

A review of speaker diarization: Recent advances with deep learning

Tae Jin Park ^{a,*}, Naoyuki Kanda ^{b,1}, Dimitrios Dimitriadis ^{b,1}, Kyu J. Han ^{c,1}, Shinji Watanabe ^{d,1}, Shrikanth Narayanan ^a

^a University of Southern California, Los Angeles, USA

^b Microsoft, Redmond, USA

^c ASAPP, Mountain View, USA

^d Johns Hopkins University, Baltimore, USA

ARTICLE INFO

Keywords:

Speaker diarization
Automatic speech recognition
Deep learning

ABSTRACT

Speaker diarization is a task to label audio or video recordings with classes that correspond to speaker identity, or in short, a task to identify “who spoke when”. In the early years, speaker diarization algorithms were developed for speech recognition on multispeaker audio recordings to enable speaker adaptive processing. These algorithms also gained their own value as a standalone application over time to provide speaker-specific meta-information for downstream tasks such as audio retrieval. More recently, with the emergence of deep learning technology, which has driven revolutionary changes in research and practices across speech application domains, rapid advancements have been made for speaker diarization. In this paper, we review not only the historical development of speaker diarization technology but also the recent advancements in neural speaker diarization approaches. Furthermore, we discuss how speaker diarization systems have been integrated with speech recognition applications and how the recent surge of deep learning is leading the way of jointly modeling these two components to be complementary to each other. By considering such exciting technical trends, we believe that this paper is a valuable contribution to the community to provide a survey work by consolidating the recent developments with neural methods and thus facilitating further progress toward a more efficient speaker diarization.

1. Introduction

“Diarize” means making a note or keeping an event in a diary. Speaker diarization, like keeping a record of events in such a diary, addresses the question of “who spoke when” (Tranter et al., 2003; Tranter and Reynolds, 2006; Anguera et al., 2012) by logging speaker-specific salient events on multiparticipant (or multispeaker) audio data. Throughout the diarization process, the audio data would be divided and clustered into groups of speech segments with the same speaker identity/label. As a result, salient events, such as non-speech/speech transition or speaker turn changes, are automatically detected. In general, this process does not require any prior knowledge of the speakers, such as their real identity or the number of participating speakers in the audio data. Thanks to its feature of separating audio streams by these speaker-specific events, speaker diarization can be effectively employed for indexing or analyzing various types of audio data, e.g., audio/video broadcasts from media stations, conversations in conferences, personal videos from online social media or hand-held devices, court proceedings, business meetings, earnings reports in a financial sector, just to name a few.

* Corresponding author.

E-mail address: taejinpa@usc.edu (T.J. Park).

¹ Authors contributed equally.

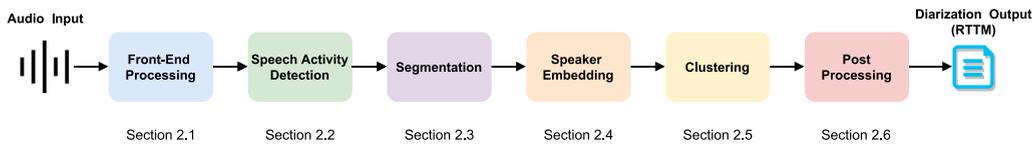


Fig. 1. Traditional speaker diarization system.

Traditionally speaker diarization systems consist of multiple, independent sub-modules as presented in Fig. 1. To mitigate any artifacts in acoustic environments, various front-end processing techniques, for example, speech enhancement, dereverberation, speech separation or target speaker extraction, are employed. Voice or speech activity detection (SAD) is then applied to separate speech from non-speech events. Raw speech signals in the selected speech portion are transformed to acoustic features or embedding vectors. In the clustering stage, the transformed speech portions are grouped and labeled by speaker classes and in the post-processing stage, the clustering results are further refined. Each of these sub-modules is optimized individually in general.

1.1. Historical development of speaker diarization

During the early years of diarization technology (in the 1990s), the research objective was to benefit automatic speech recognition (ASR) on air traffic control dialogues and broadcast news recordings, by separating each speaker's speech segments and enabling speaker-adaptive training of acoustic models (Gish et al., 1991; Siu et al., 1992; Rohlicek et al., 1992; Jain et al., 1996; Padmanabhan et al., 1996; Gauvain et al., 1998; Liu and Kubala, 1999). In this period some fundamental approaches for measuring the distance between speech segments for speaker change detection and clustering, such as generalized likelihood ratio (GLR) (Gish et al., 1991) and Bayesian information criterion (BIC) (Chen and Gopalakrishnan, 1998), were developed and quickly became the golden standard. All these efforts collectively laid out paths to consolidate activities across research groups worldwide, leading to several research consortia and challenges in the early 2000s, among which there were the Augmented Multiparty Interaction (AMI) Consortium (AMI, 2009) supported by the European Commission and the RT Evaluation (NIST, 2009) hosted by the National Institute of Standards and Technology (NIST). These organizations, spanning over from a few years to a decade, fostered further advancements on speaker diarization technologies across different data domains from broadcast news (Ajmera and Wooters, 2003; Tranter and Reynolds, 2004; Reynolds and Torres-Carrasquillo, 2005; Zhu et al., 2005; Meignier et al., 2006) and conversational telephone speech (CTS) (Rosenberg et al., 2002; Liu and Kubala, 2003; Tranter et al., 2004; Kenny et al., 2010) to meeting conversations (Ajmera et al., 2004; Jin et al., 2004; Anguera et al., 2006; Leeuwen and Konecny, 2007; Vijayasenan et al., 2009). The new approaches resulting from these advancements include, but not limited to, beamforming (Anguera et al., 2007), information bottleneck clustering (IBC) (Vijayasenan et al., 2009), variational Bayesian (VB) approaches (Valente et al., 2010), joint factor analysis (JFA) (Kenny et al., 2010).

Speaker specific representation in a total variability space derived from simplified JFA, known as i-vector (Dehak et al., 2011), found great success in speaker recognition and was quickly adopted by speaker diarization systems as feature representation for short speech segments, segmented in an unsupervised fashion. i-Vector successfully replaced its predecessors such as merely mel-frequency cepstral coefficient (MFCC) or speaker factors (or eigenvoices) (Castaldo et al., 2008) to bolster clustering performance in speaker diarization, being combined with principal component analysis (PCA) (Shum et al., 2011, 2012), variational Bayesian Gaussian mixture model (VB-GMM) (Shum et al., May 2013), mean shift (Senoussaoui et al., 2013b) and probabilistic linear discriminant analysis (PLDA) (Sell and Garcia-Romero, 2014).

Since the advent of deep learning in the 2010s, there has been a considerable amount of research to take advantage of powerful modeling capabilities of the neural networks for speaker diarization. One representative example is the extraction of the speaker embeddings using neural networks, such as the d-vectors (Variani et al., 2014; Heigold et al., 2016; Wang et al., 2018) or the x-vectors (Snyder et al., 2018), which most often are embedding vector representations based on the bottleneck layer output of a deep neural network (DNN) trained for speaker recognition. The shift from i-vector to these neural embeddings contributed to enhanced performance, easier training with more data (Zhang et al., 2019), and robustness against speaker variability and acoustic conditions. More recently, end-to-end neural diarization (EEND) where individual sub-modules in the traditional speaker diarization systems (c.f., Fig. 1) can be replaced by one neural network gets more attention with promising results (Fujita et al., 2019,b). This research direction, although not fully matured yet, could open up unprecedented opportunities to address challenges in the field of speaker diarization, such as, the joint optimization with other speech applications, with overlapping speech, if large-scale data is available for training such powerful neural network-based models.

1.2. Motivation

Till now, there are two well-rounded overview papers in the area of speaker diarization that survey the development of speaker diarization technology with different focuses. In Tranter and Reynolds (2006), various speaker diarization systems and their subtasks in the context of broadcast news and CTS data are reviewed up to till mid 2000s. Thus, the historical progress of speaker diarization technology development in the 1990s and early 2000s are covered. Contrarily, the focus of Anguera et al. (2012) was put more on speaker diarization for meeting speech and its respective challenges. This paper thus weighs more in the corresponding technologies

Table 1
Table of Taxonomy.

	Trained based on Non-diarization objective	Trained based on Diarization objective
Single-module optimization	Sections 2.1–2.6 Front-end (Hershey et al., 2016; Kolbæk et al., 2017; Luo and Mesgarani, 2019), speaker embedding (Variani et al., 2014b; Snyder et al., 2017, 2018), SAD (Drugman et al., 2015), etc.	Section 3.1 Affinity matrix refinement (Wang et al., 2020), IDEC (Dimitriadis, 2019), TS-VAD (Medennikov et al., 2020a), etc.
Joint optimization	Section 2.7 VB-HMM (Diez et al., 2018a), VBx (Diez et al., 2019b) Out of scope Joint front-end & ASR (Yu et al., 2017; Seki et al., 2018; Chang et al., 2019; Kanda et al., 2019b,d; Wang et al., 2021), joint speaker identification & speech separation (Wang et al., 2019; Han et al., 2020), etc.	Section 3.2 UIS-RNN (Zhang et al., 2019), RPN (Huang et al., 2020), online RSAN (von Neumann et al., 2019), EEND (Fujita et al., 2019,b), etc. Section 4 Joint ASR & speaker diarization. (Shafey et al., 2019; Mao et al., 2020; Kanda et al., 2019c, 2021a), etc.

to mitigate problems from the perspective of meeting environments, where there are usually more participants than broadcast news or CTS data and multi-modal data is frequently available. Since these two papers were published, speaker diarization systems have gone through a lot of notable changes, especially from the leap-frog advancements in deep learning approaches addressing technical challenges across multiple machine learning domains. We believe that this survey work is a valuable contribution to the community to consolidate the recent developments with neural methods and thus facilitate further progress toward a more efficient diarization.

1.3. Overview and taxonomy of speaker diarization

Attempting to categorize the existing, most-diverse speaker diarization technologies, both in the space of modularized speaker diarization systems before the deep learning era and those based on neural networks of recent years, a proper grouping would be helpful. The main categorization we adopt in this paper is based on two criteria, resulting total of four categories, as shown in Table 1. The first criterion is whether the model is trained based on speaker diarization-oriented objective function or not. Any trainable approaches to optimize models in a multispeaker situation and learn relations between speakers are categorized into the “diarization objective” class. The second criterion is whether multiple modules are jointly optimized toward some objective function. If a single sub-module is replaced into a trainable one, such method is categorized into the “Single-module optimization” class. Conversely, joint modeling of segmentation and clustering (Zhang et al., 2019), joint modeling of speech separation and speaker diarization (von Neumann et al., 2019) or fully end-to-end neural diarization system (Fujita et al., 2019,b) is categorized into the “Joint optimization” class.

Note that our intention of this categorization is to help readers to quickly overview the broad development in the field, and it is not our intention to divide the categories into superior–inferior.

Also, while we are aware of many techniques that fall into the category “Non-Diarization Objective” and “Joint Optimization” (e.g., joint front-end and ASR (Yu et al., 2017; Seki et al., 2018; Chang et al., 2019; Kanda et al., 2019b,d; Wang et al., 2021) and joint speaker identification and speech separation (Wang et al., 2019; Han et al., 2020)), we exclude them in the paper to focus on the review of speaker diarization techniques.

1.4. Diarization evaluation metrics

1.4.1. Diarization error rate

The accuracy of speaker diarization system is measured using diarization error rate (DER) (Fiscus et al., 2006) where DER is the sum of three different error types: False alarm (FA) of speech, missed detection of speech and confusion between speaker labels.

$$\text{DER} = \frac{\text{FA} + \text{Missed} + \text{Speaker-Confusion}}{\text{Total Duration of Time}}. \quad (1)$$

To establish a one-to-one mapping between the hypothesis outputs and the reference transcript, Hungarian algorithm (Kuhn, 1955) is employed. In the 2006 RT evaluation (Fiscus et al., 2006), 0.25 s of “no score” collar (also referred to as “score collar”) is set around every boundary of reference segment to mitigate the effect of inconsistent notation and human errors in reference transcript and this evaluation scheme has been most widely used in speaker diarization studies.

1.4.2. Jaccard error rate

The Jaccard error rate (JER) was first introduced in the DIHARD II evaluation. The goal of JER is to evaluate each speaker with equal weight. Unlike DER which is estimated for the whole utterance altogether, per-speaker error rates are first computed and then averaged to compute JER. Specifically, JER is computed as follows.

$$\text{JER} = \frac{1}{N} \sum_i^{N_{ref}} \frac{\text{FA}_i + \text{MISS}_i}{\text{TOTAL}_i}. \quad (2)$$

In Eq. (2), TOTAL_i is the union of the i th speaker's speaking time in the reference transcript and the i th speaker's speaking time in the hypotheses. N_{ref} is the number of speakers in the reference script. Note that the Speaker-Confusion in DER is reflected in the part of FA_i in the calculation of JER. Since JER is using union operation between reference and the hypotheses, JER never exceeds 100%, whereas DER can become much larger than 100%. DER and JER are highly correlated but if a subset of speakers are dominant in the given audio recording, JER tends to be higher than the ordinary case.

1.4.3. Word-level diarization error rate

While DER is based on the duration of the speaking time of each speaker, word-level DER (WDER) is designed to measure the error that is caused in the lexical (output transcription) side. The motivation of WDER is the discrepancy between DER and the accuracy of the final transcript output since DER relies on the duration of the speaking time that is not always aligned with the word boundaries. The concept of word-breakage ratio was proposed in [Silovsky et al. \(2012\)](#) where word-breakage shares similar idea with WDER. Unlike WDER, word-breakage ratio measures the number of speaker-change points occur inside a word boundary. The work in [Park and Georgiou \(2018\)](#) suggested the term WDER, evaluating the diarization output with ground-truth transcription. More recently, the joint ASR and speaker diarization system was evaluated in the WDER format in [Shafey et al. \(2019\)](#). Although the way of calculating WDER would differ over the studies, but the underlying mechanism is that the diarization error is calculated by counting the correctly or incorrectly labeled words.

1.5. Paper organization

The rest of the paper is organized as follows.

- In Section 2, we overview techniques belonging to the “Non-diarization objective” class in the proposed taxonomy, mostly those used in the traditional, modular speaker diarization systems. While there are some overlaps with the counterpart sections of the aforementioned two survey papers ([Tranter and Reynolds, 2006](#); [Anguera et al., 2012](#)) in terms of reviewing notable developments in the past, this section would add more latest schemes in the corresponding components of the speaker diarization systems.
- In Section 3, we discuss advancements mostly leveraging DNNs trained with the diarization objective where single sub-modules are independently optimized (Section 3.1) or jointly optimized (Section 3.2) toward fully end-to-end speaker diarization.
- In Section 4, we present a perspective of how speaker diarization has been investigated in the context of ASR, reviewing historical interactions between these two domains to peek into the past, present and future of speaker diarization applications.
- Section 5 provides information on speaker diarization challenges and corpora to facilitate research activities and anchor technology advances. We also discuss evaluation metrics such as DER, JER and Word-level DER (WDER) in this section.
- We share a few examples of how speaker diarization systems are employed in both research and industry practices in Section 6 and conclude this work in Section 7, providing summary and future challenges in speaker diarization.

2. Modular speaker diarization systems

This section provides an overview of algorithms for speaker diarization belonging to the “Non-diarization objective” class, as shown in [Table 1](#). Each subsection in this section corresponds to the explanation of each module in the traditional speaker diarization system, as shown in [Fig. 1](#). In addition to the introductory explanation of each module, this section also summarizes the recent techniques within the module.

2.1. Front-end processing

This section describes mostly front-end techniques, used for speech enhancement, dereverberation, speech separation, and speech extraction as part of the speaker diarization pipeline. Let $s_{i,f,t} \in \mathbb{C}$ be the short-time Fourier Transform (STFT) representation of the source speaker i on the frequency bin f at frame t . The observed noisy signal $x_{t,f}$ can be represented by a mixture of the source signals, a room impulse response $h_{i,f,t} \in \mathbb{C}$, and additive noise $n_{t,f} \in \mathbb{C}$,

$$x_{t,f} = \sum_{i=1}^K \sum_{\tau} h_{i,f,\tau} s_{i,f,t-\tau} + n_{t,f}, \quad (3)$$

where K denotes the number of speakers present in the audio signal.

The aim of the front-end techniques described in this section is to estimate the original source signal $\hat{x}_{i,t}$ given the observation $\mathbf{X} = (\{x_{t,f}\}_f)_t$ for the downstream diarization task,

$$\hat{x}_{i,t} = \text{FrontEnd}(\mathbf{X}), \quad i = 1, \dots, K, \quad (4)$$

where $\hat{x}_{i,t} \in \mathbb{C}^D$ denotes the i th speaker's estimated STFT spectrum with D frequency bins at frame t .

Although there are numerous speech enhancement, dereverberation, and separation algorithms, e.g., Haeb-Umbach et al. (2019), Vincent et al. (2018), Wang and Chen (2018), herein most of the recent techniques used in the DIHARD challenge series (Sell et al., 2018b; Ryant et al., 2019; Diez et al., 2018b), LibriCSS meeting recognition task (Chen et al., 2020; Raj et al., 2021), and the CHiME-6 challenge track 2 (Watanabe et al., 2020; Arora et al., 2020; Medennikov et al., 2020b) are covered.

2.1.1. Speech enhancement and denoising

Speech enhancement techniques focus mainly on the suppression of the noise component of the noisy speech, which has shown a significant improvement thanks to deep learning.

For example, long short-term memory (LSTM)-based speech enhancement (Gao et al., 2018; Erdogan et al., 2015) is used as a front-end technique in the DIHARD II baseline (Ryant et al., 2019), i.e.,

$$\hat{x}_t = \text{LSTM}(\mathbf{X}), \quad (5)$$

where we only consider the single source example (i.e., $K = 1$) and omit the source index i . This is a regression-based approach by minimizing the objective function,

$$\mathcal{L}_{\text{MSE}} = \|s_t - \hat{x}_t\|^2. \quad (6)$$

The log power spectrum or ideal ratio mask is often used as the target domain of the output s_t . Also, the speech enhancement used in Loizou (2013) applies this objective function in each layer based on a progressive manner.

The effectiveness of the speech enhancement techniques can be enhanced via multichannel processing, including minimum variance distortionless response (MVDR) beamforming (Haeb-Umbach et al., 2019). Raj et al. (2021) demonstrates the significant improvement of the DER from 18.3% to 13.9% in the LibriCSS meeting task based on mask-based MVDR beamforming (Heymann et al., 2016; Erdogan et al., 2016).

2.1.2. Dereverberation

Compared with other front-end techniques, the major dereverberation techniques used in various tasks are based on statistical signal processing methods. One of the most widely used techniques is the weighted prediction error (WPE) based dereverberation (Nakatani et al., 2010; Yoshioka and Nakatani, 2012; Drude et al., 2018).

The basic idea of WPE, for the case of single source, i.e., $K = 1$, without noise, is to decompose the original signal model Eq. (3) into the early reflection $x_{t,f}^{\text{early}}$ and late reverberation $x_{t,f}^{\text{late}}$ as follows:

$$x_{t,f} = \sum_{\tau} h_{f,\tau} s_{f,t-\tau} = x_{t,f}^{\text{early}} + x_{t,f}^{\text{late}}. \quad (7)$$

WPE tries to estimate filter coefficients $\hat{h}_{f,t}^{\text{wpe}} \in \mathbb{C}$, which maintain the early reflection while suppressing the late reverberation based on the maximum likelihood estimation.

$$\hat{x}_{t,f}^{\text{early}} = x_{t,f} - \sum_{\tau=\Delta}^L \hat{h}_{f,\tau}^{\text{wpe}} x_{f,t-\tau}, \quad (8)$$

where Δ denotes the number of frames to split the early reflection and late reverberation, and L denotes the filter size.

WPE is widely used as one of the gold standard front-end processing methods, e.g., it is part of the DIHARD and CHiME, which are both the baseline and the top-performing systems (Sell et al., 2018b; Ryant et al., 2019; Diez et al., 2018b; Watanabe et al., 2020; Arora et al., 2020). Although the performance improvement of WPE-based dereverberation is not significant, it provides solid performance improvement across almost all tasks. Moreover, WPE is based on linear filtering and since it does not introduce signal distortions, it can be safely combined with downstream front-end and back-end processing steps. Similar to the speech enhancement techniques, WPE-based dereverberation demonstrates additional performance improvements when applied on multichannel signals.

2.1.3. Speech separation

Speech separation is a promising family of techniques when the overlapping speech regions are significant. The effectiveness of multichannel speech separation based on beamforming has been widely confirmed (Anguera et al., 2007; Yoshioka et al., 2018; Boeddecker et al., 2018). For example, in the CHiME-6 challenge (Watanabe et al., 2020), guided source separation (GSS) (Boeddecker et al., 2018) based multichannel speech extraction techniques have been used to achieve the top result. On the other hand, single-channel speech separation techniques (Hershey et al., 2016; Kolbæk et al., 2017; Luo and Mesgarani, 2019) do not often show any significant effectiveness in realistic multispeaker scenarios, such as the LibriCSS (Chen et al., 2020) or the CHiME-6 tasks (Watanabe et al., 2020), where speech signals are continuous and contain both overlapping and overlap-free speech regions. The single-channel speech separation systems often produce a redundant non-speech or even a duplicated speech signal for the non-overlap regions, and as such the ‘‘leakage’’ of audio causes many false alarms of speech activity. A leakage filtering method was proposed in Xiao et al. (2020) tackle the problem, where a significant improvement in the diarization performance was observed after including this processing step in the top-ranked system on the VoxCeleb Speaker Recognition Challenge 2020 (Nagrani et al., 2020).

