_1, \dots, \mathbf{Y}_F) \\ &= \arg \min_{\{\gamma_1 \cdots \gamma_F\}} \left\{ \sum_{f=1}^F L \log |\mathbf{\Sigma}_{y_f}| + \text{Tr}(\mathbf{Y}_f^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{Y}_f) \right\}, \end{aligned} \quad (11)$$

where $|\cdot|$ denotes the determinant of a matrix. To obtain a minimum of this objective function, we equate its derivatives to zero,

$$\frac{\partial}{\partial \gamma_{f,m}} \left\{ \sum_{f=1}^F L \log |\mathbf{\Sigma}_{y_f}| + \text{Tr}(\mathbf{Y}_f^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{Y}_f) \right\} = 0. \quad (12)$$

This yields the fixed point update rule^{23,24}

$$\hat{\gamma}_{f,m}^{\text{new}} = \hat{\gamma}_{f,m}^{\text{old}} \frac{\|\mathbf{Y}_f^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{a}_{f,m}\|_2^2}{L \mathbf{a}_{f,m}^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{a}_{f,m}}. \quad (13)$$

The optimization is performed by iterating through the update rule in Eq. (13), which converges in practice. At each iteration, an estimate of the noise variance is required (see Sec. II B 2). The K sparse entries in $\hat{\gamma}_f$ correspond to an estimate of the powers of K sources.

We can form the multi-frequency estimate as

$$\hat{\gamma} = \sum_{f=1}^F \hat{\gamma}_f. \quad (14)$$

Similar to Eq. (6), summing F processed ambiguity surfaces with SBL [Eq. (14)] may improve localization performance if the sparsity across frequencies is the same. Since $\hat{\gamma}_f$ is the

source variance estimate at frequency f , a plot of $\hat{\gamma}$ can be interpreted as a broadband ambiguity surface.

An alternate way to *enhance* sparsity of the solution is to set $\mathbf{\Gamma} = \mathbf{\Gamma}_1 = \cdots = \mathbf{\Gamma}_F$. This assumes that \mathbf{X}_f has the same statistical distribution at each frequency. A sparse $\mathbf{\Gamma}$ would impose identical sparsity constraints on $\mathbf{X}_1, \dots, \mathbf{X}_F$. Maximizing the joint evidence, we have

$$\hat{\gamma} = \arg \min_{\gamma} \left\{ \sum_{f=1}^F L \log |\mathbf{\Sigma}_{y_f}| + \text{Tr}(\mathbf{Y}_f^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{Y}_f) \right\}, \quad (15)$$

where $\mathbf{\Sigma}_{y_f} = \sigma_f^2 \mathbf{I} + \mathbf{A}_f \mathbf{\Gamma} \mathbf{A}_f^H$. Computing the derivative of the objective function and equating it to zero gives the update rule²⁴

$$\hat{\gamma}_m^{\text{new}} = \hat{\gamma}_m^{\text{old}} \frac{\sum_{f=1}^F \|\mathbf{Y}_f^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{a}_{f,m}\|_2^2}{L \sum_{f=1}^F \mathbf{a}_{f,m}^H \mathbf{\Sigma}_{y_f}^{-1} \mathbf{a}_{f,m}}. \quad (16)$$

Similar to Eq. (14), a plot of $\hat{\gamma}$ [Eq. (16)] can be interpreted as a broadband ambiguity surface. The update rule Eq. (16) is used in our multi-frequency data analysis.

2. SBL noise estimate

Let $\mathbf{A}_{\mathcal{M}}$ denote the matrix formed by K columns of \mathbf{A} indexed by \mathcal{M} , where the set \mathcal{M} indicates the location of non-zero entries of γ with cardinality K . We estimate \mathcal{M} from $\hat{\gamma}^{\text{new}}$ by picking the strongest K peaks. The noise variance σ_f^2 estimate^{15,16,24,25} is given by

$$\hat{\sigma}_f^2 = \frac{1}{N-K} \text{Tr} \left(\left(\mathbf{I}_N - \mathbf{A}_{f,\mathcal{M}} \mathbf{A}_{f,\mathcal{M}}^+ \right) \mathbf{K}_f \right), \quad (17)$$

where $\mathbf{A}_{\mathcal{M}}^+$ denotes the Moore-Penrose pseudo-inverse of the matrix $\mathbf{A}_{\mathcal{M}}$. In scenarios including environmental mismatch, $\hat{\sigma}_f^2$ can be influenced by model mismatch and not closely represent the actual ambient noise variance.^{26,27}

3. SBL pseudo-code

SBL is implemented as shown in Table I. The unknown source variance γ_m is estimated by iterating over the fixed point update rule in Eq. (16). The iterations are performed until the error criterion (ϵ) in line 8 is met or if the maximum number of iterations (N_t) is reached. We assume a single source, i.e., $K = 1$, for estimating the noise variance. The algorithm performance was not observed to be sensitive to K in simulations. This assumption makes the algorithm flexible, as there is no requirement to know the true number of sources.

