

ROBUST BAYESIAN BEAMFORMING FOR SOURCES AT DIFFERENT DISTANCES WITH APPLICATIONS IN URBAN MONITORING

Patrick W.A. Wijnings^{*†} Sander Stuijk^{*} Bert de Vries^{*‡} Henk Corporaal^{*}

^{*} Eindhoven University of Technology, Eindhoven, The Netherlands

[†] Sorama, Eindhoven, The Netherlands

[‡] GN Hearing, Eindhoven, The Netherlands

ABSTRACT

Acoustic smart sensor networks can provide valuable actionable intelligence to authorities for managing safety in the urban environment. A spatial filter (beamformer) for localization and separation of acoustic sources is a key component of such a network. However, classical methods such as delay-and-sum beamforming fail, because sources are located at varying distances from the sensor array. This causes a regularization problem where either far-away sources are wrongly attenuated, or noise is wrongly amplified.

We solve this by considering source strength and location as random variables. The posterior distributions are approximated using Gibbs sampling. Each marginal is computed by combining importance sampling and inverse transform sampling using Chebyshev polynomial approximation. This leads to an iterative algorithm with similarities to deconvolution beamforming.

Our method is robust against deviations in manifold model, can deal with sources at different distances and power levels, and does not require an a priori known number of sources.

Index Terms—Acoustic applications, Array signal processing, Robustness, Bayes methods

1. INTRODUCTION

Smart sensor networks which provide *actionable intelligence* to authorities are increasingly important for managing safety in the urban environment [2]. Acoustic sensors are valuable in this application, because they enable capture of events that are difficult to observe with other modalities. Examples of such events are gunshots, breaking glass, car alarms, aggressive behaviour, and violation of noise regulations.

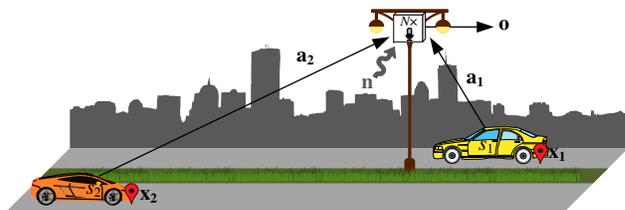
This paper focuses on a key component of such a network: a *spatial filter* that localizes and separates the sources in the environment. Fig. 1a sketches the situation. The problem is formulated in the phasor (narrowband) domain, where source m has signal $s_m [\text{Pa}] \in \mathbb{C}$. The *array manifold* $\mathbf{a}_m(\mathbf{x}_m) [\text{Pa}/\text{Pa}] \in \mathbb{C}^N$ is the transfer function from source to sensors. It depends on source locations $\mathbf{x}_m = (x_m; y_m; z_m) [\text{m}] \in \mathbb{R}^3$. N is the number of sensors and M is the (typically unknown) number of sources. $\mathbf{n} [\text{Pa}] \in \mathbb{C}^N$ is the noise.

Thus, given an observation from the sensors:

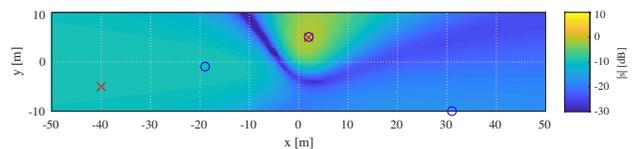
$$\mathbf{o} = \mathbf{n} + \sum_{m=1}^M \mathbf{a}_m(\mathbf{x}_m) s_m [\text{Pa}] \in \mathbb{C}^N, \quad (1)$$

the goal is to find signal and location estimators $\hat{s}_m(\mathbf{o})$ and $\hat{\mathbf{x}}_m(\mathbf{o})$ that maximize expected signal-to-noise ratio (SNR) while accounting for sensor noise, interference, and manifold perturbations.

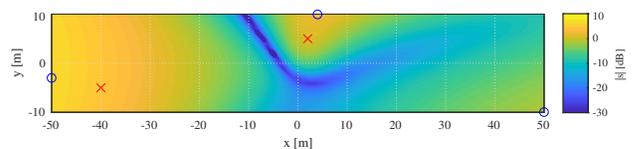
This work is funded by the NWO Perspectief program ZERO. Correspondence to: p.w.a.wijnings@tue.nl.



