

30+ Years of Source Separation Research: Achievements and Future Challenges

Araki, Shoko; Ito, Nobutaka; Haeb-Umbach, Reinhold; Wichern, Gordon; Wang, Zhong-Qiu; Mitsufuji, Yuki

TR2025-036 March 08, 2025

Abstract

Source separation (SS) of acoustic signals is a research field that emerged in the mid-1990s and has flourished ever since. On the occasion of ICASSP's 50th anniversary, we review the major contributions and advancements in the past three decades in the speech, audio, and music SS research field. We will cover both single- and multi-channel SS approaches. We will also look back on key efforts to foster a culture of scientific evaluation in the research field, including challenges, performance metrics, and datasets. We will conclude by discussing current trends and future research directions.

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2025

30+ Years of Source Separation Research: Achievements and Future Challenges

Shoko Araki¹, Nobutaka Ito², Reinhold Haeb-Umbach³, Gordon Wichern⁴, Zhong-Qiu Wang⁵, Yuki Mitsuji⁶
¹NTT Corporation, Japan, ²University of Tokyo, Japan, ³Paderborn University, Germany,
⁴Mitsubishi Electric Research Laboratories, USA, ⁵Southern University of Science and Technology, China, ⁶Sony AI, USA

Abstract—Source separation (SS) of acoustic signals is a research field that emerged in the mid-1990s and has flourished ever since. On the occasion of ICASSP’s 50th anniversary, we review the major contributions and advancements in the past three decades in the speech, audio, and music SS research field. We will cover both single- and multi-channel SS approaches. We will also look back on key efforts to foster a culture of scientific evaluation in the research field, including challenges, performance metrics, and datasets. We will conclude by discussing current trends and future research directions.

Index Terms—Speech separation, audio source separation, music source separation, 50th ICASSP

I. INTRODUCTION

Source separation (SS) is a technology that reconstructs the individual source signals from one or more mixtures of them. Given that daily acoustic environments usually contain multiple concurrent sound sources, SS is an important technology for acoustic signals, e.g., listening to each speaker’s voice when multiple people speak simultaneously or extracting vocal or specific instrumental parts from music.

SS, including blind SS (BSS), of acoustic signals [1]–[3] is an important area of research that rapidly emerged from around the mid-1990s to the early 2000s and much research is still being conducted in this field. The SS topic was added to the EDICS in 2006 as AUD-SSEN (Source Separation and Signal Enhancement), and then separated into AUD-SEP (Audio and Speech Source Separation) and AUD-SEN (Signal Enhancement and Restoration) to differentiate between those growing research areas in 2014. In recent years, ICASSP has consistently received 40–50 submissions to AUD-SEP every year. Over the past 30+ years, many researchers have entered the field and many breakthroughs have been made. In addition, since the 2000s, when BSS research based on independent component analysis (ICA) was being conducted, various attempts were made to make the SS field prosperous as described in Sec. IV. Furthermore, since deep learning was introduced in audio SS in the 2010s, many data-driven methods have been proposed, and a lot of SS research has been conducted not only in the audio and acoustic signal processing (AASP) community but also in the speech and machine learning communities.

This paper will review not only the key contributions and advancements in the last three decades, but also key efforts such as challenges, performance metrics, and datasets to foster SS research. We will also discuss current trends and future directions in SS research.

II. PROBLEM DESCRIPTION

Given N (≥ 2) source signals $s_n(\tilde{t})$, where \tilde{t} is the discrete time index, the mixture signal $y_m(\tilde{t})$ captured at the m^{th} of M (≥ 1) microphones, can be written as a convolutive mixture

$$y_m(\tilde{t}) = \sum_{n=1}^N \sum_{\tau} h_{mn}(\tau) s_n(\tilde{t} - \tau) + v_m(\tilde{t}) \quad (1)$$

$$= \sum_{n=1}^N c_{mn}(\tilde{t}) + v_m(\tilde{t}) \quad \text{with } m \in \{1, \dots, M\}, \quad (2)$$

where h_{mn} is the acoustic impulse response from source n to microphone m , and v_m is an additive noise term, which is often neglected.

The goal of SS is to estimate the source signals s_n or their images $c_{\mu n}$ at a reference microphone μ , from the observed mixtures $\{y_1, \dots, y_M\}$. If this problem is solved without leveraging any prior knowledge about the sources or mixing conditions, it is referred to as *blind* SS (BSS). The cases where $N < M$, $N = M$, and $N > M$ are called over-determined, determined, and underdetermined problems, respectively. In many studies, the number of sources N is assumed known, even for BSS.

