A Comparison and Combination of Unsupervised Blind Source Separation Techniques

Christoph Boeddeker, Frederik Rautenberg, Reinhold Haeb-Umbach

Paderborn University, Department of Communications Engineering, Paderborn, Germany
Email: {boeddeker,haeb}@nt.upb.de, frra@campus.upb.de

Abstract

Unsupervised blind source separation methods do not require a training phase and thus cannot suffer from a train-test mismatch, which is a common concern in neural network based source separation. The unsupervised techniques can be categorized in two classes, those building upon the sparsity of speech in the Short-Time Fourier transform domain and those exploiting non-Gaussianity or non-stationarity of the source signals. In this contribution, spatial mixture models which fall in the first category and independent vector analysis (IVA) as a representative of the second category are compared w.r.t. their separation performance and the performance of a downstream speech recognizer on a reverberant dataset of reasonable size. Furthermore, we introduce a serial concatenation of the two, where the result of the mixture model serves as initialization of IVA, which achieves significantly better WER performance than each algorithm individually and even approaches the performance of a much more complex neural network based technique.

1 Introduction

In the field of speech signal processing, Blind Source Separation (BSS) is concerned with separating a mixture of speech signals into the contributions of the individual speakers. Formulating BSS as a supervised learning problem, neural networks have excelled in this task, even if the mixture signal is recorded by a single microphone only. However, their performance heavily depends on the absence of a mismatch between the, typically artificially mixed, training data and the test scenario. Furthermore, neural network based single channel source separation as of today breaks down in the presence of reverberation.

On the other hand, signal processing based approaches, such as Spatial Mixture Models (SMMs) [1, 2] and the class of Independent Component Analysis (ICA) [3–7] based algorithms, are unsupervised methods, that do not require a training phase and thus cannot suffer from a train-test mismatch. These are multi-channel techniques and good separation performance has been reported even if the input signals are reverberated.

SMMs and ICA based techniques rely on quite different modeling assumptions. The core assumption of SMM based source separation is the sparsity in the Short-Time Fouier Transformation (STFT) domain: at any time-frequency (TF) bin at most one of the sources is active, while the contribution of the others is negligibly small. Based on this, the posterior probability of source activities for each TF bin and the parameters of the statistical model of the multi-channel observations can be estimated with the Expectation-Maximization (EM) algorithm. Typical models are the Watson [1] and the complex Angular Central Gaussian (cACG) [2] mixture models. The actual separation is performed with the source activity posteriors which are interpreted as masks: A source signal can be retrieved either by simply multiplying the STFT representation of the mixture signal by the speaker's mask or by beamforming. In the latter case, the mask is employed to compute source-specific spatial covariance matrices, from which the beamformer coefficients of common beamformers can be obtained [8]. Note that the number of sources K can be smaller, equal or larger than the number of microphones M. The only requirement is that at least two microphones are available ($M \geq 2$). It should also be mentioned that, with SMMs, each frequency is usually treated independently, which incurs a frequency permutation problem: the order of speakers in each frequency is undefined and needs to be aligned by a so-called permutation solver, e.g., [9], before transformation to the time domain.

Independent Vector Analysis (IVA), on the other hand, does not employ a sparsity assumption. Based on ICA, the core assumption is that sources are either non-Gaussian, non-stationary or non-white [10], or even non-proper [11]. In case of speech considered here, non-Gaussianity and non-stationarity are the most common assumptions. Most IVA algorithms assume a determined case, i.e. the number of speakers K is equal to the number of observations M. This is because IVA assumes the observation is obtained with an invertible mixing matrix from the source signals. Source separation is thus obtained by reverting this mixing process. However, an extension to the overdetermined case where $M \geq K$ can be obtained, e.g., by dimension reduction through Principal Component Analysis (PCA) [12], by introducing dummy sources [13] or by simply discarding $M - K$ microphone channels. Note that IVA treats all frequencies jointly and thus does not require a frequency permutation solver.

Surprisingly, no direct comparison of SMMs and IVA on a reasonably sized database has been done so far, at least to the best of our knowledge. This paper is meant to fill this gap. Here we use the Spatialized Multi-Speaker WallStreet Journal (SMS-WSJ) dataset proposed by [14]. This dataset offers up to $M = 6$ observations which contain mixed signals from two speakers. Additionally, the observations are reverberated and contain mild microphone noise. As evaluation metric we use the Signal-to-Distortion Ratio (SDR) [15] and the Word Error Rate (WER) of a downstream speech recognizer.

Since SMMs and IVA rely on quite different modeling assumptions they may have complementary weaknesses and strengths, which makes a combination of the two an attractive option. Here, we suggest a serial concatenation, where the output of the SMM serves as an initialization of IVA. With this the WER performance can be improved by more than 10 % relative and approaches the best WER reported on this dataset, which has been obtained by a much more complex multi-channel neural-network based source separator.