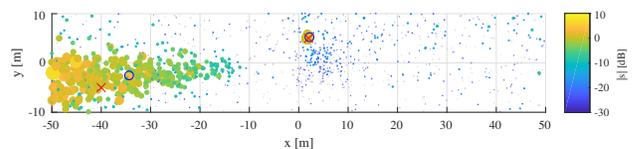
(a) Situation sketch. See main text for definition of the variables.



(b) Delay-and-sum beamformer. Color represents estimated source power for each steering direction. Local maxima marked with \circ . Source 1 is estimated as louder than the equal-power source 2.



(c) Delay-and-sum beamformer (attenuation-compensated). The sources are now estimated as equally loud (at \times), but their locations are difficult to find because point estimates based on local maxima (\circ) are unreliable; the solution is ill-conditioned. This also results in a significant ‘ghost source’ in the bottom-right corner which consists of blown-up noise and leakage from the actual sources.



(d) Our method. Dots mark samples from the posterior distribution $[\mathbf{x}_1, s_1 | \mathbf{o}]$. Color ($\propto \log|s|$) and radius ($\propto |s|^2$) represent source power. Point estimates $\hat{\mathbf{x}}_{1,2}$ (\circ) based on k-means clustering. Source power is now correctly estimated (at \circ) and no ghost sources (with power larger than the noise floor) are observed.

Fig. 1. Example case with two 3 kHz, 0 dB sources (\times) and -30 dB noise floor. The sensor array, mounted 8 m above the origin, has 64 sensors in a golden angle Fermat spiral with 14 cm aperture [1].

One property of classical methods such as delay-and-sum (DS) beamforming [3] is that a source near the array is estimated as louder than an identical source far from the array (Fig. 1b). This is because the manifold consists of both delay and *attenuation* which is neglected in DS beamforming. This is an important shortcoming in our application, where the absolute sound levels of the sources convey information about the type of event. Unfortunately, attenuation compensation inevitably comes together with amplification of noise and leakage (i.e. side-lobes) of interfering sources. Hence insufficiently accounting for observation noise results in an ill-conditioned solution, as illustrated in Fig. 1c, because the matrix $[\mathbf{a}_1 \cdots \mathbf{a}_M]$ is (near-)singular [4].

Minimum variance beamforming [5] models observation uncertainty by considering s_m and \mathbf{n} (Gaussian) random variables. Bayes' rule is then used to infer beliefs about s_m given \mathbf{o} . The expected values of these beliefs give point estimators \hat{s}_m that provably maximize SNR [6]. However, the result is still ill-conditioned with respect to perturbations in $\mathbf{a}_m(\mathbf{x}_m)$ and hence also \mathbf{x}_m , i.e. performance will degrade significantly if locations of interfering sources near the array are not precisely known. Thus, \mathbf{x}_m must also be considered a random variable. Unfortunately, this causes intractable integrals during belief inference that must be approximated.

In iterative beamforming [7], inference is approximated using Expectation Maximization (EM) [8], where iteratively expected values of the source contributions, locations and powers are updated. The method accounts for manifold attenuation, but only results in the acoustic near-field are presented. Also, because EM is (partially) a maximum likelihood method, it typically underestimates uncertainty of the posterior distribution.

Instead, Besson and Bidon [9] use Gibbs sampling [10], where iteratively samples from the source powers, locations and noise covariance are updated. By parametrizing location in terms of angle of arrival and assigning probability distributions with specific form, these samples can be efficiently drawn from standard distributions. However, they do not take manifold attenuation into account.

In this work, Gibbs sampling (like [9]) is combined with a Cartesian parametrization of location (like [7]). However, we update the sources sequentially instead of simultaneously which leads to a probabilistic version of deconvolution beamforming [11]. Due to our parametrization, the Gibbs samples must be drawn from non-standard distributions. We do this by combining importance sampling [12] and inverse transform sampling [13]. Furthermore, the number of active sources is automatically determined using a shrinking prior like relevance vector machine beamforming [14]. Finally, unlike [9], we do not consider multiple snapshots because they are not available when source locations are changing over time, and also require additional assumptions about correlations between sources.

