Sequential Sensor Placement using Bayesian Compressed Sensing for Source Localization

Milan Courcoux-Caro, Charles Vanwynsberghe, Cédric Herzet, Alexandre Baussard

HAL Id: hal-03070390
https://inria.hal.science/hal-03070390

Submitted on 15 Dec 2020

Sequential Sensor Placement using Bayesian Compressed Sensing for Source Localization

Milan Courcoux-Caro[†], Charles Vanwynsberghe[†], Cédric Herzet[‡], Alexandre Baussard[§]

[†]*Lab-STICC UMR 6285, CNRS, ENSTA Bretagne*, F-29200, Brest, France
[‡]*Inria centre Rennes-Bretagne Atlantique, Université de Rennes, CNRS, IRISA*, 35000, Rennes, France
[§]*Institut Charles Delaunay, Université de Technologie de Troyes*, F-10000, Troyes, France

Abstract—This paper deals with the sensor placement problem for an array designed for source localization. When it involves the identification of a few sources, the compressed sensing framework is known to find directions effectively thanks to sparse approximation. The present contribution intends to provide an answer to the following question: given a set of observations, how should we make the next measurement to minimize (some form of) uncertainty on the localization of the sources? More specifically, we propose a methodology for sequential sensor placement inspired from the "Bayesian compressive sensing" framework introduced by Ji *et al*. Our method alternates between a step of sparse source localization estimation, and a step to choose the sensor position that minimizes the covariance of the estimation error. Numerical results show that an array designed by the proposed procedure leads to better performance than sensors positioned at random.

Index Terms—Acoustic, source localization, sequential sensor placement, compressed sensing

I. INTRODUCTION

Sensor arrays are commonly used to achieve source localization (SL) in acoustic or radar applications. The last decade has shown promising results in the use of sparse approximation to solve this problem, when the measured field (acoustic or electromagnetic) results from a few number of sources. Moreover, when performed with an array of random sensor positions, it shows close connections with compressed sensing (CS) [1]–[3].

This paper investigates the design of the array geometry to obtain efficient compressed measurements in SL. An extensive literature exists when it is based on optimizing the directivity pattern (typically by reducing the side lobe level peak, see *e.g.*, [4], [5]) but it is rather dedicated to beamforming techniques. In sound-field control applications, a review of joint source and sensor placement methods show how to enhance the stability and accuracy of the synthesized field [6]. However, to our knowledge, only few studies are connected to the CS framework; we notice the work by Gaumond *et al* for the array geometry design by Statistical Restricted Isometry Property (StRIP), applied to underwater acoustics [7].

In this contribution we propose a sequential sensor placement to design the array geometry. The key idea stems from the following question: given a set of measurements obtained from an array of sensors, how to place a new sensor so that

the new measure minimizes the uncertainty on the angles of arrival? This idea can be identified with the so-called "Bayesian experimental design" and has recently been studied to iteratively position sensors in a stationary acoustic scene [8]. In the latter contribution, the sparse nature of the sought signal is not exploited. In the present work, we propose an algorithm for sensor placement, inspired from the "Bayesian compressive technique" introduced in [9], [10], and exploiting a sparse prior on the source locations. Similarly to [8], we assume that the measured acoustic field is stationary all along the sensing process. The proposed methodology therefore targets applications where processing time is not crucial but where the focus is on reducing the number of sensors for, *e.g.*, economic constraints.

The paper is organized as follows. In section II, we formally state the "source localization" and "sensor placement" problems. In section III, we show that a closed-form expression for the error-covariance matrix appearing in our sensor-placement problem can be obtained by using a proper Bayesian hierarchical model. In section IV, we present a low-complexity greedy procedure to find an approximate solution of the sensor-placement problem. Finally, in section V we illustrate the performance gain obtained by the proposed adaptive placement of the sensors for source localization.