III. DATA SELECTION AND PROCESSING

A. SWellEx-96 data

We use the relatively range-independent SWellEx-96 Event S5 data set^{12,28–30} recorded on a 64 element vertical

TABLE I. Algorithm 1: Multi-frequency SBL algorithm.

-
-
- 1: Parameters: $\epsilon = 10^{-8}$, $N_t = 5000$
 - 2: Input: \mathbf{K}_f , $\mathbf{A}_f \forall f$
 - 3: Initialization: $\gamma_m^{\text{old}} = 1$, $\forall m$, $\hat{\sigma}_f^2 = 0.1$, $\forall f$
 - 4: **for** $i = 1$ to N_t
 - 5: Compute: $\Sigma_{y_f} = \hat{\sigma}_f^2 \mathbf{I} + \mathbf{A}_f \Gamma^{\text{old}} \mathbf{A}_f^H$
 - 6: γ_m^{new} update $\forall m$ using Eq. (16)
 - 7: $\hat{\sigma}_f^2$ estimate $\forall f$ using Eq. (17) with $K = 1$
 - 8: If $\frac{\|\gamma^{\text{new}} - \gamma^{\text{old}}\|_1}{\|\gamma^{\text{old}}\|_1} < \epsilon$, **break**
 - 9: $\gamma^{\text{old}} = \gamma^{\text{new}}$
 - 10: **end**
 - 11: Output: γ^{new} , $\hat{\sigma}_f^2 \forall f$
-
-

line array (VLA) sampled at 1500 Hz with $N = 21$ of those elements used for processing. The surface ship R/V Sproul traveled with a radial velocity of 2.5 m/s towards the VLA with closest point of approach (CPA) at approximately 1 km (Fig. 1). The ship towed a deep and a shallow source, both projecting different multi-tonal waveforms.

For the 75 min Event S5, we use the deep source at frequencies 166 and 201 Hz for processing. The data are split into 156 segments, resulting in a single segment length of 29 s. A FFT length of 4096 samples (2.7 s) with 50% overlap results in $L = 21$ snapshots for each segment with a FFT bin width of 0.37 Hz. Our algorithm searches the adjacent ± 2 FFT bins and extracts the FFT value corresponding to the maximum bin power to accommodate Doppler shift. For the




FIG. 1. (Color online) SWellEx-96 Event S5 showing the path of the surface ship R/V Sproul in blue. The ship towed a deep source at ~ 60 m depth along roughly a 200 m isobath during the 75 min VLA recording.

Bartlett processor, each snapshot is windowed spatially with a normalized Kaiser window with $\beta = 4.7$. The adaptive processors optimize their own weights.

A delayed copy of the SWellEx-96 Event S5 data is added to the original SWellEx-96 Event S5 data to construct a two-source scenario, resulting in a source separation of 1 km. This separation (14 data segments) reduces the number of segments with two sources present from 156 to 142 and the start time of the event to approximately 7 min. Source 2 is the weaker source and located farther in range than source 1 (with respect to the VLA from 7 to 61 min). The snapshots of source 2 are added to the snapshots of source 1: $\mathbf{Y} = \mathbf{Y}_1 + \xi \mathbf{Y}_2$. The source 1 to source 2 power ratio (SSR) $\xi = 10^{-3/20} \|\mathbf{Y}_1\|_{\mathcal{F}} / \|\mathbf{Y}_2\|_{\mathcal{F}}$, hence source 2 is 3 dB below the power of source 1. \mathcal{F} denotes the Frobenius norm. The snapshots in \mathbf{Y} are used to construct the SCM.

To ensure the sources are not coherent, each snapshot in $\mathbf{Y}_i = [\mathbf{y}_1, \dots, \mathbf{y}_L]$ is multiplied by a random phase prior to adding the delayed data set to the original data set. This requirement is necessary for adaptive (MVDR and WNC) and eigenanalysis methods³¹ because signal coherence affects processor performance when inverting the SCM. SBL does not invert the SCM but requires an optimization procedure [see Eq. (16)].