Our contribution is a beamformer which can estimate an a priori unknown number of sources at different distances and power levels, as illustrated in Fig. 1d. Furthermore, our beamformer is robust against perturbations in the manifold.

Next, specification of the generative model (Sec. 2) and inference (Sec. 3–4) are presented separately. Finally, results from simulations are presented in Sec. 5 and we draw conclusions in Sec. 6.

2. MODEL SPECIFICATION

Our generative model is specified using the probability distribution:

$$[\underbrace{s_{1..M}, \mathbf{x}_{1..M}}_{\text{Joint}}, \mathbf{o}] = \underbrace{[\mathbf{o}|s_{1..M}, \mathbf{x}_{1..M}]}_{\text{Likelihood}} \prod_{m=1}^M \underbrace{[\underbrace{s_m}_{\text{Signal}} | \underbrace{\mathbf{x}_m}_{\text{Location}}]}_{\text{Prior}}, \quad (2)$$

where $[\cdot]$ is notation for probability density function (PDF).

2.1. Signal prior

Because little information about the sources is available, an *uninformative prior* is chosen. The prior is split up into a gain and phase component:

$$[s_m] = \underbrace{[|s_m|]}_{\text{Gain}} \underbrace{[\angle s_m]}_{\text{Phase}}. \quad (3)$$

For the phase, a uniform distribution is chosen:

$$[\angle s_m] = \begin{cases} \frac{1}{2\pi} & |\angle s_m| \leq \pi \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

while for the gain, the (improper) *Jeffreys prior* [15, Sec. 6.15] is chosen:

$$[\log|s_m|] \propto 1 \equiv [|s_m|] \propto 1/|s_m|, \quad (5)$$

i.e. sources of any dB(SPL) level are equally expected.

Through a change of variables this implies a distribution on $|s_m|$ (in units of Pa) with a pole at zero amplitude and slowly decaying tail. Such distributions manifest horseshoe-like shrinking behaviour [16, 17], i.e. observations which are caused by few dominant sources have much higher prior probability than sums of many small sources. Thus it is conjectured that, given an observation, this implies superfluous sources in the model will be ‘turned off’ after inference.

2.2. Location prior

The location prior is split up into independent coordinates:

$$[\mathbf{x}_m] = [x_m][y_m][z_m]. \quad (6)$$

Uniform distributions are chosen for x_m and y_m . Because a single array can hardly distinguish between sources at different ‘depth’, z_m is fixed to 0 using a Dirac delta distribution:

$$[z_m] = \delta(z_m). \quad (7)$$

This is satisfactory for our application, because sources (i.e. traffic and pedestrians) are typically close to the ground.

Because the same prior is assigned for each source, reordering the source indices m does not change the model, i.e. the model is symmetrical in m . Hence, the posteriors will necessarily also be the same for each m . Thus, before evaluating estimators \hat{s}_m and $\hat{\mathbf{x}}_m$, the ‘logical’ sources must be clustered into ‘physical’ sources. We do this using k-means (e.g. [18, Ch. 9]). Alternatively, a steering location can be implemented by limiting $[\mathbf{x}_1]$ to a small target region.

2.3. Likelihood

We assume spatially white, complex normal sensor noise:

$$[\mathbf{n}] = (\pi\sigma^2)^{-N} \exp(-\|\mathbf{n}\|^2/\sigma^2), \quad (8)$$

with known standard deviation σ and zero mean. Wind [19, 20] and diffuse ambient sources [21, Sec. 1.D] can also be modelled as (spatially low-pass) complex normal noise, but are neglected here to simplify presentation. Substituting \mathbf{n} using Eq. 1 gives the likelihood:

$$[\mathbf{o}|s_{1..M}, \mathbf{x}_{1..M}] = (\pi\sigma^2)^{-N} \exp(-\|\mathbf{o} - \sum_{m=1}^M \mathbf{a}_m(\mathbf{x}_m)s_m\|^2/\sigma^2). \quad (9)$$