II. MODEL AND GENERAL STRATEGY

Consider a linear array of aperture D which consists of M sensors. Their positions are defined in the vector $\mathbf{p} \in [-\frac{D}{2}, \frac{D}{2}]^M$. We suppose that the measured signal $\mathbf{y}(\mathbf{p}) \in \mathbb{C}^M$ results from S uncorrelated narrowband sources in the far field. In the SL literature, one classic approach is to model the forward problem on grid as:

$$\mathbf{y}(\mathbf{p}) = \mathbf{A}(\mathbf{p})\mathbf{x} + \mathbf{w}, \qquad (1)$$

where $\mathbf{w}$ is an uncorrelated noise, supposed to be a zero-mean complex Gaussian with variance α_0^{-1}. The matrix $\mathbf{A}(\mathbf{p}) \in \mathbb{C}^{M \times N}$ contains N steering vectors associated to plane waves propagating from angles $\theta_1, \cdots, \theta_N \in [-\frac{\pi}{2}, \frac{\pi}{2}]$. For convenience we define $\mathbf{A}(\mathbf{p})$ by its rows, such that $\mathbf{A}(\mathbf{p}) = [\mathbf{r}(p_1)^{\mathsf{T}}, \ldots, \mathbf{r}(p_M)^{\mathsf{T}}]^{\mathsf{T}}$ and:

$$\mathbf{r}(p_m) = [e^{-j\frac{2\pi}{\lambda}p_m \sin(\theta_1)}, \ldots, e^{-j\frac{2\pi}{\lambda}p_m \sin(\theta_N)}], \qquad (2)$$

with λ the source wavelength. Thus, finding the S sources consists in identifying an S-sparse vector $\mathbf{x} \in \mathbb{C}^N$. The

magnitude of the non-zero values gives the source amplitudes, and their indices relate to their angular positions in the far field.

The goal of this paper is to design the geometry of the sensor array (that is the choice of $\mathbf{p}$) so that the measurements minimize some form of uncertainty of the ground-truth value of $\mathbf{x}$. The main idea pursued hereafter is as follows: given some "posterior" covariance error matrix on $\mathbf{x}$, say $\boldsymbol{\Sigma}(\mathbf{p}) \in \mathbb{C}^{N \times N}$, our goal is to identify the sensor position minimizing the determinant of $\boldsymbol{\Sigma}(\mathbf{p})$:

$$\tilde{\mathbf{p}} = \underset{\mathbf{p} \in [-\frac{D}{2}, \frac{D}{2}]^M}{\operatorname{argmin}} \log |\boldsymbol{\Sigma}(\mathbf{p})|. \tag{3}$$

In the literature, this approach is known as "D-optimal" design [11]. Solving (3) entails nevertheless several problems. First of all, evaluating the covariance matrix $\boldsymbol{\Sigma}(\mathbf{p})$ may be a difficult task depending on the prior model used for $\mathbf{x}$. We address this problem in section III by relying on a Gaussian hierarchical model as in [12]. Second, solving numerically optimization problem (3) may be a difficult task. In this paper, we consider a greedy approach to find an approximate solution of this problem, see section IV.

III. SOURCE LOCALIZATION BY MAXIMUM A POSTERIORI

In order to solve the optimization problem (3), we first introduce the SL estimator to infer both an empirical estimation of the covariance matrix $\boldsymbol{\Sigma}(\mathbf{p})$, and the unknown source vector $\mathbf{x}$. It relies on the Sparse Bayesian Inference (SBI) [12], whose MAP resolution is described below.

A. Bayesian hierarchical model

First, the noise $\mathbf{w}$ follows the uncorrelated complex circular distribution with variance α_0^{-1}, that is $\mathcal{CN}(\mathbf{0}, \alpha_0^{-1} \mathbf{I}_M)$. The MAP infers α_0^{-1} since it is unknown, and uses the following non-informative Gamma prior of shape c and rate d:

$$\mathcal{P}(\alpha_0|c, d) = \Gamma(\alpha_0|c, d), \tag{4}$$

with $c, d \to 0^+$, as done in [12], [13].