For the range-independent waveguide geoacoustic model (Fig. 2), the water depth is 4 m below the deepest element of the array at 212 m. The VLA spans the lower half of the water column and the inter-element spacing is 5.6 m which corresponds to a design frequency of 133 Hz using a sound speed of 1488 m/s. The seafloor is composed of a 23.5 m thick sediment layer, overlaying a 800 m thick mudstone layer. Replica vectors are computed using the Kraken normal mode code³² with a range and depth discretization of 50 m and 10 m on a 10 km \times 200 m grid, respectively.

While many parameters may contribute to mismatch in realistic scenarios, array-tilt belongs to the set of important MFP parameters.^{29,33} Tilt generally is encountered to some degree when using vertical line arrays spanning a significant portion of a shallow water column. Hydrophone




FIG. 2. Waveguide with sound speed profile, VLA, and geo-acoustic parameters for range-independent SWellEx-96 Event S5. A single element out of the 22 element subset is excluded from processing.

displacement due to an array tilt is incorporated into the Kraken field calculation. The tilted array remains in a straight line anchored at the bottom.

To estimate the SNR at the array, we use the signal and a representative noise frequency. For an array, the conventional definition of input SNR is the average element-level signal power divided by the average noise power. When the signal is recorded in the presence of noise, we approximate this definition as

$$\text{SNR} \approx 10 \log_{10} \left(\frac{\text{Tr}(\mathbf{K}_{s+n})}{\text{Tr}(\mathbf{K}_n)} - 1 \right) (\text{dB}). \quad (18)$$

\mathbf{K}_{s+n} refers to the signal plus noise SCM and \mathbf{K}_n is the SCM computed at a neighboring noise frequency.

B. Simulations

To explore processor performance in a controlled environment subject to array tilt, we use the SWellEx-96 replica vectors to simulate a multi-frequency scenario. Source magnitudes are selected such that SSR is 3 dB. Source 2 is the weaker source and is farther in range relative to source 1 by 1 km. To simulate $L = 21$ snapshots while ensuring that both sources are incoherent, each source phase is selected independently from a uniform distribution $[0, 2\pi)$ for each snapshot. The simulated received data vectors are added,

$$\mathbf{Y} = \mathbf{a}(\theta_1)\mathbf{x}_1^T + \mathbf{a}(\theta_2)\mathbf{x}_2^T + \mathbf{N}, \quad (19)$$

where each $\mathbf{x}_i \in \mathbb{C}^L$ contains L complex amplitudes for the i th source. As with the actual data, these complex source amplitudes are modeled as deterministic sequences. The observations $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_L]$ are used to construct the SCM over L snapshots.

Processor performance at a particular frequency is evaluated for additive noise (SNR discretization is 1 dB). Here, array (or average single element) SNR is defined as the ratio of the power of the weaker source 2 to independent and identically distributed (i.i.d.) complex Gaussian noise \mathbf{n} ,

$$\text{SNR} = 10 \log_{10} \frac{\|\mathbf{a}\mathbf{x}\|_2^2}{\mathbf{E}\{\|\mathbf{n}\|_2^2\}} (\text{dB}). \quad (20)$$

Equation (20) corresponds to the single snapshot SNR, where \mathbf{a} is the source replica and x its complex amplitude.

When simulating $L = 21$ snapshots, the signals are added to i.i.d. complex Gaussian noise. When processing multiple frequencies, each frequency has the same SNR and is generated with a different noise seed.

To evaluate processor performance when the source wavefield does not correspond exactly to a dictionary entry, we simulate the stationary sources [Eq. (19)] on a more finely spaced grid of replica vectors (2 m in depth and 10 m in range). The finely spaced replica set allows each of the $L = 21$ snapshots to be drawn randomly from 24 different positions while remaining within $\pm 1/2$ cell to the grid point (10 m depth and 50 m range discretization).