Finally, the manifold vectors \mathbf{a}_{mn} are based on Green's function:

$$g(d) \propto \exp(ikd)/d [\text{Pa}], \quad (10)$$

where d [m] is the Euclidean distance to an acoustic monopole and k [rad/m] is the wavenumber [22, Sec. 6.5.1]. When expressing source strength as $s_m = g(1)$, the manifold between source m and sensor n , separated by distance d_{mn} , becomes:

$$a_{mn} = \exp(ikd_{mn})/d_{mn}. \quad (11)$$

The wavenumber is expressed in frequency f [Hz] and speed of sound c_0 [m/s] as:

$$k = 2\pi f/c_0, \quad (12)$$

where c_0 depends on temperature (12°C), relative humidity (50%) and CO2 concentration (mole fraction of $314 \cdot 10^{-6}$) according to the atmospherical model from [23, 24].

3. INFERENCE FOR A SINGLE SOURCE

When $M = 1$, applying Bayes' rule on Eq. 2 gives:

$$[s, \mathbf{x}|\mathbf{o}] \propto [s, \mathbf{x}, \mathbf{o}] = [\mathbf{o}|s, \mathbf{x}][s|\mathbf{x}], \quad (13)$$

with subscript $m = 1$ omitted for notational convenience. The goal is to sample from this posterior distribution.

3.1. Factorization of the posterior

We choose to factorize the posterior as:

$$[s, \mathbf{x}|\mathbf{o}] \propto \underbrace{[\angle s|s, \mathbf{x}, \mathbf{o}][|s|\mathbf{x}, \mathbf{o}]}_{\text{Signal}} \underbrace{[\mathbf{x}|\mathbf{o}]}_{\text{Location}}, \quad (14)$$

with (within the support of \mathbf{x} and s):

$$[\mathbf{x}|\mathbf{o}] \propto [\mathbf{o}|\mathbf{x}] \propto \int_{0^+}^{\infty} \exp((2-\zeta)\zeta q(\mathbf{x}, \mathbf{o})) \bar{I}_0(2\zeta q(\mathbf{x}, \mathbf{o}))/\zeta d\zeta \approx \exp(q(\mathbf{x}, \mathbf{o})), \quad (15)$$

$$[|s|\mathbf{x}, \mathbf{o}] \propto \exp\left(\left(2 - \frac{|s|}{p(\mathbf{x}, \mathbf{o})}\right) |s| \frac{q(\mathbf{x}, \mathbf{o})}{p(\mathbf{x}, \mathbf{o})}\right) \bar{I}_0\left(2|s| \frac{q(\mathbf{x}, \mathbf{o})}{p(\mathbf{x}, \mathbf{o})}\right) / |s|, \quad (16)$$

$$[\angle s|s, \mathbf{x}, \mathbf{o}] \propto \exp\left(2|s| \frac{q(\mathbf{x}, \mathbf{o})}{p(\mathbf{x}, \mathbf{o})} \cos(\angle s - \angle(\mathbf{a}(\mathbf{x})^* \mathbf{o}))\right), \quad (17)$$

with (scaled) modified Bessel function of the first kind:

$$\bar{I}_0(z) \stackrel{\text{def}}{=} \exp(-|\Re z|) I_0(z), \quad (18)$$

as implemented in Matlab as `besseli(0, z, 1)`, and:

$$p(\mathbf{x}, \mathbf{o}) \stackrel{\text{def}}{=} \frac{|\mathbf{a}(\mathbf{x})^* \mathbf{o}|}{|\mathbf{a}(\mathbf{x})^* \mathbf{a}(\mathbf{x})|}, \quad q(\mathbf{x}, \mathbf{o}) \stackrel{\text{def}}{=} p^2(\mathbf{x}, \mathbf{o}) \frac{\|\mathbf{a}(\mathbf{x})\|^2}{\sigma^2}. \quad (19)$$