Second, the present problem is undetermined, thus $N \gg M > S$. An intuitive choice for the prior of $\mathbf{x}$ could be the Laplace distribution to favor sparsity, but it does not lead to a tractable derivation of the MAP. Rather, we substitute this choice by a normal distribution prior on $\mathbf{x}$:

$$\mathcal{P}(\mathbf{x}|\boldsymbol{\alpha}) = \mathcal{CN}(\mathbf{0}, \operatorname{diag}(\boldsymbol{\alpha})^{-1}), \tag{5}$$

combined with a Gamma hyper-prior for each element of $\boldsymbol{\alpha}$ in the diagonal covariance matrix, that is:

$$\mathcal{P}(\boldsymbol{\alpha}|\rho) = \prod_{n=1}^{N} \Gamma(\alpha_n|1, \rho). \tag{6}$$

The shape parameter equals 1, and the rate ρ is a small positive constant as suggested in [12]. Such a hierarchical model is the basis of the well-known Relevance Vector Machine procedure [14], and appears in various works to express a simple surrogate of a sparse prior on $\mathbf{x}$ [12], [15].

B. MAP estimator by Expectation-Maximization

Consider the estimation of $\mathbf{x}$ in a MAP sense:

$$\tilde{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmax}} \, \mathcal{P}(\mathbf{x}|\mathbf{y}(\mathbf{p}), \boldsymbol{\alpha}, \alpha_0). \tag{7}$$

Taking model (4)-(6) into account, a closed-form solution of $\tilde{\mathbf{x}}$ can be found since all hyper-priors are conjugates of the Gaussian priors on $\mathbf{x}$ and $\mathbf{w}$. In fact, the posterior distribution is also Gaussian, and writes $\mathcal{P}(\mathbf{x}|\mathbf{y}(\mathbf{p}), \boldsymbol{\alpha}, \alpha_0) = \mathcal{CN}(\tilde{\mathbf{x}}_{\mathbf{p}}, \widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}})$. Its mean and covariance read:

$$\tilde{\mathbf{x}}_{\mathbf{p}} = \alpha_0 \widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}} \mathbf{A}(\mathbf{p})^{\mathsf{H}} \mathbf{y} \tag{8}$$

$$\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}} = (\alpha_0 \mathbf{A}(\mathbf{p})^{\mathsf{H}} \mathbf{A}(\mathbf{p}) + \operatorname{diag}(\boldsymbol{\alpha})^{-1})^{-1}. \tag{9}$$

The variables $\boldsymbol{\alpha}$ and α_0 are latent and must therefore be estimated as well. This can be done via an Expectation-Maximization (EM) procedure as described in [12]. Note that $\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}}$ is obtained posterior to the measurements with sensor positions fixed, and thus is not parametric. Unlike the expression of $\boldsymbol{\Sigma}(\mathbf{p})$ in (3), it is not a function of the vector $\mathbf{p}$.

The SBI algorithm is implemented as follows: after initializing $\boldsymbol{\alpha}$ and α_0, one computes $\tilde{\mathbf{x}}_{\mathbf{p}}$ and $\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}}$ from (8)-(9). Then, $\boldsymbol{\alpha}$ and α_0 are re-evaluated via an EM update. These steps are repeated until convergence of the procedure.

IV. GREEDY SENSOR PLACEMENT

In this section, we propose a greedy procedure to find an approximate solution of (3). It starts with a small number of sensors M_0, from which a first estimation of $\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}}$ is computed. Then, the most informative sensor position, named $\tilde{p}_*$, is chosen to reduce the estimation covariance error. Again, $\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}}$ is derived with the new measurement, named y_*, to continue and repeat the placement until the array consists of M sensors. Because of the incremental nature of the approach, the size of the vector $\mathbf{y}$ (resp. $\mathbf{p}$) increases with the addition of the new element y_* (resp. $\tilde{p}_*$) at each iteration.

