

MAXIMUM LIKELIHOOD ESTIMATION OF THE DIRECTION OF SOUND IN A REVERBERANT NOISY ENVIRONMENT

Mohamed F. Mansour

Amazon Inc., USA

ABSTRACT

We describe a new method for estimating the direction of sound in a reverberant environment from basic principles of sound propagation. The method utilizes SNR-adaptive features from time-delay and energy of the directional components after acoustic wave decomposition of the observed sound field to estimate the line-of-sight direction under noisy and reverberant conditions. The effectiveness of the approach is established with measured data of different microphone array configurations under various usage scenarios.

Index Terms— Maximum-Likelihood Estimation, Direction of Arrival, Reverberation, Room Acoustics, Wave Decomposition.

1. INTRODUCTION

Computing the direction of arrival (DoA) of a sound source is a classical estimation problem that is essential for sound source localization. It has many applications in robotics and speech communication systems [1, 2], and has become increasingly important with the proliferation of voice-controlled smart systems [3]. Many techniques have been proposed to address the problem including, beamforming [4, 5], subspace methods, e.g., MUSIC and ESPRIT [6, 7], time delay methods, e.g., GCC-PHAT and SRP-PHAT [8, 9], and more recently DNN methods [10, 11]. In the absence of strong reverberation/interference, existing techniques generally provide satisfactory results, and this has been studied extensively in the literature. However, a commercial-grade embedded system for sound localization, with constraints on computation/latency/memory, requires consistent localization performance under adverse reverberation and noise conditions, and this is the subject of this work.

The fundamental problem in computing the DoA of a sound source in a reverberant and noisy environment is to distinguish the line-of-sight component of the target sound source from all interfering directional components in the presence of incoherent sensor noise. These interfering directional components include acoustic reflections of the target sound source, as well as all directional components of coherent noise interference. Hence, all solutions to the DoA problem aim at finding a proper characterization of the line-of-sight component based on either a physical model or a data-driven model. Signal processing solutions to the problem, e.g., SRP-PHAT and MUSIC algorithms deploy a channel model for acoustic propagation and source/noise statistics, where it is generally assumed that the direct path component is on average stronger than acoustic reflections across a range of frequencies of interest. The direct path component is computed implicitly using inter-microphone information, e.g., generalized cross-correlation. These approaches frequently fail to accurately capture common cases, e.g., when there is a strong room reverberation and the microphone array is placed at a corner far from the source. This problem is more apparent with small microphone arrays, e.g., ≤ 4 microphones, where due to the coarse sensing resolution, the line-of-sight component might

be perceived as weaker than some reflections. Moreover, the problem gets more complicated in the presence of coherent interference with speech-like content, e.g., from TV. To resolve some of these issues, data-driven approaches that deploy variations of deep neural networks (DNN) were introduced in recent years, with the assumption that training data captures all relevant use cases. These approaches showed improvement (sometimes significant) over classical approaches on the test datasets [11] especially under noisy conditions. However, as noted in [12] that all these solutions can only work well when distance between source and microphone array is small, which is a limitation for commercial adoption. Further, this approach is not scalable to accommodate different microphone array geometries, as the training data is dependent on the microphone array and it should capture a huge number of cases that cover all usage scenarios at different kinds of rooms, noise, and sound stimuli. Unlike training of speech models for ASR, synthetic models, e.g., image source method [13], cannot replace data collections because the learning objective is the model itself and they are parameterized by a relatively small number of parameters that can be learned by the DNN model.

The work presented in this paper provides a new methodology for computing the sound direction that is based on directional decomposition of microphone array observations. The microphone array observations are mapped to directional component via acoustic wave decomposition, and these directional components are processed to compute the sound direction. The multidimensional representation of the spatial signal with directional components provides an intuitive characterization of the line-of-sight component of the sound source based on principles of acoustic propagation, which was not explored in earlier works. It utilizes a generalized acoustic propagation model that accommodates total acoustic pressure due to scattering at the mounting surface. This physical characterization is utilized to construct a statistical framework to derive the maximum-likelihood estimator of the direction of arrival. The mapping to directional components is empowered by the work in [14, 15], where a method for generalized acoustic wave decomposition of a microphone array of arbitrary geometry was described. It does not require a special microphone array geometry as in related localization work with spherical harmonics [16]. The proposed system is suited for embedded implementation and it is scalable to accommodate different microphone array geometries with minimal tuning effort. The discussion in this work is limited single-source localization. It is shown in section 4 that the proposed algorithm outperforms existing baseline solutions in mitigating large localization errors when evaluated on a large corpus of real data under diverse room conditions and different microphone array size and geometry.