Under certain assumptions, the convolutive mixture can be approximated as an instantaneous mixture in the time-frequency (TF) domain by using, e.g., the short-time Fourier transform (STFT):

$$y_{mtf} = \sum_{n=1}^N h_{mnf} s_{ntf} + v_{mtf}, \quad (3)$$

where t and f are time frame and frequency bin indices, respectively. While this is often more computationally tractable, it introduces a frequency permutation problem in narrow-band SS algorithms, since the order in which the source signals appear at the separator output at different frequencies is arbitrary.

III. SOURCE SEPARATION APPROACHES

This section introduces representative SS technologies that have been proposed over the past 30+ years.

A. Model-based SS

BSS is inherently an ill-posed problem. Thus additional assumptions must be made to arrive at a unique solution. Assumptions on the source signals and/or the mixing process are given in mathematical models to solve BSS. Approaches to BSS can be distinguished by the assumptions applied.

Multi-channel approaches:

For determined or over-determined BSS, **independent component analysis (ICA)** was applied to separate acoustic signals in the mid-1990s [4]. This approach assumes that the source signals are non-Gaussian, non-stationary, non-white, or a combination of those.

Methods for dealing with convolutive mixtures were actively studied in the TF domain (e.g., [1]–[3]). To address the permutation problem in frequency-domain ICA, **independent vector analysis (IVA)** [5]–[7] was proposed to find a separating matrix based on the independence between vectors bundling signal components of all frequencies, assuming that frequency components of the same source are statistically dependent. The permutation indeterminacy can also be addressed by employing a full-band source model. A popular choice is **non-negative matrix factorization (NMF)**, which models a

spectrogram as linear combinations of base spectra. The combination of frequency-domain ICA and NMF as a source model is known as **independent low-rank matrix analysis (ILRMA)** [8]. Solving the permutation problem in frequency-domain ICA and accelerating algorithms are still important research topics.

For an underdetermined problem ($N > M$), **TF masking** is one of the primary methods. By utilizing W-disjoint orthogonality [9] of many sound source signals in the TF domain, this method extracts the dominant sound at each TF with a binary (or soft) mask. A typical mask estimation method is clustering spatial features at each TF obtained from multi-channel mixtures in an expectation-maximization (EM) framework using a spatial mixture model (e.g., [10]). Recently, many mask estimation approaches using a neural network have been proposed, as described in Sec. III-B. The source signals can then be extracted by simply multiplying the microphone signal in the TF domain with the estimated masks. Alternatively, the TF masks are used to estimate beamformer coefficients, leading to source extraction by beamforming, which typically reduces artifacts [11]–[13].

Another approach to underdetermined BSS is **full-rank spatial covariance analysis (FCA)** [14]. FCA is capable of handling situations in which the W-disjoint orthogonality does not hold, such as reverberant environments and music SS. To this end, the method uses a multi-channel Wiener filter instead of the TF mask and necessitates spatial covariance matrices of the individual sources. These matrices are modeled as full-rank instead of rank-one to better deal with reverberant environments and are estimated from the observed mixtures using the EM algorithm. Key advancements include avoiding the permutation problem by incorporating NMF [15], [16], boosting SS performance by combining with a deep generative model [17], and reducing computation through joint diagonalization [18], [19].

In music separation, methods that work well in the underdetermined case are required since music is predominantly delivered in a stereo format and typically contains more than two instruments (i.e., sources). Panning is often used in music production such that sources have different locations in the stereo image, and by assuming that there is very little overlap between different sources in the magnitude spectrogram (i.e., the W-disjoint orthogonality), histograms of angle estimates for each time-frequency bin can be used to create separation masks. DUET [20], ADReSS [21] and PROJET [22] are well-known examples of energy vs angle separation algorithms, which, while imperfect due to the W-disjoint orthogonality assumption, are often lightweight enough to run in real-time.

Single-channel approaches:

SS has to rely on spectral cues to discriminate sources if only a single-channel input is available. An early attempt to solve this extremely hard problem was **factorial hidden Markov models (HMMs)**, where a multi-dimensional dynamic programming algorithm was employed to track and decode the time trajectory of the individual sources [23]. However, the grammar was very restrictive, and general multi-talker separation was far from being solved.