3.2. Location sampling

During initialization, L candidates $\mathbf{x}(l)$ are drawn. $L_1 = \lfloor \gamma L \rfloor$ of these are generated using rejection sampling by accepting $\mathbf{x}(l)$ only if:

$$\min_{l'=1 \dots (l-1)} \|\mathbf{a}(\mathbf{x}(l)) - \mathbf{a}(\mathbf{x}(l'))\| > \varepsilon. \quad (20)$$

The purpose of this is to sample more densely in regions where the manifold changes quickly. The corresponding PDF $[\mathbf{x}]_{\text{pilot}}$ is interpolated from the histogram for (x, y) . The remaining $L_2 = L - L_1$ candidates are drawn from a uniform distribution (i.e. the true prior) to ensure the full support of \mathbf{x} is covered. Then, the proposal PDF for these candidates is:

$$[\mathbf{x}]_{\text{prop}} = \gamma [\mathbf{x}]_{\text{pilot}} + (1-\gamma) [\mathbf{x}]. \quad (21)$$

To sample $\mathbf{x}|\mathbf{o}$, a candidate is randomly drawn with weights based on Eq. 15:

$$w(k) = [\mathbf{o}|\mathbf{x}(k)]/[\mathbf{x}(k)]_{\text{prop}}. \quad (22)$$

The denominator compensates for mismatch between true prior $[\mathbf{x}]$ and proposal $[\mathbf{x}]_{\text{prop}}$. This *importance sampling* [12] scheme allows for smaller L and hence less computational cost. The most expensive step is calculation of L N -dimensional vector in-products for each observation \mathbf{o} .

3.3. Signal sampling

Above sample for $\mathbf{x}|\mathbf{o}$ is then plugged into $[|s|\mathbf{x}, \mathbf{o}]$ (Eq. 16). Because this PDF is a smooth function of one variable, it can be efficiently represented using Chebyshev polynomials up to machine precision [25]. By integrating the polynomial to find the cumulative distribution function (CDF), a sample for $|s|\mathbf{x}, \mathbf{o}$ is drawn using *inverse transform sampling* (described in [13], implemented using [26]). This process is repeated to sample $\angle s|s, \mathbf{x}, \mathbf{y}$ using Eq. 17.

4. INFERENCE FOR MULTIPLE SOURCES

For $M > 1$, we opt for Gibbs sampling [10] which amounts to iteratively sampling a single source m while keeping all other sources $m' \neq m$ fixed, that is, from:

$$[s_m, \mathbf{x}_m | s_{1 \dots M \setminus m}, \mathbf{x}_{1 \dots M \setminus m}, \mathbf{o}] = [s, \mathbf{x} | \mathbf{o}'], \quad (23)$$

which is Eq. 14 with observation \mathbf{o} replaced with *residual*:

$$\mathbf{o}' \stackrel{\text{def}}{=} \mathbf{o} - \sum_{\substack{m'=1 \\ m' \neq m}}^M s_{m'} \mathbf{a}(\mathbf{x}_{m'}). \quad (24)$$

Note the similarity with deconvolution beamforming (e.g. [11, Eq. 24]). In the first iteration (initialization: $s_{m'} := 0$) the sources are sampled in order of decreasing observed power, because Eq. 15 is sampled globally (Sec. 3.2) and $q(\mathbf{x}, \mathbf{o})$ can be interpreted as observed SNR. This property prevents our sampler from getting stuck in local bassins of near-zero source powers and, as a consequence, wrong source locations. We did observe this behaviour in alternative sampling methods such as PyMC3's No-U-Turn Sampler [27, 28]. Subsequent iterations explore alternative explanations of \mathbf{o} and hence account for observation uncertainty.

Due to its iterative nature, the Gibbs sampler rarely swaps m -indices of two strong sources. To ensure symmetry in m of the posteriors (Sec. 2.2), each posterior includes all samples regardless of m . Then, m is used to initialize the (k-means) clusters to extract the 'physical' sources.

5. RESULTS

Two experiments illustrate behaviour of our method. Each time, $500M$ samples were drawn with $L = 3000$, $f = 500$ kHz, $\gamma = 0.3$, $\varepsilon = 0.1$, and sensor array from Fig. 1. True source power was 0 dB.

