

An experimental review of speaker diarization methods with application to two-speaker conversational telephone speech recordings

Luca Serafini ^{a,*}, Samuele Cornell ^a, Giovanni Morrone ^a, Enrico Zovato ^b, Alessio Brutti ^c, Stefano Squartini ^a

^a Università Politecnica delle Marche, Ancona, Italy

^b PerVoice S.p.A., Trento, Italy

^c Fondazione Bruno Kessler, Trento, Italy

ARTICLE INFO

Keywords:

Speaker diarization
Conversational telephone speech
Deep learning
End-to-end neural diarization
Speech separation guided diarization

ABSTRACT

We performed an experimental review of current diarization systems for the conversational telephone speech (CTS) domain. In detail, we considered a total of eight different algorithms belonging to clustering-based, end-to-end neural diarization (EEND), and speech separation guided diarization (SSGD) paradigms. We studied the inference-time computational requirements and diarization accuracy on four CTS datasets with different characteristics and languages. We found that, among all methods considered, EEND-vector clustering (EEND-VC) offers the best trade-off in terms of computing requirements and performance. More in general, EEND models have been found to be lighter and faster in inference compared to clustering-based methods. However, they also require a large amount of diarization-oriented annotated data. In particular EEND-VC performance in our experiments degraded when the dataset size was reduced, whereas self-attentive EEND (SA-EEND) was less affected. We also found that SA-EEND gives less consistent results among all the datasets compared to EEND-VC, with its performance degrading on long conversations with high speech sparsity. Clustering-based diarization systems, and in particular VBx, instead have more consistent performance compared to SA-EEND but are outperformed by EEND-VC. The gap with respect to this latter is reduced when overlap-aware clustering methods are considered. SSGD is the most computationally demanding method, but it could be convenient if speech recognition has to be performed. Its performance is close to SA-EEND but degrades significantly when the training and inference data characteristics are less matched.

1. Introduction

Speaker diarization, also often referred to as simply diarization, aims at partitioning an audio recording into temporal segments denoting the boundaries of each speaker's utterances. In other words, it addresses the problem of "who spoke when?", without any a-priori knowledge of the speakers' identities. This discipline represents an important research field in the speech processing area. Its main applications involve automatic speech recognition (ASR), speaker indexing, speaker recognition, real-time captioning, and audio analysis to name a few. Speaker diarization constitutes an important and often essential pre-processing step in most of these

* Corresponding author.

E-mail addresses: l.serafini@univpm.it (L. Serafini), s.cornell@pm.univpm.it (S. Cornell), g.morrone@univpm.it (G. Morrone), enrico.zovato@pervoice.it (E. Zovato), brutti@fbk.eu (A. Brutti), s.squartini@univpm.it (S. Squartini).

<https://doi.org/10.1016/j.csl.2023.101534>

Received 14 December 2022; Received in revised form 28 April 2023; Accepted 22 May 2023

Available online 29 May 2023

0885-2308/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

application scenarios: e.g., accurate diarization can be used effectively to drive multi-channel blind source separation algorithms to separate concurrent speakers for distant speech recognition (Boeddeker et al., 2018; Kanda et al., 2019) or speaker adaptation for ASR (Saon et al., 2013; Miao et al., 2015; Wang and Wang, 2017; Sari et al., 2020). It is also a very challenging task as speaker identity is arguably not easy to estimate especially when many speakers are present and they often overlap. In such scenarios, even humans struggle and have to resort to linguistic information from neighboring utterances or long term information to resolve speaker identities' ambiguities on short segments.

As such, it can be argued that the two-speaker conversation case is the simplest scenario when all the other characteristics that can impact diarization performance are equal: e.g., audio quality due to the registration setup, environmental noise, duration and sequence of speakers' turns, as well as the style and spontaneity of speech. On the other hand, the two-speaker scenario constitutes a particularly important use-case as it is a very common scenario and has a great commercial interest. In fact, it represents the most ordinary situation in normal conversations that take place via telephone or remotely via hands-free tele-conferencing, but it also appears in many other contexts, such as in doctor-patient meetings, interviews, air traffic control conversations, et cetera.

In light of the tremendous progress made in recent years, we present in this work an experimental review focused on conversational telephone speech (CTS), and compare current state-of-the-art methods based on clustering, separation, and end-to-end diarization on real-world CTS datasets with different languages. The main motivation is that recent research trends, such as the invention of fully end-to-end diarization methods (Fujita et al., 2019a,b, 2020; Kinoshita et al., 2021b,a, 2022b,a), seems to suggest that the technology is mature enough to attain highly accurate and reliable diarization in the CTS scenario, whereas other scenarios, such as distant-talk meetings with several speakers, still pose several challenges (Watanabe et al., 2020). Our hope is that this work will be useful to researchers and practitioners alike, as we investigate computational requirements and the performance of eight diarization algorithms on multilingual data, all belonging to the CTS domain but with diverse characteristics, such as overlapped speech ratio, average length of the conversation, and speech sparsity.

The remainder of the paper is organized as follows: we present in Sections 1.1 and 1.2 a brief overview of speaker diarization methods and a comparison of this review work with previous ones. Section 2 describes the various systems that have been considered in our experimental analysis, along with their main characteristics. In Section 3 we outline the experimental setup and describe the systems implementation and hyper-parameters employed in our experiments, as well as the diarization metrics and the CTS datasets used for training and testing. In Section 4 we analyze the results. Finally, in Section 5 we draw conclusions and discuss possible future research directions.

1.1. A brief history of speaker diarization

The first works on speaker diarization can be traced back to the 1990s (Gish et al., 1991; Siu et al., 1992; Jain et al., 1996; Chen et al., 1998; Liu and Kubala, 1999). These early works focused on applications such as radio broadcast news and communications, with the main goal of improving ASR performance. As such, some of these works relied directly on ASR outputs (Jain et al., 1996; Liu and Kubala, 1999), e.g., by performing a two-pass decoding, where the first pass was used for detecting speakers' changes. Also focused on ASR, Siu et al. (1992) can be considered one of the very first works on speech separation guided diarization (SSGD) as the authors performed diarization using blind speech separation. In these early years, the foundations for the clustering-based diarization paradigm were laid (Gish et al., 1991; Chen et al., 1998). These works formulated diarization as a clustering problem. After voice activity detection (VAD), speaker discriminative features are extracted from each speech segment. These latter are then clustered together in order to assign each of the features and their corresponding sub-segments to each speaker and thus identify the number of speakers in the recording and obtain a segmentation for each speaker. The features used in such early studies were largely "hand-crafted", with Mel-frequency cepstral coefficients (MFCC) being a popular choice together with linear frequency cepstral coefficients (LFCC), perceptual linear predictive (PLP) (Hermansky, 1990), and linear predictive coding (LPC) (O'Shaughnessy, 1988). Often a combination of more than one feature type (Yamaguchi et al., 2005) was used in order to increase the system robustness by capturing various speech characteristics possibly at different time scales. Regarding clustering, these early works focused on agglomerative hierarchical clustering (AHC) performed based on a distance measure between each feature vector belonging to a different segment. How to define such distance and what properties the feature vectors should have in order to have a well-defined distance from a speaker identity perspective is still an active research direction. The use of Bayesian information criterion (BIC) (Chen et al., 1998) to define such distance and perform AHC was found to be particularly effective at the time and remained a popular choice in subsequent years together with the generalized likelihood ratio (GLR) (Gish et al., 1991).

The first decade of the new millennium saw tremendous improvement and growth in the diarization field and research directions also included more challenging scenarios, such as meetings captured by far-field microphone arrays. For example, the still very relevant and popular AMI (Carletta et al., 2005) dataset was collected in 2005. In this period the clustering-based diarization paradigm was consolidated and became the de-facto standard approach.

One of the major improvements over previous techniques was due to the fact that researchers understood the necessity to move away from fully hand-crafted features and devise data-driven methods to obtain more robust and higher-level speaker-id discriminative features. One of the main challenges in obtaining good speaker identity features lies in the fact that these should be robust to intra-speaker variability (e.g., due to changes in intonation, background noise, etc.) while being enough discriminative to allow for differentiating between distinct speakers (i.e., inter-speaker discriminability). Diarization systems started to incorporate machine learning models such as Gaussian mixture models (GMM). A key work was the one of Reynolds et al. (2000) which introduced the speaker-independent GMM-Universal Background Model (GMM-UBM) for speaker verification. In this work, each

vector of features is derived in a data-driven fashion from a GMM (i.e., it consists of a weighted sum of a finite number of multi-dimensional Gaussian components). Since in most applications there is often not enough data to train from scratch such a GMM model without overfitting, these feature vectors are obtained by adapting some parameters of a pre-trained GMM model: e.g., the means of the Gaussian components which, when concatenated are denoted as a GMM “supervector”. The GMM-UBM model is trained on a large amount of data with a large speaker variability via a maximum a posteriori (MAP) criterion. The use of a GMM-UBM remained a de-facto standard for speaker verification and diarization till the advent of deep neural network (DNN)-based techniques.