Assuming source magnitudes are additive, **NMF** [24] decomposes the magnitude spectrum of the mixture into a product of a basis matrix of prototype spectra and an activation matrix. NMF guarantees the non-negativity of the estimated base spectra and activations, making it well suited for magnitude spectra. NMF has been predominantly applied to audio and music SS, where the source spectra of individual sounds or instruments can have quite different signatures [25].

The introduction of neural networks, which learn the spectral patterns of the source signals in a supervised learning phase brought about a real breakthrough, even for single-channel input, as described

in the next section.

In music separation, specific qualities of music signals can be exploited for designing single-channel separation algorithms. REpeating Pattern Extraction Technique (REPET) [26] identifies periodically repeating segments in music signals, which typically correspond to background musical elements, and can then be separated from non-repeating foreground elements, such as lead vocals. Another computationally efficient approach is harmonic-percussive separation [27], which uses median filtering on both the time and frequency dimensions of the magnitude spectrogram to separate harmonic elements (approximately horizontal spectrogram lines) from percussive elements (vertical lines). If a musical score is available, using it as a hint has also been an active research area [28].

B. Deep learning-based SS

For audio SS, the sound types (e.g., speech, music, and sound events) to deal with are often known as a prior. One can leverage any prior knowledge of signal patterns to address the ill-posed problem in SS, thereby improving separation. In recent years, deep learning has shown remarkable capability at learning signal patterns from massive data. In this context, *learning to separate* based on supervised learning and deep neural networks (DNNs) has attracted broad research interest and become a promising direction [29].

A major breakthrough in this direction was made by Wang and Wang [30]. Using simulated pairs of clean and noisy speech, they realized speech enhancement by training DNNs on spectral features to predict the so-called **ideal binary mask**. At run time, the estimated mask functions like a Wiener filter that can suppress noise. This concept formulates speech enhancement as a data-driven classification problem and can be approached via large-scale supervised learning on massive simulated data.

Unlike speech enhancement, where speech and noise sources exhibit different signal patterns, speaker separation has a particular difficulty in label permutation, since the sources to separate are all speech and homogeneous. This poses a major challenge for supervised learning, as the sources estimated by DNNs are often not aligned with the true sources. A major breakthrough that addresses this issue is **deep clustering** [31], which trains DNNs to embed each TF unit so that the embeddings of the TF units dominated by the same source are close to each other and far away otherwise. At run time, binary TF masks for separating sources are estimated by clustering the learned embeddings. Another key approach to address the issue is **permutation invariant training (PIT)** [31], [32], which aligns estimated sources with true sources before loss computation.

Another major direction is **target speaker extraction**, which informs DNNs about which sources to separate by inputting auxiliary information, such as speaker embeddings or visual cues [33], [34].

Even in monaural conditions, where only spectral cues can be utilized for SS, supervised deep learning has already shown remarkable effectiveness [29]. When multiple microphones are available, spatial features can complement spectral ones to improve separation [35]. Since in many applications the same microphone array is used in training and testing, a trend is to directly stack the real and the imaginary components [36] or, more simply, the waveforms [37] of input mixtures as input features. In scenarios where the array geometry could be different between training and testing, DNN modules that can model arrays with various geometries and numbers of microphones have been proposed [38], [39].

In the past decade, DNN-based SS approaches have shifted from TF masking, which estimates only the target magnitude via real-valued masking and uses the mixture phase for signal re-synthesis,

to complex spectrum estimation [40], [41], or even to time-domain waveform estimation [42]. This shift has been propelled by the rapid development of deep learning. In early work, only feed-forward DNNs with fully-connected layers were utilized [30]. Convolutional neural networks, which can better model local signal patterns, were later leveraged [43]. Subsequently, recurrent neural networks with long short-term memory (LSTM) were introduced to model temporal patterns [44]. Recently, Transformers with attention mechanisms were designed to model long-term signal dependencies for SS [45]. Modern DNN architectures in SS often combine various DNN blocks to leverage their complementary power. For example, the popular convolution recurrent network [46] sandwiches an LSTM with a U-Net so that local and longer-term signal patterns can be integrated for separation. There are also fully convolutional networks such as **Conv-TasNet** [42], which operates on short frames in the time domain. Inspired by the seminal **dual-path recurrent neural networks (DPRNN)** [47], state-of-the-art DNN architectures in SS usually employ a dual- or multi-path architecture, where signal patterns are alternately modeled along different tensor axes. For example, in **TF-GridNet** [48], a sub-band temporal module and an intra-frame full-band module are alternately stacked to model the temporal and spatial patterns within each frequency and the spectral and spatial patterns in each frame.