2.2. Speech activity detection

SAD, also known as voice activity detection (VAD), distinguishes speech from non-speech such as background noise. SAD plays a significant role not only in speaker diarization but also in speaker recognition and speech recognition systems since SAD is a pre-processing step that can create errors that propagate through the whole pipeline. A SAD system consists mostly of two major parts. The first one is a feature extraction front-end, where acoustic features such as zero crossing rate (Itu, 1996), pitch (Chengalvarayan, 1999), signal energy (Woo et al., 2000), higher order statistics in the linear predictive coding residual domain Nemer et al. (2001) or MFCCs are often used. The other part is a classifier, where a model predicts and decides whether the input-frame contains speech or not. A system based on statistical models on spectrum (Sohn et al., 1999), Gaussian mixture models (GMMs) (Ng et al., 2012) and on Hidden Markov Models (HMMs) (Pfafu et al., 2001; Sarikaya and Hansen, 1998) has been traditionally used. After the deep learning approaches gained popularity in the speech signal processing field, numerous DNN-based systems, such as the ones based on MLP (Ryant et al., 2013), convolutional neural network (CNN) (Thomas et al., 2014), LSTM (Gelly and Gauvain, 2017), have been also proposed with superior performance to the traditional methods.

The performance of SAD largely affects the overall performance of the speaker diarization system as it can create a significant amount of false positive salient events or miss speech segments (Haws et al., 2016). A common practice in speaker diarization tasks is to report DER with “oracle SAD” setup which indicates that the system output is using SAD output that is identical to the ground truth. Conversely, the system output with an actual speech activity detector is referred to as “system SAD” output.

2.3. Segmentation

In the context of speaker diarization, speech segmentation is a process of breaking the input audio stream into multiple segments to obtain speaker-uniform segments. Therefore, the output unit of the speaker diarization system is determined via a segmentation process. In general, speech segmentation methods for speaker diarization are divided into two major categories: Segmentation by speaker-change point detection and uniform segmentation.

Segmentation by the speaker-change point detection was the gold standard of the earlier speaker diarization systems, where speaker-change points are detected by comparing two hypotheses: H_0 assumes that both the left and right speech windows are from the same speaker, whereas H_1 assumes that the two speech windows are from the different speakers. To test these two hypotheses, metric-based approaches (Chen et al., 1998; Kemp et al., 2000) were most widely applied. In metric-based approaches, the distribution of the speech feature is assumed to follow a Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ with mean μ and covariance Σ . The two hypotheses H_0 and H_1 can be then represented as follows:

$$\begin{aligned} H_0 : \mathbf{x}_1 \cdots \mathbf{x}_N &\sim \mathcal{N}(\mu, \Sigma), \\ H_1 : \mathbf{x}_1 \cdots \mathbf{x}_i &\sim \mathcal{N}(\mu_1, \Sigma_1), \\ &\mathbf{x}_{i+1} \cdots \mathbf{x}_N \sim \mathcal{N}(\mu_2, \Sigma_2), \end{aligned} \quad (9)$$

where $(\mathbf{x}_i | i = 1, \dots, N)$ is a sequence of speech features in the interest of the hypothesis testing. A slew of criteria for the metric-based approach were proposed to quantify the likelihood of the two hypotheses. The examples include the Kullback Leibler (KL) distance (Siegler et al., 1997), Generalized Likelihood Ratio (GLR) (Gish et al., 1991; Bonastre et al., 2000; Gangadharaiyah et al., 2004) and BIC (Chen et al., 1998; Tritschler and Gopinath, 1999). Among these criteria, BIC has been the most widely used method followed by numerous variants (Delacourt and Wellekens, 2000; Mori and Nakagawa, 2001; Ajmera et al., 2004b; Malegaonkar et al., 2006). Thus, in this section, we introduce BIC as a representative example for a metric-based method. If we apply BIC and to the hypotheses described in Eq. (9), a BIC value between two models from two hypotheses is expressed as follows:

$$BIC(i) = N \log |\Sigma| - N_1 \log |\Sigma_1| - N_2 \log |\Sigma_2| - \lambda P, \quad (10)$$

where the sample covariance Σ is from $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, Σ_1 is from $\{\mathbf{x}_1, \dots, \mathbf{x}_i\}$ and Σ_2 is from $\{\mathbf{x}_{i+1}, \dots, \mathbf{x}_N\}$ and P is the penalty term (Chen et al., 1998) defined as

$$P = \frac{1}{2} \left(d + \frac{1}{2} d(d+1) \right) \log N, \quad (11)$$

where d denotes the dimension of the feature; N_1 and N_2 are frame lengths of each window, respectively and $N = N_1 + N_2$. The penalty weight λ is generally set to $\lambda = 1$. The change point is set when the following equation becomes true:

$$\left\{ \max_i BIC(i) \right\} > 0. \quad (12)$$

In general, if speech segmentation is done using the speaker-change point detection method, the length of each segment is not consistent. Therefore, after the advent of the i-vector (Dehak et al., 2011) and DNN-based embeddings (Variansi et al., 2014b; Snyder et al., 2018), the segmentation based on speaker-change point detection was mostly replaced with uniform segmentation (Senoussaoui et al., 2013b; Sell et al., 2018a; Wang et al., 2018), since varying lengths of the segment created an additional variability into the speaker representation and deteriorated the fidelity of the speaker representations.

In uniform segmentation schemes, the given audio stream input is segmented with a fixed window length and overlap length. Thus, the unit duration of the speaker diarization output stays constant. However, the process of uniform segmentation of the input signals for diarization poses some potential problems because it introduces a trade-off error related to the segment length. The segments created from the uniform segmentation need to be sufficiently short to safely assume that they do not contain multiple speakers. However, at the same time it is important to capture sufficient acoustic information to extract reliable speaker representations.

2.4. Speaker representations and similarity measure

Speaker representation plays a crucial role for speaker diarization systems to measure the similarity between speech segments. This section will cover such speaker representation and also the similarity measure because they are tightly connected. We first introduce metric-based approaches for similarity measures which were popular from the late 1990s to early 2000s in Section 2.4.1. We then introduce widely used speaker representations for speaker diarization systems that are usually employed together with the uniform segmentation method described in Sections 2.4.2 and 2.4.3.

2.4.1. Metric based similarity measure

From the late 1990s to early 2000s, metric-based approaches were most commonly used for the similarity measurement between speech segments for speaker diarization systems. Methods used for speaker segmentation were also applied to measure the similarity between segments, such as KL distance (Siegler et al., 1997), GLR (Gish et al., 1991; Bonastre et al., 2000; Gangadharaiah et al., 2004), and BIC (Chen et al., 1998; Tritschler and Gopinath, 1999). As with the case of the segmentation, the BIC-based method, where the similarity between two segments are computed by Eq. (10), was one of the most extensively used metrics due to its effectiveness and ease of implementation. Metric-based approaches are usually employed together with the segmentation approaches based on a speaker-change point detection. The agglomerative hierarchical clustering (AHC) is often applied to obtain the diarization result, which will be detailed in Section 2.5.1.

2.4.2. Joint factor analysis, i-vector and PLDA

Before the advent of speaker representations such as i-vector (Dehak et al., 2011) or x-vector (Snyder et al., 2018), Gaussian Mixture Model-based Universal Background Model (GMM-UBM) (Reynolds et al., 2000) applied to acoustic features demonstrated success in speaker verification tasks. A UBM consists of a large GMM (typically with 512 to 2048 mixtures) trained to represent the speaker-independent distribution of acoustic features. Thus, a GMM-UBM model can be described with the following quantities: mixture weights, mean values and covariance matrix of the mixtures. The log-likelihood ratio between a speaker-adapted GMM and the speaker-independent GMM-UBM is used for speaker verification. Despite the success on modeling the speaker identity, GMM-UBM based speaker verification systems have suffered from intersession variability (Kenny et al., 2007), which is the variability exhibited by a given speaker from one recording session to another. Such difficulty occurs because the relevance maximum a posteriori (MAP) adaptation step during the speaker enrollment process in the GMM-UBM based speaker verification systems not only captures the speaker-specific characteristics of the speech, but also unwanted channel noise and other nuisances from the acoustic environment.

Joint factor analysis (JFA) (Kenny et al., 2007, 2008) was proposed to compensate for the variability issues by separately modeling the inter-speaker variability and the channel or session variability. The JFA approach employs a GMM supervector, which is a concatenated mean of the adapted GMM. For example, suppose a F by 1 speaker-independent GMM mean vector m_c , where c is the mixture component index and F is the dimension of the feature. Then, a supervector \mathbf{M} has dimension of CF by 1 by concatenating the F -dimensional mean vector for C mixture components. Thus, the supervector \mathbf{M} can be described as follows:

$$\mathbf{M} = [m_1^t, m_2^t, \dots, m_C^t]^t. \quad (13)$$

In the JFA approach, the given GMM supervector is decomposed into speaker independent, speaker dependent, channel dependent, and residual components. Thus, the ideal speaker supervector \mathbf{M}_J can be decomposed as indicated in Eq. (14), where \mathbf{m}_J denotes a speaker independent supervector from the UBM, \mathbf{V} denotes a speaker dependent component matrix, \mathbf{U} denotes a channel dependent component matrix, and \mathbf{D} denotes a speaker-dependent residual component matrix. Along with these component matrices, vector \mathbf{y} is for the speaker factors, vector \mathbf{x} is for the channel factors, and vector \mathbf{z} is for the speaker-specific residual factors. All of these vectors have a prior distribution of $N(0, 1)$.

$$\mathbf{M}_J = \mathbf{m}_J + \mathbf{V}\mathbf{y} + \mathbf{U}\mathbf{x} + \mathbf{D}\mathbf{z}. \quad (14)$$

The JFA approach was followed by the study in Dehak et al. (2011), in which it was discovered that channel factors in the JFA also contain information regarding the speakers. Thus, Dehak et al. (2011) proposed a new method combining the channel and speaker spaces into a combined variability space through a *total variability* matrix. Thus, the total variability matrix \mathbf{T} models both the channel and the speaker variability, and the latent variable \mathbf{w} weights the column of the matrix \mathbf{T} . The variable \mathbf{w} is referred to as the *i-vector* and is also considered a speaker representation vector. Each speaker and channel in a GMM supervector \mathbf{M}_J can be modeled as follows:

$$\mathbf{M}_J = \mathbf{m}_J + \mathbf{T}\mathbf{w}, \quad (15)$$

where \mathbf{m}_J is a speaker-independent and channel-independent supervector that can be taken as a UBM supervector. The process of extracting an i-vector \mathbf{w} for the given recording is formulated as a MAP estimation problem (Kenny et al., 2005; Dehak et al., 2011) using the Baum–Welch statistics extracted using the UBM, mean supervector \mathbf{m}_J , and total variability matrix \mathbf{T} trained from the EM algorithm as parameters. The idea behind a speaker representation was greatly popularized through the use of i-vectors, where the speaker representation vector can contain a numerical feature characterizing the vocal tract of each speaker. The i-vector speaker representations have been employed in not only speaker recognition studies but also in numerous speaker diarization studies (Senoussaoui et al., 2013b; Sell and Garcia-Romero, 2014; Zhu and Pelecanos, 2016) and have shown a superior performance compared to metric-based methods such as BIC, GLR, and KL, as mentioned in the previous subsection.

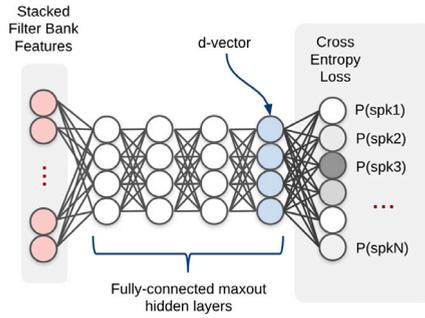


Fig. 2. Diagram of d-vector model.

Intersession variability in the i-vector approach has been further compensated using backend procedures, such as a linear discriminant analysis (LDA) (Kanagasundaram et al., 2012, 2014) and within-class covariance normalization (WCCN) (Senoussaoui et al., 2010; Kanagasundaram et al., 2011), followed by simple cosine similarity scoring. The cosine similarity scoring was later replaced with a probabilistic LDA (PLDA) model in Matějka et al. (2011). In the following studies (Garcia-Romero and Espy-Wilson, 2011; Kenny, 2010), a method applying a Gaussianization of the i-vectors and thus generating Gaussian assumptions in the PLDA, referred to as G-PLDA or simplified PLDA, was proposed for speaker verification. In general, PLDA employs the following modeling for the given speaker representation ϕ_{ij} of the i th speaker in the j th session as indicated below:

$$\phi_{ij} = \boldsymbol{\mu} + \mathbf{F}\mathbf{h}_i + \mathbf{G}\mathbf{w}_{ij} + \epsilon_{ij}. \quad (16)$$

Here, $\boldsymbol{\mu}$ is the mean vector, \mathbf{F} is the speaker variability matrix, \mathbf{G} is the channel variability matrix, and ϵ_{ij} is a residual component. In addition, \mathbf{h}_i and \mathbf{w}_{ij} are latent variables specific for the speaker and session, respectively. In G-PLDA, both latent variables, \mathbf{h}_i and \mathbf{w}_{ij} , are assumed to follow a standard Gaussian prior. During the training process of the PLDA, $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, \mathbf{F} , and \mathbf{G} are estimated using the expectation maximization (EM) algorithm. Based on the estimated statistics, two hypotheses are tested: hypothesis H_0 for a case in which two samples are from the same speaker, and hypothesis H_1 for when two samples are from different speakers. Under the hypothesis H_0 , the given speaker representations ϕ_1 and ϕ_2 are modeled as follows with a common latent variable \mathbf{h}_{12} :

$$\begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu} \end{bmatrix} + \begin{bmatrix} \mathbf{F} & \mathbf{G} & 0 \\ \mathbf{F} & 0 & \mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{h}_{12} \\ \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \end{bmatrix}. \quad (17)$$

On the other hand, under the hypothesis H_1 , ϕ_1 and ϕ_2 are modeled as follows with separate latent variable \mathbf{h}_1 and \mathbf{h}_2 :

$$\begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix} = \begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu} \end{bmatrix} + \begin{bmatrix} \mathbf{F} & \mathbf{G} & 0 & 0 \\ 0 & 0 & \mathbf{F} & \mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{h}_1 \\ \mathbf{w}_1 \\ \mathbf{h}_2 \\ \mathbf{w}_2 \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \end{bmatrix}. \quad (18)$$

In G-PLDA, it is assumed that ϕ is generated from a Gaussian distribution, which results in the following conditional density function (Jiang et al., 2014).

$$p(\phi | \mathbf{h}, \mathbf{w}) = \mathcal{N}(\phi | \boldsymbol{\mu} + \mathbf{F}\mathbf{h} + \mathbf{G}\mathbf{w}, \boldsymbol{\Sigma}). \quad (19)$$

Using Eq. (16)–(18), the log likelihood ratio can be described as follows:

$$s(\phi_1, \phi_2) = \log p(\phi_1, \phi_2 | H_0) - \log p(\phi_1, \phi_2 | H_1). \quad (20)$$

The log likelihood ratio $s(\phi_1, \phi_2)$ in the above equation was originally used for speaker verification to choose a hypothesis between H_0 and H_1 by checking whether $s(\phi_1, \phi_2)$ is positive or negative. The PLDA for speaker representations also employed in speaker diarization and the log likelihood $s(\phi_1, \phi_2)$ is used to check the similarity between clusters. Further details regarding the clustering approach using PLDA is described in Section 2.5.1.

2.4.3. Neural network based speaker representations

Speaker representations for speaker diarization have also been heavily affected by the rise of deep learning approaches. The idea behind DNN-based representation learning was first introduced for face recognition tasks (Sun et al., 2014; Taigman et al., 2014). As the fundamental idea of a neural network-based representation, we can use the deep neural network architecture to map the input signal source (an image or an audio clip) to a dense vector containing floating-point numbers. This is achieved by taking the values from a layer in the neural network model after forward-propagating the input signal to the layer that we take the values from. The mapping process from the input signal to the speaker embedding is based on the nonlinear modeling capability of multiple layers in the DNNs. In so doing, the training process of the DNNs allows the neural networks to learn the mapping without specifying

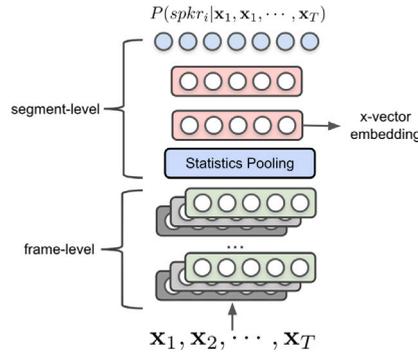


Fig. 3. Diagram of x-vector model.

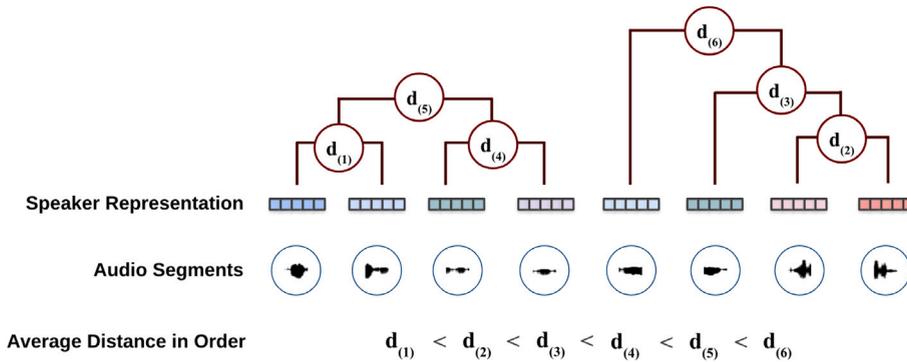


Fig. 4. Agglomerative Hierarchical Clustering.

any components or factors, which is in contrast to traditional factor analysis models based on decomposable components. In this sense, the components in JFA are more explainable than the parameters trained in DNN models trained for speaker embedding extraction. In addition, DNN-based speaker representation learning does not involve predefined probabilistic models (e.g., GMM-UBM) for the input acoustic features. In relation to this, DNN-based speaker representation achieves an improved efficiency during the inference phase because the solution used by factor-analysis based methods involves a computationally intensive matrix inversion operation (Jiang et al., 2014), whereas DNN-based embedding extractors involve fewer demanding operations, such as multiple linear transformations, with non-linear function computations for obtaining the speaker representation vector. Thus, the representation learning process has become more straight-forward and the inference speed has been improved compared to the traditional factor-analysis based methods. Among many of the neural-network based speaker representations, d-vector (Variansi et al., 2014b) remains one of the most prominent speaker representation extraction frameworks. The stacked filter bank features, which include context frames as an input feature, are employed, and multiple fully connected layers are trained with cross entropy loss. Speaker representation vectors, also referred to as d-vectors, are obtained from the last fully connected layer, as indicated in Fig. 2. The d-vector scheme appears in numerous speaker diarization papers, e.g., in Wang et al. (2018), Zhang et al. (2019).

DNN-based speaker representations are even more improved when using an x-vector (Snyder et al., 2017, 2018), which demonstrates a superior performance, winning the NIST speaker recognition challenge 2018 (Villalba et al., 2019) and the first DIHARD challenge (Sell et al., 2018b). Fig. 3 shows the structure of an x-vector extractor. The time-delay architecture and statistics pooling layer differentiate the x-vector architecture from that of a d-vector. The statistics pooling layer aggregates the frame-level outputs from the previous layer and computes its mean and standard deviation, passing them on to the following layer. Thus, it can allow the extraction of x-vectors from a variable length input. This is advantageous not only for speaker verification but also for speaker diarization because speaker diarization systems are bound to process segments that are shorter than the predetermined uniform segment length when the segment should be truncated at the end of an utterance.

2.5. Clustering

A clustering algorithm is applied to make clusters of the speech segments based on the speaker representation and similarity measure explained in the previous section. Here, we introduce the most commonly used clustering methods for speaker diarization.

2.5.1. Agglomerative hierarchical clustering

AHC is a clustering method that has been constantly employed in many speaker diarization systems with different distance metrics such as BIC (Chen et al., 1998; Han and Narayanan, 2007), KL (Rougui et al., 2006) and PLDA (Sell et al., 2018b; Arora et al., 2020; Novoselov et al., 2019). AHC is an iterative process of merging the existing clusters until the clustering process meets a criterion. The AHC process starts with the calculation of the similarity between N singleton clusters. At each step, a pair of clusters that has the highest similarity is merged. The iterative merging process of AHC is illustrated in a dendrogram, which is presented in Fig. 4.

One of the most important aspects of AHC is the stopping criterion. For the speaker diarization task, the AHC process can be stopped using either a similarity threshold or a target number of clusters. Ideally, if PLDA is used as a distance metric, the AHC process should be stopped at $s(\phi_1, \phi_2) = 0$ in Eq. (20). However, the stopping metric is adjusted to obtain an accurate number of clusters based on a development set. Conversely, if the number of speakers is known or estimated using other methods, the AHC process can be stopped when the clusters created by the AHC process reaches the pre-determined number of speakers k .

2.5.2. Spectral clustering

Spectral clustering is a widely used clustering approach for speaker diarization. While there are many variations, spectral clustering involves the following steps (see Fig. 5).

- (i) Affinity matrix calculation: There are many ways to generate affinity matrix \mathbf{A} depending on the way the affinity value is processed. The raw affinity value d is processed by kernels such as $\exp(-d^2/\sigma^2)$, where σ is a scaling parameter. On the other hand, the raw affinity value d could also be masked by zeroing the values below a threshold to only keep the prominent values.
- (ii) Laplacian matrix calculation (Von Luxburg, 2007): The graph Laplacian can be calculated in two ways: normalized and unnormalized. The degree matrix \mathbf{D} contains diagonal elements $d_i = \sum_{j=1}^n a_{ij}$ where a_{ij} is the element of the i th row and j th column in an affinity matrix \mathbf{A} .