Processor performance is measured by comparing the two dominant peaks found on the ambiguity surface to the grid point of the weaker source. Comparing both peaks (rather than deciding which peak corresponds to which source) is reasonable because processors fail in localizing the weaker source first.¹² The localization statistic P_L for the weaker source is computed by

$$P_L = \frac{C}{Q}, \quad (21)$$

where C is the number of correctly found peaks for $Q = 1000$ simulations.³⁴

IV. RESULTS

A. Simulations

First we compare processor performance using data simulated with and without array tilt mismatch. Panels in Fig. 3 introduce the two source scenario for SNR 0 dB showing normalized ambiguity surfaces using $L = 21$ snapshots at a single frequency (166 Hz). Panels on the left are mismatch free and Bartlett [Fig. 3(a)] displays many ambiguous positions competitive to source 2. This poor performance indicates that adaptive processing is required for localizing the weaker source. WNC -3 dB [panel (b)] and SBL [panel (c)] localize both sources. Panels (d)–(f) include a mismatch of 1° array tilt. The array tilt is included in the data using Eq. (19), whereas the replica vectors used to localize sources are computed for a 0° array tilt. The adaptive processors display an increased amount of ambiguity but localize the weaker source. Its power is close to -3 dB in panels (b), (c), and (f), and at -2.1 dB in panel (e). The strongest sidelobe in panel




FIG. 3. (Color online) Localization for two simulated sources at 166 Hz and SNR 0 dB. True positions are indicated by white squares and each panel is normalized by its respective peak value. Source 1 is located at 2.5 km with power 3 dB above source 2 located at 3.5 km. Left panels are mismatch free, right panels have a data mismatch of 1° array-tilt: (a), (d) Bartlett; (b), (e) WNC -3 dB; and (c), (f) SBL.




FIG. 4. (Color online) Probability of localizing (P_L) the weaker source 2 for 166 Hz [(a), (b), (d), (e)] and two-frequency (166 and 201 Hz) [(c), (f)]. $L = 21$ data snapshots are drawn randomly from a finer replica vector mesh for stationary sources in (b), (c), (e), (f). Left panels are without array-tilt mismatch, right panels include a 1° array-tilt mismatch.

(e) at approximately 70m depth and 8.6km range is at -4 dB.

Normalized plots are used to facilitate a visual comparison of ambiguity surfaces. However, they give the appearance that sidelobe level increases as mismatch increases while the mainlobe power remains unchanged. When mismatch is present, processor output power at the location of the mainlobe decreases.

To compare processor localization performance of the weaker source at different SNRs, we compute localization statistics (Fig. 4) using Eq. (21). The scenario is the same as in Fig. 3. The single frequency (166 Hz), no mismatch panel in Fig. 4(a) shows that Bartlett yields poor performance in localizing the weaker source. MVDR exhibits ideal ($P_L = 1$) performance in localizing the weaker source until a SNR of 5 dB. The WNC -3 dB and SBL have a $P_L < 1$ at a SNR below -3 and -6 dB, respectively.

Next we draw data snapshots independently from a more finely spaced grid of replica vectors while sources remain within a discretization cell. The performance of all processors degrades [Fig. 4(b)] and SBL exhibits a $P_L = 1$ at SNR above -2 dB. Adding an additional frequency (201 Hz) helps with localizing the source at reduced SNRs for all adaptive processors [Fig. 4(c)].

To investigate processor robustness to mismatch, we include a 1° array tilt in the data while the other parameters remain unchanged. Localization statistics [Fig. 4(d)–4(f)] show that the MVDR exhibits a significant loss in performance in all panels, demonstrating that it is not robust to mismatch. SBL and WNC -3 dB exhibit ideal performance at SNRs above -3 and 6 dB, respectively, in Fig. 4(d). Their relative performance is similar to the mismatch-free case. Processing an additional frequency [panel (f)] is beneficial due to the reduced performance of all processors when using only a single frequency [panel (e)] and when data snapshots do not exactly correspond to the replica vectors used for localization.

Processor robustness is further investigated for different degrees of array tilt mismatch (Fig. 5). All data snapshots are randomly drawn from a finer replica vector mesh located within respective discretization cells. Using two frequencies

(166 and 201 Hz) increases the processor output power at the location of source 2 but Bartlett and MVDR display poor performance. SBL exhibits improved performance when localizing the quiet source under increased mismatch and at lower SNRs in panels (a)–(c). Results for SNR 10 dB (not shown) yield only minor improvements in localizing the weaker source when compared to Fig. 5(a) for all processors but the MVDR.

B. SWellEx-96 data

It is useful to estimate the SNR at the VLA of the frequencies emitted by the deep source over the entire Event S5 to allow for a comparison to the simulated data. Figure 6(a) shows SNR at 166 and 201 Hz using two neighboring noise




FIG. 5. (Color online) Probability of localizing (P_L) the weaker source using multi-frequency (166 and 201 Hz) in the presence of array-tilt mismatch at SNR (a) 0 dB, (b) -5 dB, and (c) -10 dB. Data snapshots are drawn randomly from a finer replica vector mesh for stationary sources.