SS is not limited to speech but can be applied to various types of audio signals. One such application is music, which often includes vocals and instruments. The first significant success was achieved with Open-Unmix [49], which employs LSTM and data augmentation for stereo music [50]. Later, combining hybrid architectures [51], using different data representations [52], and bridging multiple instrument networks [53] showed improved performance. Another application is mixed sound events. A seminal work is universal sound separation [54], which trained an improved Conv-TasNet via supervised PIT or unsupervised **mixture invariant training** [55] to unmix mixed sound events. Another direction is target sound extraction, where the auxiliary cues informing the DNN can be a binary vector indicating a subset of a pre-defined set of sound events to separate [56] or natural language which can offer more flexible prompting [57].

C. Hybrid SS

The classical way to exploit spatial information present in multi-channel input is to use a beamformer that points a beam of increased sensitivity towards the source of interest. The issue with SS through beamforming is that the beamformer weight computation requires knowledge of the statistics of the source signals, which, naturally, is not readily available if only the mixture is given. A popular way to obtain this information is through TF mask estimation, whereby the spatial covariance matrix of a source is estimated from the TF bins that the mask estimator has classified to be dominated by that source. TF mask estimation can be done either by clustering spatial features as described earlier or by using a DNN that learns the spectro-temporal patterns of the source signals. The latter, which employs a DNN for parameter estimation and beamformers for source extraction, is an example of a hybrid approach, that blends signal processing with deep learning. This hybrid method was first introduced in the CHiME-3 challenge, where it was extremely successful [58].

Hybrid techniques can also be used for mask estimation itself, because spatial clustering and DNNs exploit different signal properties: the former takes advantage of spatial and the latter of spectral information. The former relies on unsupervised, while the latter on supervised learning. These complementary properties have been exploited in various ways, e.g., for overcoming domain mismatch

between training and test [59] or for training a DNN for SS on real mixtures, where the separated signals as a training target are not available [60], [61].

IV. INITIATIVES THAT FOSTERED SS RESEARCH

In the 2000s, there were no benchmarks in SS research, and researchers often used their own datasets and evaluation criteria. This made it difficult to discuss the advantages and disadvantages of technologies and to make objective comparisons, and thus a common language and common base were needed in the community. A number of efforts have been made to remedy this situation and foster SS research, including challenges, performance metrics, datasets, and open-source software. This section introduces these key efforts.

The International Conference on Independent Component Analysis and Blind Signal Separation (renamed the International Conference on Latent Variable Analysis and Signal Separation in 2010), held every 1.5 years since 1999, has also played an important role in developing the fundamental theory of SS and its applications to audio.

A. Challenges

In conjunction with the conference, Signal Separation Evaluation Campaign (SiSEC) was launched in 2007 to benchmark various separation models and continued until 2018 [62]. The SiSEC addressed the above issue by providing a dataset, evaluation metrics, and sample software codes. SiSEC originated from the community-based Signal Separation Evaluation Campaign (SASSEC)¹, which was organized in 2007 and provided development and test sets for speech and music. Since then, it has been organized in 2008, 2010, 2011, 2013, 2015, 2016, and 2018 (see references in [62], [63]). Although the data size of the SiSECs was small compared to today, it created a culture of benchmarking and challenges in the SS research field.

This SiSEC initiative was taken over by CHiME². The first PASCAL CHiME challenge (CHiME-1) in 2010 was on separating and recognizing speech in everyday listening conditions. Since then, CHiME has continued to provide challenge projects using data recorded in real environments and to lead the SS field.

Subsequently, the music separation (MUS) track of SiSEC was succeeded by a crowd-based competition called the Music Demixing (MDX) Challenge 2021 [64]. This was followed by an expanded edition, the Sound Demixing (SDX) Challenge 2023, featuring both music and cinematic demixing tracks [65], [66].