The greedy procedure consists in finding the sensor position p_* which maximizes the reduction of the ellipsoid error volume. According to eq. (3), it writes:

$$\tilde{p}_* = \underset{p_* \in [-\frac{D}{2}, \frac{D}{2}]}{\operatorname{argmin}} \log \left| \boldsymbol{\Sigma}\left(\begin{bmatrix} \mathbf{p} \\ p_* \end{bmatrix} \right) \right|, \tag{10}$$

by keeping the previous positions $\mathbf{p}$ constant so that the covariance is parametric to p_* only. By relying on eq. (9), one can write the hypothetical expression of $\boldsymbol{\Sigma}\left([\mathbf{p}^{\mathsf{T}}, p_*]^{\mathsf{T}}\right)$ if y_* was known. If (a) the hyper-parameters α_0 and $\boldsymbol{\alpha}$ are supposed to be approximately the same with one new measurement, and (b) by using the Matrix determinant lemma, one can evidence that:

$$|\boldsymbol{\Sigma}\left([\mathbf{p}^{\mathsf{T}}, p_*]^{\mathsf{T}}\right)| \approx \frac{|\widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}}|}{1 + \alpha_0 \mathbf{r}(p_*) \widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}} \mathbf{r}(p_*)^{\mathsf{H}}}, \tag{11}$$

with $\mathbf{r}(p_*)$ the candidate row vector computed at the position p_*. Thus, by defining $\delta h(p_*) = \log(1 + \alpha_0 \mathbf{r}(p_*) \widetilde{\boldsymbol{\Sigma}}_{\mathbf{p}} \mathbf{r}(p_*)^{\mathsf{H}})$, problem (10) can be rewritten as:

$$\tilde{p}_* = \underset{p_* \in [-\frac{D}{2}, \frac{D}{2}]}{\operatorname{argmax}} \delta h(p_*). \tag{12}$$

Algorithm 1 Sequential sensor placement for sparsity-promoting Source Localization by a greedy approach.

Require: M_0, M, $\mathbf{y} \in \mathbb{C}^{M_0}$, $\mathbf{p} = [p_1, \ldots, p_{M_0}]^\mathsf{T}$
  for $m = M_0 + 1, M_0 + 2 \ldots, M$ **do**
    - Localize sources by SBI to compute $\widetilde{\boldsymbol{\Sigma}}_\mathbf{P}$
    - Find the next optimal sensor position $\widetilde{p}_*$ – eq. (12)
    - Acquire the new measurement y_*
    - Update variables $\mathbf{y} \leftarrow [\mathbf{y}^\mathsf{T}, y_*]^\mathsf{T}$ and $\mathbf{p} \leftarrow [\mathbf{p}^\mathsf{T}, \widetilde{p}_*]^\mathsf{T}$
  end for
  return $\mathbf{p}$

The quantity maximized in the $\log$ function corresponds, in ratio, to the volume reduction of the ellipsoid error by the current sequence. Note that, in previous works, a similar result is obtained via the definition of the Shannon entropy [9], [15]. Indeed, the entropy of a multivariate normal distribution of covariance $\boldsymbol{\Sigma}$ equals $|\boldsymbol{\Sigma}|$ up to a constant. Consequently, given the assumptions, reducing the entropy of the unknown $\mathbf{x}$ amounts to accomplishing the so-called D-optimal design as described in equation (3).

Finally, we remark that finding p_* is doable by an extensive search on the line $[-\frac{D}{2}, \frac{D}{2}]$ regularly sampled on H candidate positions. In this way, the computation remains very fast, and the longest step of the strategy remains the SBI step, necessary to update the empirical covariance matrix. The proposed method is summarized in the Algorithm 1.

V. NUMERICAL EXPERIMENTS

This section presents results from numerical simulations of the SL problem in the far field. It compares the performance of the SL, when the sequential sensor placement is either from the proposed method, or done randomly as in the classical CS framework [2].

A. Simulation setup

In the considered scenario, we configure the array with an aperture $D = 10$ m and an SNR $= 12$ dB. The sensor noise variance α_0^{-1} is calculated according to this SNR. The forward model regularly samples the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$ on $N = 181$ angles. $S = 10$ far-field uncorrelated sources emit at the wavelength $\lambda = 1.0$ m. Their angular positions are sufficiently separated to respect the necessary condition for the sparse source recovery – see [16] for more details. But to our knowledge, in the SL framework, no condition on the minimal number of measurements guaranteeing the localization of S sources exists.