FIG. 6. (Color online) (a) SWellEx-96 deep source average element level SNR. (b) Data correlation for three replica vector tilts and (c) VLA tilt bearing angle with respect to the source-VLA plane. At times, the source interrupted the CW transmission over the 75 min track (most notably at 59, 56, 39, 22, and 18 min).

frequencies at 161 and 214 Hz, respectively. The theoretical array gain is 13 dB. Twenty-one data snapshots are used to estimate the SCM. As the source travels towards the array, the average element SNR increases and is a maximum of 14 dB at CPA, approximately at 61 min. To be consistent with results presented in Ref. 12, we plot the x axis in reverse time.

The tilt with respect to the source-array plane changes over time [panel (b)]. Correlation is defined as the normalized inner product of the data snapshot and best-matching replica vector [see Eqs. (3) and (4) in Ref. 35 for a detailed discussion]. This value is plotted over time for replica vectors computed with different degrees of VLA tilt in the source-array plane. Computed curves follow a similar trend as SNR in panel (a) and it can be observed that different tilts are dominant at different times throughout the 75 min event. Most notably, replica vectors computed with a 1° tilt perform well over most of the data set while replica vectors computed with a 0° tilt yield poor performance at CPA. At CPA, the tilt with respect to the source-array plane is 2° .

The azimuth of the source to the VLA (Fig. 1) as well as the bearing of the tilted VLA change over time. Source-VLA azimuth is computed using GPS data and a tilt/heading sensor mounted on the VLA recorded bearing angle with respect to magnetic north. These data are combined [panel (c)] in order to visualize the tilt bearing angle with respect to the source-VLA plane. This angle varies in time from -50° to 60° . The mismatch is greatest for a 0° relative bearing angle when using replica vectors computed with a 0° tilt. Note that the top of the array is always tilted away from the source. To reduce mismatch when processing the entire data set, we generate replica vectors having a 1° tilt in the source-array plane.

Next we investigate processor performance to localize a weaker source in a scenario subject to array tilt mismatch. Panels in Fig. 7 show two SWellEx-96 sources located at 2.5 and 3.5 km in range at approximately 48 min into the event. Left panels use replica vectors with a 0° tilt, right panels use replica vectors with a 1° tilt. The number of ambiguous positions is reduced in panels (d)–(f) when compared to panels (a)–(c). Similar to the simulations in Fig. 3, Bartlett [panels (a), (d)] displays the most ambiguity whereas WNC -3 dB [panels (b), (e)] and SBL [panels (c), (f)] display less ambiguity for source localization. Processing two frequencies (Fig. 8) slightly reduces the number of ambiguous peaks competitive in power to sources 1 and 2 to less than 5 for all but the Bartlett processor.

To demonstrate processor robustness to localizing the weaker source, we extend our processing in Fig. 8 to the entire Event S5. We extract the five highest power levels and corresponding range information for each ambiguity surface. This data then is displayed as a vertical stripe, containing only the five power levels in their respective range cells. The vertical stripes are assembled in temporal order and individually normalized range-time panels are shown in Fig. 9 for each processor.

The left panels use replica vectors computed with a 0° array tilt, whereas panels on the right are computed using replica vectors with a 1° array tilt. All processors localize the stronger source 1 such that a track is evident over the displayed 75 min event. WNC -3 dB [panel (b)] localizes source 2 sporadically. It is not evident from panel (d) using Bartlett that a second source is present due to parallel running sidelobes. When accounting for a 1° array tilt, WNC -3 dB localizes the weaker source 2 over most of the event




FIG. 7. (Color online) Localization for two SWellEx-96 sources at 166 Hz for data segment 85 of 142. True positions are indicated by white squares. Replica vectors are computed with 0° (left panels) and 1° (right panels) tilt: (a), (d) Bartlett; (b), (e) WNC -3 dB, and (c), (f) SBL. Bartlett displays the most ambiguity while all other processors exhibit fewer ambiguous localizations.




FIG. 8. (Color online) As in Fig. 7 but using two frequencies (166 and 201 Hz) reduces the ambiguous peaks competitive in power to source 1 and source 2 to approximately 5 or less for WNC -3 dB and SBL.