B. Evaluation metrics

Vincent *et al.* [67] introduced a well-known set of evaluation metrics, encompassing the source-to-distortion ratio (SDR), source-to-interference ratio (SIR), and sources-to-artifacts ratio (SAR). These metrics are based on decomposing each separated signal into components corresponding to the target sound source, residual interference from the other sources, and artifacts introduced during separation, such as musical noise. Implementations are available in the widely-used BSS Eval Matlab toolbox³ and in Python libraries^{4,5}. Additionally, a modified metric known as the scale-invariant SDR (SI-SDR) [68] was proposed, which allows only a rescaling factor, rather than a finite impulse response filter, to align the target source signal with each separated signal. The word error rate (WER) is often used to measure the effectiveness of SS as an automatic speech recognition

¹<https://www.irisa.fr/metiss/SASSEC07/>

²<https://www.chimechallenge.org/>

³https://gitlab.inria.fr/bass-db/bss_eval/

⁴https://craffel.github.io/mir_eval/

⁵<https://github.com/sigsep/sigsep-mus-eval/>

(ASR) front-end. Short-time objective intelligibility (STOI) [69] is an objective measure of speech intelligibility based on the correlation coefficient between the short-time temporal envelopes of the target and the separated speech. The mean opinion score (MOS) measures perceptual quality using human subjective evaluation scores collected in a listening test. Since obtaining subjective scores is time-consuming and costly, objective metrics, such as perceptual evaluation of speech quality (PESQ) [70], are often used to predict perceptual quality.

C. Datasets and open source

As deep learning approaches began to dominate SS research, publicly available datasets became increasingly necessary to enable reproducible scientific research. Particularly important examples include CHiME and MUSDB18 [62]⁶, which served as the basis for the challenges described in Sec. IV-A. Because SS training requires both source signals and mixtures, scripts that combine audio signals from existing datasets (e.g., speech datasets originally collected for ASR) further aided reproducibility. For example, speech SS research was facilitated by scripts for creating the wsj0-2mix dataset [31], which was later extended to WHAMR! [71] in noisy and reverberant scenarios and to LibriMix [72] with proprietary speech data replaced with open data. SMS-WSJ [73] enabled evaluation as an ASR front-end and that of multi-channel SS methods. With the recent release of the EARS dataset [74], there now exists a dataset with high sampling rate speech recordings in anechoic conditions, which should further advance the field.

In addition to publicly available datasets, the proliferation of the open source software ecosystem has aided SS research as it has in many other fields. Notable tools include pyroomacoustics [75] for simulating audio sources in reverberant environments, gpurir [76] for enabling reverberation modeling on the GPU, and Nara-WPE [60] for providing one of the first open implementations of a state-of-the-art dereverberation algorithm. For deep learning algorithms, the open-sourcing of training recipes and pre-trained models has been the key to accelerating research, such as Open-Unmix [49], Asteroid [77], and ESPnet-SE [78].

V. REMAINING CHALLENGES AND POTENTIAL DIRECTIONS

SS performance on many simulated benchmarks under constrained setups (e.g., fully overlapped speech signals, a given number of sources) has saturated. The research community has been moving toward more and more realistic conditions, such as real-recorded conversational speech in CHiME-7 and 8. DNN-based SS algorithms have shown limited success in such real-world settings, where there remain two main challenges. First, simulated data used for supervised DNN training are usually mismatched with real data, which may cause poor generalization to real data. To alleviate this, recent efforts in unsupervised, weakly supervised, and semi-supervised SS [55], [60], [61], [79], [80] aim to leverage real (“in-the-wild”) data without ground-truth sources. Generative models, such as diffusion models [81], may also handle out-of-domain data. Second, separating an unknown and time-varying number of sources necessitates source activity detection (e.g., speaker diarization, sound event detection) and needs further exploration.

Other challenges include improving the performance of SS with a single or a limited number of microphones. Additionally, the development of evaluation metrics applicable to real-world mixtures (i.e., reference-free) that strongly correlate with human perception for

not only speech signals, but also music and general sounds remains challenging. Moreover, as in other areas, developing lightweight SS models for edge devices and pursuing low-latency algorithms for real-time applications continue to be or become increasingly important issues. Given that sound sources often move in applications such as robot audition, multi-channel SS for moving sources remains a significant challenge. Another promising avenue is the synchronization of observed mixtures across independent devices to form a distributed microphone array. We look forward to further development of SS research in the following decades.

Acknowledgement We would like to thank Dr. Mike Goodwin for providing the EDICS history of SS.