One of the problems of such an approach is that it has the drawback of suffering from the aforementioned intra-speaker variability. A number of important works (Kenny et al., 2007; Castaldo et al., 2008) tried to address this problem. For example, joint factored analysis (JFA) (Kenny et al., 2007) and Eigenvoice priors (Castaldo et al., 2008) alleviate this issue by exploiting lower dimensional factorizations of the GMM supervectors to obtain more robust speaker-dependent supervectors which have lower intra-speaker variability. For example, in JFA, it is assumed that the supervector covariance matrix can be decomposed into a channel space and speaker space, where the channel space is responsible for the inter-speaker variations. This leads to the fact that each speaker and channel-dependent supervector can be decomposed into a sum of components that belongs to these two sub-spaces. As such, the two components can be disentangled and a more robust representation is obtained.

However, Dehak et al. (2010) found that the JFA channel factors also contain speaker-related information and instead propose to define just one total variability space which contains the channel and speaker factor of variations simultaneously. They then propose to find, for each speaker utterance, through Baum–Welch statistics, a projection for this total variability space which is representative of the speaker’s identity. This projection vector is called i-vector and was found to be a very effective speaker-id discriminative feature. To compensate for channel-dependent variability, Dehak et al. (2010) proposes to use linear discriminant analysis (LDA) and within-class covariance normalization (WCCN) (Hatch et al., 2006) where the i-vectors distances are defined with cosine similarity or with a support vector machine (SVM), respectively. In Matějka et al. (2011) it was shown that by using probabilistic LDA (PLDA) the performance can be improved further. This coupling of i-vectors and PLDA has been proven extremely effective and it has been a popular technique for speaker verification and diarization till the advent of deep learning based approaches.

Starting from 2014, the studies and refinements in the deep learning area, together with the increasing availability of annotated transcriptions and data, made it possible to exploit DNNs in place of GMMs for obtaining speaker-discriminative features. One of the first works in this direction is the one of Variani et al. (2014), in which DNN-based features (so-called d-vectors) were shown to be able to outperform i-vectors, the state-of-the-art approach of the time, especially in noisy conditions. As it happens with the GMM-UBM, in Variani et al. (2014) the DNN was trained on large corpora with a vast number of speakers and different acoustic conditions. The DNN is trained on fixed-length segments to try to classify the correct speaker (multi-class classification problem) among all the ones in the training set. In inference then the output of the last hidden layer is the d-vector and can be used as a speaker-discriminative feature for speaker verification or diarization. Follow-up works (Bai and Zhang, 2021) continued to improve this new paradigm based on DNNs used to extract speaker-identity discriminative features. In particular, a great improvement was due to the invention of the x-vector extractor (Snyder et al., 2018) which employs a time-delay neural network (TDNN) and a statistical pooling layer to obtain a low-dimensional speaker-id representation. This trend of designing better DNN architectures to improve speaker-id discriminative features is continuing today. Some recent improvements include the use of ResNet-based designs (Landini et al., 2021a), TitaNet (Koluguri et al., 2022), ECAPA-TDNN (Desplanques et al., 2020) and the use of self-supervised learning pre-trained models such as WavLM (Chen et al., 2022). Other works (Bredin, 2017; Li et al., 2018; Chung et al., 2020) instead have focused on the loss function and training strategy to use to train such a DNN speaker-id feature extractor. For example, Bredin (2017), Chung et al. (2020) proposed to use metric learning approaches, whereas (Li et al., 2018) the use of angular-softmax loss to improve performance. This is also a very active research direction. In all these works, clustering continued to be, as said, the main approach with DNN-based speaker-id representation used often in conjunction with PLDA to reduce dimensionality and intra-speaker variability before clustering, as done previously with i-vectors.

Indeed many advances also regarded other components of the diarization pipeline, such as the clustering step or the post-processing step.

For example, spectral clustering based on the Ng-Jordan-Weiss (NJW) algorithm (Ng et al., 2001), used in early works such as Ning et al. (2006), was then replaced by spectral clustering with maximum eigengap. Indeed, Wang et al. (2018) improved over the at the time state-of-the-art PLDA followed by AHC approach, by using this type of clustering procedure to aggregate DNN-extracted speaker embeddings, i.e., d-vectors. Afterward, Park et al. (2019) also added an auto-tuning strategy to overcome the need for a development set for the hyper-parameter tuning.

Another notable work, this time regarding post-processing, is variational Bayesian (VB) resegmentation, initially proposed for an i-vector-based system (Diez et al., 2018) and later adapted to an x-vector-based system (Landini et al., 2020). This latter approach called VBx has been proven extremely effective on a wide number of datasets (Landini et al., 2022) and challenges (Ryant et al., 2019). Among the post-processing works it is worth mentioning ensembling or fusion methods that allow combining the output of multiple heterogeneous diarization systems in order to improve the performance. Such works include diarization output voting error reduction (DOVER) (Stolcke and Yoshioka, 2019) and DOVER-Lap, a recent improvement over DOVER that allows handling also overlapped speech regions (Raj et al., 2021a; Raj and Khudanpur, 2021).

A number of recent works also explored how to improve clustering-based methods with deep learning based techniques in order to let them deal also with overlapped speech. For example, Bullock et al. (2020) proposed to use an overlap detector to mask the speaker posterior matrix in the VBx method. Raj et al. (2021b) instead devised a way to handle speaker-id discriminative features from overlapped speech regions during the clustering step.

Due to the greater availability of annotated data and possibilities opened up by deep learning recently a new line of research arose, involving other alternative approaches to diarization that try a tighter integration with DNNs. Among such works are region-proposal networks diarization (RPND) (Huang et al., 2020), unbounded interleaved-state recurrent neural network (UIS-RNN) (Zhang et al., 2019), discriminative neural clustering (DNC) (Li et al., 2021), deep learning SSGD-based approaches (Fang et al., 2021; Morrone et al., 2022) and target-speaker VAD (TS-VAD) (Medennikov et al., 2020). Some of these approaches and in particular SSGD and TS-VAD have proven to be particularly effective: for example, the system winning the recent third DIHARD Challenge (Ryant et al., 2021), was based on a combination of SSGD and TS-VAD based approaches (Wang et al., 2021b).

Indeed TS-VAD has been, in the recent three years, the de-facto best approach in terms of raw diarization performance; moreover, all the top ranking submissions in DIHARD-III (Ryant et al., 2021), VoxSRC-21/22 (Brown et al., 2022; Huh et al., 2023), CHiME-6 (Watanabe et al., 2020) and Alimeeting (Yu et al., 2022) employ such approach at least as part of a larger ensemble model (Wang et al., 2021a, 2022c,a,d). TS-VAD aims at producing speech probabilities for each speaker given a sequence of speaker-discriminative features e.g., in the original proposal (Medennikov et al., 2020), i-vectors of each speaker. As such, it can be considered more as a post-processing technique rather than a standalone diarization system since, to gather i-vectors or other speaker-id discriminative features, it is necessary first to perform diarization via e.g., a clustering-based approach. After this initial diarization (which can be rather rough) the TS-VAD can be run iteratively to refine the speaker-id discriminative features and predictions simultaneously. Despite the obvious advantages, it can be argued that most TS-VAD methods have also some drawbacks, especially when actual industrial applications are considered. To be fair, some disadvantages are being addressed at the time of this writing and this is a very active research direction. For example, a recent work (Wang and Li, 2022) expanded the TS-VAD approach to also incorporate the initial diarization step, obtaining thus a full-fledged standalone end-to-end diarization system (E2E TS-VAD), with impressive results on DIHARD-III. Again, an issue is that most proposed TS-VAD approaches are offline, as the requirement for an initial diarization to gather speaker identities across the whole recording is best done offline. A notable exception is Wang et al. (2022b) where a fully streamable TS-VAD method is proposed. Here the initial diarization is done in a chunk-wise fashion and the whole system becomes streamable, with competitive results with respect to a fully offline model (Wang et al., 2022d), especially for the two-speaker case. Another possible drawback is the fact that many of the TS-VAD approaches that perform best are not flexible regarding the number of output speakers. In fact, for example in the original TS-VAD work (Medennikov et al., 2020), the best results are achieved for the system that models jointly all four speakers (by concatenating the i-vectors of each speaker) and the single, per-speaker system performs rather poorly compared to this one. Again, recent works are addressing this issue, for example, in Cheng et al. (2022), a sequence-to-sequence TS-VAD approach is proposed, that can seamlessly generalize to arbitrarily output speakers and also different output resolutions.