(a) Normalized Graph Laplacian:

$$\mathbf{L} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}. \quad (21)$$

(b) Unnormalized Graph Laplacian:

$$\mathbf{L} = \mathbf{D} - \mathbf{A}. \quad (22)$$

- (iii) Eigen decomposition: The graph Laplacian matrix \mathbf{L} is decomposed into the eigenvector matrix \mathbf{X} and the diagonal matrix that contains eigenvalues. Thus, $\mathbf{L} = \mathbf{X} \mathbf{\Lambda} \mathbf{X}^T$.
- (iv) Re-normalization (optional): the rows of \mathbf{X} is normalized so that $y_{ij} = x_{ij} / \left(\sum_j x_{ij}^2 \right)^{1/2}$ where x_{ij} and y_{ij} are the elements of the i th row and j th column in matrices \mathbf{X} and \mathbf{Y} , respectively.
- (v) Speaker counting: Speaker number is estimated by finding the maximum eigengap (Von Luxburg, 2007; Park et al., 2019b).
- (vi) Spectral embedding clustering: The k -smallest eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ and the corresponding k eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ are stacked to construct a matrix $\mathbf{U} \in \mathbb{R}^{n \times k}$. The row vectors of \mathbf{U} are referred to as k -dimensional spectral embeddings. Finally, the spectral embeddings are clustered using a clustering algorithm. In general k -means clustering (MacQueen et al., 1967) is employed for clustering the spectral embeddings.

Among many variations of spectral clustering algorithm, the Ng–Jordan–Weiss (NJW) algorithm (Ng et al., 2001) is often employed for the speaker diarization task with variation in the kernel for the calculation of the affinity values (Ning et al., 2006; Luque and Hernando, 2012; Shum et al., 2012). Unlike the AHC approach, spectral clustering is mostly used with cosine distance (Ning et al., 2006; Luque and Hernando, 2012; Shum et al., 2012; Wang et al., 2018; Park et al., 2019b). In addition, the LSTM based similarity measurement with spectral clustering (Lin et al., 2019) also exhibited competitive performance. Depending on the datasets, the spectral clustering approach with cosine distance measurement outperforms AHC with PLDA (Park et al., 2019b; Medennikov et al., 2020b) while using the same speaker representation for both clustering methods.

2.5.3. Other clustering algorithms

The k -means algorithm is often employed in studies on speaker diarization (Zajíc et al., 2016; Dimitriadis and Fousek, 2017a; Wang et al., 2018; Zhang et al., 2019; Huang et al., 2020) due to its simplicity and ease of implementation. However, the k -means algorithm generally underperforms (Wang et al., 2018; Zhang et al., 2019) the well-known clustering algorithms such as spectral clustering and AHC. In addition, there are a few speaker diarization studies employed the mean-shift (Comaniciu and Meer, 2002) clustering algorithm, which assigns the given data points to the clusters iteratively by finding the modes in a non-parametric distribution. The Mean-shift clustering algorithm was employed in the speaker diarization task with KL distance in Stafylakis et al. (2010), i-vector and cosine distance in Senoussaoui et al. (2013b,a), and i-vector and PLDA in Salmun et al. (2017).

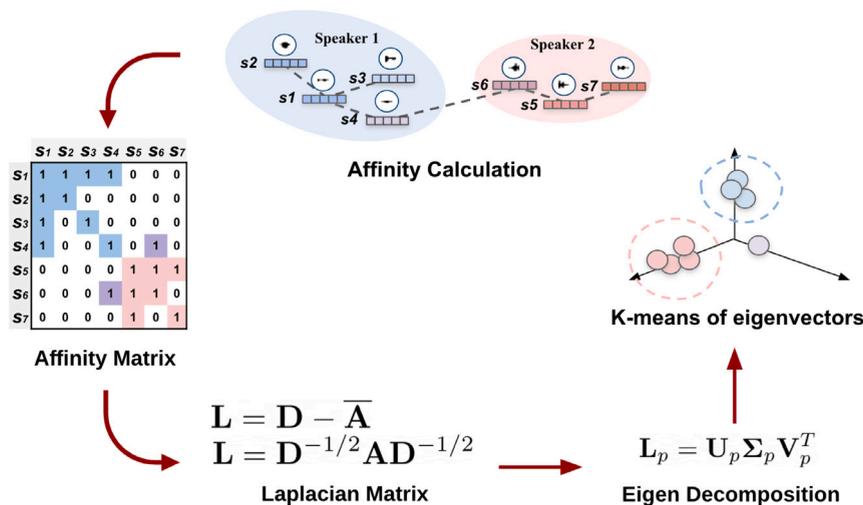


Fig. 5. General steps of spectral clustering.

2.6. Post-processing

2.6.1. Resegmentation

Resegmentation is a process of refining the speaker boundary that is roughly estimated using the clustering procedure. In [Kenny et al. \(2010\)](#), the Viterbi resegmentation method based on the Baum–Welch algorithm was introduced. In this method, the estimation of Gaussian mixture model corresponding to each speaker and Viterbi-algorithm-based ressegmentation using the estimated speaker GMM are alternately applied.

A method for representing the diarization process based on the variational Bayesian hidden Markov model (VB-HMM) was proposed, and was shown to be superior to Viterbi ressegmentation ([Diez et al., 2018a, 2019a](#)). The VB-HMM-based diarization can be seen as a joint optimization of segmentation and clustering, which will be separately introduced in Section 2.7.

2.6.2. System fusion

As another direction of post processing, there have been a series of studies on the fusion method of multiple diarization results to improve the diarization accuracy. While it is widely known that the system combination generally yields better result for various systems (e.g., speech recognition [Fiscus, 1997](#) or speaker recognition [Brummer et al., 2007](#)), the combination of multiple diarization hypotheses poses several unique problems. First, speaker labeling is not standardized among different diarization systems. Second, the estimated number of speakers may differ among different diarization systems. Finally, the estimated time boundaries may also be different among multiple diarization systems. System combination methods for speaker diarization systems need to handle these problems during the fusion process of multiple hypotheses.

In [Huijbregts et al. \(2009\)](#), a method for selecting the best diarization result among many diarization systems was proposed. In this method, a whole sequence of diarization result for a recording from each diarization system is treated as one object to be clustered. AHC is applied to the set of diarization results, in which the distance of two clusters is measured using the symmetric DER between the diarization results belonging to the two clusters. The iterative merging process of AHC is executed until the number of clusters becomes two. Finally, in the bigger cluster among the two final clusters according to the number of elements in each cluster, the diarization result that has the smallest distance to all other diarization results is selected as the final result.

In [Bozonnet et al. \(2010\)](#), two diarization systems are combined by finding the matching between two speaker clusters, and then performing ressegmentation based on the matching result.

More recently, the diarization output voting error reduction (DOVER) method ([Stolcke and Yoshioka, 2019](#)) was proposed to combine multiple diarization results based on the voting scheme. In the DOVER method, the speaker labels among different diarization systems are aligned one by one to minimize the DER between the hypotheses (processes 2 and 3 of [Fig. 6](#)). After aligning all hypotheses, each system votes its speaker label to each segmented region (each system may have different weights for voting), and the speaker label that gains the highest voting weight is selected for each segmented region (the process 4 of [Fig. 6](#)). In case multiple speaker labels get the same voting weight, a heuristic approach is employed to break the ties (such as selecting the result from the first system) is used.

The DOVER method has an implicit assumption that there is no overlapping speech, i.e., at most only one speaker is assigned for each time index. To combine the diarization hypotheses with overlapping speakers, two methods were recently proposed. In [Xiao et al. \(2020\)](#), the authors proposed the modified DOVER method, in which the speaker labels in different diarization results are first aligned with a root hypothesis, and the speech activity of each speaker is estimated based on the weighted voting score for each speaker for each small segment. The speakers of multiple hypotheses are aligned via a weighted k-partite graph matching, and the

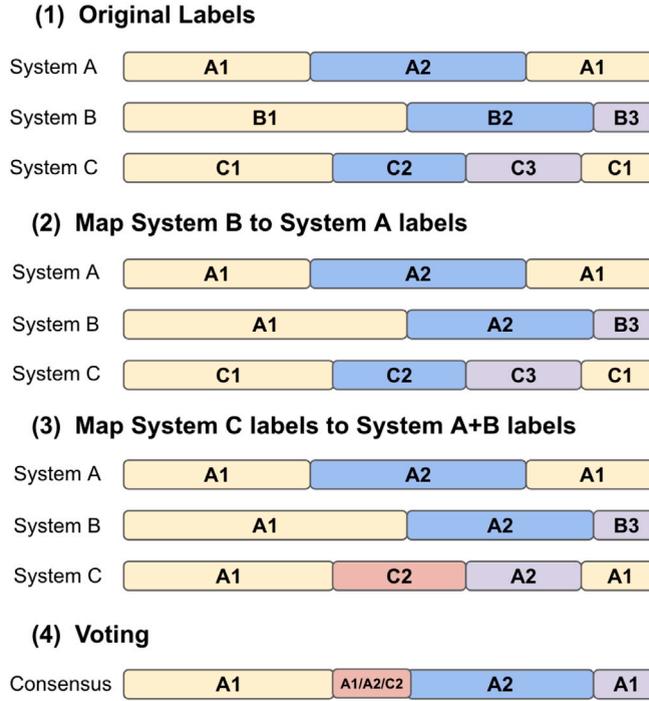


Fig. 6. Example of DOVER system.

number of speakers K for each small segment is estimated based on the weighted average of multiple systems to select the top- K voted speaker labels. Both the modified DOVER and DOVER-Lap showed DER improvement for the speaker diarization result with speaker overlaps.

2.7. Joint optimization of segmentation and clustering

This subsection introduces a VB-HMM-based diarization technique, which can be regarded as a joint optimization of segmentation and clustering, and thus cannot be well categorized in Sections 2.1–2.6. The VB-HMM framework was proposed as an extension of the VB-based speaker clustering (Valente, 2005; Kenny, 2008) by introducing HMM to constrain the speaker transitions. In the VB-HMM framework (Diez et al., 2018a), the speech feature $\mathbf{X} = (x_t | t = 1, \dots, T)$ is assumed to be generated from HMM where each HMM state corresponds to one of K possible speakers. Suppose that we have M HMM states, M -dimensional variable $\mathbf{Z} = (z_t | t = 1, \dots, T)$ is introduced where k th element of z_t is 1 if the k th speaker is speaking at the time index t , and 0 otherwise. At the same time, the distribution of x_t is modeled based on a hidden variable $\mathbf{Y} = \{y_k | k = 1, \dots, K\}$, where y_k denotes a low dimensional vector for the k th speaker. Given these notations, the joint probability of \mathbf{X} , \mathbf{Y} , and \mathbf{Z} is decomposed as follows:

$$P(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) = P(\mathbf{X}|\mathbf{Z}, \mathbf{Y})P(\mathbf{Z})P(\mathbf{Y}), \quad (23)$$

where $P(\mathbf{X}|\mathbf{Z}, \mathbf{Y})$ is the emission probability modeled by GMM whose mean vector is represented by \mathbf{Y} , $P(\mathbf{Z})$ is the transition probability of the HMM, and $P(\mathbf{Y})$ is the prior distribution of \mathbf{Y} . Because \mathbf{Z} represents the trajectory of speakers, the diarization problem can be expressed as the inference problem of \mathbf{Z} that maximizes the posterior distribution $P(\mathbf{Z}|\mathbf{X}) = \int P(\mathbf{Z}, \mathbf{Y}|\mathbf{X})d\mathbf{Y}$. Since it is intractable to directly solve this problem, the VB method is used to estimate the model parameters that approximate $P(\mathbf{Z}, \mathbf{Y}|\mathbf{X})$.

Recently, a simplified version of VB-HMM that works on the x-vector, known as VBx, was proposed (Diez et al., 2019b; Landini et al., 2020). In VBx, $P(\mathbf{X}|\mathbf{Z}, \mathbf{Y})$ is calculated using the x-vector based on the PLDA model. While the original VB-HMM works on the granularity of the frame-level feature, VBx works on the granularity of the x-vector, and thus can be seen as a clustering method that jointly models speaker turn and speaker duration.

The VB-HMM diarization was originally designed as a standalone diarization framework. However, it requires parameter initialization to start the VB estimation, and the parameters are usually initialized based on the result of another speaker clustering. In that context, VB-HMM is widely employed as the final step of speaker diarization (e.g., Sell and Garcia-Romero, 2015; Sell et al., 2018a; Diez et al., 2020). For example in Diez et al. (2020), AHC was first performed to under-cluster the x-vector, and VBx was then applied to obtain a better cluster given the AHC-based result as the initial parameter. Finally, VB-HMM was further applied to refine the boundary obtained by VBx.

Table 2

Overview of speaker diarization techniques using deep learning. Note that there are also a lot of studies that use deep learning for front-end, SAD, segmentation, and speaker embedding extraction, which are introduced in Section 2.

Section	Front-end	SAD	Seg.	Speaker embed.	Clust.	Post-proc.	Joint w/ ASR	Speaker profile	Overlap handling	Inference mode	Max # of speakers	References
3.1.1	-	-	-	-	√	-	-	-	no	offline	Variable	Affinity matrix (Wang et al., 2020; Lin et al., 2020; Park et al., 2021), IDEC (Dimitriadis, 2019)
3.1.2	-	-	-	-	√	-	-	required	no	online	Fixed	RRNN (Flemotomos and Dimitriadis, 2020)
3.1.3	- ^c	-	-	-	-	√	-	required	yes	offline	Fixed	TS-VAD ^a (Medennikov et al., 2020b,a)
3.1.3	- ^c	-	-	-	-	√	-	-	yes	offline	Variable	EEND-based post-processing ^b (Horiguchi et al., 2020b)
3.2.1	-	-	√	-	√	-	-	-	no	both	Variable	UIS-RNN (Zhang et al., 2019)
3.2.2	- ^c	-	√	√	√	√	-	-	yes	offline	Variable	RPN (Huang et al., 2020)
3.2.3	√	-	√	√	√	-	-	-	yes	online	Variable	Online-RSAN (von Neumann et al., 2019; Kinoshita et al., 2020)
3.2.4	- ^c	√	√	√	√	-	-	-	yes	offline	Fixed	EEND (Fujita et al., 2019,b)
3.2.4	- ^c	√	√	√	√	-	-	-	yes	offline	Variable	EEND-EDA (Horiguchi et al., 2020), SC-EEND (Fujita et al., 2020)
3.2.4	- ^c	√	√	√	√	-	-	-	yes	online	Fixed	EEND with speaker tracing buffer (Xue et al., 2020)
3.2.4	- ^c	√	√	√	√	-	-	-	yes	online	Variable	BW-EDA-EEND (Han et al., 2021)
3.2.4	- ^c	√	√	√	-	-	-	-	yes	offline	Variable	EEND-vector clustering (Kinoshita et al., 2021)
4.3	-	-	√	-	√	-	√	-	no	offline	Fixed	Speaker-tag insertion (Shafey et al., 2019; Mao et al., 2020)
4.3	- ^c	-	√	-	√	-	√	-	yes	offline	Variable	MAP-decoding (Kanda et al., 2019c)
4.3	- ^c	-	√	√	√	-	√	required	yes	offline	Variable	End-to-end SA-ASR (Kanda et al., 2020a, 2021b)
4.3	- ^c	-	√	√	-	-	√	-	yes	offline	Variable	End-to-end SA-ASR+Clustering (Kanda et al., 2021a)

^aTS-VAD can also be interpreted as a joint model of SAD, segmentation, and speaker identification.

^bEEND (Fujita et al., 2019,b) itself is a joint model of SAD, segmentation, speaker embedding, and clustering.

^cIt can also be seen that a speech separation module is implicitly embedded in the model to cope with speaker overlaps.

3. Recent advances in speaker diarization using deep learning

This section introduces various recent efforts toward deep learning-based speaker diarization techniques. The methods that incorporate deep learning into a single component of speaker diarization, such as clustering or post-processing, are introduced in Section 3.1. The methods that unify several components of speaker diarization into a single neural network are introduced in Section 3.2. For the overview of speaker diarization techniques using deep learning, refer to Table 2. It should be noted that there are some works that take additional input of speaker profiles. These methods may not be categorized as a diarization technique in a traditional definition. Nevertheless, we introduce them as they are optimized in a multispeaker situation to learn the relations between speakers and hence categorized as “Trained Based on the Diarization Objective” in Table 1.

3.1. Single-module optimization

3.1.1. Speaker clustering enhanced by deep learning

Enhancing the clustering procedure based on the deep learning is an active research area and several methods have been proposed for speaker diarization. This section will cover the representative works in such a direction.

An approach based on the graph neural network (GNN) was proposed in Wang et al. (2020). As shown in Fig. 7, this method aims at purifying the similarity matrix used in the spectral clustering (Section 2.5.2). Assuming a sequence of speaker embeddings $\{\mathbf{e}_1, \dots, \mathbf{e}_N\}$, where N is the sequence length, the input \mathbf{x}_i^0 to the GNN is $\{\mathbf{x}_i^0 = \mathbf{e}_i | i = 1, \dots, N\}$. The output $\mathbf{x}_i^{(p)}$ of the p th layer of the GNN is now:

$$\mathbf{x}_i^{(p)} = \sigma(\mathbf{W} \sum_j \mathbf{L}_{i,j} \mathbf{x}_j^{(p-1)}), \quad (24)$$

where \mathbf{L} is a normalized affinity matrix added by self-connection, \mathbf{W} is a trainable weight matrix for the p th layer, and $\sigma(\cdot)$ is a nonlinear function. The GNN is trained by minimizing the distance between the reference and the estimated affinity matrices. The distance is calculated using a combination of histogram loss (Ustinova and Lempitsky, 2016) and nuclear norm (Recht et al., 2010). The GNN-based speaker diarization method was evaluated on the CALLHOME dataset and an in-house meeting dataset, and significantly outperformed any of the conventional clustering methods.

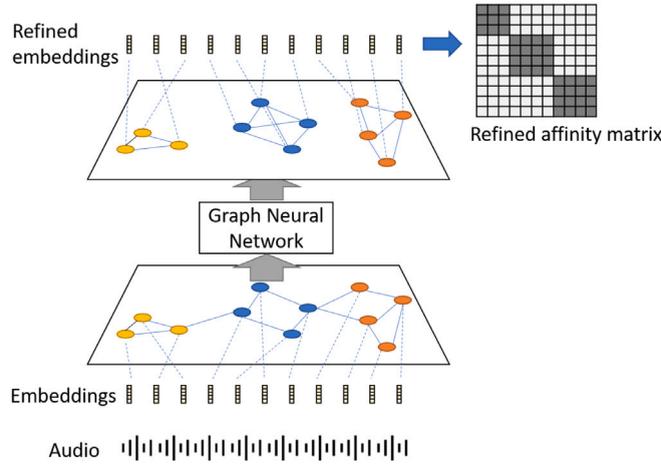


Fig. 7. Speaker diarization with graph neural network.

Besides that, different approaches have been proposed to generate the affinity matrix. In [Lin et al. \(2020\)](#), a self-attention-based neural network model was introduced to directly generate a similarity matrix from a sequence of speaker embeddings. In [Park et al. \(2021\)](#), several affinity matrices with different temporal resolutions were fused into a single affinity matrix based on a neural network.

A different approach aiming at improving clustering was proposed in [Xie et al. \(2016\)](#), called deep embedded clustering (DEC). The goal of DEC was to transform the input features (herein referred to as speaker embeddings) making them more separable in a given number of clusters/speakers. In order to make cluster differentiable, each embedding is provided with a probability of “belonging” to each of the available speaker clusters, i.e. q_{ij} can be interpreted as the probability of assigning sample i to cluster j (i.e., a soft assignment):

$$q_{ij} = \frac{(1 + \|z_i - \mu_j\|^2/a)^{-\frac{a+1}{a}}}{\sum_l (1 + \|z_i - \mu_l\|^2/a)^{-\frac{a+1}{a}}}, \quad p_{ij} = \frac{q_{ij}^2/f_i}{\sum_l q_{il}^2/f_l}, \quad (25)$$

where z_i is the bottleneck feature, a is the degree of freedom of the Student’s t -distribution, μ_i is the centroid of i th cluster and f_i is the soft cluster frequency with $f_i = \sum q_{ij}$. The clusters are iteratively refined based on the target distribution according to the bottleneck features estimated using an autoencoder.

The initial version of DEC had some problems, and refined algorithm called improved DEC (IDEC) was later proposed with better accuracy on speaker diarization ([Guo et al., 2017](#); [Dimitriadis, 2019](#)). Firstly, there was a potential risk that the neural network is converged to a trivial solution to generates corrupted embeddings. To avoid this risk, [Guo et al. \(2017\)](#) proposed to explicitly preserve the local structure of the data by adding a reconstruction loss between the output of the autoencoder and the input feature. [Dimitriadis \(2019\)](#) further addressed the issue by introducing the loss function to enforce the distribution of speaker turns being uniform across all speakers, i.e., all speakers contribute equally to the session. This assumption is not always valid for real recordings but it constrains the solution space enough to avoid the empty clusters without affecting the overall performance. Finally, [Dimitriadis \(2019\)](#) also proposed an additional loss term that penalizes the distance from the centroid μ_i , bringing the behavior of the algorithm closer to k -means.

Overall, the loss function of the IDEC consists of four loss terms, i.e., L_c , the clustering error term that is originally proposed in DEC; L_r , the reconstruction error term ([Guo et al., 2017](#)); L_u is the uniform “speaker airtime” distribution loss ([Dimitriadis, 2019](#)); and L_{MSE} , the loss to measure the distance of the bottleneck features from the centroids ([Dimitriadis, 2019](#)),

$$L = \alpha L_c + \beta L_r + \gamma L_u + \delta L_{MSE}, \quad (26)$$

where α, β, γ , and δ are the weight on the loss functions that is fine-tuned on some held-out data.

3.1.2. Learning the distance estimator

In this section a novel approach using a trainable distance function is presented. The basic idea is based on the relational recurrent neural networks (RRNNs). RRNNs were introduced by [Santoro et al. \(2018, 2016\)](#), [Sukhbaatar et al. \(2015\)](#) to address “relational information learning” problems. Such models learn relations between a sequence of input features like the notion of “closer” or “further”, e.g., two points in space are closer than a third one, etc. Speaker diarization can be seen as part of this class of problems, since the final decision depends on the distance between speech segments and speaker profiles or centroids.