The rest of this paper is organized as follows, Section 3 describes the model and the two algorithms and the way we initialize the IVA algorithm. The setup of the experiments and the discussion of the results are in Section 4. Section 5 summarize this paper.

2 Notation

In this paper we use the following notation, small characters x, bold characters $\mathbf{x}$ and bold capital characters $\mathbf{X}$ defines scalars, column vectors and matrices, respectively. $\mathbf{e}_k$ defines a unit vector, which is the k'th column of the identity matrix $\mathbf{I}_K \in \mathbb{R}^{K \times K}$. The superscript T and H denote the matrix transpose and the conjugate transpose of a matrix, respectively. All signals in this paper are in the STFT domain, with t and f being the time frame and the frequency bin index, respectively.

3 Model

Let $\mathbf{s}_{f,t} = [s_{f,t,1} \quad \cdots \quad s_{f,t,K}]^{\mathrm{T}} \in \mathbb{C}^K$ be the vector of K source signals at frequency bin index f and time frame t. The "image" of this mixture at the M microphones is given by

$$\mathbf{x}_{f,t,k} \approx \sum_{\tau=0}^{\infty} \mathbf{a}_{f,\tau,k} s_{f,t-\tau,k} \qquad (1)$$

where $\mathbf{x}_{f,t,k} \in \mathbb{C}^M$ and where $\mathbf{a}_{f,\tau,k} \in \mathbb{C}^M$ is the vector of relative transfer functions from the source k to the microphones.

The microphone signals further contain noise resulting in the observation vector $\mathbf{y}_{f,t} \in \mathbb{C}^M$:

$$\mathbf{y}_{f,t} = \sum_k \mathbf{x}_{f,t,k} + \tilde{\mathbf{n}}_{f,t} \tag{2}$$

$$= \sum_k \mathbf{h}_{f,k} d_{f,t,k} + \mathbf{n}_{f,t} = \mathbf{H}_f \mathbf{d}_{f,t} + \mathbf{n}_{f,t}. \tag{3}$$

Here, we introduced the early part of the reverberated signal $d_{f,t,k}$ and $\mathbf{h}_{f,k}$ represents the early part of the Relative Transfer Function (RTF), while the distortion $\mathbf{n}_{f,t}$ captures both the noise $\tilde{\mathbf{n}}_{f,t}$ and the late reverberation [16]. In Eq. (3), $\mathbf{H}_f \in \mathbb{C}^{M \times K}$ and $\mathbf{d}_{f,t} \in \mathbb{C}^K$ collect the early part of the RTFs and desired signals, respectively.

In the following, we describe two separation algorithms: SMM with source extraction by beamforming, and IVA. Both conduct a linear estimation, but the core assumptions are quite different as we will see.

3.1 Spatial Mixture Model

SMMs rely on the sparsity assumption of speech in the STFT domain: At most one source is active at any time frequency bin. Let $z_{f,t} \in \{0, 1, \ldots, K\}$ be a hidden variable, where $z_{f,t} = k$ indicates that source k is dominant in TF bin (f,t), and $k = 0$ is meant to indicated absence of a speech source. Then sparsity implies that Eq. (3) can be rewritten as

$$\mathbf{y}_{f,t} = \begin{cases} \mathbf{h}_{f,z_{f,t}} d_{f,t,z_{f,t}} & z_{f,t} \in \{1, \ldots K\} \\ \mathbf{n}_{f,t} & z_{f,t} = 0. \end{cases} \tag{4}$$

With this assumption, a mixture model is appropriate to represent the distribution of $\mathbf{y}_{f,t}$ [1]:

$$p(\widetilde{\mathbf{y}}_{f,t}; \pi_{f,k}, \mathbf{B}_{f,k}) = \sum_k \pi_{t,k} \cdot \mathcal{A}\left(\widetilde{\mathbf{y}}_{f,t}; \mathbf{B}_{f,k}\right), \tag{5}$$

where $\mathcal{A}(\cdot)$ is the component distribution, for which we used the cACG [2] in the following experiments. The input for the SMMs is the observation normalised to unit length

$$\widetilde{\mathbf{y}}_{f,t} = \frac{\mathbf{y}_{f,t}}{\|\mathbf{y}_{f,t}\|} \in \mathbb{C}^M. \tag{6}$$

The rationale for the removal of the vector length is the fact that the signal amplitude is mainly determined by the source signals, while the above model is meant to capture the spatial arrangement of the sources, represented by the orientation of $\widetilde{\mathbf{y}}_{f,t}$, since the spatial diversity of the sources is exploited for separation.