2. PROBLEM DEFINITION

The underlying physical model of the estimation problem is the generalized Acoustic Wave Decomposition (AWD) as described in

[14, 15], where the observed sound field, $\mathbf{p}(\omega; t)$, at the microphone array is expressed as

$$\mathbf{p}(\omega; t) = \sum_l \alpha_l(\omega; t) \boldsymbol{\psi}(\omega; \theta_l(t), \phi_l(t)) \quad (1)$$

where θ_l and ϕ_l denote respectively the elevation and azimuth (in polar coordinates) of the direction of propagation of the l -th acoustic wave, and $\boldsymbol{\psi}(\omega; \theta, \phi)$ denotes the total acoustic pressure at the microphone array when a free-field acoustic plane wave with direction (θ, ϕ) impinges on the device. The total acoustic pressure is the superposition of the incident free-field plane wave and the scattered component at the device surface. At each ω , $\boldsymbol{\psi}(\omega; \cdot)$ is a vector whose length equals the number of microphones in the microphone array. The ensemble of all vectors that span the three-dimensional space at all ω defines the *acoustic dictionary* of the device, and it is computed offline with standard acoustic simulation techniques [14, 15]. Note that, even though the elevation is not reported in the direction of sound, it is important to include it in the signal model, as acoustic waves with the same azimuth but different elevation might have different impact at the microphone array when surface scattering is accommodated.

This model generalizes the free-field plane wave decomposition to accommodate scattering at the device surface which is modeled as a hard boundary. This scattering component partially resolves spatial aliasing due to phase ambiguity at high frequencies. Each directional component of the acoustic wave expansion in (1) at frame t is characterized by its direction $(\theta_l(t), \phi_l(t))$ (which are frequency-independent), and the corresponding complex-valued weight $\alpha_l(\omega; t)$.

The objective of sound direction estimation is to compute the azimuth angle $\hat{\phi}$ that corresponds to the line-of-sight direction of the sound source, given the observed sound field $\{\mathbf{p}(\omega; t)\}_{\omega, t}$ at successive time frames that span the duration of the source signal. In this work, we assume only a single sound source, though background coherent or incoherent noise can be present.

In the absence of other sound sources, the line-of-sight component of a sound source is usually contaminated by other directional components due to acoustic reflections at nearby surfaces, as well as incoherent noise at the microphone array. Nevertheless, the line-of-sight component has two distinct features:

1. The energy of the line-of-sight component of a sound source is higher than the energy of any of its individual reflections.
2. The line-of-sight component of a sound source arrives first at the microphone array before any other reflection of the same sound source.

The following section describes a statistical framework that utilizes these two features to design a maximum-likelihood estimator of the sound direction

3. MAXIMUM-LIKELIHOOD ESTIMATION

3.1. Estimation Procedure

The estimation procedure exploits the true direction features as described in the previous section to compute the maximum-likelihood estimate of the user direction. It computes from microphone array observations two likelihood functions for the time delay and the signal energy; then applies late fusion to compute the total likelihood at each time frame. The two likelihood functions are computed from the directional components in (1) at each time frame. Finally, the

total likelihood values at different frames are smoothed over the duration of the sound signal to produce the aggregate likelihood function that is used to find the maximum-likelihood estimate. Hence, the estimation flow is as follows:

1. At each time step t , process the observed sound field, $\{\mathbf{p}(\omega; t)\}_{\omega}$, as follows
 - (a) Compute the acoustic wave decomposition at all ω in (1) using the multistage solver as in [15].
 - (b) Compute the time-delay likelihood function of each directional component (as described in section 3.2).
 - (c) Compute the energy-based likelihood function of each directional component (as described in section 3.3).
 - (d) Combine the two likelihood functions on a one-dimensional grid of possible azimuth angles (as described in section 3.4).
2. Compute the aggregate likelihood of each angle candidate over the whole signal period, and choose the angle that corresponds to the maximum-likelihood value as the sound direction estimate (as outlined in section 3.4).