[panel (e)]. Panels (c), (f) demonstrate that SBL localizes the weaker source at most times using either set of replica vectors when compared to WNC panels (b), (e), respectively.

It is apparent that processor performance differs in localizing the weaker source in Figs. 9(a)–9(c). To quantify this difference, we turn to GPS data recorded at the surface ship. We overlaid the track of the weaker source in panels (a)–(c) with the groundtruth GPS track and allowed a variation of ± 2 range cells at each time step. A processor successfully detected the source if it had a single non-zero entry within these five candidate range cells. Bartlett, WNC -3 dB, and SBL have 40, 72, and 98 localizations, respectively, over the entire track. SBL localizes the weaker source at closer distance to CPA when compared to WNC -3 dB, most notably between 65 and 72 min.

Representative CPU times for WNC -3 dB, SBL and basis pursuit are shown in Fig. 10. The BP method is implemented using CVX (Ref. 36) for single and multiple frequencies.¹² The benchmark ambiguity surface (using $20 \times 200 = 4000$ replica vectors) is the same as in Fig. 3 but snapshots are drawn from the more finely spaced grid of replica vectors. Mismatch increases the time required for SBL and BP to converge. BP uses a pre-determined regularizer to compute two sparse solutions rather than determining the

desired regularizer with a walk along the LASSO path. We average over 100 realizations to estimate CPU time for each number of snapshots. Single frequency Bartlett results (not shown) are constant across snapshots with 0.002 s. WNC -3 dB and SBL CPU time roughly are independent of the number of snapshots and scale approximately linearly with the number of frequencies. SBL is slower by a factor of 10 when compared to WNC -3 dB, using either 1 or 2 frequencies.

V. DISCUSSION

Simulation results demonstrate that SBL localizes a weaker source at lower SNRs than WNC -3 dB in ideal scenarios [Fig. 4(a)] and in more realistic single- and multi-frequency scenarios when snapshots do not correspond exactly to replica entries [Figs. 4(b) and 4(c)]. These observations are consistent with CS results implemented using basis pursuit.¹² Results in Figs. 4(d)–4(f) and Fig. 5 indicate that SBL offers a degree of robustness against mismatch in the form of array tilt including snapshots drawn from a finer replica vector mesh while exhibiting improved localization performance when compared to other processors at low SNR.




FIG. 9. Multi-frequency (166 and 201 Hz) range localization of two sources for the SWellEx-96 Event S5. For each of the 142 processed segments/ambiguity surfaces (Fig. 8 shows No. 85), five peaks corresponding to the highest power levels are plotted. Replica vectors are computed with 0° (left panels) and 1° (right panels) tilt: (a), (d) Bartlett; (b), (e) WNC -3 dB; and (c), (f) SBL.




FIG. 10. (Color online) CPU time required to compute a single ambiguity surface similar to Figs. 3(e) and 3(f). Basis pursuit uses a single, pre-determined regularizer.

The SWellEx-96 data set includes mismatch in the form of array tilt (Fig. 6) which is exploited to demonstrate processor localization robustness. SBL [Fig. 9(c)] exhibits improved performance over the entire track when compared to WNC [Fig. 9(b)] in localizing a weaker source in the presence of a stronger source. SBL also offers improved localization performance for the weaker source close to CPA at 61 min when compared to panels (a), (b). This observation is supported by simulations in Fig. 5, showing that SBL can accommodate mismatch in the form of array tilt up to 2° . The adaptive processors perform well in localizing the weaker source when this mismatch is within $\pm 1^\circ$ [Figs. 9(e) and 9(f)]. Of course, in addition to tilt, the data contains a modest amount of unknown environmental mismatch. It is noteworthy that ambiguous solutions in panel (f) are not as connected when compared to Bartlett [panel (d)] and WNC [panel (e)], which further improves visualization of the track. A similar argument can be made when comparing individual panels on the left, but SBL displays a small amount of ambiguous, connected peaks.

Snapshots used to construct the SCM cover a time interval of 29 s. In this interval, the source moves approximately 70 m which is more than one range discretization cell (50 m). Previous results using the basis pursuit method¹² indicate that processing snapshots and or multiple frequencies corresponding to adjacent cells yield adjacent sparse solutions. Panels in Figs. 8(b), 8(c), 8(e), and 8(f) display significant processor output power in two adjacent cells corresponding to the location of the stronger source. Using BP, the number of sparse solutions required to localize all sources is unknown since multiple solutions might correspond to a single source. This problem is amplified in scenarios with mismatch: the processor identifies false locations before both sources are localized [Fig. 8(c) displays many ambiguous positions similar in power to both sources]. SBL requires no *a priori* knowledge of number of sparse solutions, automatically determines sparsity, and therefore sidesteps this problem entirely.