3.1.1 Parameter estimation

Maximum likelihood estimation of the SMM parameters is achieved by the iterative EM algorithm [2]:

$$\gamma_{f,t,k} = \frac{\pi_{t,k} \cdot \mathcal{A}\left(\widetilde{\mathbf{y}}_{f,t}; \mathbf{B}_{f,k}\right)}{\sum_{\check{k}} \pi_{t,\check{k}} \cdot \mathcal{A}\left(\widetilde{\mathbf{y}}_{f,t}; \mathbf{B}_{f,\check{k}}\right)}, \tag{7}$$

$$\mathbf{B}_{f,k} = M \cdot \frac{\sum_t \gamma_{f,t,k} \frac{\widetilde{\mathbf{y}}_{f,t} \cdot \widetilde{\mathbf{y}}_{f,t}^{\mathrm{H}}}{\widetilde{\mathbf{y}}_{f,t}^{\mathrm{H}} \mathbf{B}_{f,k}^{-1} \widetilde{\mathbf{y}}_{f,t}}}{\sum_t \gamma_{f,t,k}} \tag{8}$$

$$\pi_{t,k} = \frac{1}{F} \sum_f \gamma_{f,t,k}. \tag{9}$$

Here, $\gamma_{f,t,k} = \Pr(z_{f,t} = k | \widetilde{\mathbf{y}}_{f,t})$ is the posterior probability that TF bin (f,t) is dominated by source k. Further, $\mathbf{B}_{f,k}$ is a parameter matrix. Please note that the mixture weight $\pi_{t,k}$ has been chosen here to be time dependent [17]. While, in theory, the time dependent mixture weight [17] avoids the frequency permutation problem, experiments showed that more reliable estimates are obtained, if a permutation solver is nevertheless integrated in the EM algorithm, i.e., applied after each EM step [18]. We use an unpublished[1] similarity based permutation solver from Tran Vu [1] similar to [9] and apply it after Eq. (7).

---

[1] Code is published: https://github.com/fgnt/pb_bss

3.1.2 Beamforming

The SMM does not yield an estimate of the source signals. Signal extraction is done by interpreting the posterior probabilities $\gamma_{f,t,k}^{\mathrm{ML}}$ as a mask. The k-th source can be recovered simply by multiplying the observed signal with the posterior probability. But superior performance is achieved by employing the mask to compute spatial covariance matrices, which are then used to calculate a Minimum Variance Distortionless Response (MVDR) beamformer [19] with an SDR based reference channel selection [8]. Finally, we can obtain the estimate with

$$\hat{d}_{f,t,k} = \check{\mathbf{w}}_{f,k}^{\mathrm{H}} \cdot \mathbf{y}_{f,t}, \tag{10}$$

where $\check{\mathbf{w}}_{f,k}^{\mathrm{H}}$ are the beamformer coefficients.

3.2 Independent Vector Analysis

IVA [4, 5] is an extension of ICA [3, 7]. In ICA we assume that the observation $\mathbf{y}_{f,t}$ is obtained by a linear combination of the "source" signals $\mathbf{d}_{f,t}$

$$\mathbf{y}_{f,t} = \mathbf{H}_f \mathbf{d}_{f,t}, \tag{11}$$

where it is assumed, that the mixing matrix $\mathbf{H}_f \in \mathbb{C}^{M \times K}$ is invertible. This implies that the number of sources K is equal to the number of microphones M. The goal is then to find a separation matrix $\mathbf{W}_f \in \mathbb{C}^{K \times M}$ that estimates the source signals $\mathbf{d}_{f,t}$

$$\hat{\mathbf{d}}_{f,t} = \mathbf{W}_f \cdot \mathbf{y}_{f,t}, \tag{12}$$

where the key assumption is, that the source signals are independent.

The difference between ICA and IVA is that ICA assumes an independent source model for each frequency, while in IVA the source models are coupled between the frequencies. Here, the coupling is given by the time-varying, however frequency-independent variance $r_{t,k}$ of the source models:

$$p(\hat{\mathbf{d}}_{f,t}) = \prod_k \mathcal{N}\left(\hat{d}_{f,t,k}; 0, r_{t,k}\right). \tag{13}$$

Source separation thus exploits the nonstationarity of the sources, rather than the non-Gaussianity.[2]

The separation matrix $\mathbf{W}_f$ is commonly estimated with the maximum likelihood approach.