There are several issues that potentially limits the accuracy of speaker diarization systems. Firstly, as mentioned in Section 2.3, the duration of segments when extracting speaker embeddings poses a trade-off between the time resolution and the robustness

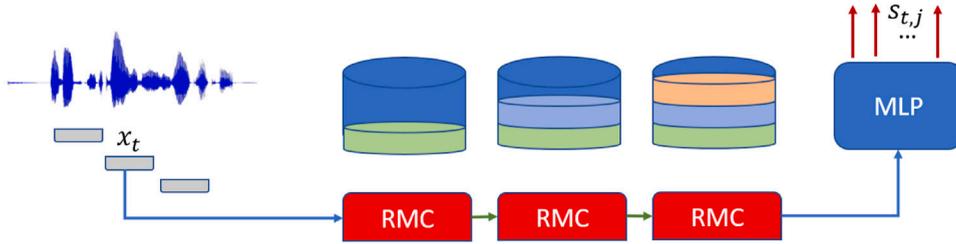


Fig. 8. Continuous speaker identification system based on RMC. The speech signal is segmented uniformly and each segment x_t is compared against all the available speaker profiles according to a distance metric $d(\cdot, \cdot)$. A speaker label $s_{t,j}$ is assigned to each x_t minimizing this metric.

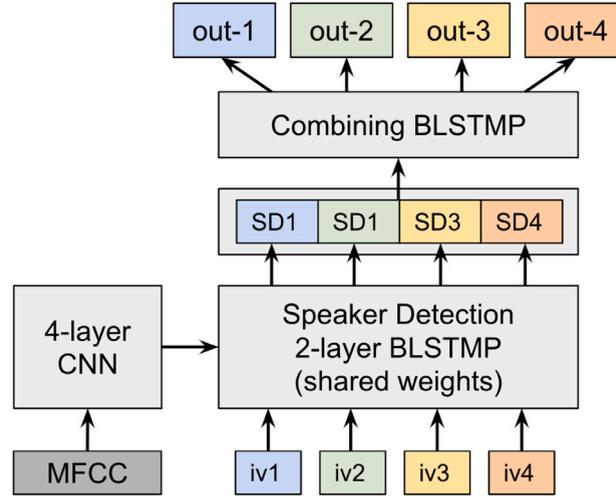


Fig. 9. Target Speaker Voice Activity Detector.

of the extracted speaker representations. Secondly, speaker embedding extractors are not explicitly trained to provide optimal representations for speaker diarization, despite the fact these invariant, discriminative representations are used to separate thousands of speakers (Snyder et al., 2018). Thirdly, the distance metric is often based on a heuristic approach and/or dependent on certain assumptions that do not necessarily hold, e.g., assuming Gaussianity in the case of PLDA (Garcia-Romero and Espy-Wilson, 2011). Finally, the audio chunks are treated independently and any temporal context is simply ignored in conventional clustering methods as described in Section 2.5. These issues can be attributed to the distance metric function, and most of them can be addressed with RRNNs, where a data-driven, memory-based approach bridges the performance gap between the heuristic and the trainable distance estimation approaches.

In this context, an approach for learning the relationship between the speaker cluster centroids (or speaker profiles) and the embeddings is proposed in Flemotomos and Dimitriadis (2020) (Fig. 8). In this work, the diarization process is considered to be a classification task on an already segmented audio, as in Section 2.3, either uniformly (Zajic et al., 2016) or based on estimated speaker-change salient points (Yoshioka et al., 2019b). The speaker embeddings x_j for each segment, which are assumed to be speaker-homogeneous, are extracted and then compared with all the available speaker profiles or speaker centroids. The most suitable speaker label is assigned to each segment by minimizing the distance-based loss function, i.e., the relationship between embeddings and profiles. As discussed in Flemotomos and Dimitriadis (2020), the RRNN-based distance estimation exhibits consistent improvements in its performance when compared with the more traditional distance estimation approaches such as the cosine distance (Dehak et al., 2011) or the PLDA-based (Garcia-Romero and Espy-Wilson, 2011) distance. Note that although the task in Flemotomos and Dimitriadis (2020) is speaker identification, an extension to the speaker diarization is rather straightforward when the speaker profiles are pre-estimated, either as centroids using any of the traditional clustering algorithm (Sections 2.5 and 3.1.1), or using prior knowledge.

3.1.3. Post processing based on deep learning

There are a few recent studies on the neural network-based diarization method that is applied on top of the result from a traditional clustering-based speaker diarization. These methods can be categorized as an extension of the post-processing. Medennikov et al. (2020b,a) proposed the target-speaker voice activity detection (TS-VAD) to achieve accurate speaker diarization even under noisy conditions with many speaker overlaps. TS-VAD assumes that a set of i-vectors $\mathcal{E} = \{e_k \in \mathbb{R}^f | k = 1, \dots, K\}$ are

available for each speaker in the audio, where f is the dimension of i-vector and K is the number of speakers. As presented in Fig. 9, TS-VAD takes not only a sequence of MFCC, $\mathbf{X} = (\mathbf{x}_t \in \mathbb{R}^d | t = 1, \dots, T)$, where d is the dimension of MFCC and T the length of the sequence, but also a set of i-vectors \mathcal{E} . Given \mathbf{X} and \mathcal{E} , the model outputs a sequence of k -dimensional vector $\mathbf{O} = (\mathbf{o}_t \in \mathbb{R}^K | t = 1, \dots, T)$ where the k th element of \mathbf{o}_t represents the probability of the speech activity of the speaker corresponding to \mathbf{e}_k at the time frame t . In other words, the k th element of \mathbf{o}_t is expected to be 1 if the speaker of \mathbf{e}_k is speaking at time t , and 0 otherwise.

Because TS-VAD requires the i-vectors of speakers, pre-processing to obtain the i-vectors is necessary. The procedure proposed in Medennikov et al. (2020b,a) is as follows:

1. Apply clustering-based diarization.
2. Estimate i-vectors for each speaker given the diarization result.
3. Repeat (a) and (b).
 - (a) Apply TS-VAD given the estimated i-vectors.
 - (b) Refine i-vectors given the TS-VAD result.

TS-VAD was proposed as a part of the winning system of CHiME-6 Challenge (Watanabe et al., 2020), and showed a significantly better DER compared with the conventional clustering based approach (Medennikov et al., 2020b). However, it has a drawback, i.e., the maximum number of speakers that the model can handle is limited by the dimension of the output vector.

As a different approach, Horiguchi et al. proposed the application of the EEND model (detailed in Section 3.2.4) to refine the result of a clustering-based speaker diarization (Horiguchi et al., 2020b). A clustering-based speaker diarization method can handle a large number of speakers but unable to handle overlapped speech. Conversely, EEND has opposite characteristics. To complementarily use the two methods, The authors in Medennikov et al. (2020b) first applied a conventional clustering method. Then, the two-speaker EEND model was iteratively applied for each pair of detected speakers to refine the time boundary of overlapped regions.

3.2. Joint optimization for speaker diarization

3.2.1. Joint segmentation and clustering

A model called unbounded interleaved-state recurrent neural networks (UIS-RNN) was proposed, which replaced the segmentation and clustering methods with a trainable model (Zhang et al., 2019). Given the input sequence of embeddings $\mathbf{X} = (\mathbf{x}_t \in \mathbb{R}^d | t = 1, \dots, T)$, UIS-RNN generates the diarization result $\mathbf{Y} = (y_t \in \mathbb{N} | t = 1, \dots, T)$ as a sequence of speaker index for each time frame. The joint probability of \mathbf{X} and \mathbf{Y} can be decomposed by the chain rule as follows.

$$P(\mathbf{X}, \mathbf{Y}) = P(\mathbf{x}_1, y_1) \prod_{t=2}^T P(\mathbf{x}_t, y_t | \mathbf{x}_{1:t-1}, y_{1:t-1}). \quad (27)$$

To model the distribution of the speaker change, UIS-RNN then introduces a latent variable $\mathbf{Z} = (z_t \in \{0, 1\} | t = 2, \dots, T)$, where z_t becomes 1 if the speaker indices at time $t-1$ and t are different, and 0 otherwise. The joint probability including \mathbf{Z} is then decomposed as follows.

$$P(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = P(\mathbf{x}_1, y_1) \prod_{t=2}^T P(\mathbf{x}_t, y_t, z_t | \mathbf{x}_{1:t-1}, y_{1:t-1}, z_{1:t-1}). \quad (28)$$

Finally, the term $P(\mathbf{x}_t, y_t, z_t | \mathbf{x}_{1:t-1}, y_{1:t-1}, z_{1:t-1})$ is further decomposed into three components.

$$P(\mathbf{x}_t, y_t, z_t | \mathbf{x}_{1:t-1}, y_{1:t-1}, z_{1:t-1}) = P(\mathbf{x}_t | \mathbf{x}_{1:t-1}, y_{1:t}) P(y_t | z_t, y_{1:t-1}) P(z_t | z_{1:t-1}). \quad (29)$$

Here, $P(\mathbf{x}_t | \mathbf{x}_{1:t-1}, y_{1:t})$ denotes the sequence generation probability, and modeled by gated recurrent unit (GRU)-based recurrent neural networks. $P(y_t | z_t, y_{1:t-1})$ denotes the speaker assignment probability and modeled by a distance-dependent Chinese restaurant process (Blei and Frazier, 2011), which can model the distribution of unbounded number of speakers. Finally, $P(z_t | z_{1:t-1})$ represents the speaker change probability and is modeled by the Bernoulli distribution. Since all models are represented by trainable ones, UIS-RNN can be trained in a supervised fashion by finding the parameters that maximize $\log P(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ over training data. The inference can be conducted by finding \mathbf{Y} that maximizes $\log P(\mathbf{X}, \mathbf{Y})$ given \mathbf{X} based on the beam search in an online fashion. While UIS-RNN works in an online fashion, it demonstrated better DER than that of the offline system based on spectral clustering.

3.2.2. Joint segmentation, embedding extraction, and re-segmentation

A speaker diarization method based on the region proposal networks (RPN) was proposed to jointly perform segmentation, speaker embedding extraction, and resegmentation (Huang et al., 2020). The RPN was originally proposed to detect multiple objects from a two-dimensional image (Ren et al., 2016), and one-dimensional variant along with the time-axis is used for speaker diarization.

As can be seen from Fig. 10(a), the STFT features with a size of time and frequency bin is first converted to the feature map with a size of time, frequency and channels using CNNs. Then, other three types of neural networks are applied on various sizes of sliding windows (named ‘‘anchor’’) along with the time axis. For each anchor, the three neural networks perform SAD, speaker embedding extraction, and region refinement, respectively. Here, SAD is the task to estimate the probability of speech activity for the anchor region. Speaker embedding extraction is the task to generate an embedding to represent the speaker characteristics of

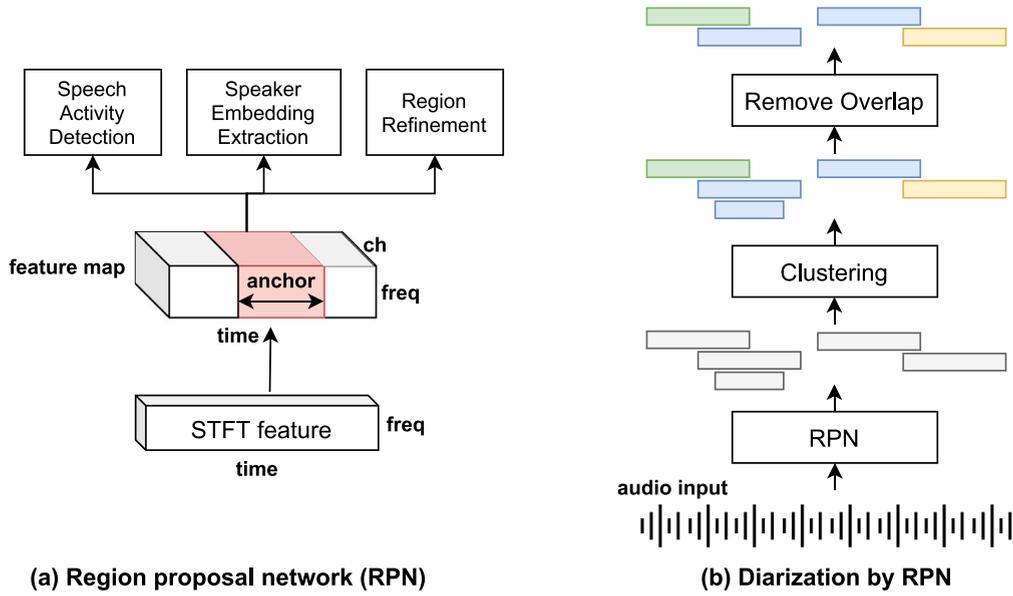


Fig. 10. (a) RPN for speaker diarization, (b) diarization procedure based on RPN.

the audio corresponding to the anchor region. Finally, region refinement is the task to estimate the difference between the shape (i.e. duration and center position) of the anchor and that of the corresponding reference region.

The inference procedure by RPN is presented in Fig. 10(b). RPN is first applied to list the anchors with speech activity probability higher than the pre-determined threshold. The anchors are then clustered using a conventional clustering method (e.g., k-means) based on the speaker embeddings estimated for each anchor. Finally, highly overlapped anchors after region refinement are removed, a method known as the non-maximum suppression.

The RPN-based speaker diarization system has the advantage of handling overlapped speech with possibly any number of speakers. Also, it is much simpler than the conventional speaker diarization system. It was shown in multiple datasets that this system achieved significantly better DER than the conventional clustering-based speaker diarization system (Huang et al., 2020; Raj et al., 2021).

3.2.3. Joint speech separation and diarization

There are also recent researches on the joint modeling of speech separation and speaker diarization. Kounades-Bastian et al. (2017a,b) proposed the incorporation of a speech activity model into speech separation based on the spatial covariance model with non-negative matrix factorization. They derived the EM algorithm to estimate separated speech and the speech activity of each speaker from the multichannel overlapped speech. While their method jointly performs speaker diarization and speech separation, it is based on a statistical modeling, and the estimation was conducted solely based on the observation, i.e. without any model training.

von Neumann et al. (2019), Kinoshita et al. (2020) later proposed a trainable model, namely online Recurrent Selective Attention Network (online RSAN), for joint speech separation, speaker counting, and speaker diarization based on a single neural network (Fig. 11). Their neural network takes the input of spectrogram $\mathbf{X}_b \in \mathbb{R}^{T \times F}$, a residual mask $\mathbf{R}_{b,i-1} \in \mathbb{R}^{T \times F}$, and a speaker embedding $\mathbf{e}_{b-1,i} \in \mathbb{R}^d$, where b is the index of the audio block; i , the index of the speaker; T , the length of the audio block; and F , the maximum frequency bin of the spectrogram. It outputs the speech mask $\mathbf{M}_{b,i} \in \mathbb{R}^{T \times F}$ and an updated speaker embedding for the speaker corresponding to $\mathbf{e}_{b,i}$. The neural network is applied in an iterative fashion for each audio block b , and for each speaker i as follows:

1. Repeat (a) and (b) for $b = 1, 2, \dots$

(a) $\mathbf{R}_{b,0} = \mathbf{1}$

(b) Repeat (i)–(iii) for $i = 1, 2, \dots$ until being stopped at (iii).

(i) $\mathbf{M}_{b,i}, \mathbf{e}_{b,i} = \text{NN}(\mathbf{X}_b, \mathbf{R}_{b,i-1}, \mathbf{e}_{b-1,i})$ ($\mathbf{e}_{b-1,i}$ is set to $\mathbf{0}$ if it was not calculated previously)

(ii) $\mathbf{R}_{b,i} = \max(\mathbf{R}_{b,i-1} - \mathbf{M}_{b,i}, \mathbf{0})$

(iii) If $\frac{1}{TF} \sum_{t,f} \mathbf{R}_{b,i}(t, f) < \text{threshold}$, stop iteration.

A separated speech for speaker i at audio block b can be obtained by $\mathbf{M}_{b,i} \odot \mathbf{X}_b$ where \odot is the element-wise multiplication. The speaker embedding $\mathbf{e}_{b,i}$ is used to keep track of the speaker of adjacent blocks. Thanks to the iterative approach, this neural network

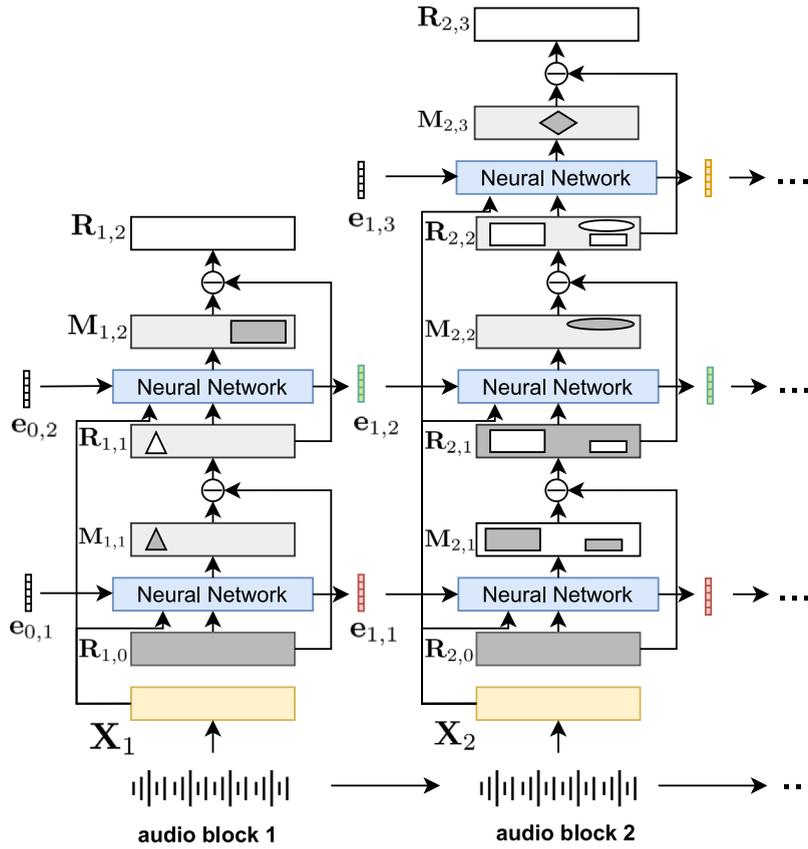


Fig. 11. Joint speech separation, speaker counting, and speaker diarization model.

can cope with the variable number of speakers while jointly performing speech separation and speaker diarization. The online RSAN was evaluated by using real meeting dataset with up to six speakers, and showed better results than the clustering-based method (Kinoshita et al., 2020).

3.2.4. Fully end-to-end neural diarization

Recently, the framework called EEND was proposed (Fujita et al., 2019,b), which performs all the speaker diarization procedures based on a single neural network. The architecture of EEND is shown in Fig. 12. An input to the EEND model is a T -length sequence of acoustic features (e.g., log Mel-filterbank), $\mathbf{X} = (\mathbf{x}_t \in \mathbb{R}^F | t = 1, \dots, T)$. A neural network then outputs the corresponding speaker label sequence $\mathbf{Y} = (y_t | t = 1, \dots, T)$ where $y_t = [y_{t,k} \in \{0, 1\} | k = 1, \dots, K]$. Here, $y_{t,k} = 1$ represents the speech activity of the speaker k at the time frame t , and K is the maximum number of speakers that the neural network can output. Importantly, $y_{t,k}$ and $y_{t,k'}$ can be both 1 for different speakers k and k' , indicating that these two speakers k and k' are speaking simultaneously (i.e. overlapping speech). The neural network is trained to maximize $\log P(\mathbf{Y}|\mathbf{X}) \sim \sum_t \sum_k \log P(y_{t,k}|\mathbf{X})$ over the training data by assuming the conditional independence of the output $y_{t,k}$. Because there can be multiple candidates of the reference label \mathbf{Y} by swapping the speaker index k , the loss function is calculated for all possible reference labels and the reference label that has the minimum loss is used for the error back-propagation, which is inspired by the permutation free objective used in speech separation (Kolbæk et al., 2017). EEND was initially proposed using a bidirectional long short-term memory (BLSTM) network (Fujita et al., 2019), and was soon extended to the self-attention-based network (Fujita et al., 2019b) by showing the state-of-the-art DER for two-speaker data such as the two-speaker excerpt from the CALLHOME dataset (LDC2001S97) and the dialogue audio in the corpus of Spontaneous Japanese (Maekawa, 2003).

EEND has multiple advantages. First, EEND can handle overlapping speech in a sound way. Second, the network is directly optimized toward the maximization of diarization accuracy, by which we can expect a high accuracy. Third, it can be retrained by a real data (i.e. not synthetic data) just by feeding a reference diarization label while it is often not straightforward for the prior works. However, EEND also has several limitations. First, the model architecture limits the maximum number of speakers that the model can cope with. Second, EEND consists of BLSTM or self-attention based neural networks, making it difficult to do online processing. Third, it was empirically suggested that EEND tends to overfit to the distribution of the training data (Fujita et al., 2019).

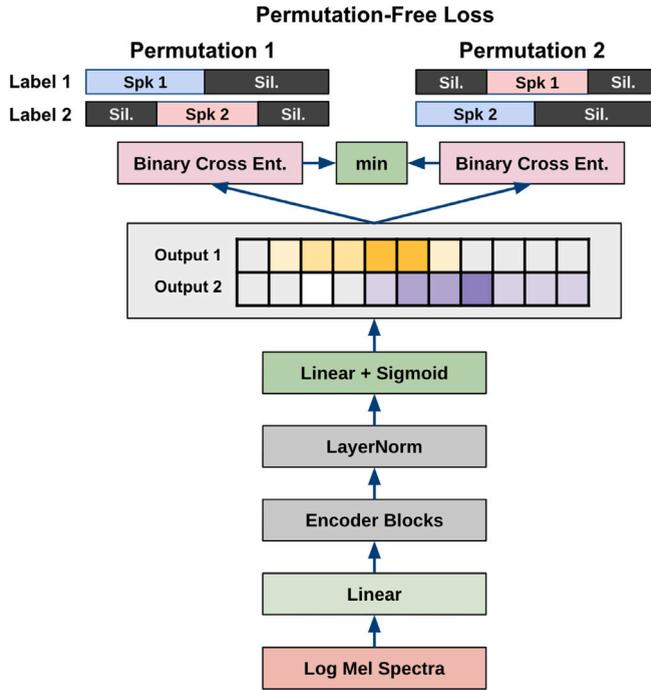


Fig. 12. Two-speaker end-to-end neural diarization model.