REFERENCES

- [1] S. Makino *et al.*, Eds., *Blind Speech Separation*. Springer, 2007.
- [2] E. Vincent *et al.*, Eds., *Audio Source Separation and Speech Enhancement*. Wiley, 2018.
- [3] S. Makino, Ed., *Audio Source Separation*. Springer, 2018.
- [4] A. J. Bell and T. J. Sejnowski, “An information-maximization approach to blind separation and blind deconvolution,” *Neural Computation*, vol. 7, pp. 1129–1159, 1995.
- [5] T. Kim *et al.*, “Independent vector analysis: An extension of ICA to multivariate components,” in *Proc. ICA*, 2006, pp. 165–172.
- [6] A. Hiroe, “Solution of permutation problem in frequency domain ICA, using multivariate probability density functions,” in *Proc. ICA*, 2006, pp. 601–608.
- [7] N. Ono, “Stable and fast update rules for independent vector analysis based on auxiliary function technique,” in *Proc. WASPAA*, 2011, pp. 189–192.
- [8] D. Kitamura, N. Ono, H. Sawada, H. Kameoka, and H. Saruwatari, “Determined blind source separation unifying independent vector analysis and nonnegative matrix factorization,” *IEEE/ACM TASLP*, vol. 24, no. 9, pp. 1626–1641, 2016.
- [9] O. Yilmaz and S. Rickard, “Blind separation of speech mixtures via time-frequency masking,” *IEEE Transactions on Signal Processing*, vol. 52, no. 7, pp. 1830–1847, 2004.
- [10] N. Ito *et al.*, “Complex angular central Gaussian mixture model for directional statistics in mask-based microphone array signal processing,” in *Proc. EUSIPCO*, 2016, pp. 1153–1157.
- [11] M. Souden *et al.*, “A multichannel MMSE-based framework for speech source separation and noise reduction,” *IEEE TASLP*, vol. 21, no. 9, pp. 1913–1928, 2013.
- [12] T. Yoshioka *et al.*, “The NTT CHiME-3 system: Advances in speech enhancement and recognition for mobile multi-microphone devices,” in *Proc. ASRU*, 2015, pp. 436–443.
- [13] C. Boeddecker *et al.*, “Front-end processing for the CHiME-5 dinner party scenario,” in *Proc. CHiME*, 2018, pp. 35–40.
- [14] N. Q. K. Duong *et al.*, “Under-determined reverberant audio source separation using a full-rank spatial covariance model,” *IEEE TASLP*, vol. 18, no. 7, pp. 1830–1840, 2010.
- [15] S. Arberet *et al.*, “Nonnegative matrix factorization and spatial covariance model for under-determined reverberant audio source separation,” in *Proc. ISSPA*, 2010.
- [16] H. Sawada *et al.*, “Multichannel extensions of non-negative matrix factorization with complex-valued data,” *IEEE TASLP*, vol. 21, no. 5, pp. 971–982, 2013.
- [17] Y. Bando *et al.*, “Neural full-rank spatial covariance analysis for blind source separation,” *IEEE Signal Processing Letters*, vol. 28, pp. 1670–1674, 2021.
- [18] N. Ito *et al.*, “A joint diagonalization based efficient approach to underdetermined blind audio source separation using the multichannel Wiener filter,” *IEEE/ACM TASLP*, vol. 29, pp. 1950–1965, 2021.
- [19] K. Sekiguchi *et al.*, “Fast multichannel nonnegative matrix factorization with directivity-aware jointly-diagonalizable spatial covariance matrices for blind source separation,” *IEEE/ACM TASLP*, vol. 28, pp. 2610–2625, 2020.
- [20] S. Rickard, “The DUET blind source separation algorithm,” in *Blind speech separation*. Springer, 2007, pp. 217–241.
- [21] D. Barry, B. Lawlor, and E. Coyle, “Sound source separation: Azimuth discrimination and resynthesis,” in *Proc. DAFX*, 2004, pp. 240–244.