To accommodate the overdetermined case, where the number of microphones exceeds the number of sources ($M > K$), we here follow the approach of Overdetermined IVA (OverIVA) in [13]. They proposed to add $M - K$ dummy estimates $\mathbf{v}_{f,t} = \begin{bmatrix} v_{f,t,1} & \cdots & v_{f,t,M-K} \end{bmatrix}^{\mathrm{T}}$, so that no dimension reduction technique is necessary. For the $M - K$ estimates, they proposed to use a Gaussian model with a time-invariant but frequency dependent covariance, which stands in contrast to the time-varying frequency independent variances $r_{t,k}$ used for the true sources. In [13] it is shown that the dummy sources $\mathbf{v}_{f,t}$ can be modeled as dependent variables without negatively affecting the estimation of the sources of interest $\mathbf{d}_{f,t}$, while at the same time allowing to find a simpler solution. Hence we use a full covariance matrix Σ_f in:

$$p(\mathbf{v}_{f,t}) = \mathcal{N}\left(\mathbf{v}_{f,t}; 0, \Sigma_f\right). \tag{14}$$

Now we can define the square separation matrix

$$\widetilde{\mathbf{W}}_f = \begin{bmatrix} \mathbf{W}_f \\ \mathbf{U}_f \end{bmatrix} \in \mathbb{C}^{M \times M}, \tag{15}$$

and use complex linear random variable transformation [20] for $\mathbf{y}_{f,t} = \widetilde{\mathbf{W}}_f^{-1} \begin{bmatrix} \hat{\mathbf{d}}_{f,t}^{\mathrm{T}} & \mathbf{v}_{f,t}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$ to obtain the likelihood:

$$L = p(\mathbf{y}_{f,t}; \forall f \in \{1, \ldots, F\} \text{ and } t \in \{1, \ldots, T\}; \boldsymbol{\theta})$$

$$= \prod_f |\det(\widetilde{\mathbf{W}}_f)|^2 \cdot \prod_t \mathcal{N}\left(\mathbf{U}_f \mathbf{y}_{f,t}; \mathbf{0}, \Sigma_f\right)$$

$$\cdot \prod_k \mathcal{N}\left(\mathbf{w}_{f,k}^{\mathrm{H}} \mathbf{y}_{f,t}; 0, r_{t,k}\right), \tag{16}$$

---

[2] Note, though, that when averaging over the distribution of the variances (typically a Gamma distribution), the resulting predictive distribution is student-t, i.e., is supergaussian.

Algorithm 1 OverIVA inspired by [13]

1: Initialize $\widetilde{\mathbf{W}}_f$
2: **while** not converged **do**
3:     **for** $k \in \{1, \dots, K\}$ **do**
4:         $r_{t,k} \leftarrow \frac{1}{F}\sum_f |\mathbf{w}_{f,k}^{\mathrm{H}}\mathbf{y}_{f,t}|^2 \quad \forall t$
5:         **for** $f \in \{1, \dots, F\}$ **do**
6:             $\mathbf{R}_{f,k} \leftarrow \frac{1}{T}\sum_t \frac{1}{r_{t,k}}\mathbf{y}_{f,t}\mathbf{y}_{f,t}^{\mathrm{H}}$
7:             $\mathbf{w}_{f,k} \leftarrow \left(\widetilde{\mathbf{W}}_f \mathbf{R}_{f,k}\right)^{-1}\mathbf{e}_k$
8:             $\mathbf{w}_{f,k} \leftarrow \mathbf{w}_{f,k}/\sqrt{\mathbf{w}_{f,k}^{\mathrm{H}}\mathbf{R}_{f,k}\mathbf{w}_{f,k}}$
9:             update k'th row of $\mathbf{W}_f$ with $\mathbf{w}_{f,k}^{\mathrm{H}}$
10:     $\mathbf{J}_f \leftarrow \left(\mathbf{E}_2\mathbf{\Sigma}_{\mathbf{y},f}\mathbf{W}_f^{\mathrm{H}}\right)\left(\mathbf{E}_1\mathbf{\Sigma}_{\mathbf{y},f}\mathbf{W}_f^{\mathrm{H}}\right)^{-1}$
11:     $\mathbf{U}_f \leftarrow [\mathbf{J}_f \quad -\mathbf{I}_{M-K}]$
12:     $\widetilde{\mathbf{W}}_f \leftarrow \begin{bmatrix}\mathbf{W}_f \\ \mathbf{U}_f\end{bmatrix}$

where $\mathbf{w}_{f,k}$ is an entry of $\mathbf{W}_f = [\mathbf{w}_{f,1} \quad \dots \quad \mathbf{w}_{f,K}]^{\mathrm{H}}$ and $\boldsymbol{\theta}$ are all parameters, i.e. $\mathbf{W}_f$ and $r_{t,k}$ for every f, t and k.