3.2. Delay-based Likelihood

Assume that a source signal, $X(\omega)$, experiences multiple reflections in the acoustic path towards a microphone array. Denote the l -th reflection at a receiving microphone by $X_l(\omega)$, which can be expressed as

$$X_l(\omega) \approx e^{-j\omega\tau_l} \delta_l(\omega) X(\omega) \quad (2)$$

where $\tau_l > 0$ is the corresponding delay, and δ_l is a real-valued propagation loss. Note that, $X_l(\omega)$ refers to the l -th AWD component in (1) due to sound source $X(\omega)$. To eliminate the nuisance parameters δ_l and $X(\omega)$, we introduce the auxiliary parameter

$$\begin{aligned} Q_{lk}(\omega) &\triangleq \frac{X_l(\omega)/\|X_l(\omega)\|}{X_k(\omega)/\|X_k(\omega)\|} \\ &\approx e^{j\omega(\tau_k - \tau_l)} \end{aligned} \quad (3)$$

This parameter is utilized to find the time delay between the two components. However, it is susceptible to phase wrapping at large ω , and one extra step is needed to mitigate its impact. Define for a frequency shift Δ

$$\begin{aligned} R_{lk}(\omega) &\triangleq \frac{Q_{lk}(\omega + \Delta)}{Q_{lk}(\omega)} \\ &\approx e^{j\Delta(\tau_k - \tau_l)} \end{aligned} \quad (4)$$

which eliminates the dependence on ω , and if Δ is chosen small enough, then phase wrapping is eliminated. Then, the estimated delay between components l and k , $\bar{\rho}_{lk} \triangleq \tau_k - \tau_l$, is computed as

$$\bar{\rho}_{lk} = \frac{1}{\sum_{\omega \in \Omega} W(\omega)} \sum_{\omega \in \Omega} W(\omega) \frac{\angle R_{lk}(\omega)}{\Delta} \quad (5)$$

where $\angle R$ denotes the angle of R , and $W(\omega)$ is a sigmoid weighting function that depends on SNR at ω . Note that, the procedure does not require computing the inverse FFT as in common generalized cross-correlation schemes [17], which significantly reduces the overall complexity. If $\bar{\rho}_{lk} > 0$, then the k -th reflection is delayed from

the l -th reflection and vice versa. Hence, the probability that the k -th component is delayed from the l -th component is $P(\rho_{lk} > 0)$ (where ρ_{lk} is the true value of $\tau_k - \tau_l$). If $\rho_{lk} \sim \mathcal{N}(\bar{\rho}_{lk}, \sigma^2)$, then

$$P(\tau_k > \tau_l) \equiv P(\rho_{lk} > 0) \approx \frac{1}{2} \operatorname{erfc}\left(\frac{-\bar{\rho}_{lk}}{\sigma\sqrt{2}}\right) \quad (6)$$

where $\operatorname{erfc}(\cdot)$ is the complementary error function. Note that, $P(\tau_k < \tau_l) = 1 - P(\tau_k > \tau_l)$, hence, $\bar{\rho}_{lk}$ is computed once for each pair of components.

The acoustic reflections $\{X_l(\omega)\}$ are approximated by $\{\alpha_l(\omega; t)\}_\omega$ in (1). Denote the probability that the l -th component is the first to arrive at the microphone array by β_l , which can be expressed as

$$\begin{aligned} \beta_l &\triangleq P(\tau_l < \min\{\tau_k\}_{k \neq l}) \\ &= \prod_{k \neq l} P(\rho_{lk} > 0) \end{aligned} \quad (7)$$

which, using (6), can be expressed in the log-domain as

$$\bar{\beta}_l \approx \sum_{k \neq l} \log\left(\operatorname{erfc}\left(\frac{-\bar{\rho}_{lk}}{\sigma\sqrt{2}}\right)\right) \quad (8)$$

This is a good approximation of the time-delay likelihood function as long as the AWD components follow the signal model in (2). A simple test to validate this assumption is to compute the pair-wise correlation coefficient between components, and run the computation only if it is above a predetermined threshold.