UIS-RNN (Zhang et al., 2019) and DNC (Li et al., 2021) consist of supervised neural models based on speaker embeddings, where the first is clustering-free, whereas the other relies on neural-based clustering. They are both not overlap-aware, but they are able to manage a variable number of speakers.

Building on the invention of permutation invariant training (PIT) for DNN-based speech separation in 2019 (Kolbæk et al., 2017), Fujita et al. developed the first fully end-to-end DNN-based system (Fujita et al., 2019a) which was later improved using self-attention (Fujita et al., 2019b). This again sparked another research direction on systems based on end-to-end neural diarization (EEND) which is very active nowadays. Horiguchi et al. (2020) proposed to extend EEND with an encoder–decoder-based attractor architecture (EEND-EDA) able to handle a flexible number of output speakers thanks to an autoregressive decoder. It embeds a speaker counting functionality consisting of an LSTM-based encoder–decoder applied on the EEND's output to generate a variable number of attractors from a speech embedding sequence (which is the EEND's output). It also estimates the probability that the attractor is really associated with an active speaker. It generates attractors until a stopping criterion is met (i.e., when the existence probability associated with an attractor is smaller than a threshold), thus it can deal with an arbitrary number of speakers. Diarization results are obtained computing the dot product of the frame-wise speech embeddings with each speaker-wise attractor, followed by sigmoid activation. Speaker overlap is detected when multiple higher dot products between embeddings and attractors are present in the same frame (Horiguchi et al., 2022a). Kinoshita et al. proposed several improvements to the initial EEND approach by integrating speaker-id embeddings extraction so that the strengths of EEND and clustering-based approaches can be combined in a framework called EEND-vector clustering (EEND-VC) (Kinoshita et al., 2021b,a, 2022b,a). As EEND-EDA, EEND-VC is able to estimate the number of speakers, which can be arbitrary. After splitting the input into chunks, it applies end-to-end diarization on each of them and derives speaker embeddings. For each chunk the upper limit of the number of speakers is defined by the neural architecture, as for SA-EEND. Chunk size has to be sufficiently long to obtain good speaker embeddings, whereas it has to be also sufficiently short to contain a number of speakers that does not exceed the maximum number allowed by the architecture. Finally, the diarization output is obtained by applying a constrained-clustering algorithm between embeddings (Kinoshita et al., 2021b). Other recent works focused on streaming processing: (Xue et al., 2021) proposed to extend EEND with a speaker-tracing buffer to solve the permutation ambiguity caused by PIT when the model inference is performed via sliding windows. EEND-EDA streaming versions have been also recently proposed (Han et al., 2021; Horiguchi et al., 2022b). Also focusing on online processing, Coria et al. (2021), combined the use of EEND with an x-vector extractor and online clustering, where the EEND model is used to gate the representation before the x-vector statistical pooling layer, to extract per-speaker embeddings even in overlap regions. Recent results in the most popular diarization challenges indicate that EEND-based systems are increasingly competitive (Wang et al., 2021b; Horiguchi et al., 2021; Landini et al., 2021b) and nowadays surpass clustering-based methods.

Finally, other end-to-end approaches, that do not rely on PIT, such as DIVE (Zeghidour et al., 2021) and the joint speaker diarization/ASR system based on a recurrent neural network transducer (RNN-T) proposed in Shafey et al. (2019), have been also recently proposed.

1.2. Motivation and previous review works

Many works attempted to make a comprehensive “snapshot” of the current advances and previous, up to the time of their writing, algorithms proposed for speaker diarization. In 2006, [Tranter and Reynolds \(2006\)](#) wrote one of the very first reviews of diarization systems developed from the 1990s until the early 2000s, in the broadcast news and CTS domains. Following, in 2012, the work of [Anguera et al. \(2012\)](#) reviewed the new advancements, focusing mainly on meeting conversations, such as ones from AMI. [Park et al. \(2022\)](#) made recently a very comprehensive review of speaker diarization including more up-to-time methods, which are largely based on DNNs. They proposed a taxonomy of speaker diarization which divided the algorithms into various groups. These aforementioned works presented a comprehensive review of the speaker diarization field but did not analyze in depth the performance and computational aspects of different algorithms on a common dataset.

Here, on the other hand, we aim at providing quantitative guidance on the performance of various recent diarization systems in a common scenario. To the best of our knowledge, this work is the first experimental review on speaker diarization. As said, we limit our scope to the CTS domain because current techniques seem particularly promising in this commercially relevant domain.

We consider eight diarization systems which belong to clustering-based, end-to-end, and hybrid approaches and evaluate them on four multi-language datasets. In addition, we study the computational cost of different systems (i.e., processing time, memory footprint). This aspect was neglected by previous review works, but it is of paramount importance as computational requirements highly impact the financial and environmental ([Schwartz et al., 2020](#)) costs in production environments.

2. Diarization systems under study

Most diarization systems can be divided into two categories: clustering-based and end-to-end methods. There are, however, as explained, some notable exceptions (e.g., TS-VAD, SSGD, etc.), which are not fully end-to-end methods as involve usually a two-step procedure to obtain diarization (whilst clustering-based methods involve many different steps which are separately optimized) and make extensive use of deep learning based approaches.

In this paper, we perform a comparative study of algorithms that belong to clustering-based, end-to-end, and SSGD methods.

Regarding clustering-based approaches, we include a total of five different algorithms which come from the state-of-the-art such as VBx ([Landini et al., 2020](#)), x-vector extractor with spectral clustering ([Park et al., 2019](#)) and their overlapped-aware variants ([Bullock et al., 2020](#); [Raj et al., 2021b](#)). We also included a method based on AHC and BIC which is representative of early diarization approaches.

Regarding end-to-end methods, we considered self-attentive EEND (SA-EEND) ([Fujita et al., 2019b](#)) and EEND-vector clustering (EEND-VC) ([Kinoshita et al., 2021b,a](#)). Among the other methods that do not fall into EEND or clustering-based frameworks, we consider SSGD. The choice of SSGD ([Fang et al., 2021](#); [Morrone et al., 2022](#)) over TS-VAD ([Medennikov et al., 2020](#)) is considered as it is more straightforward to use (does not require initialization via clustering-based diarization) and has a greater commercial interest as the separated signals can be fed to an ASR back-end directly. For example, it has been shown to be particularly effective in the CTS domain recently ([Morrone et al., 2022](#)) with near state-of-the-art diarization and ASR results close to ones obtained with oracle speech separation.

We describe these different diarization paradigms in more detail thereafter.

2.1. Clustering-based methods

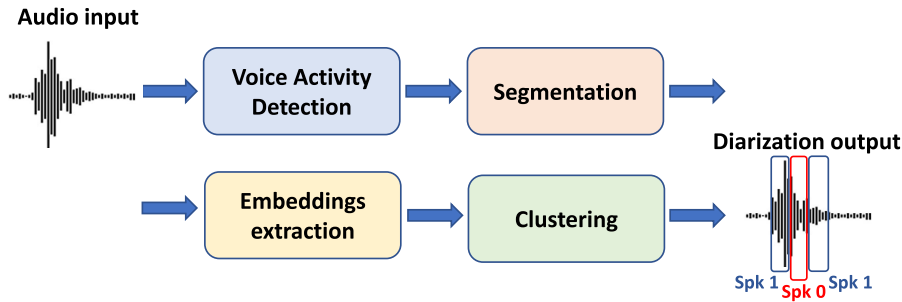
As explained in Section 1.1, clustering-based methods have been the mainstream approach for diarization for several decades and continue to be relevant and competitive today. Clustering-based methods typically consist of four modules arranged in a pipeline structure (see [Fig. 1\(a\)](#)): VAD, segmentation, speaker embedding extraction, and clustering.