3.2.1 Parameter estimation

Since no closed form solution is known, the likelihood is maximized in an iterative fashion, similar to [21, 22]. Algorithm 1 summarizes the update rules for the OverIVA parameter estimation presented by [13]. The maximization regarding the separation matrix leads to the Hybrid Exact-Approximate Joint Diagonalization (HEAD) Problem [23], which also has no closed form solution if the number of sources is greater than two [21, 24, 25]. So we update the separation matrix in another iterative update scheme proposed by [22]. In the algorithm we used the following notation

$$\mathbf{E}_1 = [\mathbf{I}_K \quad \mathbf{0}_{K\times M-K}] \in \mathbb{C}^{K\times M}, \tag{17}$$

$$\mathbf{E}_2 = [\mathbf{0}_{M-K\times K} \quad \mathbf{I}_{M-K}] \in \mathbb{C}^{(M-K)\times M}. \tag{18}$$

Further,

$$\mathbf{\Sigma}_{\mathbf{y},f} = \frac{1}{T}\sum_t \mathbf{y}_{f,t}\mathbf{y}_{f,t}^{\mathrm{H}} \in \mathbb{C}^{M\times M}, \tag{19}$$

is the covariance matrix of the observations.

After the separation has been performed with Eq. (12), the scaling ambiguity of the sources is resolved with the minimum distortion principle [26]

$$\hat{\beta}_{f,k} = \underset{\beta_{f,k}}{\mathrm{argmin}} \sum_t \left\| y_{f,t,r} - \beta_{f,k}\hat{d}_{f,t,k} \right\|^2, \tag{20}$$

where r is a reference microphone index and $\hat{\beta}_{f,k}$ a scalar to fix the scaling ambiguity. This leads to

$$\hat{d}_{f,t,k} \leftarrow \frac{\sum_{\check{t}} y_{f,\check{t},r}\hat{d}_{f,\check{t},k}^*}{\sum_{\check{t}} |\hat{d}_{f,\check{t},k}|^2}\hat{d}_{f,t,k}. \tag{21}$$

where $(\cdot)^*$ is the complex conjugate operation.

3.2.2 Parameter initialisation

To start the estimation iterations for IVA an initialisation of the separation matrix $\mathbf{W}_f$ is necessary. The simplest option is to employ the identity matrix for $\widetilde{\mathbf{W}}_f$. A more elaborated way is to use the eigenvectors of $\mathbf{\Sigma}_{\mathbf{y},f}$ that belong to the largest eigenvalues [13] and then calculate $\widetilde{\mathbf{W}}_f$ with the help of lines 10 to 12 from Algorithm 1.

However, initialization of an iterative algorithm turns often times to be critical for overall performance. Therefore, we see here an opportunity for combining the two approaches for source separation: using the result of the SMM to initialize IVA.

Let $\hat{d}_{f,t,k}$ be the estimates of the separated signals obtained by the SMM. The separation coefficients $\mathbf{w}_{f,k}$ of IVA can then be initialized by solving the following least squares (LS) problem

$$\mathbf{w}_{f,k} = \underset{\mathbf{w}_{f,k}}{\mathrm{argmin}} \left\{ \sum_t \|\hat{d}_{f,t,k} - \mathbf{w}_{f,k}^{\mathrm{H}}\mathbf{y}_{f,t}\|^2 \right\}, \tag{22}$$

whose solution is well-known:

$$\mathbf{w}_f = \left(\sum_{\check{t}}\mathbf{y}_{f,\check{t}}\mathbf{y}_{f,\check{t}}^{\mathrm{H}}\right)^{-1}\sum_t \mathbf{y}_{f,t}\hat{d}_{f,t,k}^*. \tag{23}$$

We mentioned earlier that source extraction based on SMMs is done by beamforming. Therefore, it appears natural to use the beamforming vectors, such as $\mathbf{w}_{f,k}^{\mathrm{H}}$, Eq. (10), directly as the initial values of the separation matrix. Indeed, this is a valid option, since the beamformer coefficients have been derived as linear estimates that optimize some objective function, such as mean squared error or minimum variance under a distortionless constraint. In the following we will, for reasons to be explained below, use different STFT sizes for SMM and IVA based separation. While transforming the coefficients from one STFT size to another is in principle possible, it is easier to obtain the coefficients in the target STFT window size by solving the above LS problem. Note, in case of the same STFT window size and shift in SMM and IVA, the beamformer is equal to the least squares estimate from Eq. (23), which can be seen, when we use Eq. (10) in Eq. (22):

$$\mathbf{w}_{f,k} = \underset{\mathbf{w}_{f,k}}{\mathrm{argmin}} \left\{ \sum_t \|\check{\mathbf{w}}_{f,k}^{\mathrm{H}}\cdot\mathbf{y}_{f,t} - \mathbf{w}_{f,k}^{\mathrm{H}}\mathbf{y}_{f,t}\|^2 \right\}. \tag{24}$$

4 Experiments

For evaluating the models, we use the SMS-WSJ dataset proposed in [14]. This dataset uses the audio data from the WSJ Database [27], resampled to 8 kHz sampling rate. Utterances are artificially reverberated by convolving with simulated Room Impulse Responses (RIRs) with the sound decay time T_{60} sampled uniformly in the range from 200 to 500 ms, to create observations for 6 microphones. Two utterances are added at an average SDR of 0 dB to create a mixture. Further, to simulate the microphone noise, Gaussian noise with an average Signal-to-Noise Ratio (SNR) of 25 dB is added.