3.3. Energy-based Likelihood

The true energy of the line-of-sight component is theoretically higher than the energy of each individual reflection. However, due to the finite number of microphones, the true directional component might be diluted in the AWD computation. Nevertheless, the line-of-sight energy is usually among the highest energy components. The energy of each component is computed as

$$E_l = \sum_{\omega} \|\alpha_l(\omega; t)\|^2 W(\omega) \quad (9)$$

where $W(\omega)$ is a weighting function as in (5). An AWD component is a candidate to be the line-of-sight component if $E_l > \nu \max\{E_k\}$, where ν is a predetermined threshold. All directional components above the energy threshold are considered equally likely to be the line-of-sight component. Hence, if the number of AWD components that satisfy this condition is M , then the energy-based likelihood is computed as

$$\gamma_l = \begin{cases} -\log M & \text{if } E_l > \nu \max\{E_k\} \\ \varepsilon & \text{otherwise} \end{cases} \quad (10)$$

where $\varepsilon \ll -\log M$ corresponds to a small probability value to account for measurement/computation errors.

3.4. Aggregate Likelihood

At each time frame, the log-likelihoods $\bar{\beta}_l$ and γ_l are computed for the l -th AWD component as in (8) and (10) respectively. This corresponds to azimuth angle ϕ_l of the corresponding entry of the device dictionary, and the total likelihood at ϕ_l is the sum of the two likelihoods. Due to the finite dictionary size and the finite precision of the computation, the true angle of the l -th component can be an angle

adjacent to ϕ_l . If we assume normal distribution (with variance κ) of the true azimuth angle around ϕ_l , then the likelihood for azimuth angles adjacent to ϕ_l is approximated as

$$\chi(\phi; t) = \max\left(\chi(\phi; t); \beta_l + \gamma_l - \frac{(\phi - \phi_l)^2}{2\kappa}\right). \quad (11)$$

and the likelihood function, $\chi(\cdot)$, of all azimuth angles is updated according to (11) with every new AWD component. Note that, the azimuth likelihood is smoothed with the azimuth angle of each AWD component, ϕ_l , whereas the elevation component, θ_l , is treated as nuisance parameter that is averaged out. If joint estimation of azimuth and elevation is required, then a two-dimensional likelihood function $\chi(\theta, \phi; t)$ is utilized rather than the one-dimensional likelihood in (11).

The final step is to compute the maximum-likelihood estimate of the azimuth angle by aggregating the local likelihood values in (11) over the duration of the sound event. The final likelihood aggregates the local likelihood at different time frames after proper weighting by the total SNR at each frame.

$$\bar{\chi}(\phi) = \sum_t \chi(\phi; t) \eta(t) \quad (12)$$

where $\eta(t)$ is a weighting sigmoid function that is proportional to the total SNR at frame t . This temporal weighting is necessary to mitigate errors in the sound event time boundaries. The maximum-likelihood estimate of the angle of arrival is computed from (12) as

$$\hat{\phi} = \operatorname{argmax} \bar{\chi}(\phi) \quad (13)$$

3.5. Discussion

The signal measurement model for the maximum-likelihood estimation is the general physical model in (1), and the properties of the line-of-sight component as described in section 2. The formulation of a statistical model, in the form of the combined likelihood function in (11) from this physical model is the key contribution of this work.

The aggregate likelihood function combines time-delay likelihood and energy likelihood to capture the physical properties of the line-of-sight component. These likelihood computations utilize the directional components from the acoustic wave decomposition in (1) as described in [15]. The reverberation impact is mitigated by incorporating the time-delay component, while the noise impact is mitigated by incorporated SNR-dependent weighting. The algorithm is fundamentally different from existing model-based algorithms in few aspects:

- Incorporating both time-delay and energy to compute the direction of sound.
- Incorporating magnitude component in steering vectors to reduce spatial aliasing.
- Utilizing sparse techniques to find relevant directions in 3D space, rather than deploying exhaustive beamforming that unavoidably has spatial leakage from adjacent directions,
- It scales properly with minimal tuning to other microphone array geometries through the acoustic dictionary.