The sequential array design starts from $M_0 = 2$ sensors, located at $\pm D/2$ to respect the aperture. Note that there is an ambiguity in the SL with a low number of sensors. Nevertheless, we arbitrarily limit M_0 to 2 in order to reduce the effect of the initial condition imposed by the user. 43 sensors are added one by one to reach $M = 45$ in total. At each of the 43 iterations, the optimum criterion seeks the next sensor into the interval $[-\frac{D}{2}, \frac{D}{2}]$ regularly sampled on $H = 10001$ points.

B. Performance score

In this paper, the SL is assessed as a source detection tool. To do so, the chosen performance score is the Jaccard index J, a measure of similarity by calculating the ratio between the intersection and the union of 2 sample sets. Here, the sample sets are the detected sources and the ground truth. The Jaccard index is a function of its true positive (TP), false positive (FP) and false negative (FN) rates. As in [17], it writes:

$$J = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}}. \tag{13}$$

Note that $0 \leq J \leq 1$, and $J = 1$ reveals an exact detection, whereas $J = 0$ if none of the sources is correctly detected. To compute the TP, FP and FN rates, we pair the spikes from the estimation $\widetilde{\mathbf{x}}$ with the ground truth positions if the distance between them is lower than a tolerance radius r. The number of reconstructed pairs is TP, the unpaired spikes from $\widetilde{\mathbf{x}}$ give the FP, and the unpaired ground truth positions give the FN. Knowing that the angular grid equally samples $[-\frac{\pi}{2}, \frac{\pi}{2}]$, the tolerance radius r arbitrarily equals the angular step $|\theta_{n+1} - \theta_n|$. In all the results below, the score J is averaged over several realizations to obtain statistical trends.

C. Relevance of the sequential design for SL

In this section, we analyze how much the score δh is relevant to sequentially improve the source detection score J. To do this, we test sensor positioning as a function of δh values. First, $\delta h(p_*)$ is linearly scaled between 0 and 1, to obtain $\overline{\delta h}(p_*)$. Accordingly, $\overline{\delta h}(p_*) = 1$ (resp. 0) is the most (resp. least) informative position candidate. Then, we run the sequential placement and choose randomly the M sensor positions $\widetilde{p}_*$ such that $\overline{\delta h}(\widetilde{p}_*) \in [0, 0.25]$. The same process is repeated 4 times, from the least to the most informative intervals $[0.25, 0.5]$, $[0.5, 0.75]$, and $[0.75, 1]$. Note that a fully random placement corresponds to $\overline{\delta h}(\widetilde{p}_*) \in [0, 1]$, and the optimal choice is obtained for $\overline{\delta h}(\widetilde{p}_*) = 1$.
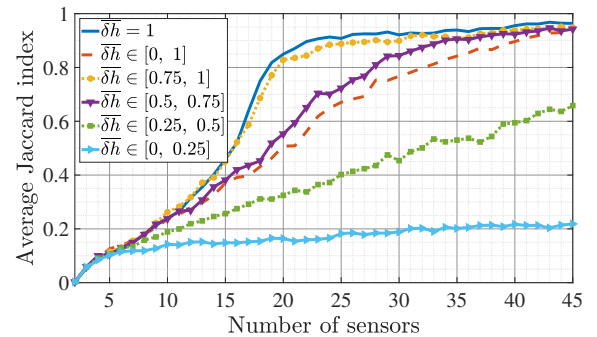

Fig. 1. SL performance for 10 sources, in a sequential placement of 43 sensors (with $M_0 = 2$ initial sensors). Jaccard index averaged on 100 realizations, with fixed sources and different noise $\mathbf{w}$.