The models in this paper don't require a training, so only the test set is used. This set contains 1 332 mixtures, 45 144 spoken words giving rise to 201 minutes of test data.

As performance metrics we use the SDR [15] and the WER. For the WER calculation we used the Kaldi [28] Automatic Speech Recognition (ASR) model for this database [14].

4.1 Baseline and Topline

The first row of Table 1 shows the SDR and the WER if the observations are taken as they are, i.e., without any separation. Furthermore, for each speaker separately, the early image signals[3] are used [14, 16]. This corresponds to perfect dereverberation and source separation, serving as topline for our experiments.

Next, those image signals are used as initialization for the IVA separation matrix, i.e., replacing $\hat{d}_{f,t}$ by $\mathbf{d}_{f,t}$ in Eq. (22), which can serve as another indication of which performance is best possible. Note that the performance strongly depends on the STFT window size and the STFT advance, with longer windows leading to improved performance up to 2048. This was to be expected because the approximation of a convolution by a simple multiplication in the STFT domain, Eq. (3), (so-called multiplicative transfer function approximation [29]) is better justified.

---

[3]The early image is the speech source convolved with the initial 50 ms of the RIR

Table 1: Reference scores on SMS-WSJ. Early image is the speech source convolved with the initial part of the RIR [14].

	Size / Shift	WER [%]	SDR [dB]
Observation	- / -	78.12	−0.01
Early image	- / -	7.24	57.10
+ linear constraint: Eq. (22)	1024 / 128	11.64	16.33
+ linear constraint: Eq. (22)	2048 / 256	9.29	20.07

Table 2: Scores for SMM and IVA with different STFT parameters.

	Size / Shift	WER [%]	SDR [dB]
SMM [14]	512 / 128	18.70	12.30
SMM	1024 / 128	14.33	12.80
SMM	2048 / 256	17.53	11.15
IVA	1024 / 128	13.84	13.53
IVA	2048 / 256	11.80	13.72

Table 3: Results of the IVA algorithm, using the mixture model as initialization

Initialization	SMM Size / Shift	IVA Size / Shift	WER [%]	SDR [dB]
$\mathbf{w}_{f,k}^{\mathrm{BF}}$	1024 / 128	1024 / 128	12.57	13.72
	2048 / 256	2048 / 256	10.84	13.84
$\hat{\mathbf{d}}_{f,t,k}$	1024 / 128	1024 / 128	12.56	13.71
	1024 / 128	2048 / 256	10.69	13.88
	2048 / 256	2048 / 256	10.84	13.85

4.2 STFT Sizes

In all following simulations, $M = 6$ microphones are used for separation and the number of considered sources are $K = 3$.[4]

Following up on the effect of the STFT size, Table 2 shows the performance of the source separation algorithms SMM and IVA for different STFT window sizes and shifts. The result in the first row is the baseline result from [14]. It can be observed that SMM obtains its best results for an STFT size of 1024, while IVA can be further improved by increasing the size to 2048. This can be explained by the fact that SMM builds upon the sparsity of speech in STFT domain, which is lost if the STFT size is chosen too large, while IVA does not have to bother with sparsity and therefore can afford larger sizes.

Further, it is striking that IVA clearly outperforms SMM based source separation. It should also be mentioned that the SDR is not always a good indicator of WER performance. Comparing the results in the last two rows, although the SDR improves by only 0.19 dB, the WER decreases by more than two percentage points! We will come back to this discrepancy below.

4.3 Model Chaining

Table 3 displays results obtained with the initialization of the IVA separation matrix using the SMM output. Using the initialization improves the results of IVA. While the SDR is improved by only 0.15 dB, the WER improved by more than 1 percentage point compared to the best results in Table 2.

In order to shed some light on the significance of the SDR on WER, Fig. 1 displays the cumulative distribution function of the SDR for different models. Here we can see that IVA achieves very low SDR values for some examples: Without the SMM-based initialization, 3.2 % of the examples separated by IVA achieved a SDR of ≤ 7 dB. With the SMM-based initialization, the number of poorly separated mixtures is reduced to 2 %.

Mixtures with such low separation performance will cause recognition errors in the ASR engine. Thus, reducing the percentage of poorly separated mixtures may not have a large impact on the average SDR performance, but nevertheless have significant effect on the average WER.

4.4 Dereverberation and Comparison with State-of-the-Art

Since the data is reverberated and since the source separation algorithms by themselves are not meant to carry out dereverberation, we experimented with Weighted Prediction Error (WPE) [30, 31] as a preprocessing step in front of source separation. WPE is a powerful dereverberation algorithm which has led to improved ASR performance on many databases.