4. EXPERIMENTAL EVALUATION

The proposed algorithm is evaluated using two different microphone arrays. The first microphone array is a circular array with 8 microphones mounted atop a cylindrical surface. The second microphone

array is a star-shaped 3D array with 4 microphones that are mounted on an LCD screen. The geometry and mounting surface of the two arrays are quite different to illustrate the generality of the proposed method. The test dataset has approximately 55k utterances recorded in different rooms at different SNR levels. The dataset covers all angles around the microphone array, and covers possible placements of the microphone array inside a room, i.e., center, wall, and corner. For each microphone array, the device acoustic dictionary is computed offline as described in section 2. The noise was recorded at the microphone array separately and added to clean speech for evaluation. It covers a wide set of household noises, e.g., fan, vacuum, TV, microwave, ... etc.

Fig. 1 shows the average performance of the proposed algorithm at different SNR values. In both microphone array cases, the mean absolute error is around 6° at high SNR and it degrades gracefully with lower SNR. The 8-mic configuration provides 5 to 9 dB advantage over the 4-mic configuration depending on the operating SNR. Note that, the proposed algorithm does not explicitly deploy a denoising mechanism. Rather, it deploys a noise mitigation mechanism through the SNR-dependent weighting as discussed in section 3. The performance at low SNR can be improved with a denoising procedure prior to estimation but this is outside the scope of this work.




Fig. 1. Mean Absolute Error in degrees of the proposed algorithm for microphone arrays of size 8 and 4

In Fig. 2, the cumulative density function (CDF) of the absolute error for the proposed algorithm is shown. It is compared to the CDF of the SRP-PHAT and state-of-the-art DNN solution. In both cases, the compared algorithms are fully tuned to the respective microphone array by subject matter experts for the respective hardware. The DNN solution utilizes an enhanced implementation of CRNN-SSL algorithm in [18] with few architecture changes to match state of the art performance. It is trained with data from the device with 4-mic microphone array at numerous room configurations with a combination of synthetic and real-data. The SRP-PHAT solution utilizes heuristics to increase robustness to strong reflections and interfering noise. As shown in the figure, the proposed algorithm provides improvement in both cases especially for the high-error case, which corresponds to low-SNR cases and cases with strong reverberation. For example, both the 90-percentile and the 95-percentile errors are reduced by more than 50% as compared to SRP-PHAT as illustrated in Fig. 2a. Hence, the proposed algorithm is more effective in mitigating large estimation errors, which usually have big negative impact on the user experience.

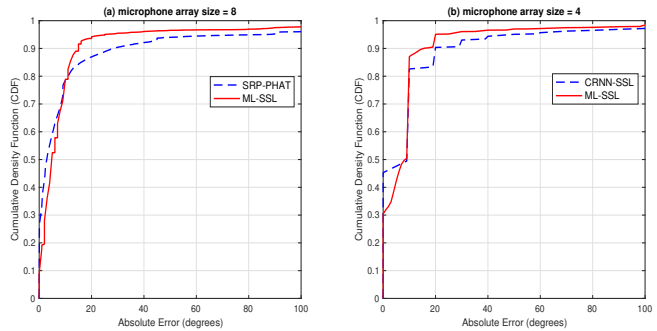


Fig. 2. Cumulative density function of the absolute error for proposed algorithm vs (a.) SRP-PHAT with 8-mic, (b) CRNN-SSL with 4-mic

5. CONCLUSION

The proposed algorithm addresses the two fundamental problems in computing sound source direction, namely reverberation and noise interference. It is founded on a rigorous and general physical model for sound propagation, which is mapped to a statistical model that is solved by standard estimation techniques to compute the maximum-likelihood estimation. The proposed algorithm is shown to outperform existing solutions in the literature when evaluated with a large dataset of real data.