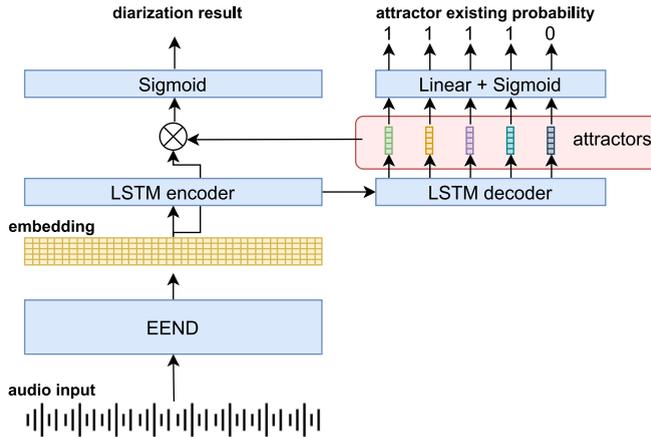


Fig. 13. End-to-end neural diarization with encoder-decoder-based attractor (EDA).

To cope with an unbounded number of speakers, several extensions of EEND have been investigated. Horiguchi et al. (2020) proposed an extension of EEND with the encoder-decoder-based attractor (EDA) (Fig. 13). This method applies an LSTM-based encoder-decoder on the output of EEND to generate multiple attractors. Attractors are generated until the attractor existing probability becomes less than the threshold. Then, each attractor is multiplied by the embeddings generated from EEND to calculate the speech activity for each speaker. EEND-EDA was evaluated on CALLHOME (two to six speakers) and DIHARD 2 (one to nine speakers) dataset and showed better performance than the clustering-based baseline system.

On the other hand, Fujita et al. (2020) proposed another approach to output the speech activity one after another by using a conditional speaker chain rule. In this method, a neural network is trained to produce a posterior probability $P(y_k|y_1, \dots, y_{k-1}, \mathbf{X})$, where $y_k = (y_{t,k} \in \{0, 1\} | t = 1, \dots, T)$ is the speech activity for the k th speaker. Then, the joint speech activity probability of all speakers can be estimated from the following speaker-wise conditional chain rule as:

$$P(y_1, \dots, y_K | \mathbf{X}) = \prod_{k=1}^K P(y_k | y_1, \dots, y_{k-1}, \mathbf{X}). \tag{30}$$

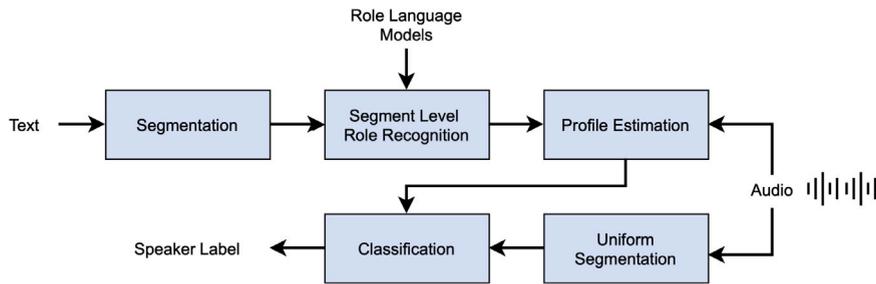


Fig. 14. Integration of lexical information and acoustic information.

During inference, the neural network is repeatedly applied until the speech activity y_k for the last estimated speaker approaches zero. Kinoshita et al. (2021) proposed a different approach that combines EEND and speaker clustering. In their method, a neural network is trained to generate speaker embeddings and the speech activity probability. Speaker clustering constrained by the estimated speech activity by EEND is applied to align the estimated speakers among the different processing blocks.

There are also a few recent trials to extend the EEND for online processing. Xue et al. (2020) proposed a method using a speaker tracing buffer to better align the speaker labels of adjacent processing blocks. Han et al. (2021) proposed a block online version of EEND-EDA (Horiguchi et al., 2020) by carrying the hidden state of the LSTM-encoder to generate the attractors block by block.

4. Speaker diarization in the context of ASR

From a conventional perspective, speaker diarization is considered a pre-processing step for ASR. In the traditional system structures for speaker diarization, presented in Fig. 1, speech inputs are processed sequentially across the diarization components without considering the ASR performance, which is usually measured using the word error rate (WER). WER is the number of misrecognized words (substitution error, insertion error, and deletion error) divided by the number of reference words. One issue is that the tight boundaries of speech segments as the outcomes of speaker diarization have a high chance of causing unexpected word truncation or deletion errors in ASR decoding. In this section we discuss how the speaker diarization systems have been developed in the context of ASR, not only resulting in better WER by preventing speaker diarization from affecting the ASR performance, but also benefiting from ASR artifacts to enhance diarization performance. More recently, there have been a few pioneering proposals made for the joint modeling of speaker diarization and ASR, which will also be introduced in this section.

4.1. Early works

The lexical information from the ASR output has been employed for the speaker diarization system in a few different ways. First, the earliest approach was the RT03 evaluation (Tranter et al., 2003) which used word boundary information for the purpose of segmentation. In Tranter et al. (2003), a general ASR system for broadcast news data was built, in which the basic components are segmentation, speaker clustering, speaker adaptation and system combination after ASR decoding from the two sub-systems with the different adaptation methods. The authors used the word boundary information from the ASR system for speech segmentation, and compared it with the BIC-based speech segmentation. While the performance gain by the ASR-based segmentation was insignificant, this was the first attempt to take advantage of ASR output to enhance the diarization performance. In addition, the ASR result was used to refine SAD in IBM's submission (Huang et al., 2007) for RT07 evaluation. The system that appeared in Huang et al. (2007) incorporates word alignments from the speaker independent ASR module and refines the SAD result to reduce false alarms so that the speaker diarization system can have better clustering quality. The segmentation system in Silovsky et al. (2012) also takes advantage of word alignments from ASR. The authors in Silovsky et al. (2012) focused on the word-breakage problem, in which the words from the ASR output are truncated by segmentation results since the segmentation results and the decoded word sequences are not aligned. Therefore, word-breakage ratio was proposed to measure the rate of change points detected inside intervals corresponding to words. The DER and word-breakage ratio were used to measure the influence of the word truncation problem. While the aforementioned early works of speaker diarization systems that leverage the ASR output focus on the word alignment information to refine the SAD or segmentation result, the speaker diarization system in Canseco-Rodriguez et al. (2004) created a dictionary for the phrases commonly appearing in broadcast news. The phrases in this dictionary provide the identity of who is speaking, who will speak and who spoke in the broadcast news scenario. For example, "This is [name]" indicates who was the speaker of the broadcast news section. Although the early studies on speaker diarization did not fully leverage the lexical information to drastically improve the DER, the idea of integrating the information from ASR output has been adopted by many studies to refine or improve the speaker diarization output.

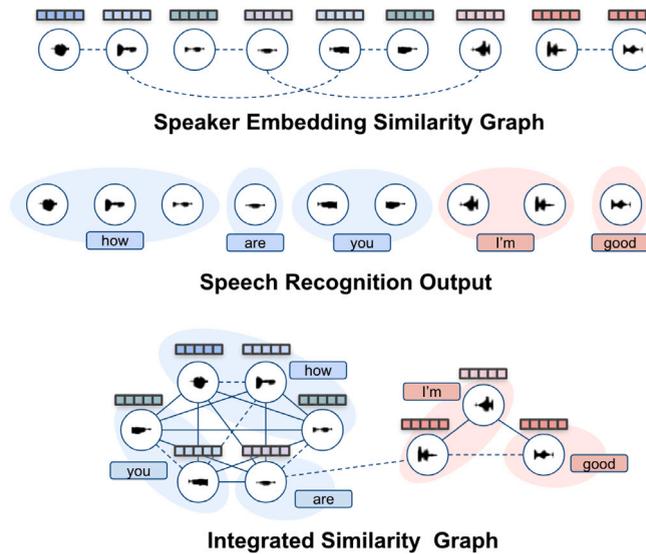


Fig. 15. Integration of lexical information and acoustic information.

4.2. Using lexical information from ASR

The more recent speaker diarization systems that take advantage of the ASR transcript have employed a DNN model to capture the linguistic pattern in the given ASR output to enhance the speaker diarization result. The authors in [Flemotomos et al. \(2020\)](#) proposed a way of using the linguistic information for the speaker diarization task where participants have distinct roles that are known to the speaker diarization system. [Fig. 14](#) shows the diagram of the speaker diarization system discussed in [Flemotomos et al. \(2020\)](#). In this system, a neural text-based speaker change detector and a text-based role recognizer are employed. By using both linguistic and acoustic information, DER was significantly improved compared with the acoustic only system.

Lexical information from the ASR output was also used for speaker segmentation ([Park and Georgiou, 2018](#)) by employing a sequence-to-sequence model that outputs speaker turn tokens. Based on the estimated speaker turn, the input utterance is segmented accordingly. The experimental results in [Park and Georgiou \(2018\)](#) indicate that using both acoustic and lexical information can be exploited and an extra advantage can be obtained owing to the word boundaries we get from the ASR output.

The authors of [Park et al. \(2019\)](#) presented follow-up research within the above thread. Unlike the system in [Park and Georgiou \(2018\)](#), the lexical information from the ASR module was integrated with the speech segment clustering process by employing an integrated adjacency matrix. The adjacency matrix is obtained from the max operation between the acoustic information created from affinities among audio segments and lexical information matrix created by segmenting the word sequence into word chunks that are likely to be spoken by the same speaker. [Fig. 15](#) presents a diagram that explains how lexical information is integrated in an affinity matrix with acoustic information. The integrated adjacency matrix leads to an improved speaker diarization performance for the CALLHOME American English dataset.

4.3. Joint ASR and speaker diarization with deep learning

Motivated by the recent success of deep learning and end-to-end modeling, several models have been proposed to jointly perform ASR and speaker diarization. As with the previous section, the ASR results contain a strong cue to improve speaker diarization. On the other hand, speaker diarization results can be used to improve the accuracy of ASR, for example, by adapting the ASR model toward each estimated speaker. Joint modeling can leverage such inter-dependency to improve both ASR and speaker diarization. In the evaluation, a WER metric that counts word hypotheses with speaker-attribution errors as misrecognized words, such as speaker-attributed WER ([Fiscus et al., 2007](#)) or concatenated minimum-permutation WER (cpWER) ([Watanabe et al., 2020](#)), is often used. ASR-specific metrics (e.g., speaker-agnostic WER) or diarization-specific metrics (e.g., DER mentioned in [Section 1.4.1](#)) are also used complementarily.

The first approach is the introduction of a speaker tag in the transcription of end-to-end ASR models ([Fig. 16](#)). [Shafey et al. \(2019\)](#) proposed to insert a speaker role tag (e.g., <doctor> and <patient>) into the output of a recurrent neural network-transducer (RNN-T)-based ASR system. This method was evaluated by using doctor-patient conversation, and a significant reduction in WDER was reported with a marginal degradation of WER.

Similarly, [Mao et al. \(2020\)](#) proposed the insertion of a speaker identity tag into the output of an attention-based encoder-decoder ASR system, and showed an improvement of DER especially when the oracle utterance boundaries were not given. The works by [Shafey et al.](#) and [Mao et al.](#) showed that the insertion of speaker tags is a simple and promising way to jointly perform ASR and

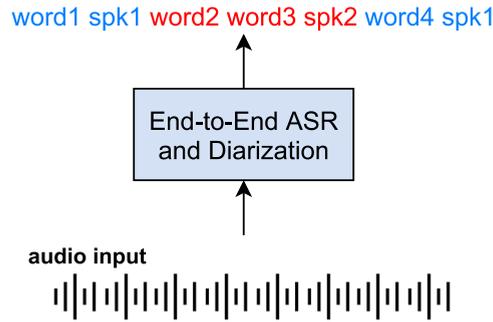


Fig. 16. Joint ASR and diarization by inserting a speaker tag in the transcription.

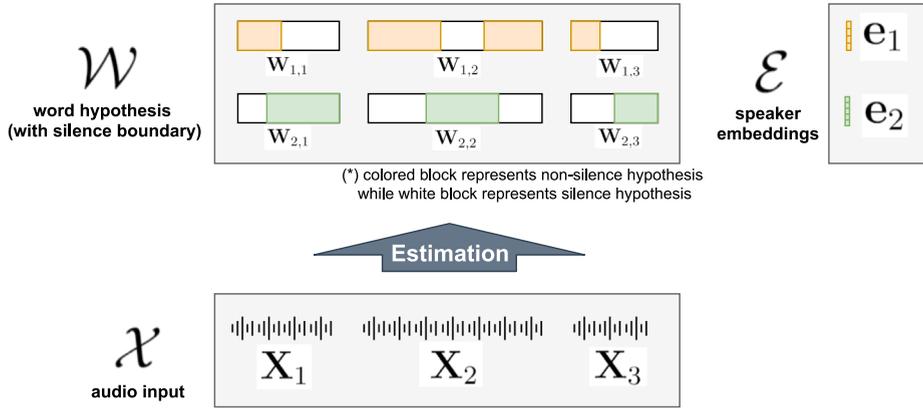


Fig. 17. Joint decoding framework for ASR and speaker diarization.

speaker diarization. On the other hand, the speaker roles or speaker identity tags need to be determined and fixed during training. Thus, it is difficult to cope with an arbitrary number of speakers using this approach.

The second approach is a MAP-based joint decoding framework. Kanda et al. (2019c) formulated the joint decoding of ASR and speaker diarization as follows (see also Fig. 17). Assume that a sequence of observations is represented by $\mathcal{X} = \{X_1, \dots, X_U\}$, where U denotes the number of segments (e.g., generated by applying SAD on a long audio) and X_u denotes the acoustic feature sequence of the u th segment. Further assume that word hypotheses with time boundary information are represented by $\mathcal{W} = \{W_1, \dots, W_U\}$ where W_u is the speech recognition hypothesis corresponding to the segment u . Here, $W_u = (W_{1,u}, \dots, W_{K,u})$ contains all the speakers' hypotheses in the segment u where K denotes the number of speakers, and $W_{k,u}$ represents the speech recognition hypothesis of the k th speaker of the segment u . Finally, a tuple of speaker embeddings $\mathcal{E} = (e_1, \dots, e_K)$, where $e_j \in \mathbb{R}^d$ is the d -dimensional speaker embedding of the k th speaker, is also assumed. With all these notations, the joint decoding framework of multispeaker ASR and diarization can be formulated as a problem to find most likely $\hat{\mathcal{W}}$ as follows:

$$\hat{\mathcal{W}} = \underset{\mathcal{W}}{\operatorname{argmax}} P(\mathcal{W}|\mathcal{X}) \quad (31)$$

$$= \underset{\mathcal{W}}{\operatorname{argmax}} \left\{ \sum_{\mathcal{E}} P(\mathcal{W}, \mathcal{E}|\mathcal{X}) \right\} \quad (32)$$

$$\approx \underset{\mathcal{W}}{\operatorname{argmax}} \{ \max_{\mathcal{E}} P(\mathcal{W}, \mathcal{E}|\mathcal{X}) \}, \quad (33)$$

where the Viterbi approximation is applied to obtain the final equation. This maximization problem is further decomposed into two iterative problems as follows:

$$\hat{\mathcal{W}}^{(i)} = \underset{\mathcal{W}}{\operatorname{argmax}} P(\mathcal{W}|\hat{\mathcal{E}}^{(i-1)}, \mathcal{X}), \quad (34)$$

$$\hat{\mathcal{E}}^{(i)} = \underset{\mathcal{E}}{\operatorname{argmax}} P(\mathcal{E}|\hat{\mathcal{W}}^{(i)}, \mathcal{X}), \quad (35)$$

where i is the iteration index of the procedure. In Kanda et al. (2019c), Eq. (34) is modeled by the target speaker ASR (Zmolikova et al., 2017; Delcroix et al., 2018, 2019; Kanda et al., 2019d) and Eq. (35) is modeled by the overlap-aware speaker embedding estimation. This method obtains a speaker-attributed WER similar to that of the target-speaker ASR with oracle speaker embeddings for two-speaker conversation data of the Corpus of Spontaneous Japanese (Maekawa, 2003). On the other hand, it requires an

- **CALLHOME: NIST SRE 2000 (LDC2001S97)**

NIST SRE 2000 (Disk-8), often referred to as the CALLHOME dataset, is the most widely used dataset for speaker diarization in recent papers. This dataset contains 500 sessions of multilingual telephonic speech. Each session has two to seven speakers with two dominant speakers in each conversation.

- **AMI Corpus**

The AMI database (Carletta et al., 2005) includes 100 h of meeting recordings from multiple sites in 171 meeting sessions. The AMI database provides an audio source recorded using lapel microphones that are separately recorded and amplified for each speaker. Another audio source is recorded using omnidirectional microphone arrays mounted on the table while meeting. The AMI database is a suitable dataset for the evaluation of speaker diarization systems integrated with the ASR module since AMI provides forced alignment data that contains word and phoneme level timings along with the transcript and speaker label. Each meeting session has three to five speakers.

- **ICSI Meeting Corpus**

The ICSI meeting corpus (Janin et al., 2003) contains 75 meeting corpus with four meeting types. The ICSI meeting corpus provides word level timing along with the transcript and speaker label. The audio source is recorded using close-talking individual microphone and six tabletop microphones to provide speaker-specific channel and multichannel recording. Each meeting has 3 to 10 participants.

- **CHiME-5/6 challenge and its dataset**

The CHiME-5 challenge (Barker et al., 2018) and CHiME-6 challenge (Watanabe et al., 2020) were designed as series of ASR competitions for the daily conversation of multiple speakers. The dataset was provided at the CHiME-5 challenge, and it contains 50 h of multiparty real conversations in the everyday home environment. It contains speaker labels, segmentation, and corresponding transcriptions. The audio source is recorded using six four-channel microphone arrays located in the kitchen and dining/living rooms in a house and also binaural microphones worn by participants. The number of participants is fixed at four. While the oracle diarization results were allowed to be used for the ASR task in the CHiME-5 challenge, CHiME-6 challenge track 2 requires the result of both ASR and diarization. The primary evaluation metric for such a track was cpWER, which counts both the speaker-attributed errors and word recognition errors in the WER calculation. DER and JER were also evaluated as secondary metrics without “score collar” and with overlapped regions. The CHiME-5/6 corpus was also used as one track in the DIHARD 2 challenge.

- **VoxSRC Challenge and VoxConverse corpus**

The VoxCeleb Speaker Recognition Challenge (VoxSRC) is the recent evaluation series for speaker recognition systems (Chung et al., 2019; Nagrani et al., 2020). The goal of VoxSRC is to test how well the current technology can cope with the speech “in the wild”. This evaluation series initially started with a pure speaker verification task (Chung et al., 2019), and the diarization task was added as track 4 at the latest evaluation at the VoxCeleb Speaker Recognition Challenge 2020 (VoxSRC-20) (Nagrani et al., 2020). The VoxConverse dataset (Chung et al., 2020) was used for the speaker diarization task with DER as the primary metric, and JER as the secondary metric. The VoxConverse dataset contains 74 h of human conversation extracted from YouTube videos. The dataset is divided into development set (20.3 h, 216 recordings), and test set (53.5 h, 310 recordings). The number of speakers in each recording has a wide range of variety from 1 speaker to 21 speakers. The audio includes various types of noises such as background music, laughter, etc. It also contains a significant portion of overlapping speech from 0% to 30.1% depending on the recording. While the dataset contains the visual information as well as audio, as of June 2021, only the audio of the development set was released under a Creative Commons Attribution 4.0 International License for research purposes. The audio of the evaluation set was used as a blind test set.

- **LibriCSS**

The LibriCSS corpus (Chen et al., 2020) contains 10 h of multichannel recordings and was designed for the research of speech separation, speech recognition, and speaker diarization. It was created by playing the audio in the LibriSpeech corpus (Panayotov et al., 2015) in a real meeting room, and recorded using a 7-channel microphone array. It consists of 10 sessions, each of which is further decomposed to six 10-min mini-sessions. Each mini-session was made by audio of eight speakers and designed to have different overlap ratios from 0% to 40%. To facilitate the research, the baseline system for speech separation and ASR (Chen et al., 2020) and the baseline system that integrates speech separation, speaker diarization and ASR (Raj et al., 2021) have been developed and released.

- **DIHARD Challenge and its dataset**

DIHARD evaluation (Ryant et al., 2018, 2019) focuses on the performance gap of state-of-the-art diarization systems on challenging domains. The first DIHARD challenge, DIHARD 1, started with track 1 (oracle SAD) and track 2 (system SAD). The evaluation data was a collection of various corpus. It includes very challenging datasets such as clinical interviews, web videos, and speech in the wild (e.g., recordings in restaurants), as well as relatively less challenging datasets, such as CTS and audio books to diversify the domains. DIHARD 2 additionally included multichannel speaker diarization task in track 3 (oracle SAD) and track 4 (system SAD) using the recordings from the CHiME-5 corpus (Barker et al., 2018). In the latest DIHARD challenge, DIHARD 3, the CTS dataset was added, whereas multichannel tracks 3 and 4 were excluded. The DIHARD challenge employs DER and JER for the evaluation metric without “score collar” and with overlap region.