---

[4]We observed better performance, if we used a separate class for the noise: 2 speakers plus noise amount to a total of 3 classes.

Table 4: Results of the proposed chain and comparison with literature

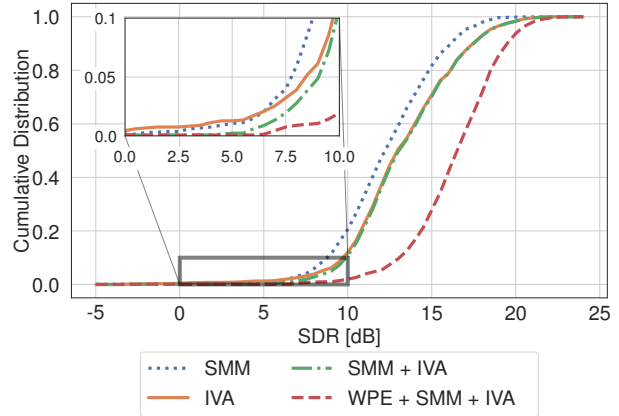|	Array agnostic	WER [%]	SDR [dB]
WPE + SMM	✓	11.52	15.58
WPE + IVA	✓	10.91	16.60
WPE + SMM + IVA	✓	9.55	16.84
NN + BF + NN [32]	✗	8.52	-
NN + BF [32]	✗	28.71	-
NN + BF [33]	✓	16.84	14.20

Figure 1: Cumulative distribution of the achieved SDR values on the testset

Here, we apply WPE using an STFT with a window size of 512 and a shift of 128. Those parameter values have been proposed in earlier studies and can be motivated by the fact that WPE explicitly aims at extracting the direct signal and early reflections, up to 50 ms of the RIR. Therefore, relatively short STFT windows should be used such that the late reverberation is captured by the windows following the window containing only the direct signal and early reflections.

With WPE preprocessing and initializing IVA with SMM we obtain the overall best results with 9.55% WER and 16.84 dB SDR.

In Table 4 we have added the results reported in [32]. They tried different combinations of Neural Networks (NNs) and beamforming. The best configuration consisted of first a NN, then beamforming and then again a NN, where both NNs had access to all microphones and the objective was to do dereverberation and source separation. The system which achieved the currently best WER on SMS-WSJ is tuned to the specific array configuration and is a huge system with 14 million parameters.

We note that our approach proposed here is worse by about one percentage point, but achieves this result with an unsupervised approach and much less parameters. Further, SMM and IVA are array geometry independent and have already proven to work with real recordings.

Since beamforming alone didn't worked well in [32], see entry NN+BF [32] in Table 4, we applied the NN system from [33] to this dataset (entry NN+BF [33]). When comparing to our proposal that uses a linear enhancement (i.e. beamforming), we see an advantage of ours compared to the NNs.

5 Conclusion

In this paper, two unsupervised blind source separation techniques are compared on the SMS-WSJ dataset. It is shown that IVA outperforms SMM. Furthermore, initializing IVA with the output of SMM further improved the WER performance of IVA by about 1 percentage point or 10 % relative. In a comparison with results from the literature, the unsupervised techniques have shown to get close to the best NNs, without being array dependent and outperform array agnostic NN approaches.

References

[1] D. H. T. Vu and R. Haeb-Umbach, "Blind speech separation employing directional statistics in an expectation maximization framework," in *ICASSP*, pp. 241–244, IEEE, 2010.

[2] N. Ito, S. Araki, and T. Nakatani, "Complex angular central Gaussian mixture model for directional statistics in mask-based microphone array signal processing," in *EUSIPCO*, pp. 1153–1157, IEEE, 2016.

[3] S.-i. Amari, A. Cichocki, H. H. Yang, *et al.*, "A new learning algorithm for blind signal separation," in *Advances in neural information processing systems*, pp. 757–763, Morgan Kaufmann Publishers, 1996.

[4] T. Kim, T. Eltoft, and T.-W. Lee, "Independent vector analysis: An extension of ica to multivariate components," in *International Conference on Independent Component Analysis and Signal Separation*, pp. 165–172, Springer, 2006.

[5] A. Hiroe, "Solution of permutation problem in frequency domain ICA, using multivariate probability density functions," in *International Conference on Independent Component Analysis and Signal Separation*, pp. 601–608, Springer, 2006.

[6] H. Sawada, N. Ono, H. Kameoka, D. Kitamura, and H. Saruwatari, "A review of blind source separation methods: two converging routes to ILRMA originating from ICA and NMF," *APSIPA Transactions on Signal and Information Processing*, vol. 8, 2019.

[7] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," *Neural networks*, vol. 13, no. 4-5, pp. 411–430, 2000.