The VAD module detects portions of the audio signal containing voice only, by discarding non-speech parts, such as silence, music, background noise, etc. The objective of the segmentation module is to divide the speech portions into homogeneous single-speaker segments. Early segmentation modules were based on the detection of speaker-change points, generating variable-length single-speaker segments. In recent years, uniform segmentation gained success. It divides the input recording into overlapped fixed-length speech segments, which should be long enough to catch representative speaker information, without containing more than one speaker. The extraction module’s goal is to capture the speaker’s characteristics within a speech segment and encapsulates them into a speaker embedding: a dense high-dimensional representation. The clustering step groups the embeddings that belong to the same speaker. Hereafter we describe the five clustering-based systems we considered in this experimental review.

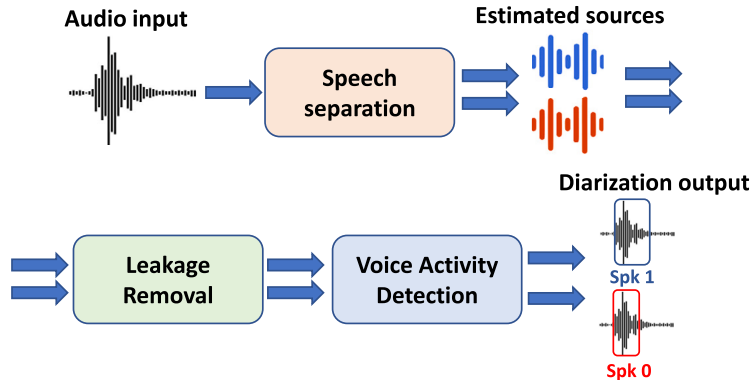
2.1.1. Variational Bayesian-hidden Markov model clustering of x-vector sequences (VBx)

As mentioned in Section 1.1, the VBx method proposed by [Landini et al. \(2020\)](#) combines variational Bayesian resegmentation ([Diez et al., 2019](#)) with x-vector clustering diarization. It has been proven a highly effective diarization technique and this is attested by its wide use in popular diarization challenges such as DIHARD.


VBx employs a standard clustering-based diarization pipeline based on x-vectors ([Snyder et al., 2018](#)) and PLDA+AHC clustering to derive initial diarization hypotheses. These are then refined via a Bayesian hidden Markov model (BHMM). In detail, the AHC step firstly derives hard assignments between speakers and the x-vectors sequence extracted from audio segments containing speech with a VAD. In VBx it is assumed that this sequence can be modeled by a hidden Markov model (HMM) with speaker-specific state distributions that are derived from the PLDA model. In detail, an ergodic HMM is used and each HMM state is associated



(a) General architecture of a clustering-based diarization system.



(b) Architecture of the SSGD system.



(c) Diarization system based on an end-to-end approach.

Fig. 1. General architectures of different categories of diarization systems.

with a single speaker, thus this approach cannot model overlapped speech. The hard assignment derived from AHC clustering is refined by finding the most probable sequence of HMM-associated latent variables that defines the hard alignment between the HMM states (and thus speakers) and x-vectors. This most probable sequence constitutes the diarization output of VBx and is found by firstly defining a hidden variable for the x-vector sequence distribution. This hidden variable is used to exploit the variational Bayes approach and the solution is found via evidence lower bound objective (ELBO) maximization.

2.1.2. Spectral clustering based on x-vectors (SCx)

The most used spectral clustering in early works on diarization was based on the NJW algorithm (Ng et al., 2001).

In NJW-based spectral clustering, the speakers' clusters are estimated after the construction of a similarity matrix, a multi-dimensional vector whose elements represent pairwise similarities between speaker embeddings of different speakers. This special matrix is transformed into spectral embeddings, which are then clustered through the k-means algorithm (Lloyd, 1982).

Its main drawback is represented by the limited accuracy since this technique is very sensitive to the noisiness of the similarity matrix due to intra-speaker variability of the speaker embeddings. In order to overcome this issue, a scaling parameter or a threshold should be introduced to assign different weights to pairwise similarities (Park et al., 2019). This approach was followed, for example, by Wang et al. (2018), which first exploited spectral clustering with maximum eigengap, in place of the popular PLDA+AHC pipeline, to aggregate and cluster together DNN-extracted speaker embeddings, i.e., d-vectors, and demonstrated it as state-of-the-art. The

optimal values of these hyper-parameters have to be found through a tuning on the development set. This, however, can increase the computational cost and makes generalization more challenging.

Shortly after, [Park et al. \(2019\)](#) propose to overcome this issue by deriving the optimal hyper-parameters through an auto-tuning strategy that avoids the need for a development set.

For simplicity, in this paper, since we consider this last approach applied to x-vectors, we will call it, for convenience here, SCx (in assonance to VBx).

2.1.3. *Overlap-aware versions of clustering-based diarization approaches based on x-vectors (VBx+OVL and SCx+OVL)*

Conventional clustering-based approaches are unable to model overlap-speech and thus incur in severe diarization accuracy degradation in conversational speech. Recent works by [Bullock et al. \(2020\)](#) and [Raj et al. \(2021b\)](#) proposed two different ways to make clustering approaches more “overlap friendly” by leveraging an overlap detection front-end. In detail, [Bullock et al. \(2020\)](#) makes use of the VBx pipeline and combines it with an overlap detector. The output of VB resegmentation is a sequence of probabilities of each x-vector belonging to each speaker. In this overlap-aware algorithm, the most likely speaker sequence is multiplied by a mask obtained from the output of VAD module, whereas the second most likely sequence is masked with the overlap detection output. In regions in which overlap has not been detected the output sequence is equal to the most likely speaker sequence. Instead, when overlap is detected, the output sequence is the union of the two most likely sequences.

The approach described in Raj et al. instead relies on the spectral clustering approach described previously in Section 2.1.2. Firstly, the authors propose to cast the discrete spectral clustering problem to continuous space via convex optimization. The continuous space then is discretized again, to obtain the speakers’ assignments, but by using the output of an overlap detector to relax cluster exclusivity constraints.

2.1.4. *Clustering-based diarization algorithm using agglomerative hierarchical clustering and the Bayesian information criterion (AHC-BIC)*

The BIC can be used to perform clustering in a diarization system as outlined in [Chen et al. \(1998\)](#), [Tritschler and Gopinath \(1999\)](#). The underlying assumption for BIC-based clustering is that the speaker-id discriminative features can be described by a multivariate Gaussian distribution. Early works such as [Chen et al. \(1998\)](#), [Tritschler and Gopinath \(1999\)](#) employed hand-crafted acoustic features such as MFCC which only approximately satisfy this assumption. Once the features are collected these are merged iteratively ([Tritschler and Gopinath, 1999](#)) in an AHC way by computing the BIC distance, ΔBIC , between each pair of clusters i and j as:

$$\Delta BIC = n \log |\Sigma| - n_i \log |\Sigma_i| - n_j \log |\Sigma_j| - \frac{\lambda}{2} \left\{ d + \frac{d}{2}(d+1) \right\} \log n, \quad (1)$$

where $|\Sigma|$ is the determinant of the covariance matrix estimated on acoustic data, $n = n_i + n_j$ (where n_i and n_j are, respectively, the number of segments belonging to cluster i and j), λ is a penalty factor that controls the aggressiveness of clustering (higher values lead to fewer clusters), and d is the size of the acoustic vectors (the number of acoustic spectral features).

This ΔBIC is a sort of penalized likelihood ratio expressing the difference between the assumptions that all the vectors are split into two Gaussian populations or belong all to the same Gaussian population, where the penalty is related to the different numbers of model parameters needed to represent the data in the two cases. The smaller this value is, the closer the two clusters are, meaning that they are good candidates for merging.

2.2. *Speech separation guided diarization (SSGD)*

In SSGD, diarization is performed through speech separation. As said, it was firstly proposed in [Siu et al. \(1992\)](#), but has recently gained more attention ([Fang et al., 2021](#); [Morrone et al., 2022](#)), as DNN-based methods for speech separation had become more and more effective. The classical SSGD pipeline is illustrated in [Fig. 1\(b\)](#). A speech separation algorithm is firstly used to obtain from the input conversation audio estimated audio streams for each speaker. These audio streams are then fed to VAD module which then estimates speech boundaries for each separated speaker, thus obtaining diarization annotation. As it happens with the VAD component in clustering-based systems, also here, usually post-processing, such as median filtering, is performed to smooth out the VAD predictions. In [Morrone et al. \(2022\)](#) an additional “leakage removal” post-processing module was proposed to reduce false alarms due to imperfect separation, occurring mainly on long single-speaker segments.