- **Rich Transcription Evaluation Series**

The RT Evaluation (NIST, 2009) is the pioneering evaluation series of initiating deeper investigation on speaker diarization in relation to ASR. The main purpose of this effort was to create ASR technologies that would produce transcriptions with descriptive metadata, like who said when, where speaker diarization plays in. Thus, the main tasks in the evaluation were

ASR and speaker diarization. The domains of the data of interest were broadcast news, CTS and meeting recordings with multiple participants. Throughout the period of 2002–2009, the RT evaluation series promoted and gauged advances in speaker diarization and ASR technology. The evaluations among this period are named as RT evaluations (RT-02, RT-03S, RT-03F, and RT-05F) and RT Meeting Recognition (RT-06S, RT-07S and RT-09). These evaluations and their datasets include speaker diarization evaluation as a part of automatic metadata extraction (MDE).

- **Other datasets**

There are also several corpora that have been used for the diarization research but not covered in the list above. The Corpus of Spontaneous Japanese (Maekawa, 2003) contains about 12 h of two-speaker dialogue recorded using headset microphones. AISHELL-4 (Fu et al., 2021) is a relatively new Mandarin Chinese dataset containing 118 h of four to eight speakers in a conference scenario. It is recorded by 8-ch circular microphone array as well as headset microphones for each participant. The ESTER-1 (Gravier et al., 2004) and ESTER-2 (Galliano et al., 2009) evaluation campaign datasets are a set of French recordings designed for three task category: Segmentation (S), Transcription (T) and Information Extraction (E). In the ESTER-1 and ESTER-2 evaluation campaign, speaker diarization was evaluated as one of the core tasks among other tasks including speaker tracking, sound event tracking, and transcriptions. The datasets for ESTER-1 and ESTER-2 include 100 h and 150 h of manually transcribed French radio broadcast news, respectively. ETAPE (Gravier et al., 2012) is also a French speech processing evaluation dataset that contains 36 h of TV and radio shows with both prepared and spontaneous speech. Unlike the ESTER evaluation series, ETAPE targets cross-show speaker diarization.

6. Applications

6.1. Meeting transcription

The goal of meeting transcription is to automatically generate speaker-attributed transcripts during real-life meetings based on their audio and optional video recordings. Accurate meeting transcription is among the processing steps in a pipeline for several tasks, such as, summarization and topic extraction. Similarly, the same transcription system can be used in other domains such as healthcare (Chiu et al., 2017).

Although this task was introduced by NIST in the RT evaluation series back in 2003 (Fiscus et al., 2007; Janin et al., 2003; Carletta et al., 2006), the initial systems had very poor performance, and consequently commercialization of the technology was not possible. However, recent advances in the areas of speech recognition (Xiong et al., 2016; Saon et al., 2017), far-field speech processing (Yoshioka et al., 2015; Du et al., 2016; Li et al., 2017), speaker ID and diarization (Dimitriadis and Fousek, 2017b; Zhang et al., 2019; Sell et al., 2018b), have greatly improved the speaker-attributed transcription accuracy, enabling such commercialization. Bimodal processing combining cameras with microphone arrays has further improved the overall performance (He et al., 2016, 2017).

The variety of application scenarios, customer needs, and business scope, different constraints may be imposed on meeting transcription systems. For example, it is most often required to provide the resulting transcriptions in low latency, making the diarization and recognition even more challenging. However, the architecture of the transcription system can substantially improve the overall performance, e.g., by using microphone arrays of known geometry as the input device. Also, in the case where the expected meeting attendees are known beforehand, the transcription system can further improve speaker attribution, all while providing the exact name of the speaker, instead of randomly generated discrete speaker labels.

Two different scenarios in this space are presented: first, a fixed-geometry microphone array combined with a fish-eye camera system. Second, an ad-hoc geometry microphone array system without a camera. In both scenarios, a “non-binding” list of participants and their corresponding speaker profiles are considered to be known. In particular, the transcription system has access to the invitees’ names and profiles; however, the actual attendees may not accurately match those invited. As such, there is an option to include “unannounced” participants. In addition, some of the invitees may not have profiles. In both scenarios, there is a constraint of low-latency transcriptions, where initial results need to be shown with low latency. The finalized results can be updated later in an offline mode. Some of the technical challenges to overcome are Yoshioka et al. (2019a):

1. Although ASR on overlapping speech is one of the main challenges in meeting transcription, limited progress has been made over the years. Numerous multichannel speech separation methods have been proposed based on independent component analysis (ICA) or spatial clustering (Buchner et al., 2005; Sawada et al., 2007; Nesta et al., 2011; Sawada et al., 2011; Ito et al., 2014; Drude and Haeb-Umbach, 2017), but their application to a meeting setup had limited success. In addition, neural network-based separation methods such as permutation invariant training (PIT) (Kolbæk et al., 2017) or deep clustering (DC) (Hershey et al., 2016) cannot adequately address reverberation and background noise (Maciejewski et al., 2018).
2. Flexible framework: It is desirable that the transcription system is capable of processing all the available information, such as the multichannel audio and visual cues. The system needs to process a dynamically changing number of audio channels without loss of performance. As such, the architecture needs to be modular enough to encompass the different settings.
3. The speaker-attributed ASR of natural meetings requires online or streaming ASR, audio pre-processing such as dereverberation, and accurate diarization and speaker identification. These multiple processing steps are usually optimized separately and thus, the overall pipeline is most frequently inefficient.

4. The use of multiple, unsynchronized audio streams, e.g., audio capturing using mobile devices, adds complexity to the meeting setup and processing. In return, we gain a potentially better spatial coverage since the devices are usually distributed around the room and near the speakers. As part of the application scenario, the meeting participants bring their personal devices, which can be repurposed to improve the overall quality of meeting transcription quality. On the other hand, while there are several pioneering studies (Araki et al., 2018), it is unclear what the best strategies are for consolidating multiple asynchronous audio streams and to what extent they work for natural meetings in online and offline setups.

Based on these considerations, an architecture of a meeting transcription system with asynchronous distant microphones has been proposed in Yoshioka et al. (2019b). In this work, various fusion strategies have been investigated: from early fusion beamforming of the audio signals, to mid-fusion combination of senones per channel, to late fusion combination of the diarization and ASR results (Stolcke and Yoshioka, 2019). The resulting system performance was benchmarked on real-world meeting recordings against fixed geometry systems. As aforementioned, the requirement of speaker-attributed transcriptions with low latency was also adhered to. In addition to the end-to-end system analysis, the paper (Yoshioka et al., 2019b) proposed the idea of “leave-one-out beamforming” in the asynchronous multi-microphone setup, enriching the “diversity” of the resulting signals, as proposed in Stolcke (2011). Finally, it is described how an online, incremental version of recognizer output voting error reduction (ROVER) (Fiscus, 1997) can process both the ASR and diarization outputs, enhancing the overall speaker-attributed ASR performance.

6.2. Conversational interaction analysis and behavioral modeling

Speech and spoken language are central to conversational interactions. They carry crucial information about a speaker’s intent, emotions, identity, age and other individual and interpersonal traits and state variables including health state. Computational advances are increasingly allowing access to such rich information (Narayanan and Georgiou, 2013; Bone et al., 2017). For example, knowing how much, and how, a child speaks in an interaction contains critical information about the child’s developmental state, and offers clues to clinicians in diagnosing disorders such as autism (Kumar et al., 2020). Such analyses are made possible by capturing and processing the audio recordings of the interactions, which often involve two or more people. An important foundational step is the identification and association of the speech portions belonging to specific individuals involved in the conversation. The technologies providing these capabilities are SAD and speaker diarization. Speech portions segmented with speaker-specific information provided by speaker diarization, by itself without any explicit lexical transcription, can offer important information to domain experts who can take advantage of speaker diarization results for quantitative turn-taking analysis.

A domain that is the most relevant in such analyses of spoken conversational interactions relates to behavioral signal processing (BSP) (Georgiou et al., 2011; Narayanan and Georgiou, 2013), which refers to the technology and algorithms for modeling and understanding human communicative, affective and social behaviors. For example, these may include analyzing how positive or negative a person is, how empathic an individual is toward another, what the behavior patterns reveal about the relationship status, and the health condition of an individual (Bone et al., 2017). BSP involves addressing all the complexities of spontaneous interactions in conversations with additional challenges involved in handling and understanding emotional, social and interpersonal behavioral dynamics revealed through verbal and nonverbal cues of the interaction participants. Therefore, the knowledge of speaker specific vocal information plays a significant role in BSP, requiring highly accurate speaker diarization performance. For example, the speaker diarization module is employed as a pre-processing module for analyzing psychotherapy mechanisms and quality (Xiao et al., 2016) and suicide risk assessment (Chakravarthula et al., 2020).

Another popular application of speaker diarization for conversation interaction analysis is the medical doctor-patient interactions. In the system described in Mirheidari et al. (2017), the nature of memory problems of a patient is detected from the conversations between neurologists and patients. The speech and language features extracted from the ASR transcripts combined with the speaker diarization results are used to predict the type of disorder. An automated assistant system for medical domain transcription is proposed in Finley et al. (2018), which includes the speaker diarization module, ASR module and natural language generation module. The automated assistant module accepts the audio clip and outputs grammatically correct sentences describing the topic of the conversation, subject and subject’s symptom.

6.3. Audio indexing

Content-based audio indexing is a well known application domain for speaker diarization. It can provide meta-information such as the content or data type of a given audio data to make information retrieval efficient since search query by machines would be limited by such metadata. The more diverse information was available, the better efficiency could be achieved in retrieving audio contents from a database.

One useful piece of information for the audio indexing would be ASR transcripts to understand the content of speech portions in the audio data. Speaker diarization can augment those transcripts in terms of “who spoke when”, which was the main purpose of the RT evaluation series (NIST, 2009), as discussed in Sections 4.1 and 5.3. The aggregated spoken utterances from speakers by a speaker diarization system also enable per-speaker summary or keyword list-up, which can be used for other query values to retrieve relevant contents from the database. Guo et al. (2016), we can get a view of how speaker diarization outputs can be linked for information searching in consumer-facing applications.

6.4. Conversational AI

Thanks to the advances of ASR technology, the applications of ASR have evolved from simple voice command recognition systems to conversational AI systems. Conversational AI systems, as opposed to voice command recognition systems, have features that are lacking in voice command recognition systems. The fundamental idea of conversational AI is to build a machine that humans can talk to and interact with. In this sense, focusing on an interested speaker in a multiparty setting is one of the most important features of conversational AI. Moreover, speaker diarization becomes an essential feature for conversational AI. For example, conversational AI equipped in a car can pay attention to a specific speaker that is demanding a piece of information from the navigation system by applying speaker diarization along with ASR.

Smart speakers and voice assistants are the most popular products in which speaker diarization plays a significant role for conversational AI. Since the response time and online processing are the crucial factors in real-life settings, the demand for end-to-end speaker diarization systems integrated into the ASR pipeline is growing. The performance of incremental (online) ASR and speaker diarization of the commercial ASR services are evaluated and compared in [Addlesee et al. \(2020\)](#). It is expected that the real-time and low latency aspect of speaker diarization will be more emphasized in the speaker diarization systems in the future since the performance of online diarization and online ASR still have much room for improvement.

7. Challenges and the future of speaker diarization

This paper has provided a comprehensive overview of speaker diarization techniques, highlighting the recent development of deep learning-based diarization approaches. In the early days, a speaker diarization system was developed as a pipeline of sub-modules including front-end processing, SAD, segmentation, speaker embedding extraction, clustering, and post-processing, leading to a standalone system without much connection to other components in a given speech application. With the emergence of the deep learning technology, more and more advancements have been made for speaker diarization, from a method that replaces a single module into a deep-learning-based method, to a fully EEND. Furthermore, as the speech recognition technology becomes more accessible, a trend to tightly integrate speaker diarization into the ASR systems has emerged, such as benefiting from the ASR output to improve the accuracy of speaker diarization. Recently, joint modeling for speaker diarization and speech recognition is investigated in an attempt to enhance the overall performance of speaker diarization. Thanks to these great achievements, speaker diarization systems have already been used in many applications, including meeting transcription, conversational interaction analysis, audio indexing, and conversational AI systems.

As we have seen, tremendous progress has been made for speaker diarization systems. Nevertheless, there is still much room for improvement. As the final remark, we conclude this paper by listing the remaining challenges for speaker diarization toward future research and development.

Online processing of speaker diarization. Most speaker diarization methods assume that an entire recording can be observed to execute speaker diarization. However, numerous applications such as meeting transcription systems or smart agents require very short latency for assigning the speaker. While there have been several attempts to make online speaker diarization system both for clustering-based systems (e.g., [Dimitriadis and Fousek, 2017b](#)) and neural network-based diarization systems (e.g., [Zhang et al., 2019](#); [Xue et al., 2020](#); [Han et al., 2021](#)), it still remains as a challenging problem.

Domain mismatch. A model that is trained on a data in a specific domain often works poorly on data in another domain. For example, it is experimentally known that the EEND model tends to overfit to the distribution of the speaker overlaps of the training data ([Fujita et al., 2019](#)). Such a domain mismatch issue is universal for any training-based method. Given the growing interest for trainable speaker diarization systems, it will become more important to evaluate the ability for handling the variety of inputs. The international evaluation efforts for speaker diarization such as the DIHARD challenge ([Ryant et al., 2018, 2019, 2020](#)) or VoxSRC ([Chung et al., 2019](#); [Nagrani et al., 2020](#)) also have great importance in this direction.

Speaker overlap. Overlap of multitalker speech is the inevitable nature of conversation. For example, an average of 12% to 15% of speaker overlap was observed in meeting recordings ([Cetin and Shriberg, 2006](#); [Yoshioka et al., 2018](#)), and it could even increase in daily conversations ([Kanda et al., 2019a](#); [Barker et al., 2018](#); [Watanabe et al., 2020](#)). Nevertheless, many traditional speaker diarization systems, especially clustering-based systems, have only focused on non-overlapping regions and even the overlapping regions are excluded in the evaluation metric ([Garofolo et al., 2004](#)). While the topic has been studied for long years (e.g. early works [Otterson and Ostendorf, 2007](#); [Boakye et al., 2008](#)), there is a growing interest for handling the speaker overlaps toward better speaker diarization, including the application of speech separation ([Xiao et al., 2020](#)), post-processing ([Bullock et al., 2020](#); [Horiguchi et al., 2020b](#)), and joint modeling of speech separation and speaker diarization ([von Neumann et al., 2019](#); [Kanda et al., 2020a](#)).

Integration with ASR. Many applications require ASR results along with speaker diarization results. In the modular combination of speaker diarization and ASR, some systems locate a speaker diarization system before ASR ([Medennikov et al., 2020b](#)) while some systems locate a speaker diarization system after ASR ([Yoshioka et al., 2019a](#); [Park et al., 2019](#)). Both types of systems showed a strong performance for a specific task, and determining the best kind of system architecture for the speaker diarization and ASR tasks is still an open problem ([Raj et al., 2021](#)). Furthermore, there is another line of research to jointly perform speaker diarization and ASR ([Shafey et al., 2019](#); [Mao et al., 2020](#); [Kanda et al., 2019c, 2020a](#)), which was introduced in Section 4. The joint modeling approach could leverage the inter-dependency between speaker diarization and ASR to better perform both tasks. However, it has not yet been fully investigated whether such joint frameworks perform better than the well-tuned modular systems. Overall, the integration of speaker diarization and ASR is one of the hottest topics that is still being investigated by many researchers.

Audiovisual modeling. Visual information contains a strong clue for the identification of speakers. For example, the video captured by a fish-eye camera was used to improve the accuracy of speaker diarization in a meeting transcription task (Yoshioka et al., 2019a). The visual information was also used to significantly improve the accuracy of speaker diarization for speaker diarization on YouTube video (Chung et al., 2020). While these studies showed the effectiveness of visual information, the audiovisual speaker diarization has yet been rarely investigated compared with audio-only speaker diarization, and there will be many rooms for improvement.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Addlesee, A., Yu, Y., Eshghi, A., 2020. A comprehensive evaluation of incremental speech recognition and diarization for conversational AI. In: Proceedings of the International Conference on Computational Linguistics. pp. 3492–3503.
- Ajmera, J., Lathoud, G., McCowan, L., 2004. Clustering and segmenting speakers and their locations in meetings. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 605–608.
- Ajmera, J., McCowan, L., Bourlard, H., 2004b. Robust speaker change detection. *IEEE Signal Process. Lett.* 11 (8), 649–651.
- Ajmera, J., Wooters, C., 2003. A robust speaker clustering algorithm. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. pp. 411–416.
- AMI, 2009. AMI consortium. <http://www.amiproject.org/index.html>.
- Anguera, X., Bozonnet, S., Evans, N., Fredouille, C., Friedland, G., Vinyals, O., 2012. Speaker Diarization: a Review of Recent Research. Vol. 2. No. 2. IEEE, pp. 356–370.
- Anguera, X., Wooters, C., Hernando, J., 2006. Purity algorithms for speaker diarization of meetings data. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 1. pp. 1025–1028.
- Anguera, X., Wooters, C., Hernando, J., 2007. Acoustic beamforming for speaker diarization of meetings. *IEEE Trans. Audio Speech Lang. Process.* 15 (7), 2011–2023.
- Araki, S., Ono, N., Kinoshita, K., Delcroix, M., 2018. Meeting recognition with asynchronous distributed microphone array using block-wise refinement of mask-based MVDR beamformer. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 5694–5698.
- Arora, A., Raj, D., Subramanian, A.S., Li, K., Ben-Yair, B., Maciejewski, M., Żelasko, P., Garcia, P., Watanabe, S., Khudanpur, S., 2020. The JHU multi-microphone multi-speaker ASR system for the CHiME-6 challenge. arXiv preprint [arXiv:2006.07898](https://arxiv.org/abs/2006.07898).
- Barker, J., Watanabe, S., Vincent, E., Trmal, J., 2018. The fifth 'CHiME' speech separation and recognition challenge: dataset, task and baselines. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1561–1565.
- Blei, D.M., Frazier, P.I., 2011. Distance dependent Chinese restaurant processes. *J. Mach. Learn. Res.* 12 (8).
- Boakye, K., Trueba-Hornero, B., Vinyals, O., Friedland, G., 2008. Overlapped speech detection for improved speaker diarization in multiparty meetings. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 4353–4356.
- Boeddecker, C., Heitkaemper, J., Schmalenstroer, J., Drude, L., Heymann, J., Haeb-Umbach, R., 2018. Front-end processing for the CHiME-5 dinner party scenario. In: Proceedings of CHiME 2018 Workshop on Speech Processing in Everyday Environments. pp. 35–40.
- Bonastre, J.-F., Delacourt, P., Fredouille, C., Merlin, T., Wellekens, C., 2000. A speaker tracking system based on speaker turn detection for NIST evaluation. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 2. IEEE, pp. 1177–1180.
- Bone, D., Lee, C.-C., Chaspari, T., Gibson, J., Narayanan, S., 2017. Signal processing and machine learning for mental health research and clinical applications. *IEEE Signal Process. Mag.* 34 (5), 189–196.
- Bozonnet, S., Evans, N., Anguera, X., Vinyals, O., Friedland, G., Fredouille, C., 2010. System output combination for improved speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association. ISCA, pp. 2642–2645.
- Brunner, N., Burget, L., Cernocky, J., Glombek, O., Grezl, F., Karafiat, M., A. van Leeuwen, D., Matejka, P., Schwarz, P., Strasheim, A., 2007. Fusion of Heterogeneous Speaker Recognition Systems in the STBU Submission for the NIST Speaker Recognition Evaluation 2006. Vol. 15. No. 7. pp. 2072–2084.
- Buchner, H., Aichner, R., Kellermann, W., 2005. A Generalization of Blind Source Separation Algorithms for Convolutional Mixtures Based on Second-Order Statistics. Vol. 13. No. 1. Pp. 120–134.
- Bullock, L., Bredin, H., Garcia-Perera, L.P., 2020. Overlap-aware diarization: Resegmentation using neural end-to-end overlapped speech detection. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 7114–7118.
- Canseco-Rodriguez, L., Lamel, L., Gauvain, J.-L., 2004. Speaker diarization from speech transcripts. In: Proceedings of the International Conference on Spoken Language Processing. Vol. 4. pp. 3–7.
- Carletta, J., Ashby, S., Bourban, S., Flynn, M., Guillemot, M., Hain, T., Kadlec, J., Karaiskos, V., Kraaij, W., Kronenthal, M., Lathoud, G., Lincoln, M., Lisowska, A., McCowan, L., and D. Reidsma, W.P., Wellner, P., 2006. The AMI meeting corpus: a pre-announcement. In: Proceedings of Int. Worksh. Machine Learning for Multimodal Interaction. pp. 28–39.
- Carletta, J., Ashby, S., Bourban, S., Flynn, M., Guillemot, M., Hain, T., Kadlec, J., Karaiskos, V., Kraaij, W., Kronenthal, M., et al., 2005. The AMI meeting corpus: A pre-announcement. In: International Workshop on Machine Learning for Multimodal Interaction. Springer, pp. 28–39.
- Castaldo, F., Colibro, D., Dalmasso, E., Laface, P., Vair, C., 2008. Stream-based speaker segmentation using speaker factors and eigenvoices. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 4133–4136.
- Cetin, O., Shriberg, E., 2006. Speaker overlaps and ASR errors in meetings: Effects before, during, and after the overlap. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 1. IEEE, pp. 357–360.
- Chakravarthula, S.N., Nasir, M., Tseng, S.-Y., Li, H., Park, T.J., Baucom, B., Bryan, C.J., Narayanan, S., Georgiou, P., 2020. Automatic prediction of suicidal risk in military couples using multimodal interaction cues from couples conversations. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 6539–6543.
- Chang, X., Qian, Y., Yu, K., Watanabe, S., 2019. End-to-end monaural multi-speaker ASR system without pretraining. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 6256–6260.
- Chen, S.S., Gopalakrishnan, P.S., 1998. Speaker, Environment and Channel Change Detection and Clustering Via the Bayesian Information Criterion. Tech. Rep., IBM T. J. Watson Research Center, pp. 127–132.
- Chen, S., Gopalakrishnan, P., et al., 1998. Speaker, environment and channel change detection and clustering via the bayesian information criterion. In: Proceedings DARPA Broadcast News Transcription and Understanding Workshop. Vol. 8. Virginia, USA, pp. 127–132.
- Chen, Z., Yoshioka, T., Lu, L., Zhou, T., Meng, Z., Luo, Y., Wu, J., Xiao, X., Li, J., 2020. Continuous speech separation: Dataset and analysis. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 7284–7288.