It is noteworthy that some SBL ambiguous solutions correspond to ambiguous WNC solutions [e.g., Figs. 7(b)

and 7(c)]. WNC yielded good performance using a constant tuning parameter of -3 dB and WNC generally performs well when selecting a value within $[-6 -2]$ dB. However, it is difficult to assess how to select the *optimum* white noise gain constraint to localize the weaker source.¹² In contrast, SBL explains the data by use of an optimization without being constrained by a similar tuning parameter.

On average, Bartlett, WNC, and SBL require 3, 25, and 375 s, respectively, to compute all multi-frequency ambiguity surfaces used to construct a respective panel in Fig. 9. Our MFP toolbox is implemented using MATLAB whereas the 142 data segments are parallelized on 25 processors. The SBL algorithm can be parallelized over frequency except for the update rule in Eq. (16) to offset SBL's frequency dependent CPU time (Fig. 10). BP's execution time increases quadratically with the number of snapshots while SBL's execution time largely is snapshot independent.

VI. CONCLUSIONS

We demonstrated that sparse Bayesian learning (SBL) behaves similar to an adaptive processor and outperforms WNC -3 dB when localizing a weaker source in the presence of a stronger source in simulations and with the SWellEx-96 data set. Furthermore, SBL is robust to mismatch in the form of array tilt and additional, modest (unknown) data-replica mismatch. Unlike other sparse methods, SBL automatically determines sparsity. SBL's processing time is independent of the number of snapshots and it is a factor 10 slower when compared to the WNC.

ACKNOWLEDGMENTS

The authors thank David Ensberg (MPL) for providing help with the SWellEx-96 data set. This work was supported by the Office of Naval Research under Grant No. N00014-13-1-0632.

- ¹A. B. Baggeroer, W. A. Kuperman, and P. N. Mikhalevsky, "An overview of matched field methods in ocean acoustics," *IEEE J. Ocean. Eng.* **18**, 401–424 (1993).
- ²H. P. Buckner, "Use of calculated sound fields and matched-field detection to locate sound sources in shallow water," *J. Acoust. Soc. Am.* **59**, 368–373 (1976).
- ³J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory* **50**, 2231–2242 (2004).
- ⁴S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.* **20**, 33–61 (1998).
- ⁵P. A. Forero and P. A. Baxley, "Shallow-water sparsity-cognizant source-location mapping," *J. Acoust. Soc. Am.* **135**, 3483–3501 (2014).
- ⁶G. F. Edelmann and C. F. Gaumont, "Beamforming using compressive sensing," *J. Acoust. Soc. Am.* **130**, EL232–EL237 (2011).
- ⁷A. Xenaki, P. Gerstoft, and K. Mosegaard, "Compressive beamforming," *J. Acoust. Soc. Am.* **136**, 260–271 (2014).
- ⁸J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *IEEE Proc.* **57**, 1408–1418 (1969).
- ⁹R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Ant. Propag.* **34**, 276–280 (1986).
- ¹⁰P. Gerstoft, A. Xenaki, and C. F. Mecklenbräuker, "Multiple and single snapshot compressive beamforming," *J. Acoust. Soc. Am.* **138**, 2003–2014 (2015).
- ¹¹H. Cox, R. Zeskind, and M. Owen, "Robust adaptive beamforming," *IEEE Trans. Acoust. Speech Sign. Process.* **35**, 1365–1376 (1987).
- ¹²K. L. Gemba, W. S. Hodgkiss, and P. Gerstoft, "Adaptive and compressive matched field processing," *J. Acoust. Soc. Am.* **141**, 92–103 (2017).