It can be argued that, in SSGD, diarization almost comes “for free”, as the most computationally expensive operation is speech separation and the VAD component can be very simple. A simple energy-based VAD can already achieve decent performance ([Morrone et al., 2022](#)). This makes it very appealing for use-cases that require downstream applications such as ASR or natural language understanding, for which we can use the available separated outputs.

SSGD, as SA-EEND, needs a-priori knowledge of the maximum number of speakers in a conversation; therefore the system will work as desired even in single speaker situations, but not when the number of speakers is greater than the one considered during the training phase. As SA-EEND, SSGD is able of operating in situations in which overlap is not present (overlap-free chunks).

2.3. End-to-end methods

Clustering-based methods implicitly assume the presence of single-speaker segments only. Therefore, they are not able to handle situations where speaker overlap is present. As explained, overlap-aware assignment techniques such as [Bullock et al. \(2020\)](#), [Raj et al. \(2021b\)](#), try to solve this problem but have intrinsic weaknesses, since the speaker embedding extractor module is still optimized without considering speaker overlap. Thus embeddings may not be very accurate on overlapped speech especially when overlap is significant.

In contrast, end-to-end approaches can natively manage speech segments where multiple speakers are active at the same time. This makes them extremely appealing especially on domains such as CTS, where, as explained, the amount of overlapped speech can be quite significant.

Their general structure is depicted in [Fig. 1\(c\)](#). These methods have three main advantages over the clustering-based approaches. First, they are conceptually simpler. Clustering-based algorithms need separate training for all the trainable models in the pipeline (i.e., VAD module, x-vector extraction module, PLDA, overlap detector). End-to-end systems require just one single model to be trained. Moreover, end-to-end systems can be trained to directly minimize diarization errors, in contrast to clustering which is usually performed in an unsupervised fashion. Finally, whereas clustering-based systems rely on embedding extractors such as x-vectors that are trained for speaker verification on single-speaker segments only, end-to-end can be trained with real conversations, and thus naturally learn how to handle overlap and speaker turns.

However, it can be argued that this last point could be also a weakness, as the cost of annotating real-world conversations, can constitute an obstacle to the adoption of end-to-end methods ([Landini et al., 2022](#)); an evaluation study carried out by [Xia et al. \(2022\)](#) quantifies this effort as approximately two hours of work to annotate 10 min of audio, for a human annotator. This cost can be reduced by using synthetically generated noisy mixtures as in [Fujita et al. \(2019b\)](#), [Kinoshita et al. \(2021a\)](#), however, care is needed during the training phase in order to avoid overfitting on a specific overlap ratio, and adaptation on a real dataset is then crucial to mitigate potential mismatches ([Fujita et al., 2019b](#); [Horiguchi et al., 2020](#); [Kinoshita et al., 2021a](#)).

In this work, we experiment with two end-to-end systems: SA-EEND and EEND-VC, both based on the PIT framework. We describe them thereafter.

2.3.1. Self-attentive end-to-end neural diarization (SA-EEND)

SA-EEND ([Fujita et al., 2019b](#)) tackles the speaker diarization task by formulating it as a multi-label classification problem. Diarization is performed within a single DNN, which is fed in with a sequence of acoustic features as input, and outputs the joint speech activities of multiple speakers i.e., VAD logits for each speaker.

SA-EEND is heavily inspired by DNN-based speech separation with PIT ([Kolbæk et al., 2017](#)). The only difference from PIT-based speech separation is the fact that the output in EEND is the frame-wise speech presence posteriors for each speaker. As such, the standard binary cross entropy (BCE) loss is used, and the permutation ambiguity is tackled by selecting the speakers' permutation that minimizes this loss with respect to the current estimate. As it happens with speech separation it is necessary for SA-EEND that the number of output speakers is equal to or higher than the maximum number of speakers expected in one conversation. Also, the use of PIT means that the inference of the model should be performed in one pass over the conversation and not, sequentially on sliding windows as there is no guarantee that permutation will be consistent on different windows. A technique similar to continuous speech separation (CSS) ([Chen et al., 2020](#)) can also be employed in practice, where outputs are re-ordered based on cosine-similarity with previous window predictions that overlap with current ones.

2.3.2. End-to-end neural diarization-vector clustering (EEND-VC)

This architecture ([Kinoshita et al., 2021b](#)) combines the EEND approach with clustering-based diarization principles in order to exploit the advantages of both methods. It represents an improvement of the SA-EEND, as it is able to handle an arbitrary number of speakers and achieves better performance especially when input recordings are long. As shown in [Fig. 2](#), it splits the audio input in chunks, and for each of them derives, as SA-EEND, local speakers' speech posteriors as well as their speaker embeddings associated with the current segment. In this way, the inter-chunk permutation problem of SA-EEND is avoided as the identity of speakers across a long conversation can be recovered via the clustering of speaker embeddings. This means that the model in inference can be run on chunks of several seconds, whose length is similar to the one employed in training, without the use of CSS-like stitching or a speaker buffer ([Xue et al., 2021](#)), which instead is necessary for SA-EEND. Another considerable advantage over SA-EEND is that EEND-VC only needs to know the maximum number of speakers within a single inference window which can be much smaller (e.g., 2, 3) than the total number of speakers in a conversation.

EEND-VC can also leverage constrained clustering algorithms for preventing speaker embeddings belonging to the same chunk from being associated with the same speaker. This is in practice implemented by setting a cannot-link constraint between these embeddings. For example, original EEND-VC work ([Kinoshita et al., 2021b](#)) used the COP-k-means ([Wagstaff et al., 2001](#)) algorithm. In a follow-up work ([Kinoshita et al., 2021a](#)) they demonstrated that the constrained AHC outperforms COP-k-means and other clustering algorithms, especially when the number of speakers is high. When in one input chunk the number of speakers is less than the preset number of local speakers, it is necessary to identify whose speaker embeddings belong to silence. This is achieved in inference by looking at the local speakers' speech posteriors and seeing if their mean is below a certain "activity threshold". If this applies the corresponding embedding is considered as belonging to silence. During the learning phase, the network tries to minimize a loss function that consists of the weighted sum of two losses: the same PIT BCE loss as used in SA-EEND, and a speaker embedding loss. The former is formulated in the same way as the SA-EEND loss function, but for local speakers, the latter is inspired by [Zeghidour and Grangier \(2021\)](#) and is designed to drive the speaker embeddings to have high inter-speaker distance and low intra-speaker variability.

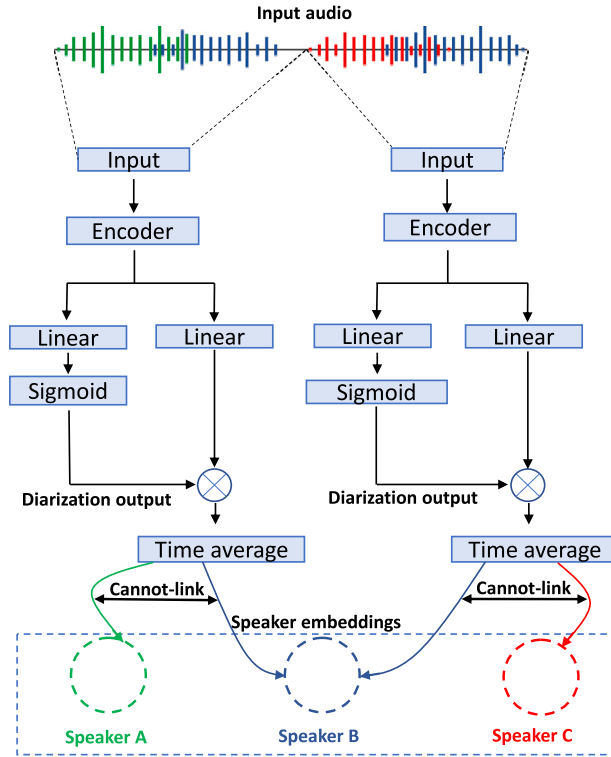


Fig. 2. Architecture of the EEND-VC.

Table 1

Main characteristics of the datasets. This table reports the average total audio time, active speech time, speech sparsity, overlap ratio (computed with respect to the active speech time duration), and DNSMOS of each dataset's test set.