[8] H. Erdogan, J. R. Hershey, S. Watanabe, M. I. Mandel, and J. Le Roux, "Improved MVDR beamforming using single-channel mask prediction networks.," in *Interspeech*, pp. 1981–1985, 2016.

[9] H. Sawada, S. Araki, and S. Makino, "Underdetermined convolutive blind source separation via frequency bin-wise clustering and permutation alignment," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 3, pp. 516–527, 2010.

[10] H. Buchner, R. Aichner, and W. Kellermann, "TRINICON: A versatile framework for multichannel blind signal processing," in *ICASSP*, vol. 3, IEEE, 2004.

[11] B. Loesch and B. Yang, "Cramér-rao bound for circular and noncircular complex independent component analysis," *IEEE transactions on signal processing*, vol. 61, no. 2, pp. 365–379, 2012.

[12] F. Asano, Y. Motomura, H. Asoh, and T. Matsui, "Effect of pca filter in blind source separation," in *Proc. ICA*, pp. 57–62, 2000.

[13] R. Scheibler and N. Ono, "Independent vector analysis with more microphones than sources," in *WASPAA*, pp. 185–189, IEEE, 2019.

[14] L. Drude, J. Heitkaemper, C. Boeddeker, and R. Haeb-Umbach, "SMS-WSJ: Database, performance measures, and baseline recipe for multi-channel source separation and recognition," *arXiv preprint arXiv:1910.13934*, 2019.

[15] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE transactions on audio, speech, and language processing*, vol. 14, no. 4, pp. 1462–1469, 2006.

[16] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, "Far-Field Automatic Speech Recognition," *Proceedings of the IEEE*, 2020.

[17] N. Ito, S. Araki, and T. Nakatani, "Permutation-free convolutive blind source separation via full-band clustering based on frequency-independent source presence priors," in *ICASSP*, pp. 3238–3242, IEEE, 2013.

[18] L. Drude, *Integration of neural networks and probabilistic spatial models for acoustic blind source separation*. PhD thesis, Paderborn University, 2020.

[19] M. Souden, J. Benesty, and S. Affes, "On optimal frequency-domain multichannel linear filtering for noise reduction," *IEEE Transactions on audio, speech, and language processing*, vol. 18, no. 2, pp. 260–276, 2009.

[20] F. D. Neeser and J. L. Massey, "Proper complex random processes with applications to information theory," *IEEE transactions on information theory*, vol. 39, no. 4, pp. 1293–1302, 1993.

[21] N. Ono and S. Miyabe, "Auxiliary-function-based independent component analysis for super-Gaussian sources," in *International Conference on Latent Variable Analysis and Signal Separation*, pp. 165–172, Springer, 2010.

[22] N. Ono, "Stable and fast update rules for independent vector analysis based on auxiliary function technique," in *WASPAA*, pp. 189–192, IEEE, 2011.

[23] A. Yeredor, "On hybrid exact-approximate joint diagonalization," in *3rd International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 312–315, IEEE, 2009.

[24] H. Sawada, N. Ono, H. Kameoka, and D. Kitamura, "Blind audio source separation on tensor representation," *ICASSP, Apr*, 2018.

[25] N. Ono, "Fast stereo independent vector analysis and its implementation on mobile phone," in *IWAENC*, IEEE, 2012.

[26] K. Matsuoka, "Minimal distortion principle for blind source separation," in *Proceedings of the 41st SICE Annual Conference*, vol. 4, pp. 2138–2143, IEEE, 2002.

[27] D. B. Paul and J. Baker, "The design for the wall street journal-based CSR corpus," in *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York*, 1992.

[28] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, *et al.*, "The kaldi speech recognition toolkit," in *Workshop on Automatic Speech Recognition and Understanding (ASRU)*, IEEE, 2011.

[29] Y. Avargel and I. Cohen, "On multiplicative transfer function approximation in the short-time fourier transform domain," *IEEE Signal Processing Letters*, vol. 14, no. 5, pp. 337–340, 2007.

[30] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1717–1731, 2010.

[31] L. Drude, J. Heymann, C. Boeddeker, and R. Haeb-Umbach, "NARA-WPE: A Python package for weighted prediction error dereverberation in Numpy and Tensorflow for online and offline processing," in *Speech Communication; 13th ITG-Symposium*, IEEE, 2018.

[32] Z.-Q. Wang, P. Wang, and D. Wang, "Multi-microphone complex spectral mapping for utterance-wise and continuous speaker separation," *arXiv preprint arXiv:2010.01703*, 2020.

[33] C. Boeddeker, W. Zhang, T. Nakatani, K. Kinoshita, T. Ochiai, M. Delcroix, N. Kamo, Y. Qian, and R. Haeb-Umbach, "Convolutive transfer function invariant sdr training criteria for multi-channel reverberant speech separation," in *ICASSP*, pp. 8428–8432, IEEE, 2021.