- Chengalvarayan, R., 1999. Robust energy normalization using speech/nonspeech discriminator for German connected digit recognition. In: Sixth European Conference on Speech Communication and Technology.
- Chiu, C.-C., Tripathi, A., Chou, K., Co, C., Jaitly, N., Jaunzeikare, D., Kannan, A., Nguyen, P., Sak, H., Sankar, A., et al., 2017. Speech recognition for medical conversations. arXiv preprint [arXiv:1711.07274](https://arxiv.org/abs/1711.07274).
- Chung, J.S., Huh, J., Nagrani, A., Afouras, T., Zisserman, A., 2020. Spot the conversation: speaker diarisation in the wild. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 299–303.
- Chung, J.S., Nagrani, A., Coto, E., Xie, W., McLaren, M., Reynolds, D.A., Zisserman, A., 2019. VoxSRC 2019: The first VoxCeleb speaker recognition challenge. arXiv preprint [arXiv:1912.02522](https://arxiv.org/abs/1912.02522).
- Comaniciu, D., Meer, P., 2002. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (5), 603–619.
- Dehak, N., Kenny, P., Dehak, R., Dumouchel, P., Ouellet, P., 2011. Front-End Factor Analysis for Speaker Verification. Vol. 19. No. 4. *IEEE*.
- Delacourt, P., Wellekens, C.J., 2000. DISTBIC: A speaker-based segmentation for audio data indexing. *Speech Commun.* 32 (1–2), 111–126.
- Delcroix, M., Watanabe, S., Ochiai, T., Kinoshita, K., Karita, S., Ogawa, A., Nakatani, T., 2019. End-to-end speakerbeam for single channel target speech recognition. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 451–455.
- Delcroix, M., Zmolikova, K., Kinoshita, K., Ogawa, A., Nakatani, T., 2018. Single channel target speaker extraction and recognition with speaker beam. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 5554–5558.
- Diez, M., Burget, L., Landini, F., Černocký, J., 2019a. Analysis of Speaker Diarization Based on Bayesian HMM with Eigenvoice Priors. Vol. 28. *IEEE*, pp. 355–368.
- Diez, M., Burget, L., Landini, F., Wang, S., Černocký, H., 2020. Optimizing Bayesian HMM based x-vector clustering for the second DIHARD speech diarization challenge. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 6519–6523.
- Diez, M., Burget, L., Matejka, P., 2018a. Speaker diarization based on Bayesian HMM with eigenvoice priors. In: *Odyssey*. pp. 147–154.
- Diez, M., Burget, L., Wang, S., Rohdin, J., Černocký, J., 2019. Bayesian HMM based x-vector clustering for speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 346–350.
- Diez, M., Landini, F., Burget, L., Rohdin, J., Silnova, A., Zmoliková, K., Novotný, O., Veselý, K., Glembek, O., Plchot, O., et al., 2018b. BUT system for DIHARD speech diarization challenge 2018. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2798–2802.
- Dimitriadis, D., 2019. Enhancements for audio-only diarization systems. arXiv preprint [arXiv:1909.00082](https://arxiv.org/abs/1909.00082).
- Dimitriadis, D., Fousek, P., 2017. Developing on-line speaker diarization system. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2739–2743.
- Dimitriadis, D., Fousek, P., 2017. Developing on-line speaker diarization system. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2739–2743.
- Drude, L., Haeb-Umbach, R., 2017. Tight integration of spatial and spectral features for BSS with deep clustering embeddings. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2650–2654.
- Drude, L., Heymann, J., Boeddeker, C., Haeb-Umbach, R., 2018. 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*. VDE, pp. 1–5.
- Drugman, T., Stylianou, Y., Kida, Y., Akamine, M., 2015. Voice activity detection: Merging source and filter-based information. *IEEE Signal Process. Lett.* 23 (2), 252–256.
- Du, J., Tu, Y., Sun, L., Ma, F., Wang, H., Pan, J., Liu, C., Chen, J., Lee, C., 2016. The USTC-iFlytek system for CHiME-4 challenge. In: Proceedings of CHiME-4 Workshop. pp. 36–38.
- Erdogan, H., Hershey, J.R., Watanabe, S., Le Roux, J., 2015. Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 708–712.
- Erdogan, H., Hershey, J.R., Watanabe, S., Mandel, M.I., Le Roux, J., 2016. Improved MVDR beamforming using single-channel mask prediction networks. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1981–1985.
- Finley, G.P., Edwards, E., Robinson, A., Sadoughi, N., Fone, J., Miller, M., Suendermann-Oeft, D., Brenndorfer, M., Axtmann, N., 2018. An automated assistant for medical scribes. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 3212–3213.
- Fiscus, J.G., 1997. A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (ROVER). In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. *IEEE*, pp. 347–354.
- Fiscus, J.G., Ajot, J., Garofolo, J.S., 2007. The rich transcription 2007 meeting recognition evaluation. pp. 373–389.
- Fiscus, J.G., Ajot, J., Michel, M., Garofolo, J.S., 2006. The rich transcription 2006 spring meeting recognition evaluation. In: Proceedings of International Workshop on Machine Learning and Multimodal Interaction. NIST, pp. 309–322.
- Flemtomos, N., Dimitriadis, D., 2020. A memory augmented architecture for continuous speaker identification in meetings. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 6524–6528.
- Flemtomos, N., Georgiou, P., Narayanan, S., 2020. Linguistically aided speaker diarization using speaker role information. In: *Odyssey*. pp. 117–124.
- Fu, Y., Cheng, L., Lv, S., Jv, Y., Kong, Y., Chen, Z., Hu, Y., Xie, L., Wu, J., Bu, H., et al., 2021. AISHELL-4: An open source dataset for speech enhancement, separation, recognition and speaker diarization in conference scenario. arXiv preprint [arXiv:2104.03603](https://arxiv.org/abs/2104.03603).
- Fujita, Y., Kanda, N., Horiguchi, S., Nagamatsu, K., Watanabe, S., 2019. End-to-end neural speaker diarization with permutation-free objectives. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 4300–4304.
- Fujita, Y., Kanda, N., Horiguchi, S., Xue, Y., Nagamatsu, K., Watanabe, S., 2019b. End-to-end neural speaker diarization with self-attention. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. *IEEE*, pp. 296–303.
- Fujita, Y., Watanabe, S., Horiguchi, S., Xue, Y., Shi, J., Nagamatsu, K., 2020. Neural speaker diarization with speaker-wise chain rule. arXiv preprint [arXiv:2006.01796](https://arxiv.org/abs/2006.01796).
- Galliano, S., Gravier, G., Chaubard, L., 2009. The ESTER 2 evaluation campaign for the rich transcription of French radio broadcasts. In: Tenth Annual Conference of the International Speech Communication Association.
- Gangadharaiah, R., Narayanaswamy, B., Balakrishnan, N., 2004. A novel method for two-speaker segmentation. In: Eighth International Conference on Spoken Language Processing.
- Gao, T., Du, J., Dai, L.-R., Lee, C.-H., 2018. Densely connected progressive learning for lstm-based speech enhancement. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 5054–5058.
- Garcia-Romero, D., Espy-Wilson, C.Y., 2011. Analysis of i-vector Length Normalization in Speaker Recognition Systems. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 249–252.
- Garofolo, J.S., Laprun, C.D., Fiscus, J.G., 2004. The Rich Transcription 2004 Spring Meeting Recognition Evaluation. NIST.
- Gauvain, J.-L., Lamel, L., Adda, G., 1998. Partitioning and transcription of broadcast news data. In: Proceedings of the International Conference on Spoken Language Processing. pp. 1335–1338.
- Gelly, G., Gauvain, J.-L., 2017. Optimization of RNN-based speech activity detection. *IEEE/ACM Trans. Audio Speech Lang Process.* 26 (3), 646–656.
- Georgiou, P.G., Black, M.P., Narayanan, S.S., 2011. Behavioral signal processing for understanding (distressed) dyadic interactions: some recent developments. In: Proceedings of the Joint ACM Workshop on Human Gesture and Behavior Understanding. pp. 7–12.
- Gish, H., Siu, M., Rohlicek, R., 1991. Segregation of speakers for speech recognition and speaker identification. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 873–876.

- Gish, H., Siu, M.-H., Rohlicek, J.R., 1991. Segregation of speakers for speech recognition and speaker identification. In: *Icassp*. Vol. 91. pp. 873–876.
- Gravier, G., Adda, G., Paulson, N., Carré, M., Giraudel, A., Galibert, O., 2012. The ETAPE corpus for the evaluation of speech-based TV content processing in the French language. In: *LREC-Eighth International Conference on Language Resources and Evaluation*. pp. na.
- Gravier, G., Bonastre, J.-F., Geoffrois, E., Galliano, S., McTait, K., Choukri, K., 2004. The ESTER evaluation campaign for the rich transcription of french broadcast news. In: *LREC*.
- Guo, A., Faria, A., Riedhammer, J., 2016. Remeeting – Deep insights to conversations. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 1964–1965.
- Guo, X., Gao, L., Liu, X., Yin, J., 2017. Improved deep embedded clustering with local structure preservation. In: *Proceedings of International Joint Conference on Artificial Intelligence*. pp. 1753–1759.
- Haeb-Umbach, R., Watanabe, S., Nakatani, T., Bacchiani, M., Hoffmeister, B., Seltzer, M.L., Zen, H., Souden, M., 2019. Speech processing for digital home assistants: Combining signal processing with deep-learning techniques. *36* (6), 111–124.
- Han, E., Lee, C., Stolcke, A., 2021. BW-EDA-EEND: Streaming End-to-End Neural Speaker Diarization for a Variable Number of Speakers. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 7193–7197.
- Han, C., Luo, Y., Li, C., Zhou, T., Kinoshita, K., Watanabe, S., Delcroix, M., Erdogan, H., Hershey, J.R., Mesgarani, N., et al., 2020. Continuous speech separation using speaker inventory for long multi-talker recording. arXiv preprint [arXiv:2012.09727](https://arxiv.org/abs/2012.09727).
- Han, K.J., Narayanan, S.S., 2007. A robust stopping criterion for agglomerative hierarchical clustering in a speaker diarization system. In: *Proceedings of the Annual Conference of the International Speech Communication Association*.
- Haws, D., Dimitriadis, D., Saon, G., Thomas, S., Picheny, M., 2016. On the importance of event detection for ASR. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2961–2969.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Heigold, G., Moreno, I., Bengio, S., Shazeer, N., 2016. End-to-end text-dependent speaker verification. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 5115–5119.
- Hershey, J.R., Chen, Z., Le Roux, J., Watanabe, S., 2016. Deep clustering: Discriminative embeddings for segmentation and separation. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, pp. 31–35.
- Heymann, J., Drude, L., Haeb-Umbach, R., 2016. Neural network based spectral mask estimation for acoustic beamforming. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, pp. 196–200.
- Horiguchi, S., Fujita, Y., Watanabe, S., Xue, Y., Nagamatsu, K., 2020. End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 269–273.
- Horiguchi, S., Garcia, P., Fujita, Y., Watanabe, S., Nagamatsu, K., 2020b. End-to-end speaker diarization as post-processing. arXiv preprint [arXiv:2012.10055](https://arxiv.org/abs/2012.10055).
- Huang, J., Marcheret, E., Viswesvariah, K., Potamianos, G., 2007. The IBM RT07 evaluation systems for speaker diarization on lecture meetings. In: *Multimodal Technologies for Perception of Humans*. Springer, pp. 497–508.
- Huang, Z., Watanabe, S., Fujita, Y., Garcia, P., Shao, Y., Povey, D., Khudanpur, S., 2020. Speaker diarization with region proposal network. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, pp. 6514–6518.
- Huijbregts, M., van Leeuwen, D., Jong, F., 2009. The majority wins: a method for combining speaker diarization systems. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. ISCA, pp. 924–927.
- Ito, N., Araki, S., Yoshioka, T., Nakatani, T., 2014. Relaxed disjointness based clustering for joint blind source separation and dereverberation. In: *Proceedings of International Workshop on Acoustic Echo and Noise Control*. pp. 268–272.
- ITU, A., 1996. Silence compression scheme for G. 729 optimized for terminals conforming to recommendation V. 70. ITU-T Recommendation G. 729.
- Jain, U., Sieglar, M.A., Doh, S.-J., Gouvea, E., Huerta, J., Moreno, P.J., Raj, B., Stern, R.M., 1996. Recognition of continuous broadcast news with multiple unknown speakers and environments. In: *Proceedings of ARPA Spoken Language Technology Workshop*. pp. 61–66.
- Janin, A., Baron, D., Edwards, J., Ellis, D., Gelbart, D., Morgan, N., Peskin, B., Pfau, T., Shriberg, E., Stolcke, A., Wooters, C., 2003. The ICSI meeting corpus. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 1–364–1–367.
- Jiang, Y., Lee, K.A., Wang, L., 2014. PLDA In the i-supervector space for text-independent speaker verification. *EURASIP J. Audio Speech Music Process.* 2014 (1), 1–13.
- Jin, Q., Laskowski, K., Schultz, T., Waibel, A., 2004. Speaker segmentation and clustering in meetings. In: *Proceedings of the International Conference on Spoken Language Processing*. pp. 597–600.
- Kanagasundaram, A., Dean, D., Sridharan, S., McLaren, M., Vogt, R., 2014. i-vector Based speaker recognition using advanced channel compensation techniques. *Comput. Speech Lang.* 28 (1), 121–140.
- Kanagasundaram, A., Dean, D., Vogt, R., McLaren, M., Sridharan, S., Mason, M., 2012. Weighted LDA techniques for i-vector based speaker verification. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 4781–4784.
- Kanagasundaram, A., Vogt, R., Dean, D., Sridharan, S., Mason, M., 2011. i-vector Based speaker recognition on short utterances. In: *Proceedings of the 12th Annual Conference of the International Speech Communication Association*. International Speech Communication Association, pp. 2341–2344.
- Kanda, N., Boeddeker, C., Heitkaemper, J., Fujita, Y., Horiguchi, S., Nagamatsu, K., Haeb-Umbach, R., 2019. Guided source separation meets a strong ASR backend: hitachi/paderborn university joint investigation for dinner Party ASR. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 1248–1252.
- Kanda, N., Chang, X., Gaur, Y., Wang, X., Meng, Z., Chen, Z., Yoshioka, T., 2021. Investigation of end-to-end speaker-attributed ASR for continuous multi-talker recordings. In: *Proceedings of IEEE Spoken Language Technology Workshop*.
- Kanda, N., Fujita, Y., Horiguchi, S., Ikeshita, R., Nagamatsu, K., Watanabe, S., 2019. Acoustic modeling for distant multi-talker speech recognition with single-and multi-channel branches. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 6630–6634.
- Kanda, N., Gaur, Y., Wang, X., Meng, Z., Chen, Z., Zhou, T., Yoshioka, T., 2020. Joint speaker counting, speech recognition, and speaker identification for overlapped speech of any number of speakers. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 36–40.
- Kanda, N., Gaur, Y., Wang, X., Meng, Z., Yoshioka, T., 2020. Serialized output training for end-to-end overlapped speech recognition. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 2797–2801.
- Kanda, N., Horiguchi, S., Fujita, Y., Xue, Y., Nagamatsu, K., Watanabe, S., 2019. Simultaneous speech recognition and speaker diarization for monaural dialogue recordings with target-speaker acoustic models. In: *Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding*. pp. 31–38.
- Kanda, N., Horiguchi, S., Takashima, R., Fujita, Y., Nagamatsu, K., Watanabe, S., 2019. Auxiliary interference speaker loss for target-speaker speech recognition. In: *Proceedings of the Annual Conference of the International Speech Communication Association*. pp. 236–240.
- Kanda, N., Meng, Z., Lu, L., Gaur, Y., Wang, X., Chen, Z., Yoshioka, T., 2021. Minimum Bayes training for end-to-end speaker-attributed ASR. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 6503–6507.
- Kanda, N., Ye, G., Gaur, Y., Wang, X., Meng, Z., Chen, Z., Yoshioka, T., 2021c. End-to-end speaker-attributed ASR with transformer. arXiv preprint [arXiv:2104.02128](https://arxiv.org/abs/2104.02128).

- Kemp, T., Schmidt, M., Westphal, M., Waibel, A., 2000. Strategies for automatic segmentation of audio data. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 3. pp. 1423–1426.
- Kenny, P., 2008. Bayesian Analysis of Speaker Diarization with Eigenvoice Priors. CRIM, Montreal, Technical report.
- Kenny, P., 2010. Bayesian speaker verification with heavy-tailed priors. In: *Odyssey*. Vol. 14.
- Kenny, P., Boulianne, G., Dumouchel, P., 2005. Eigenvoice modeling with sparse training data. *IEEE Trans. Speech Audio Process.* 13 (3), 345–354.
- Kenny, P., Boulianne, G., Ouellet, P., Dumouchel, P., 2007. Speaker and Session Variability in GMM-Based Speaker Verification. Vol. 15. No. 4. *IEEE*, pp. 1448–1460.
- Kenny, P., Ouellet, P., Dehak, N., Gupta, V., Dumouchel, P., 2008. A Study of Interspeaker Variability in Speaker Verification. Vol. 16. No. 5. *IEEE*, pp. 980–988.
- Kenny, P., Reynolds, D., Castaldo, F., 2010. Diarization of Telephone Conversations Using Factor Analysis. Vol. 4. No. 6. *IEEE*, pp. 1059–1070.
- Kenny, P., Reynolds, D., Castaldo, F., 2010. Diarization of telephone conversations using factor analysis. *IEEE J. Sel. Top. Sign. Proces.* 4 (6), 1059–1070.
- Kinoshita, K., Delcroix, M., Araki, S., Nakatani, T., 2020. Tackling real noisy reverberant meetings with all-neural source separation, counting, and diarization system. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 381–385.
- Kinoshita, K., Delcroix, M., Tawara, N., 2021. Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 7198–7202.
- Kolbæk, M., Yu, D., Tan, Z.-H., Jensen, J., 2017. Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks. 25 (10), 1901–1913.
- Kounades-Bastian, D., Girin, L., Alameda-Pineda, X., Gannot, S., Horaud, R., 2017a. An EM algorithm for joint source separation and diarization of multichannel convolutive speech mixtures. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 16–20.
- Kounades-Bastian, D., Girin, L., Alameda-Pineda, X., Horaud, R., Gannot, S., 2017b. Exploiting the intermittency of speech for joint separation and diarization. In: Proceedings of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics. *IEEE*, pp. 41–45.
- Kuhn, H.W., 1955. The hungarian method for the assignment problem. *Nav. Res. Logist. Q.* 2 (1–2), 83–97.
- Kumar, M., Kim, S.H., Lord, C., Narayanan, S., 2020. Speaker diarization for naturalistic child-adult conversational interactions using contextual information. *J. Acoust. Soc. Am.* 147 (2), EL196–EL200. <http://dx.doi.org/10.1121/10.0000736>.
- Landini, F., Profant, J., Diez, M., Burget, L., 2020. Bayesian HMM clustering of x-vector sequences (VBx) in speaker diarization: theory, implementation and analysis on standard tasks. arXiv preprint [arXiv:2006.07898](https://arxiv.org/abs/2006.07898).
- Leeuwen, D.A.V., Konecny, M., 2007. Progress in the AMIDA speaker diarization system for meeting data. In: Proceedings of International Evaluation Workshops CLEAR 2007 and RT 2007. pp. 475–483.
- Li, B., Sainath, T.N., Narayanan, A., Caroselli, J., Bacchiani, M., Misra, A., Shafran, I., Sak, H., Punduk, G., Chin, K., Sim, K.C., Weiss, R.J., Wilson, K.W., Variiani, E., Kim, C., Siohan, O., Weintraub, M., McDermott, E., Rose, R., Shannon, M., 2017. Acoustic modeling for Google Home. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 399–403.
- Lin, Q., Hou, Y., Li, M., 2020. Self-attentive similarity measurement strategies in speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 284–288.
- Lin, Q., Yin, R., Li, M., Bredin, H., Barras, C., 2019. LSTM based similarity measurement with spectral clustering for speaker diarization. In: Proc. Interspeech 2019. pp. 366–370.
- Liu, D., Kubala, F., 1999. Fast speaker change detection for broadcast news transcription and indexing. In: Proceedings of the International Conference on Spoken Language Processing. pp. 1031–1034.
- Liu, D., Kubala, F., 2003. A cross-channel modeling approach for automatic segmentation of conversational telephone speech. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. pp. 333–338.
- Loizou, P.C., 2013. *Speech Enhancement: Theory and Practice*. CRC Press.
- Luo, Y., Mesgarani, N., 2019. Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation. 27 (8), 1256–1266.
- Luque, J., Hernando, J., 2012. On the use of agglomerative and spectral clustering in speaker diarization of meetings. In: *Odyssey*. pp. 130–137.
- Maciejewski, M., Sell, G., Garcia-Perera, L.P., Watanabe, S., Khudanpur, S., 2018. Building corpora for single-channel speech separation across multiple domains. arXiv preprint [arXiv:1811.02641](https://arxiv.org/abs/1811.02641).
- MacQueen, J., et al., 1967. Some methods for classification and analysis of multivariate observations. In: Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability. Vol. 14. Oakland, CA, USA, pp. 281–297.
- Maekawa, K., 2003. Corpus of spontaneous Japanese: Its design and evaluation. In: ISCA & IEEE Workshop on Spontaneous Speech Processing and Recognition. pp. 7–12.
- Malegaonkar, A., Ariyaeeinia, A., Sivakumaran, P., Fortuna, J., 2006. Unsupervised speaker change detection using probabilistic pattern matching. *IEEE Signal Process. Lett.* 13 (8), 509–512.
- Mao, H.H., Li, S., McAuley, J., Cottrell, G., 2020. Speech recognition and multi-speaker diarization of long conversations. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 691–695.
- Matějka, P., Glembek, O., Castaldo, F., Alam, M.J., Plchot, O., Kenny, P., Burget, L., Černocký, J., 2011. Full-covariance UBM and heavy-tailed PLDA in i-vector speaker verification. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. *IEEE*, pp. 4828–4831.
- Medennikov, I., Korenevsky, M., Prisyach, T., Khokhlov, Y., Korenevskaya, M., Sorokin, I., Timofeeva, T., Mitrofanov, A., Andrusenko, A., Podluzhny, I., Laptev, A., Romanenko, A., 2020. Target-speaker voice activity detection: a novel approach for multi-speaker diarization in a dinner party scenario. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 274–278.
- Medennikov, I., Korenevsky, M., Prisyach, T., Khokhlov, Y., Korenevskaya, M., Sorokin, I., Timofeeva, T., Mitrofanov, A., Andrusenko, A., Podluzhny, I., et al., 2020. The STC system for the CHiME-6 challenge. In: CHiME 2020 Workshop on Speech Processing in Everyday Environments.
- Meignier, S., Moraru, D., Fredouille, C., Bonastre, J.-F., Besacier, L., 2006. Step-by-step and integrated approaches in broadcast news speaker diarization. *Comput. Speech Lang.* 20 (2–3), 303–330.
- Mirheidari, B., Blackburn, D., Harkness, K., Walker, T., Venneri, A., Reuber, M., Christensen, H., 2017. Toward the automation of diagnostic conversation analysis in patients with memory complaints. *J. Alzheimer's Dis.* 58 (2), 373–387.
- Mori, K., Nakagawa, S., 2001. Speaker change detection and speaker clustering using VQ distortion for broadcast news speech recognition. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 1. pp. 413–416.
- Nagrani, A., Chung, J.S., Huh, J., Brown, A., Coto, E., Xie, W., McLaren, M., Reynolds, D.A., Zisserman, A., 2020. VoxSRC 2020: The second VoxCeleb speaker recognition challenge. arXiv preprint [arXiv:2012.06867](https://arxiv.org/abs/2012.06867).
- Nakatani, T., Yoshioka, T., Kinoshita, K., Miyoshi, M., Juang, B.-H., 2010. Speech dereverberation based on variance-normalized delayed linear prediction. 18 (7), 1717–1731.
- Narayanan, S., Georgiou, P.G., 2013. Behavioral signal processing: Deriving human behavioral informatics from speech and language. *Proc. IEEE* 101 (5), 1203–1233.
- Nemer, E., Goubran, R., Mahmoud, S., 2001. Robust voice activity detection using higher-order statistics in the LPC residual domain. *IEEE Trans. Speech Audio Process.* 9 (3), 217–231.
- Nesta, F., Svaizer, P., Omologo, M., 2011. Convolutional BSS of short mixtures by ICA recursively regularized across frequencies. Vol. 19. No. 3. pp. 624–639.
- Ng, A., Jordan, M., Weiss, Y., 2001. On spectral clustering: Analysis and an algorithm. *Adv. Neural Inf. Process. Syst.* 14, 849–856.