Metric	Unit	CALLHOME	Fisher	CallCntrITA	CallCntrPOR
Total audio time	h	2.97	10.11	1.70	1.90
Active speech time	h	2.67	9.74	1.22	1.07
Speech sparsity	%	10.10	3.66	28.24	43.68
Audio file duration	s	72.14	596.37	255.20	325.20
Overlap ratio	%	13.08	14.59	5.14	5.35
DNSMOS	–	2.91	2.96	2.52	2.32

3. Experimental setup

3.1. Datasets

In this experimental review, we consider four different two-speaker CTS datasets, all with 8 kHz sample rate. Two of these are well-known, widely used datasets: CALLHOME (Przybocki and Alvin, 2001), Fisher (Cieri et al., 2004). Another two datasets, i.e., CallCntrITA and CallCntrPOR, are novel and feature Italian and Portuguese languages, respectively. The datasets have very different characteristics which are described in the following subsections. Statistics are shown in Table 1. We report the average total audio time, active speech time, speech sparsity (computed as the amount of time where speech is not present divided by the total audio amount), overlap ratio, and audio quality using DNSMOS (Reddy et al., 2021) which is a blind mean opinion score (MOS) estimation method.

3.1.1. CALLHOME

CALLHOME (Przybocki and Alvin, 2001) is one of the most commonly used data corpora containing telephone conversations for speaker diarization evaluation. It contains spontaneous conversations, distributed across six different languages: Arabic, English, German, Japanese, Mandarin, and Spanish. Each conversation can have from two to a maximum of seven different speakers.

In this work, we use a two-speaker sub-portion of CALLHOME. In particular, we use the subset of the 2000 NIST Speaker Recognition Evaluation (SRE) (Przybocki and Alvin, 2001) contained in the R65_8_1 folder, as widely done in the literature, for example in EEND (Fujita et al., 2019a,b). To allow for direct comparison, following such previous works, we select two-speaker

conversations, and perform a split, following the Kaldi CALLHOME diarization v2 recipe,¹ in order to obtain a validation set and a test set of 155 and 148 recordings, respectively. As reported in Table 1, the CALLHOME test set is characterized by a relatively high overlap ratio, with respect to the other datasets in this study.

3.1.2. Fisher Corpus Part I

Fisher Corpus Part I (Cieri et al., 2004) is an English-based dataset, containing telephone conversations between strangers on assigned topics. As such, dialogues are more formal than CALLHOME's ones. Fisher provides separate streams for each of the two speakers; as separate streams are not available in many applications and also would make the diarization problem trivial, we sum these streams obtaining a mixture signal. In our experiments, we will perform diarization on these two-speaker mixture signals. We split it in training (5 000 files), validation (61 files), and test (61 files) partitions. As reported in Table 1, the test set of this corpora is the one with the largest overlap ratio and the smallest speech sparsity. Therefore, its conversations are characterized by relatively few moments of silence and speakers often overlap.

3.1.3. CallCntrITA

CallCntrITA is a proprietary data corpus that consists of real-world telephone conversations in a call-center scenario between a customer and the operator, in the Italian language. Audio files are recorded by a call center company that offers customer care oriented services. The customer and operator are recorded in two different audio streams thus, as in Fisher, oracle speech sources are available. Again, as with Fisher, we sum them to obtain a mixture signal to be used as input to the diarization systems. CallCntrITA is composed of 48 conversations which we randomly split in two equal validation and test partitions (24 each then). As shown in Table 1, the associated test set presents the lowest overlap ratio among the datasets considered. Compared to CALLHOME and Fisher the DNSMOS values are lower because of the presence of babble noise on the operator side. CALLHOME and Fisher are instead relatively clean in most conversations as the high DNSMOS overall value suggests.

3.1.4. CallCntrPOR

It is the Portuguese counterpart of CallCntrITA, thus it is recorded in the same manner. It is composed of 42 files and as for CallCntrITA, we randomly split it in two equal 21 validation and test parts. As shown in Table 1, the associated test set is denoted by a speech sparsity equal to 43%. thus, on average, almost half of the duration of each of its files consists of non-speech. It also presents one of the lowest overlap ratios. It is the dataset with the lowest DNSMOS predicted audio quality: as CallCntrITA in the operator side there is babble noise but here it is more severe.

3.2. Systems implementation details

Thereafter we report implementation details and hyper-parameters for the methods under study described in Section 2: e.g., the SSGD's separation and speaker embedding extraction models employed in our experiments. Table 2 summarizes the major aspects regarding the implementation of all diarization systems considered in this manuscript (e.g., feature front-end, the ability to handle speakers overlap, training data) to highlight similarities and differences between them. More details are reported in each specific subsection, as follows.

3.2.1. Variational Bayesian-hidden Markov model clustering of x-vector sequences (VBx)

For the VBx algorithm, we use an open source implementation² which was also used in Morrone et al. (2022). This implementation is derived directly from the official one made available by the VBx authors (Landini et al., 2020). For speech segmentation, we use the publicly available Kaldi ASPIRE VAD model (Peddinti et al., 2015)³, which deals with 40 MFCC as input. Since these are computed without dimensionality reduction, they contain the same information as filterbank features. As the x-vector extractor model we use a pre-trained 8 kHz ResNet101 model⁴ (He et al., 2016), and x-vectors are extracted every 0.24 s, by considering a frame size equal to 1.44 s.

The ResNet101 x-vector extraction module is depicted in Fig. 3. It consists of a neural network that takes 64 MFCC without dimensionality reduction for T frames as input. Mel filters span the range 20–3700 Hz, and features are extracted every 10 ms by considering a window size equal to 25 ms. This model is trained for a speaker recognition task as the original x-vector architecture (Snyder et al., 2018) on a large amount of data consisting of *VoxCeleb1* (323 h, 1 211 speakers) (Nagrani et al., 2017), *VoxCeleb2* (2 290 h, 5 994 speakers) (Chung et al., 2018), *CN-CELEB* (264 h, 973 speakers) (Fan et al., 2020), *NIST-SRE 2004–2010* (3805 h, 4 254 speakers), *Switchboard (SWBD)* (1 170 h, 2 591 speakers), *DeepMine* (688 h, 1 858 speakers) (Zeinali et al., 2018). Data augmentation consisting of artificial reverberation and addition of music and noise from MUSAN (Snyder et al., 2015) is also employed. PLDA was trained on the same data.

This same pipeline consisting of the Kaldi ASPIRE VAD and the ResNet101 x-vector extractor, with their respective hyper-parameters (e.g., x-vector extractor hop and window) is also used in the other clustering-based techniques except for the AHC-BIC.

¹ https://github.com/kaldi-asr/kaldi/blob/master/egs/callhome_diarization/v2/run.sh

² D. Raj, diarizer, <https://github.com/desh2608/diarizer/tree/is2022>.

³ <https://kaldi-asr.org/models/m4>

⁴ https://github.com/desh2608/diarizer/tree/is2022/diarizer/models/ResNet101_8kHz




Fig. 3. Architecture of the x-vector extraction model used in the latest VBx system. The output size of each block is reported inside or near the respective box. T indicates the number of input frames.

Table 2

This table summarizes the major aspects regarding the implementation of all diarization systems considered in this work. Dagger, double dagger, and star (\dagger , \ddagger , \star) are used for clustering-based systems based on x-vectors to associate, respectively, the pair PLDA and x-vector extractor, the VAD, and the overlap (OVL) detector with datasets used to train them.