⁶<https://zenodo.org/records/1117372>

- [22] D. FitzGerald *et al.*, "Projection-based demixing of spatial audio," *IEEE/ACM TASLP*, vol. 24, no. 9, pp. 1560–1572, 2016.
- [23] J. R. Hershey *et al.*, "Super-human multi-talker speech recognition: A graphical modeling approach," *Computer Speech and Language*, vol. 24, no. 1, pp. 45–66, 2010.
- [24] D. Lee and H. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, 1999.
- [25] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria," *IEEE TASLP*, vol. 15, no. 3, pp. 1066–1074, 2007.
- [26] Z. Rafii and B. Pardo, "Repeating pattern extraction technique (REPET): A simple method for music/voice separation," *IEEE TASLP*, vol. 21, no. 1, pp. 73–84, 2013.
- [27] D. FitzGerald, "Harmonic/percussive separation using median filtering," in *Proc. DAFx*, 2010.
- [28] S. Ewert *et al.*, "Score-informed source separation for musical audio recordings: An overview," *IEEE Signal Processing Magazine*, vol. 31, no. 3, pp. 116–124, 2014.
- [29] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM TASLP*, vol. 26, pp. 1702–1726, 2017.
- [30] Y. Wang and D. Wang, "Towards scaling up classification-based speech separation," *IEEE TASLP*, no. 7, pp. 1381–1390.
- [31] J. R. Hershey *et al.*, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. ICASSP*, 2016.
- [32] D. Yu, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *Proc. ICASSP*, 2017, pp. 241–245.
- [33] A. Ephrat *et al.*, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *ACM TOG*, vol. 37, no. 4, apr 2018.
- [34] K. Zmolikova *et al.*, "Neural target speech extraction: An overview," *IEEE Signal Processing Magazine*, vol. 40, no. 3, pp. 8–29, 2023.
- [35] Z.-Q. Wang and D. Wang, "Combining spectral and spatial features for deep learning based blind speaker separation," *IEEE/ACM TASLP*, vol. 27, no. 2, pp. 457–468, 2019.
- [36] Z.-Q. Wang *et al.*, "Multi-microphone complex spectral mapping for utterance-wise and continuous speaker separation," *IEEE/ACM TASLP*, vol. 29, pp. 2001–2014, 2021.
- [37] C. L. Liu *et al.*, "Multichannel speech enhancement by raw waveform-mapping using fully convolutional networks," *IEEE/ACM TASLP*, vol. 28, pp. 1888–1900, 2020.
- [38] Y. Luo *et al.*, "End-to-end microphone permutation and number invariant multi-channel speech separation," in *Proc. ICASSP*, 2020, pp. 6394–6398.
- [39] T. Yoshioka *et al.*, "VarArray: Array-geometry-agnostic continuous speech separation," in *Proc. ICASSP*, 2022, pp. 6027–6031.
- [40] D. S. Williamson *et al.*, "Complex ratio masking for monaural speech separation," *IEEE/ACM TASLP*, pp. 483–492, 2016.
- [41] K. Tan and D. Wang, "Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement," *IEEE/ACM TASLP*, vol. 28, pp. 380–390, 2020.
- [42] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM TASLP*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [43] D. Stoller *et al.*, "Wave-U-Net: A multi-scale neural network for end-to-end audio source separation," in *Proc. ISMIR*, 2018, pp. 334–340.
- [44] H. Erdogan *et al.*, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. ICASSP*, 2015, pp. 708–712.
- [45] C. Subakan *et al.*, "Attention is all you need in speech separation," in *Proc. ICASSP*, 2021, pp. 21–25.
- [46] K. Tan and D. Wang, "A convolutional recurrent neural network for real-time speech enhancement," in *Interspeech*, 2018, pp. 3229–3233.
- [47] Y. Luo *et al.*, "Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation," in *Proc. ICASSP*, 2020, pp. 46–50.
- [48] Z.-Q. Wang *et al.*, "TF-GridNet: Integrating full- and sub-band modeling for speech separation," *IEEE/ACM TASLP*, vol. 31, pp. 3221–3236, 2023.
- [49] F.-R. Stöter *et al.*, "Open-Unmix - a reference implementation for music source separation," *Journal of Open Source Software*, 2019.
- [50] S. Uhlich *et al.*, "Improving music source separation based on deep neural networks through data augmentation and network blending," in *Proc. ICASSP*, 2017, pp. 261–265.
- [51] N. Takahashi *et al.*, "MMDenseLSTM: An efficient combination of convolutional and recurrent neural networks for audio source separation," in *Proc. IWAENC*, 2018, pp. 106–110.
- [52] A. Défossez, "Hybrid spectrogram and waveform source separation," *arXiv:2111.03600*, 2022.