- Ng, T., Zhang, B., Nguyen, L., Matsoukas, S., Zhou, X., Mesgarani, N., Vesely, K., Matějka, P., 2012. Developing a speech activity detection system for the DARPA RATS program. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1969–1972.
- Ning, H., Liu, M., Tang, H., Huang, T.S., 2006. A spectral clustering approach to speaker diarization. In: Proceedings of the International Conference on Spoken Language Processing. pp. 2178–2181.
- NIST, 2009. Rich transcription evaluation. <https://www.nist.gov/itl/iad/mig/rich-transcription-evaluation>.
- Novoselov, S., Gusev, A., Ivanov, A., Pekhovsky, T., Shulipa, A., Avdeeva, A., Gorlanov, A., Kozlov, A., 2019. Speaker diarization with deep speaker embeddings for DIHARD challenge II. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1003–1007.
- Otterson, S., Ostendorf, M., 2007. Efficient use of overlap information in speaker diarization. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. IEEE, pp. 683–686.
- Padmanabhan, M., Bahl, L.R., Nahamoo, D., Picheny, M.A., 1996. Speaker clustering and transformation for speaker adaptation in large-vocabulary speech recognition systems. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 701–704.
- Panayotov, V., Chen, G., Povey, D., Khudanpur, S., 2015. LibriSpeech: an ASR corpus based on public domain audio books. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 5206–5210.
- Park, T.J., Georgiou, P., 2018. Multimodal speaker segmentation and diarization using lexical and acoustic cues via sequence to sequence neural networks. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1373–1377.
- Park, T.J., Han, K.J., Huang, J., He, X., Zhou, B., Georgiou, P., Narayanan, S., 2019. Speaker diarization with lexical information. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 391–395.
- Park, T.J., Han, K.J., Kumar, M., Narayanan, S., 2019b. Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap. *IEEE Signal Process. Lett.* 27, 381–385.
- Park, T.J., Kumar, M., Narayanan, S., 2021. Multi-scale speaker diarization with neural affinity score fusion. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 7173–7177.
- Pfau, T., Ellis, D.P., Stolcke, A., 2001. Multispeaker speech activity detection for the ICSI meeting recorder. In: IEEE Workshop on Automatic Speech Recognition and Understanding, 2001. ASRU'01. IEEE, pp. 107–110.
- Raj, D., Denisov, P., Chen, Z., Erdogan, H., Huang, Z., He, M., Watanabe, S., Du, J., Yoshioka, T., Luo, Y., Kanda, N., Li, J., Wisdom, S., R. Hershey, J., 2021. Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis. In: Proceedings of IEEE Spoken Language Technology Workshop.
- Recht, B., Fazel, M., Parrilo, P.A., 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.* 52 (3), 471–501.
- Ren, S., He, K., Girshick, R., Sun, J., 2016. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Vol. 39. No. 6. IEEE, pp. 1137–1149.
- Reynolds, D.A., Quatieri, T.F., Dunn, R.B., 2000. Speaker verification using adapted Gaussian mixture models. *Digit. Signal Process.* 10 (1–3), 19–41.
- Reynolds, D.A., Torres-Carrasquillo, P., 2005. Approaches and applications of audio diarization. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 953–956.
- Rohlicek, J.R., Ayuso, D., Bates, M., Bobrow, R., Boulanger, A., Gish, H., Jeanrenaud, P., Meteer, M., Siu, M., 1992. Gisting conversational speech. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 113–116.
- Rosenberg, A.E., Gorin, A., Liu, Z., Parthasarathy, P., 2002. Unsupervised speaker segmentation of telephone conversations. In: Proceedings of the International Conference on Spoken Language Processing. pp. 565–568.
- Rougui, J.E., Rziza, M., Aboutajdine, D., Gelgon, M., Martinez, J., 2006. Fast incremental clustering of gaussian mixture speaker models for scaling up retrieval in on-line broadcast. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. Vol. 5. IEEE.
- Ryant, N., Church, K., Cieri, C., Cristia, A., Du, J., Ganapathy, S., Liberman, M., 2018. The first DIHARD speech diarization challenge. In: Proceedings of the Annual Conference of the International Speech Communication Association.
- Ryant, N., Church, K., Cieri, C., Cristia, A., Du, J., Ganapathy, S., Liberman, M., 2019. The second DIHARD diarization challenge: dataset, task, and baselines. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 978–982.
- Ryant, N., Church, K., Cieri, C., Du, J., Ganapathy, S., Liberman, M., 2020. Third DIHARD challenge evaluation plan. arXiv preprint [arXiv:2006.05815](https://arxiv.org/abs/2006.05815).
- Ryant, N., Liberman, M., Yuan, J., 2013. Speech activity detection on youtube using deep neural networks. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 728–731.
- Salmun, I., Shapiro, I., Opher, I., Lapidot, I., 2017. PLDA-based mean shift speakers' short segments clustering. *Comput. Speech Lang.* 45, 411–436.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T., 2016. Meta-learning with memory-augmented neural networks. In: Proceedings of International Conference on Machine Learning. pp. 1842–1850.
- Santoro, A., Faulkner, R., Raposo, D., Rae, J., Chrzanowski, M., Weber, T., Wierstra, D., Vinyals, O., Pascanu, R., Lillicrap, T., 2018. Grelational recurrent neural networks. *Proc. Adv. Neural Inf. Process. Syst.* 7299–7310.
- Saon, G., Kurata, G., Sercu, T., Audhkhasi, K., Thomas, S., Dimitriadis, D., Cui, X., Ramabhadran, B., Picheny, M., Lim, L.-L., et al., 2017. English conversational telephone speech recognition by humans and machines. arXiv preprint [arXiv:1703.02136](https://arxiv.org/abs/1703.02136).
- Sarikaya, R., Hansen, J.H., 1998. Robust detection of speech activity in the presence of noise. In: Proceedings of the International Conference on Spoken Language Processing. Vol. 4. Citeseer, pp. 1455–1458.
- Sawada, H., Araki, S., Makino, S., 2007. Measuring dependence of bin-wise separated signals for permutation alignment in frequency-domain BSS. *Int. Symp. Circ. Syst.* 3247–3250.
- Sawada, H., Araki, S., Makino, S., 2011. Underdetermined Convolutional Blind Source Separation Via Frequency Bin-Wise Clustering and Permutation Alignment. Vol. 19. No. 3. pp. 516–527.
- Seki, H., Hori, T., Watanabe, S., Le Roux, J., Hershey, J.R., 2018. A purely end-to-end system for multi-speaker speech recognition. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics. Vol. 1. pp. 2620–2630.
- Sell, G., Garcia-Romero, D., 2014. Speaker diarization with PLDA i-vector scoring and unsupervised calibration. In: Proceedings of IEEE Spoken Language Technology Workshop. IEEE, pp. 413–417.
- Sell, G., Garcia-Romero, D., 2015. Diarization resegmentation in the factor analysis subspace. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 4794–4798.
- Sell, G., Snyder, D., McCree, A., Garcia-Romero, D., Villalba, J., Maciejewski, M., Manohar, V., Dehak, N., Povey, D., Watanabe, S., Khudanpur, S., 2018. Diarization is hard: some experiences and lessons learned for the JHU team in the inaugural DIHARD challenge. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2808–2812.
- Sell, G., Snyder, D., McCree, A., Garcia-Romero, D., Villalba, J., Maciejewski, M., Manohar, V., Dehak, N., Povey, D., Watanabe, S., et al., 2018. Diarization is hard: some experiences and lessons learned for the JHU team in the inaugural DIHARD challenge. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2808–2812.
- Senoussoui, M., Kenny, P., Dehak, N., Dumouchel, P., et al., 2010. An i-vector extractor suitable for speaker recognition with both microphone and telephone speech. In: *Odyssey*. p. 6.

- Senoussaoui, M., Kenny, P., Dumouchel, P., Stafylakis, T., 2013a. Efficient iterative mean shift based cosine dissimilarity for multi-recording speaker clustering. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 7712–7715.
- Senoussaoui, M., Kenny, P., Stafylakis, T., Dumouchel, P., 2013b. A study of the cosine distance-based mean shift for telephone speech diarization. *IEEE/ACM Trans. Audio Speech Lang. Process.* 22 (1), 217–227.
- Shafey, L.E., Soltau, H., Shafraan, I., 2019. Joint speech recognition and speaker diarization via sequence transduction. In: Proceedings of the Annual Conference of the International Speech Communication Association. ISCA, pp. 396–400.
- Shum, S., Dehak, N., Chuangsuwanich, E., Reynolds, D., Glass, J., 2011. Exploiting intra-conversation variability for speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association.
- Shum, S.H., Dehak, N., Dehak, R., Glass, J.R., May 2013. *Unsupervised Methods for Speaker Diarization: an Integrated and Iterative Approach*. Vol. 21. No. 10. IEEE, pp. 2015–2028.
- Shum, S., Dehak, N., Glass, J., 2012. On the use of spectral and iterative methods for speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 482–485.
- Siegler, M.A., Jain, U., Raj, B., Stern, R.M., 1997. Automatic segmentation, classification and clustering of broadcast news audio. In: Proc. DARPA Speech Recognition Workshop.
- Silovsky, J., Zdansky, J., Nouza, J., Cerva, P., Prazak, J., 2012. Incorporation of the ASR output in speaker segmentation and clustering within the task of speaker diarization of broadcast streams. In: *International Workshop on Multimedia Signal Processing*. IEEE, pp. 118–123.
- Siu, M.-H., George, Y., Gish, H., 1992. An unsupervised, sequential learning algorithm for segmentation for speech waveforms with multiple speakers. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 189–192.
- Snyder, D., Garcia-Romero, D., Povey, D., Khudanpur, S., 2017. Deep neural network embeddings for text-independent speaker verification. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 999–1003.
- Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., Khudanpur, S., 2018. X-vectors: Robust DNN embeddings for speaker recognition. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 5329–5333.
- Sohn, J., Kim, N.S., Sung, W., 1999. A statistical model-based voice activity detection. *IEEE Signal Process. Lett.* 6 (1), 1–3.
- Stafylakis, T., Katsouros, V., Carayannis, G., 2010. Speaker clustering via the mean shift algorithm. *Recall* 2, 7.
- Stolcke, A., 2011. Making the most from multiple microphones in meeting recordings. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 4992–4995.
- Stolcke, A., Yoshioka, T., 2019. DOVER: A method for combining diarization outputs. In: *Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE, pp. 757–763.
- Sukhbaatar, S., Weston, J., Fergus, R., et al., 2015. End-to-end memory networks. *Proc. Adv. Neural Inf. Process. Syst.* 2440–2448.
- Sun, Y., Wang, X., Tang, X., 2014. Deep learning face representation from predicting 10,000 classes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1891–1898.
- Taigman, Y., Yang, M., Ranzato, M., Wolf, L., 2014. Deepface: Closing the gap to human-level performance in face verification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1701–1708.
- Thomas, S., Ganapathy, S., Saon, G., Soltau, H., 2014. Analyzing convolutional neural networks for speech activity detection in mismatched acoustic conditions. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 2519–2523.
- Tranter, S.E., Reynolds, D.A., 2004. Speaker diarisation for broadcast news. In: *Odyssey*. pp. 337–344.
- Tranter, S.E., Reynolds, D.A., 2006. An Overview of Automatic Speaker Diarization Systems. Vol. 14. No. 5. pp. 1557–1565.
- Tranter, S.E., Yu, K., Evermann, G., Woodland, P.C., 2004. Generating and evaluating for automatic speech recognition of conversational telephone speech. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 753–756.
- Tranter, S.E., Yu, K., Reynolds, D.A., Evermann, G., Kim, D.Y., Woodland, P.C., 2003. An investigation into the interactions between speaker diarisation systems and automatic speech transcription. *CUED/F-INFENG/TR-464*.
- Tritschler, A., Gopinath, R.A., 1999. Improved speaker segmentation and segments clustering using the bayesian information criterion. In: Sixth European Conference on Speech Communication and Technology.
- Ustinova, E., Lempitsky, V., 2016. Learning deep embeddings with histogram loss. *Proc. Adv. Neural Inf. Process. Syst.* 29, 4170–4178.
- Valente, F., 2005. *Variational Bayesian Methods for Audio Indexing* (Ph.D. thesis).
- Valente, F., Motlicek, P., Vijayasenan, D., 2010. Variational Bayesian speaker diarization of meeting recordings. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 4954–4957.
- Variani, E., Lei, X., McDermott, E., Moreno, I.L., G-Dominguez, J., 2014. Deep neural networks for small footprint text-dependent speaker verification. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 4052–4056.
- Variani, E., Lei, X., McDermott, E., Moreno, I.L., Gonzalez-Dominguez, J., 2014b. Deep neural networks for small footprint text-dependent speaker verification. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 4052–4056.
- Vijayasenan, D., Valente, F., Boulard, H., 2009. An information theoretic approach to speaker diarization of meeting data. *IEEE Trans. Audio Speech Lang. Process.* 17 (7), 1382–1393.
- Villalba, J., Chen, N., Snyder, D., Garcia-Romero, D., McCree, A., Sell, G., Borgstrom, J., Richardson, F., Shon, S., Grondin, F., et al., 2019. State-of-the-art speaker recognition for telephone and video speech: the JHU-MIT submission for NIST SRE18. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 1488–1492.
- Vincent, E., Virtanen, T., Gannot, S., 2018. *Audio Source Separation and Speech Enhancement*. John Wiley & Sons.
- Von Luxburg, U., 2007. A tutorial on spectral clustering. *Stat. Comput.* 17 (4), 395–416.
- von Neumann, T., Kinoshita, K., Delcroix, M., Araki, S., Nakatani, T., Haeb-Umbach, R., 2019. All-neural online source separation, counting, and diarization for meeting analysis. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 91–95.
- Wang, D., Chen, J., 2018. Supervised Speech Separation Based on Deep Learning: an Overview. Vol. 26. No. 10. IEEE, pp. 1702–1726.
- Wang, P., Chen, Z., Xiao, X., Meng, Z., Yoshioka, T., Zhou, T., Lu, L., Li, J., 2019. Speech separation using speaker inventory. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. pp. 230–236.
- Wang, Q., Downey, C., Wan, L., Mansfield, P.A., Moreno, I.L., 2018. Speaker diarization with LSTM0. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 5239–5243.
- Wang, X., Kanda, N., Gaur, Y., Chen, Z., Meng, Z., Yoshioka, T., 2021. Exploring end-to-end multi-channel ASR with bias information for meeting transcription. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding.
- Wang, J., Xiao, X., Wu, J., Ramamurthy, R., Rudzicz, F., Brudno, M., 2020. Speaker diarization with session-level speaker embedding refinement using graph neural networks. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 7109–7113.
- Watanabe, S., Mandel, M., Barker, J., Vincent, E., Arora, A., Chang, X., Khudanpur, S., Manohar, V., Povey, D., Raj, D., et al., 2020. CHiME-6 Challenge: Tackling multispeaker speech recognition for unsegmented recordings. In: 6th International Workshop on Speech Processing in Everyday Environments (CHiME 2020).
- Woo, K.-H., Yang, T.-Y., Park, K.-J., Lee, C., 2000. Robust voice activity detection algorithm for estimating noise spectrum. *Electron. Lett.* 36 (2), 180–181.
- Xiao, B., Huang, C., Imel, Z.E., Atkins, D.C., Georgiou, P., Narayanan, S.S., 2016. A technology prototype system for rating therapist empathy from audio recordings in addiction counseling. *PeerJ Comput. Sci.* 2, e59.

- Xiao, X., Kanda, N., Chen, Z., Zhou, T., Yoshioka, T., Zhao, Y., Liu, G., Wu, J., Li, J., Gong, Y., 2020. Microsoft speaker diarization system for the voxceleb speaker recognition challenge 2020. arXiv preprint [arXiv:2010.11458](https://arxiv.org/abs/2010.11458).
- Xie, J., Girshick, R., Farhadi, A., 2016. Unsupervised deep embedding for clustering analysis. In: Proceedings of International Conference on Machine Learning. pp. 478–487.
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M., Stolcke, A., Yu, D., Zweig, G., 2016. Achieving human parity in conversational speech recognition. arXiv preprint [arXiv:1610.05256](https://arxiv.org/abs/1610.05256).
- Xue, Y., Horiguchi, S., Fujita, Y., Watanabe, S., Nagamatsu, K., 2020. Online end-to-end neural diarization with speaker-tracing buffer. arXiv preprint [arXiv:2006.02616](https://arxiv.org/abs/2006.02616).
- Yoshioka, T., Abramovski, I., Aksoylar, C., Chen, Z., David, M., Dimitriadis, D., Gong, Y., Gurvich, I., Huang, X., Huang, Y., Hurvitz, A., Jiang, L., Koubi, S., Krupka, E., Leichter, I., Liu, C., Parthasarathy, P., Vinnikov, A., Wu, L., Xiao, X., Xiong, W., Wang, H., Wang, Z., Zhang, J., Zhao, Y., Zhou, T., 2019. Advances in online audio-visual meeting transcription. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. pp. 276–283.
- Yoshioka, T., Dimitriadis, D., Stolcke, A., Hinthorn, W., Chen, Z., Zeng, M., Xuedong, H., 2019. Meeting transcription using asynchronous distant microphones. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2968–2972.
- Yoshioka, T., Erdogan, H., Chen, Z., Xiao, X., Alleva, F., 2018. Recognizing overlapped speech in meetings: a multichannel separation approach using neural networks. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 3038–3042.
- Yoshioka, T., Ito, N., Delcroix, M., Ogawa, A., Kinoshita, K., Fujimoto, M., Yu, C., Fabian, W.J., Espi, M., Higuchi, T., Araki, S., Nakatani, T., 2015. The NTT CHiME-3 system: advances in speech enhancement and recognition for mobile multi-microphone devices. In: Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding. pp. 436–443.
- Yoshioka, T., Nakatani, T., 2012. Generalization of multi-channel linear prediction methods for blind MIMO impulse response shortening. *20 (10)*, 2707–2720.
- Yu, D., Chang, X., Qian, Y., 2017. Recognizing multi-talker speech with permutation invariant training. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2456–2460.
- Zajíc, Z., Kunešová, M., Radová, V., 2016. Investigation of segmentation in i-vector based speaker diarization of telephone speech. In: International Conference on Speech and Computer. pp. 411–418.
- Zhang, A., Wang, Q., Zhu, Z., Paisley, J., Wang, C., 2019. Fully supervised speaker diarization. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 6301–6305.
- Zhu, X., Barras, C., Meignier, S., Gauvain, J.-L., 2005. Combining speaker identification and BIC for speaker diarization. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2441–2444.
- Zhu, W., Pelecanos, J., 2016. Online speaker diarization using adapted i-vector transforms. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 5045–5049.
- Zmolikova, K., Delcroix, M., Kinoshita, K., Higuchi, T., Ogawa, A., Nakatani, T., 2017. Speaker-aware neural network based beamformer for speaker extraction in speech mixtures. In: Proceedings of the Annual Conference of the International Speech Communication Association. pp. 2655–2659.