System	Modules					Overlap handling	Trainable modules	Training data
	Feature Front-End	VAD	Spk Embed. extractor	Clustering algorithm	Post-Processing			
VBx	MFCC	Kaldi ASpIRE ^a	ResNet101 ^b	PLDA+ AHC init.	VB-HMM	\times	VAD [‡] , PLDA [†] , x-vector extractor [†]	VoxCeleb 1&2 [†] , CN-CELEB [†] , NIST-SRE 2004–2010 [†] , SWBD [†] , DeepMine [‡] , Fisher [‡]
SCx	MFCC	Kaldi ASpIRE ^a	ResNet101 ^b	Spectral Clustering	–	\times	VAD [‡] , x-vector extractor [†]	VoxCeleb 1&2 [†] , CN-CELEB [†] , NIST-SRE 2004–2010 [†] , SWBD [†] , DeepMine [‡] , Fisher [‡]
VBx+OVL	MFCC	Kaldi ASpIRE ^a	ResNet101 ^b	PLDA+ AHC init.	VB-HMM	\checkmark	VAD [‡] , PLDA [†] , x-vector extractor [†] , OVL Detector ^{*,c}	VoxCeleb 1&2 [†] , CN-CELEB [†] , NIST-SRE 2004–2010 [†] , SWBD [†] , DeepMine [‡] , Fisher [‡] , AMI [*]
SCx+OVL	MFCC	Kaldi ASpIRE ^a	ResNet101 ^b	Spectral Clustering	–	\checkmark	VAD [‡] , x-vector extractor [†] , OVL Detector ^{*,c}	VoxCeleb 1&2 [†] , CN-CELEB [†] , NIST-SRE 2004–2010 [†] , SWBD [†] , DeepMine [‡] , Fisher [‡] , AMI [*]
AHC-BIC	MFCC+ log-energy	Proprietary energy-based	–	AHC	–	\times	–	–
SSGD	Learned Filterbanks, LFBE	TCN-based	–	–	Leakage Removal Algorithm	\checkmark	DPRNN Speech Separator, VAD	Fisher Corpus Part I
SA-EEND	LFBE	Intrinsic	–	–	Median filter, output threshold	\checkmark	Whole model	SWBD-2 (II,III), NIST-SRE, SWBD Cellular 1&2
EEND-VC	LFBE	Intrinsic	–	Constrained AHC	Median filter, output threshold	\checkmark	Whole model	SWBD-2 (I,II,III), NIST-SRE, SWBD Cellular 1&2

^a<https://kaldi-asr.org/models/m4>

^bhttps://github.com/desh2608/diarizer/tree/is2022/diarizer/models/ResNet101_8kHx

^chttps://github.com/desh2608/diarizer/blob/is2022/diarizer/models/pyannote/callhome_epoch4_step974.ckpt

The VBx hyper-parameters are tuned, for each dataset on the dataset development set. In detail, we tune VBx VB-HMM hyper-parameters: F_A , F_B , and P_{loop} . The VBx implementation considered here uses PLDA+AHC for the initialization of the VB-HMM posterior matrix. The PLDA model is trained with the same data as the x-vector extraction module.

3.2.2. Spectral clustering based on x-vectors (SCx)

As previously reported in Section 2, we also consider another clustering-based pipeline based on spectral clustering with maximum eigengap, using an auto-tuning strategy as proposed in Park et al. (2019)^{3.2.1}. Contrary to VBx, the resegmentation step

here is not performed. We employ the same VAD model and ResNet-based x-vector extractor as the VBx one described in the previous section.

3.2.3. *Overlap-aware versions of clustering-based diarization approaches based on x-vectors (VBx+OVL and SCx+OVL)*

As explained in Section 2.1.3, the overlap-aware counterparts^{3,2.1} of VBx and SCx are able to handle overlapped speech regions by leveraging an overlapped speech detector. Here we use a pre-trained Pyannote segmentation model⁵ (Bredin and Laurent, 2021; Bredin et al., 2020). This model is trained on AMI dataset and, following Morrone et al. (2022), fine-tuned on the CALLHOME development set. It uses 40 MFCC without dimensionality reduction, and the hyper-parameters for this segmentation model are tuned on each dataset corresponding validation set. These include onset, offset, min_duration_on and min_duration_off. Regarding the VAD module, the x-vector extractor, and the clustering module, we use the same implementations of VBx and SCx baseline versions.

3.2.4. *Clustering-based diarization algorithm using agglomerative hierarchical clustering and the Bayesian information criterion (AHC-BIC)*

We used a proprietary diarization algorithm developed for call-center applications. It is highly optimized to run on resource-constrained devices.

As speaker-id discriminative features, it uses 12 MFCC plus log-energy, extracted every 10 ms by using a Hamming window 20 ms long. A proprietary energy-based VAD based on Junqua et al. (1994) is used to detect speech regions and reject silence. The VAD has been tuned to minimize false alarms by adjusting the threshold and putting a constraint on the minimum duration of segments labeled as speech. If this minimum duration is not met, the segment is labeled as non-speech.

3.2.5. *Speech separation guided diarization (SSGD)*

In this study, we use the SSGD implementation from Morrone et al. (2022). We use as the speech separator the dual-path recurrent neural network (DPRNN) (Luo et al., 2020) and, as the VAD component, a temporal convolutional network (TCN) based one presented in Cornell et al. (2022), Morrone et al. (2022). The former uses learned filterbanks (Luo and Mesgarani, 2019) as input, whereas the latter exploits 40 logarithmic Mel-scaled filterbank energies features (LFBE); both models are trained on the Fisher Corpus Part I dataset using the same procedure and hyper-parameters as described in Morrone et al. (2022) which also employed synthetically mixed, fully overlapped speech examples from Fisher to pre-train the speech separation model. We consider a non-causal DPRNN and include in the pipeline also the leakage removal algorithm proposed in Morrone et al. (2022). This is also the configuration that in Morrone et al. (2022) gives the best overall performance. Following this latter work and SA-EEND, a median filter is also employed to smooth the TCN VAD predictions. The VAD threshold value and median filter length are tuned for each dataset on their respective development set.

3.2.6. *Self-attentive end-to-end neural diarization (SA-EEND)*

In this work, we used an open-source SA-EEND Pytorch (Paszke et al., 2019) implementation.⁶ We used the same features and model parameters as Fujita et al. (2019b) except that we employed 4 transformer encoder layers instead of 2. The dataset used to train this model consists of a combination of SWBD-2 (Phase II, III), SWBD Cellular (Part 1, Part 2), and NIST Speaker Recognition Evaluation (NIST-SRE) datasets. The overall amount of data consists of 26 172 mixtures (2 215 h with an overlap ratio equal to 3.7%). The training was performed as in Fujita et al. (2020) with Adam optimizer (Kingma and Ba, 2015) using the learning rate scheduler presented in Vaswani et al. (2017), where the decay phase is preceded by 25 000 warm-up steps. The batch size was set to 64 and the length of each example fed to the model had a length of 50 s. The final model is achieved by averaging the model parameters obtained during the last 10 epochs and by fine-tuning on the CALLHOME adaptation set. As in Fujita et al. (2019b) median filtering is used to smooth the output predictions. Inference is performed not on chunks but by using all the conversation, as this model is trained with PIT. For each dataset, in this study, we tune the logits threshold and the median filter length using their development set.

3.2.7. *End-to-end neural diarization-vector clustering (EEND-VC)*

We employ the implementation⁷ from Kinoshita et al. (2021a). In order to carry out the experiments, we use a pre-trained model from Kinoshita et al. (2021a). This model is comprised of 6 transformer encoder blocks and uses 8 attention heads in each multi-head self-attention layer. It uses the same input features (23 LFBE) as SA-EEND, and uses Constrained AHC as a clustering algorithm. It was trained on simulated noisy three-speaker mixtures generated from: SWBD-2 (Phase I & II & III), SWBD Cellular (Part 1 & 2), and the NIST-SRE datasets; then, it was adapted on CALLHOME. These synthetic mixtures are created following the procedure considered in Fujita et al. (2019a) with $\beta = 10$. As the mixtures contain three speakers, the model is configured to have a maximum of three output local speakers. As it can be noted, EEND-VC pre-trained model is trained with mixtures generated by using the same datasets considered for the pre-trained SA-EEND, plus SWBD-2 Phase I; therefore their training data differ. In Section 4.1.1 we compare them when they are trained on the same dataset.

⁵ https://github.com/desh2608/diarizer/blob/is2022/diarizer/models/pyannote/callhome_epoch4_step974.ckpt

⁶ Xflick, EEND PyTorch, https://github.com/Xflick/EEND_PyTorch.

⁷ nttcslab-sp, EEND-vector-clustering, <https://github.com/nttcslab-sp/EEND-vector-clustering>.

Training is performed with the same batch size and learning rate scheduler employed by SA-EEND, and Adam optimizer. During training, the model is fed with chunks of length 15 s, instead of 50 s-chunks used in the case of SA-EEND. In inference, we use chunks of 30 s.

As done for SA-EEND, also for EEND-VC we tune the output threshold and median filter length for each dataset. The “activity threshold” to detect “silent speakers” (Kinoshita et al., 2021b) is set equal to 0.05 for all datasets and is not tuned.

3.3. Evaluation metrics

This subsection presents the evaluation metrics we used to assess the performance of the diarization systems under study and their computational costs.