Fig. 1 corresponds to the averaged score J obtained from 100 realizations when the SNR and the source positions are fixed. The results show that the considered criterion is consistent for choosing informative sensor positions. Indeed,
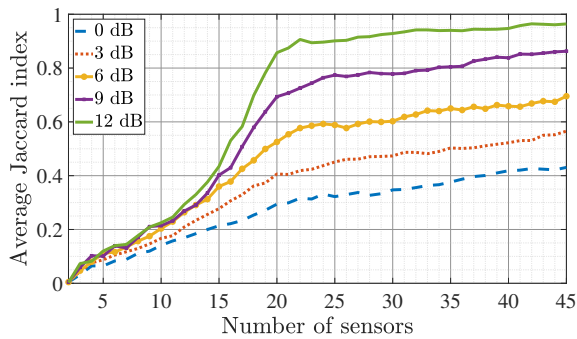
Fig. 2. SL performance for 10 sources at different SNRs. Sequential placement of 43 sensors (with $M_0 = 2$ initial sensors). Jaccard index averaged on 100 realizations, with fixed sources and different noise **w**.
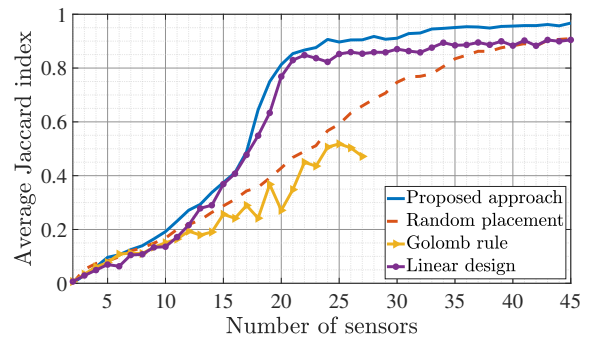

Fig. 3. SL performance for 10 sources. Sequential selection of 45 sensors (with $K_0 = 2$ initial sensors). Jaccard index averaged on 100 realizations, with different sources (position and phase).

the performance increases quicker if the interval is closer to $\overline{\delta h} = 1$. The best average Jaccard index is given with the choice "$\overline{\delta h} = 1$", and a fully random choice improves the performance at a rate between $\overline{\delta h} \in [0.25, 0.5]$ and $\overline{\delta h} \in [0.5, 0.75]$. Note that the choice "$\overline{\delta h} = 1$" only needs 25 sensors to reach the best detection performance $J \approx 0.95$, against 40 when fully-random positioning is considered. Note that J does not reach 1 in average, since a small portion of the realizations never detect all sources by SBI.

In order to evaluate the performances (or robustness) of the method according to the noise level, it is proposed in Fig. 2 to plot the optimal (that is "$\overline{\delta h} = 1$") sequential placement results when the SNR range from 0 to 12 dB. The curves follow a similar trend and the best detection performance, corresponding to $J \approx 0.9$, is only reached for SNR = 12 dB (for the considered range of sensors). It appears that the method is quite affected by the noise level. One way to overcome this problem could probably be to deal with multi-snapshot measurements.

D. Comparison with other methods

In this section, we propose to compare the proposed method with 3 state of the art approaches: random, uniform linear and the Golomb ruler [7]. The Golomb geometry is the only method found by the authors, that use the CS framework in source localization. Note that the Golomb array size does not exist for $M > 27$. Also note that the optimal and random geometries are designed by an iterative addition of sensors, however linear and Golomb arrays can not follow this iterative procedure while keeping a fixed aperture D. It means that these two geometries redefine all sensor positions at each iteration.

Fig. 3 shows the results obtained by averaging 100 realizations for a fixed SNR = 12 dB. Note that, this time, for each realization, the position and the initial phase of the sources are randomly fixed. The proposed method and the linear design have similar results. To reach, for example, $J = 0.9$, the random placement needs in our simulations 25 additional sensors compare to our method. On can note that 21 sensors are needed to reach the best detection performance.

This is related with the Shannon-Nyquist conditions, saying that $M \geqslant \frac{2D}{\lambda} - 1 = 19$ to avoid spatial aliasing. Finally, despite having a better Statistical Restricted Isometry Property (StRIP) than random geometries [7], the Golomb array poorly succeeds in detecting the sources in the presented simulation. Consequently, with enough sensors (*e.g.* $M = 40$ to 45 in Fig. 3) all methods (but Golomb) can finally detect the S sources.

Thus, the proposed approach is of interest when the number of sensors is limited. Note that the uniform linear array is highly dependent of the highest frequency of the signal contrary to our method.