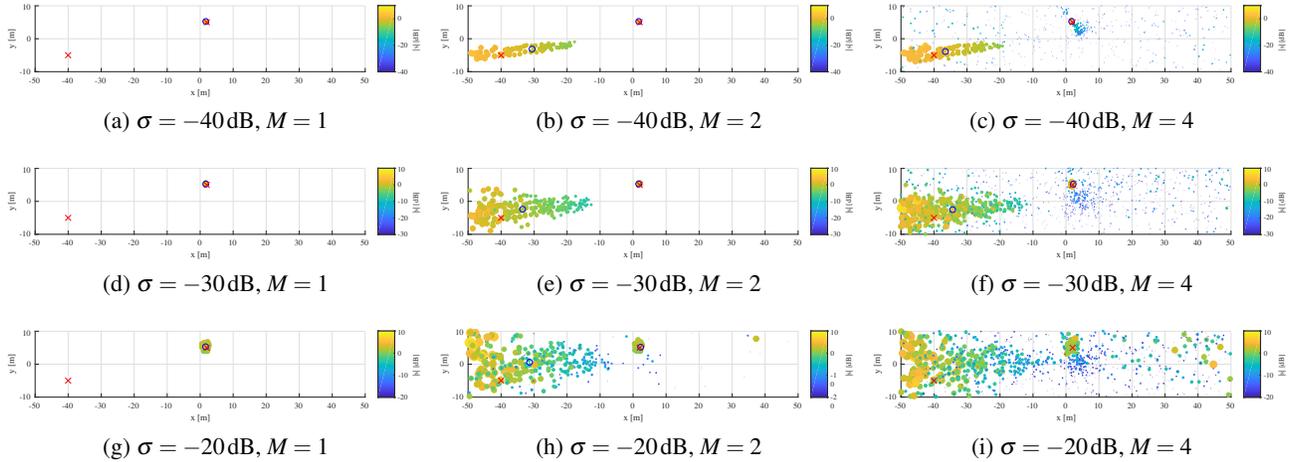


Fig. 2. Beamforming results for varying M and σ . Source locations are denoted with \times (true) and \circ (estimated; omitted if clustering failed).

First, noise level (σ) and number of model sources (M) were varied, while keeping source power and location fixed (Fig. 2). When $M = 1$, only the source with largest observed power is detected (Fig. 2a, d, g) due to the decreasing power property (Sec. 4). When $M = 4$, two model sources are ‘turned off’ (Fig. 2c, f, i) due to the shrinking signal prior (Sec. 2.1). Posterior location then becomes flat, with slight preference near the array (due to a local basin, Sec. 4). Thus, the beamformer automatically infers the number of active sources.

In Fig. 2g-i, observed SNR is negative. However, by combining sensor signals, information could still be recovered. Some blown-up samples appear in Fig. 2i near $x = 40$ m, but they are compensated by a large quantity of (less visible) zero-power samples. Thus, the beamformer can compensate for attenuation without creating ‘ghost sources’, unlike delay-and-sum beamforming (Fig. 1c).

Second, robustness was investigated. An observation was generated using perturbed sensor locations with normally distributed 3D offsets, while sensor locations in the model were unchanged. To eliminate influence of clustering errors from the experiment, SNR was high and only one source was active. Although error increases with distance from the array due to increasing attenuation and decreasing angular resolution (i.e. Abbe diffraction limit [29, sec. 12.2.5]), the beamformer is robust for sensor deviations of up to about $1/5$ wavelength (Fig. 3).

6. CONCLUSION

The experiments support our claims of a beamformer that is robust against perturbations in the manifold, can deal with sources at different distances and power levels, and does not require the number of active sources to be known a priori. The presented algorithm can be interpreted as probabilistic deconvolution beamforming. In this work we have compared our method with classical delay-and-sum beamforming. Future work will include a comparison with more literature results cf. [30].

An inherent weakness of our model is the requirement for clustering to extract point estimators from the visual results. This tends to fail in low SNR conditions or when sources are close together. One direction for future research is exploiting multiple snapshots over time, so that location prior symmetry can be eliminated by exploiting temporal correlations.