- [53] R. Sawata *et al.*, "The whole is greater than the sum of its parts: improving music source separation by bridging networks," *EURASIP J. Audio Speech Music. Process.*, vol. 2024, no. 1, p. 39, 2024.
- [54] I. Kavalerov *et al.*, "Universal sound separation," in *Proc. WASPAA2019*, 2019, pp. 175–179.
- [55] S. Wisdom *et al.*, "Unsupervised sound separation using mixture invariant training," in *Proc. NeurIPS*, vol. 33, 2020.
- [56] T. Ochiai *et al.*, "Listen to what you want: Neural network-based universal sound selector," in *Interspeech*, 2020, pp. 1441–1445.
- [57] X. Liu *et al.*, "Separate anything you describe," *arXiv preprint arXiv:2308.05037*, 2023.
- [58] J. Heymann *et al.*, "BLSTM supported GEV beamformer front-end for the 3rd CHiME challenge," in *Proc. ASRU*, 2015.
- [59] T. Nakatani *et al.*, "Integrating DNN-based and spatial clustering-based mask estimation for robust MVDR beamforming," in *Proc. ICASSP*, 2017, pp. 286–290.
- [60] L. Drude *et al.*, "Unsupervised training of a deep clustering model for multichannel blind source separation," in *Proc. ICASSP*, 2019.
- [61] E. Tzinis *et al.*, "Unsupervised deep clustering for source separation: Direct learning from mixtures using spatial information," in *Proc. ICASSP*, 2019, pp. 81–85.
- [62] F. R. Stöter *et al.*, "The 2018 signal separation evaluation campaign," in *Proc. LVA/ICA*, 2018, pp. 293–305.
- [63] E. Vincent *et al.*, "The signal separation evaluation campaign (2007–2010): Achievements and remaining challenges," *Signal Processing*, vol. 92, no. 8, pp. 1928–1936, 2012.
- [64] Y. Mitsufuji *et al.*, "Music demixing challenge 2021," *Frontiers in Signal Processing*, vol. 1, 2022.
- [65] G. Fabbro *et al.*, "The sound demixing challenge 2023 - music demixing track," *TISMIR*, vol. 7, no. 1, pp. 63–84, 2024.
- [66] S. Uhlich *et al.*, "The sound demixing challenge 2023 - cinematic demixing track," *TISMIR*, vol. 7, no. 1, pp. 44–62, 2024.
- [67] E. Vincent *et al.*, "Performance measurement in blind audio source separation," *IEEE TASLP*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [68] J. Le Roux *et al.*, "SDR – half-baked or well done?" in *Proc. ICASSP*, 2019, pp. 626–630.
- [69] C. H. Taal *et al.*, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE TASLP*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [70] A. W. Rix, M. P. Hollier, A. P. Hekstra, and J. G. Beerends, "Perceptual evaluation of speech quality (PESQ) the new ITU standard for end-to-end speech quality assessment part I–time-delay compensation," *Journal of the Audio Engineering Society*, vol. 50, no. 10, pp. 755–764, 2002.
- [71] M. Maciejewski *et al.*, "WHAMR!: Noisy and reverberant single-channel speech separation," in *Proc. ICASSP*, 2020.
- [72] J. Cosentino *et al.*, "LibriMix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [73] L. Drude *et al.*, "SMS-WSJ: Database, performance measures, and baseline recipe for multi-channel source separation and recognition," *arXiv preprint arXiv:1910.13934*, 2019.
- [74] J. Richter *et al.*, "EARS: An anechoic fullband speech dataset benchmarked for speech enhancement and dereverberation," in *Interspeech*, 2024.
- [75] R. Scheibler *et al.*, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *Proc. ICASSP*, 2018.
- [76] D. Diaz-Guerra *et al.*, "gpuRIR: A python library for room impulse response simulation with GPU acceleration," *Multimedia Tools and Applications*, vol. 80, no. 4, pp. 5653–5671, 2021.
- [77] M. Pariente *et al.*, "Asteroid: the PyTorch-based audio source separation toolkit for researchers," in *Interspeech*, 2020.
- [78] C. Li *et al.*, "ESPnet-SE: End-to-end speech enhancement and separation toolkit designed for asr integration," in *Proc. SLT*, 2021, pp. 785–792.
- [79] Z. Huang *et al.*, "Investigating self-supervised learning for speech enhancement and separation," in *Proc. ICASSP*, 2022, pp. 6837–6841.
- [80] Z.-Q. Wang, "SuperM2M: Supervised and mixture-to-mixture co-learning for speech enhancement and robust ASR," *arXiv preprint arXiv:2403.10271v2*, 2024.
- [81] R. Scheibler *et al.*, "Diffusion-based generative speech source separation," in *Proc. ICASSP 2023*, 2023, pp. 1–5.