VI. CONCLUSIONS

This paper proposes a sequential sensor placement using Bayesian experimental design to localize sources. It is based on a greedy optimization procedure, and relies on the inference of the a posteriori error covariance matrix. To our knowledge, the proposed placement method is the first in SL to be driven by both the model and the measured data. Since it relies on a sparsity-promoting prior, it is strongly connected to Bayesian Compressive Sensing introduced in [9], [10] although, unlike these contributions, the present paper deals with a parametric sensing matrix of steering vectors.

The numerical results show that the proposed approach is relevant and gives better results than the state of the art (uniform linear array and CS geometries), in the context of source detection and localization. In practice, it is useful if one needs to minimize the number of sensors (*e.g.* for economic purposes) to find all the sources. For applications, future research directions could extend the present study to arbitrary array dimensions (*i.e.*, planar or volumetric), or to the use of "off-the-grid" localization methods.

REFERENCES

[1] S. Fortunati, R. Grasso, F. Gini, M. S. Greco, and K. LePage, "Single-snapshot doa estimation by using compressed sensing," EURASIP Journal on Advances in Signal Processing, vol. 2014, no. 1, p. 120, 2014.

[2] A. Xenaki, P. Gerstoft, and K. Mosegaard, "Compressive beamforming," The Journal of the Acoustical Society of America, vol. 136, no. 1, pp. 260–271, 2014.

[3] G. F. Edelmann and C. F. Gaumond, "Beamforming using compressive sensing," The Journal of the Acoustical Society of America, vol. 130, no. 4, pp. EL232–EL237, 2011.

[4] D. W. Boeringer and D. H. Werner, "Particle swarm optimization versus genetic algorithms for phased array synthesis," IEEE Transactions on antennas and propagation, vol. 52, no. 3, pp. 771–779, 2004.

[5] K.-K. Yan and Y. Lu, "Sidelobe reduction in array-pattern synthesis using genetic algorithm," IEEE Transactions on Antennas and Propagation, vol. 45, no. 7, pp. 1117–1122, 1997.

[6] S. Koyama, G. Chardon, and L. Daudet, "Optimizing source and sensor placement for sound field control: An overview," IEEE/ACM Transactions on Audio, Speech, and Language Processing, 2020.

[7] C. F. Gaumond and G. F. Edelmann, "Sparse array design using statistical restricted isometry property," The Journal of the Acoustical Society of America, vol. 134, no. 2, p. EL191, 2013.

[8] L. Gilquin, P. Lecomte, J. Antoni, T. Le Magueresse, and C. Marteau, "Iterative positioning of microphone arrays for acoustic imaging," Journal of Sound and Vibration, vol. 469, p. 115116, 2020.

[9] S. Ji, Y. Xue, L. Carin et al., "Bayesian compressive sensing," IEEE Transactions on signal processing, vol. 56, no. 6, p. 2346, 2008.

[10] M. W. Seeger and H. Nickisch, "Compressed sensing and bayesian experimental design," in Proceedings of the 25th international conference on Machine learning. ACM, 2008, pp. 912–919.

[11] V. V. Fedorov, Theory of optimal experiments. Elsevier, 2013.

[12] Z. Yang, L. Xie, and C. Zhang, "Off-grid direction of arrival estimation using sparse bayesian inference," IEEE Transactions on Signal Processing, vol. 61, no. 1, pp. 38–43, 2012.

[13] Q. Wu, Y. D. Zhang, M. G. Amin, and B. Himed, "Complex multitask bayesian compressive sensing," in 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2014, pp. 3375–3379.

[14] M. E. Tipping, "Sparse bayesian learning and the relevance vector machine," Journal of machine learning research, vol. 1, no. Jun, pp. 211–244, 2001.

[15] M. W. Seeger, "Bayesian inference and optimal design for the sparse linear model," Journal of Machine Learning Research, vol. 9, no. Apr, pp. 759–813, 2008.

[16] A. Xenaki and P. Gerstoft, "Grid-free compressive beamforming," The Journal of the Acoustical Society of America, vol. 137, no. 4, pp. 1923–1935, 2015.

[17] Q. Denoyelle, V. Duval, G. Peyré, and E. Soubies, "The sliding frank-wolfe algorithm and its application to super-resolution microscopy," Inverse Problems, 2019.