- ¹³M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.* **1**, 211–244 (2001).
- ¹⁴D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Sign. Process.* **55**, 3704–3716 (2007).
- ¹⁵P. Gerstoft, C. F. Mecklenbräuker, A. Xenaki, and S. Nannuru, "Multisnapshot sparse Bayesian learning for DOA," *IEEE Sign. Process. Lett.* **23**, 1469–1473 (2016).
- ¹⁶Z.-M. Liu, Z.-T. Huang, and Y.-Y. Zhou, "An efficient maximum likelihood method for direction-of-arrival estimation via sparse Bayesian learning," *IEEE Trans. Wireless Commun.* **11**, 1–11 (2012).
- ¹⁷D. Wipf and S. Nagarajan, "Iterative reweighted L1 and L2 methods for finding sparse solutions," *IEEE J. Sel. Top. Sign. Process.* **4**, 317–329 (2010).
- ¹⁸R. Giri and B. D. Rao, "Type I and type II Bayesian methods for sparse signal recovery using scale mixtures," *IEEE Trans. Sign. Process.* **64**, 3418–3428 (2016).
- ¹⁹M. A. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing," *IEEE J. Select. Top. Sign. Process.* **4**, 342–349 (2010).
- ²⁰H. Xu, C. Caramanis, and S. Mannor, "Robust regression and lasso," *IEEE Trans. Inf. Theory* **56**, 3561–3574 (2010).
- ²¹J. Maksym, "A robust formulation of an optimum cross-spectral beamformer for line arrays," *J. Acoust. Soc. Am.* **65**, 971–975 (1979).
- ²²S. Nannuru, P. Gerstoft, and K. L. Gemba, "Sparse Bayesian learning with uncertain sensing matrix," in *International Conference on Acoustics, Speech, and Signal Processing*, IEEE, New Orleans (2017).
- ²³P. Gerstoft and C. F. Mecklenbräuker, "Wideband sparse Bayesian learning for DOA estimation from multiple snapshots," in *2016 IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)*, IEEE (2016), Vol. 2, pp. 1–5.
- ²⁴S. Nannuru, K. L. Gemba, P. Gerstoft, W. S. Hodgkiss, and C. F. Mecklenbräuker, "Multi-frequency sparse Bayesian learning with uncertainty models," arXiv:1704.00436 (2017).
- ²⁵P. Stoica and A. Nehorai, "On the concentrated stochastic likelihood function in array signal processing," *Circ. Syst. Sign. Process.* **14**, 669–674 (1995).
- ²⁶S. E. Dosso, P. L. Nielsen, and M. J. Wilmut, "Data error covariance in matched-field geoaoustic inversion," *J. Acoust. Soc. Am.* **119**, 208–219 (2006).
- ²⁷C. F. Huang, P. Gerstoft, and W. S. Hodgkiss, "On the effect of error correlation on matched-field geoaoustic inversion," *J. Acoust. Soc. Am.* **121**, EL64–EL69 (2007).
- ²⁸N. O. Booth, A. T. Abawi, P. W. Schey, and W. S. Hodgkiss, "Detectability of low-level broad-band signals using adaptive matched-field processing with vertical aperture arrays," *IEEE J. Ocean. Eng.* **25**, 296–313 (2000).
- ²⁹P. Hursky, W. S. Hodgkiss, and W. A. Kuperman, "Matched field processing with data-derived modes," *J. Acoust. Soc. Am.* **109**, 1355–1366 (2001).
- ³⁰D. Tollefsen, P. Gerstoft, and W. S. Hodgkiss, "Multiple-array passive acoustic source localization in shallow water," *J. Acoust. Soc. Am.* **141**, 1501–1513 (2017).
- ³¹D. Johnson and D. Dudgeon, *Array Signal Processing* (Prentice-Hall, Upper Saddle River, NJ, 1993), p. 385.
- ³²M. B. Porter, "The KRAKEN normal mode program," SACLANT Undersea Research Centre Memorandum (SM-245) and Naval Research Laboratory Memorandum Report No. 6920 (1991).
- ³³H. Schmidt, A. B. Baggeroer, W. A. Kuperman, and E. K. Scheer, "Environmentally tolerant beamforming for high-resolution matched field processing: Deterministic mismatch," *J. Acoust. Soc. Am.* **88**, 1851–1862 (1990).
- ³⁴F. Jensen, W. A. Kuperman, M. Porter, and H. Schmidt, *Computational Ocean Acoustics*, 2nd ed. (Springer Science, New York, 2011), p. 709.
- ³⁵N. O. Booth, P. A. Baxley, J. Rice, P. Schey, W. S. Hodgkiss, G. D'Spain, and J. Murray, "Source localization with broad-band matched-field processing in shallow water," *IEEE J. Ocean. Eng.* **21**, 402–412 (1996).
- ³⁶M. Grant and S. Boyd, "CVX: MATLAB software for disciplined convex programming," Version 2.1, <http://cvxr.com/cvx/> (Last viewed June 15, 2016).