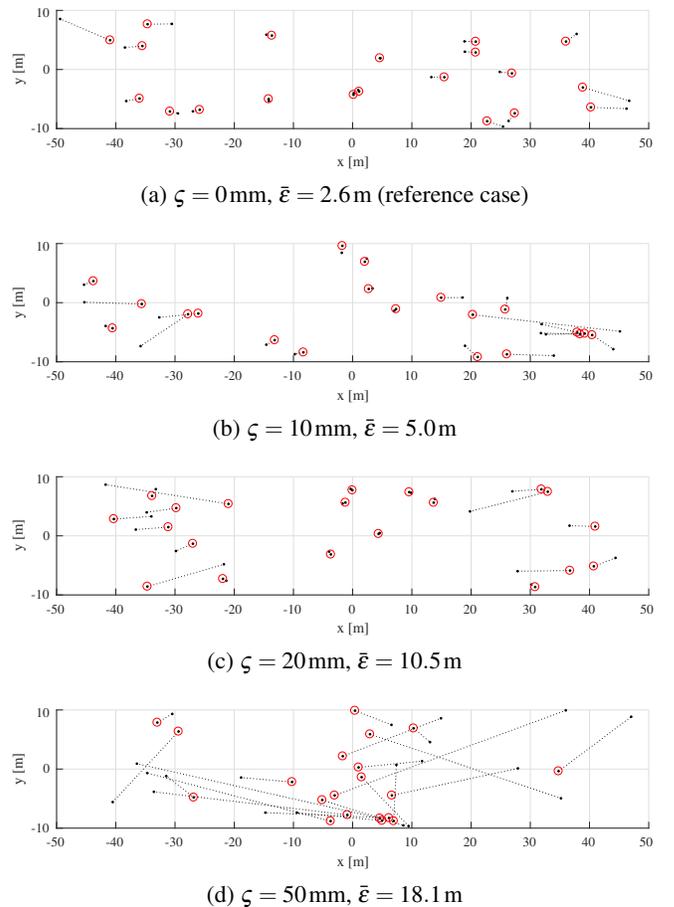


Fig. 3. Differences between true source location (\bullet) and nearest k-means estimate (\circ) for 20 random trials with perturbed sensor locations. ζ is the perturbation std. dev. in each dimension. $\bar{\epsilon}$ is average error distance. Wavelength $\lambda = 113$ mm. $\sigma = -40$ dB, $M = 2$.