Further, the proposed algorithm has two practically important advantages over prior art:

1. It is agnostic to the geometry of the microphone array and mounting surface because the input to the estimation procedure is the directional components after wave decomposition rather than microphone array observations. The device dependent part is captured in the device acoustic dictionary, which does not contribute to the algorithm hyper parameters. This enhances scalability and reduces migration effort to new hardware designs.
2. It generalizes the acoustic propagation model to accommodate scattering at the device surface. This scattering is viewed as distortion if free-field propagation model is utilized, whereas it is leveraged in the proposed system to enhance estimation. The incorporation of the magnitude components, due to scattering, in addition to phase components enhances robustness to spatial aliasing.

The proposed algorithm does not deploy a noise enhancement procedure prior to estimation. A multichannel signal enhancement system can improve the performance at low SNR if it preserves the coherence between microphones, and this is a subject of future work. Future work also includes utilizing of directional components for joint source localization and separation.

6. REFERENCES

- [1] S. Argentieri, P. Danes, and P. Souères, “A survey on sound source localization in robotics: From binaural to array processing methods,” *Computer Speech & Language*, vol. 34, no. 1, pp. 87–112, 2015.
- [2] C. Rascon and I. Meza, “Localization of sound sources in robotics: A review,” *Robotics and Autonomous Systems*, vol. 96, pp. 184–210, 2017.

- [3] A. Chhetri, P. Hilmes, T. Kristjansson, W. Chu, M. Mansour, X. Li, and X. Zhang, "Multichannel Audio Front-End for Far-Field Automatic Speech Recognition," in *2018 European Signal Processing Conference (EUSIPCO)*, 2018, pp. 1527–1531.
- [4] J. Dmochowski, J. Benesty, and S. Affes, "A generalized steered response power method for computationally viable source localization," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2510–2526, 2007.
- [5] J. Daniel and S. Kitić, "Time domain velocity vector for retracing the multipath propagation," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 421–425.
- [6] S. Argentieri and P. Danes, "Broadband variations of the music high-resolution method for sound source localization in robotics," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2007, pp. 2009–2014.
- [7] A. Hogg, V. Neo, S. Weiss, C. Evers, and P. Naylor, "A polynomial eigenvalue decomposition music approach for broadband sound source localization," in *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2021, pp. 326–330.
- [8] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE transactions on acoustics, speech, and signal processing*, vol. 24, no. 4, pp. 320–327, 1976.
- [9] J. DiBiase, H. Silverman, and M. Brandstein, "Robust localization in reverberant rooms," in *Microphone arrays*. Springer, 2001, pp. 157–180.
- [10] N. Yalta, K. Nakadai, and T. Ogata, "Sound source localization using deep learning models," *Journal of Robotics and Mechatronics*, vol. 29, no. 1, pp. 37–48, 2017.
- [11] P. Grumiaux, S. Kitić, L. Girin, and A. Guérin, "A survey of sound source localization with deep learning methods," *The Journal of the Acoustical Society of America*, vol. 152, no. 1, pp. 107–151, 2022.
- [12] Y. Wu, R. Ayyalasomayajula, M. Bianco, D. Bharadia, and P. Gerstoft, "Sslide: Sound source localization for indoors based on deep learning," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 4680–4684.
- [13] J. Allen and D. Berkley, "Image method for efficiently simulating small-room acoustics," *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, 1979.
- [14] A. Chhetri, M. Mansour, W. Kim, and G. Pan, "On acoustic modeling for broadband beamforming," in *27th European Signal Processing Conference (EUSIPCO)*, 2019, pp. 1–5.
- [15] M. Mansour, "Sparse recovery of acoustic waves," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 5418–5422.
- [16] D. Jarrett, E. Habets, and P. Naylor, "3d source localization in the spherical harmonic domain using a pseudointensity vector," in *2010 18th European Signal Processing Conference*. IEEE, 2010, pp. 442–446.
- [17] A. Brutti, M. Omologo, and P. Svaizer, "Comparison between different sound source localization techniques based on a real data collection," in *2008 hands-free speech communication and microphone arrays*. IEEE, 2008, pp. 69–72.
- [18] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 34–48, 2018.