7. REFERENCES

- [1] Oscar Martínez-Graullera, Carlos J. Martín, Gregorio Godoy, and Luis G. Ullate, “2D array design based on Fermat spiral for ultrasound imaging,” *Ultrasonics*, vol. 50, no. 2, pp. 280–289, Feb. 2010.
- [2] Jon Cropley, “Top Video Surveillance Trends for 2016,” White Paper, IHS, 2016.
- [3] B. D. van Veen and K. M. Buckley, “Beamforming: a versatile approach to spatial filtering,” *IEEE ASSP Magazine*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
- [4] P. A. Gauthier, C. Camier, Y. Pasco, A. Berry, E. Chambatte, R. Lapointe, and M. A. Delalay, “Beamforming regularization matrix and inverse problems applied to sound field measurement and extrapolation using microphone array,” *Journal of Sound and Vibration*, vol. 330, no. 24, pp. 5852–5877, Nov. 2011.
- [5] J. Capon, “High-resolution frequency-wavenumber spectrum analysis,” *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, Aug. 1969.
- [6] L. Rugini and P. Banelli, “On the Equivalence of Maximum SNR and MMSE Estimation: Applications to Additive Non-Gaussian Channels and Quantized Observations,” *IEEE Transactions on Signal Processing*, vol. 64, no. 23, pp. 6190–6199, Dec. 2016.
- [7] Xun Wang, Benjamin Quost, Jean-Daniel Chazot, and Jérôme Antoni, “Iterative beamforming for identification of multiple broadband sound sources,” *Journal of Sound and Vibration*, vol. 365, pp. 260–275, Mar. 2016.
- [8] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum Likelihood from Incomplete Data via the EM Algorithm,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.
- [9] Olivier Besson and Stéphanie Bidon, “Robust adaptive beamforming using a Bayesian steering vector error model,” *Signal Processing*, vol. 93, no. 12, pp. 3290–3299, Dec. 2013.
- [10] S. Geman and D. Geman, “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
- [11] Thomas F. Brooks and William M. Humphreys, “A deconvolution approach for the mapping of acoustic sources (DAMAS) determined from phased microphone arrays,” *Journal of Sound and Vibration*, vol. 294, no. 4, pp. 856–879, July 2006.
- [12] Herman Kahn, “Use of Different Monte Carlo Sampling Techniques,” Nov. 1955.
- [13] Sheehan Olver and Alex Townsend, “Fast inverse transform sampling in one and two dimensions,” *arXiv:1307.1223 [math, stat]*, July 2013, arXiv: 1307.1223.
- [14] David Wipf and Srikantan Nagarajan, “Beamforming Using the Relevance Vector Machine,” in *Proceedings of the 24th International Conference on Machine Learning*, New York, NY, USA, 2007, ICML ’07, pp. 1023–1030, ACM.
- [15] E. T. Jaynes, *Probability theory: the logic of science*, Cambridge University Press, Cambridge, UK ; New York, NY, 2003.
- [16] Carlos M. Carvalho, Nicholas G. Polson, and James G. Scott, “The horseshoe estimator for sparse signals,” *Biometrika*, vol. 97, no. 2, pp. 465–480, 2010.
- [17] Stéphanie van der Pas, Jean-Bernard Salomond, and Johannes Schmidt-Hieber, “Conditions for Posterior Contraction in the Sparse Normal Means Problem,” *Electronic Journal of Statistics*, vol. 10, no. 1, pp. 976–1000, 2016, arXiv: 1510.02232.
- [18] Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer-Verlag New York, Inc., 2006.
- [19] C. M. Nelke, N. Chatlani, C. Beaugeant, and P. Vary, “Single microphone wind noise PSD estimation using signal centroids,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 7063–7067.
- [20] F. Douglas Shields, “Low-frequency wind noise correlation in microphone arrays,” *The Journal of the Acoustical Society of America*, vol. 117, no. 6, pp. 3489–3496, May 2005.
- [21] Patrick W.A. Wijnings, “Near-Field Acoustical Holography: Regularization and Calibration in a Bayesian framework,” M.S. thesis, Eindhoven University of Technology, Eindhoven, July 2015.
- [22] Earl G. Williams, *Fourier acoustics: sound radiation and nearfield acoustical holography*, Academic Press, San Diego, Calif., 1999, OCLC: 634091155.
- [23] Owen Cramer, “The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity, and CO2 concentration,” *The Journal of the Acoustical Society of America*, vol. 93, no. 5, pp. 2510–2516, May 1993.
- [24] P. Giacomo, “Equation for the Determination of the Density of Moist Air,” *Metrologia*, vol. 18, no. 1, pp. 33, 1982.
- [25] Lloyd N. Trefethen, *Approximation theory and approximation practice*, Applied mathematics. Society for Industrial and Applied Mathematics, Philadelphia, 2013.
- [26] T. A Driscoll, N. Hale, and L. N. Trefethen, *Chebfun Guide*, Pafnuty Publications, 2014.
- [27] Matthew D. Hoffman and Andrew Gelman, “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo,” *arXiv:1111.4246 [cs, stat]*, Nov. 2011, arXiv: 1111.4246.
- [28] John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck, “Probabilistic programming in Python using PyMC3,” *PeerJ Computer Science*, vol. 2, pp. e55, Apr. 2016.
- [29] A. Lipson, S. G. Lipson, and H. Lipson, *Optical physics*, Cambridge University Press, Cambridge ; New York, 4th ed edition, 2011, OCLC: ocn637708967.
- [30] Lin Du, Tarik Yardibi, Jian Li, and Petre Stoica, “Review of user parameter-free robust adaptive beamforming algorithms,” *Digital Signal Processing*, vol. 19, no. 4, pp. 567–582, July 2009.