3.3.1. Diarization accuracy

We measure the speaker diarization systems’ accuracy by using the Diarization Error Rate (DER) which is the most common metric used to evaluate the difference between the output of a diarization system and a reference oracle segmentation. It quantifies the fraction of the time not correctly attributed to a speaker or non-speech and is calculated as the ratio:

$$DER = \frac{MS + FA + SC}{TOT} \quad (2)$$

where the denominator TOT is the total speech time, obtained by summing the speech time of each speaker. The numerator is the sum of the following quantities:

- *Missed Speech (MS)*: it is represented by the amount of scored time that a system wrongly indicates as non-speech.
- *False Alarm (FA)*: as opposed to MS, it represents the portion of scored time that a system hypothesizes as speech, while it is labeled as non-speech in the reference.
- *Speaker Confusion (SC)*: it is the portion of scored time assigned to the wrong speaker. It does not take into account the portion of overlapped speech not detected. Contrary to the other two errors, this is not related to the VAD stage, but it is strictly a diarization error due to segmentation and clustering steps.

In our experiments, we consider a “fair” evaluation setup, which takes into account all the speech (even segments containing overlap) and involves the introduction of a 0.25 s “no score collar” placed at each segment’s extremity. This collar-based evaluation is a common practice and allows to mitigate annotation inconsistencies which are inevitable. Precise utterance boundaries are difficult to define unambiguously: e.g., oracle segmentation can be “tighter” or “looser” depending on the dataset and the instructions given to annotators. The DER computation is performed using the NIST md-eval scoring tool⁸ version 22.

3.3.2. Computational cost

We assess the computational cost of each diarization system under study, by performing inference on the test set, for each dataset, and on a single-file basis. This mimics an actual deployment scenario where diarization is usually performed on each conversation separately so the output is readily available.

In detail, inference was performed on the same machine on an Intel i9-10920X CPU @ 3.50 GHz. In this work, we focus on CPU inference as it is the most common deployment platform for a commercial application that has to run on the customer side (it may be a powerful personal computer or an edge device) for privacy reasons. In such instances, access to a hardware accelerator such as a GPU is not guaranteed. We performed inference by ensuring that no other users were using the machine and no other processes were running. Moreover, in order to carry out the experiments under a standard condition, simulating client-side deployment, we limited the number of cores to eight units. All the diarization systems evaluated in this paper are implemented in Python, in particular by exploiting PyTorch, except AHC-BIC which is developed in C++. Anyway, it is important to mention that Pytorch also relies on calls to high-performance routines developed in C++, C or Fortran. For diarization systems with a pipeline structure, we calculated the inference time of each module net of the model’s loading time; then, we summed up the intermediate inference times together in order to compute the final result associated with the whole pipeline. For each algorithm and dataset, we calculated the average inference time, the Real-Time Factor (RTF), defined as the inference time divided by the audio file duration, and the maximum memory (RAM) occupation, during the diarization of each audio file. To get a reliable estimate for these figures, inference is repeated ten times for each audio file and in the following Section 4 we will report both the average value and the standard deviation.

4. Results and discussion

4.1. Performance evaluation

Table 3 reports diarization performance in terms of DER, obtained by performing inference with all the diarization algorithms under study on the test set of each dataset. To better understand which algorithm performs better across the four different scenarios

⁸ N. Ryant, dscore, <https://github.com/nryant/dscore>.

Table 3

DER obtained by the different diarization algorithms on each test set. For each diarization system, the table also reports the macro-average DER, which consists of the arithmetic mean (reported along with the associated standard deviation in the last row) of the DERs obtained on the different datasets. The best value in each row is highlighted in bold.

Dataset	DER							
	VBx	SCx	VBx+OVL	SCx+OVL	AHC-BIC	SSGD	SA-EEND	EEND-VC
CALLHOME	11.74	14.69	10.13	14.24	22.92	9.69	10.98	8.14
Fisher	9.75	9.15	6.91	7.01	26.13	7.14	11.04	7.60
CallCntrITA	10.53	8.87	10.47	8.79	12.86	12.19	7.85	7.65
CallCntrPOR	14.26	16.13	14.18	16.12	26.08	25.15	21.94	14.25
	macro-average DER							
	11.57	12.21	10.42	11.54	22.00	13.54	12.95	9.41
	± 1.71	± 3.24	± 2.57	± 3.75	± 5.43	± 6.94	± 5.35	± 2.80

we also computed the macro-average DER obtained by averaging the DER values across the four datasets. For this latter, we report also the standard deviation as it could outline which system is more consistent across all scenarios and thus may be preferable.

We can see that, overall, the best performing algorithm is EEND-VC, with VBx+OVL coming second. AHC-BIC achieves lower accuracy than the other approaches. Its best performance is on CallCntrITA where it is rather close to SSGD. Overlap-aware clustering is able to considerably boost performance for VBx and SCx. This is evident on the datasets with high speaker overlap, i.e., Fisher and CALLHOME: VBx+OVL is the algorithm that performs best on Fisher among all. As already reported in Table 2, the diarization systems based on DNN-based speaker embeddings employ the same x-vector extractor which is also trained with the same data. Therefore the difference in their performance is not related to the speaker embedding extraction stage.

SSGD has the least consistent performance among all algorithms under study and its performance degrades considerably on CallCntrITA and CallCntrPOR. This degradation may be due to the fact that, as it relies on speech separation, it is very sensitive to training/inference domain mismatch. Since it has been trained on Fisher, it thus performs well on CALLHOME as it has similar characteristics (e.g., speech sparsity, overlap ratio, high DNSMOS). Instead, as explained in Section 3.1, CallCntrITA and CallCntrPOR have slightly different acoustical conditions and speech statistics and thus the separation component is more prone to fail, degrading the final diarization accuracy.

SA-EEND performs slightly better than SSGD overall. The difference is very pronounced in CallCntrITA where SA-EEND seems better able to deal with sparse speech than SSGD. Compared to the other methods, it does not perform well however on long average length datasets such as Fisher, because of the mismatch between training sequence and inference sequence lengths. In fact, it appears that SA-EEND methods are more sensitive to this mismatch than separation-based ones, which instead suffer from sparsity but can generalize better to long-lasting sequences in inference (Morrone et al., 2022). EEND-VC is less prone to these problems, as it performs inference on fixed-size windows and it instead relies on clustering of locally extracted embeddings to solve the permutation ambiguity. As with the other models, the DER for SA-EEND increases significantly on CallCntrPOR, both due to challenging acoustic conditions, high average length, and sparsity.

VBx is overall the algorithm that has the most consistent performance (i.e., lowest standard deviation) across the scenarios owing to its simplicity and strong prior assumptions of its resegmentation process.

For a more complete analysis, in Table 4 we report all the DER components: missed speech (MS), false alarms (FA), and speaker confusions (SC) for each dataset and each system under study. As expected, since they rely on the same VAD model, VBx and SCx have always the same MS and FA and differ only in terms of SC, as the clustering step is different and SCx lacks the VB-resegmentation step. We can also see that their overlap-aware modifications (i.e., VBx+OVL and SCx+OVL) improve the DER by reducing in particular MS and SC on CALLHOME and Fisher (the datasets with considerable overlap-ratio) while they slightly increase FA. The AHC-BIC pipeline always detects zero FA due to the fact that its energy-based VAD implements a strict constraint on the minimal utterance duration as said in Section 3.2.4. Most of the errors come from SC as it employs hand-crafted speaker-id discriminative features which are likely not robust enough to intra-speaker variability.

SSGD obtains the worst (compared to other datasets) FA and SC on both CallCntr datasets due to the speech sparsity and the presence of background babble noise both of which are more pronounced in CallCntrPOR. The noise degrades the speech separation algorithm output which then causes the VAD to give false alarms.

Regarding EEND methods, SA-EEND always generates more SC than EEND-VC. This latter can rely on clustering for speaker tracking whereas the former cannot. In fact, the SC difference is more noticeable on the datasets with high average conversations length.


Fig. 4 reports the Pearson correlation coefficient (PCC) between DER and speech sparsity (Fig. 4(a)) and between DER and overlap ratio (Fig. 4(b)) for each dataset and for each diarization method under study. The PCC was obtained by considering each recording in each dataset a different data-point: we computed the DER separately for each recording and used the speech statistics for each recording to compute the correlation.

Regarding correlation with speech sparsity, we can see that SA-EEND is the system that overall correlates the most. This is expected as it can struggle to correctly track speakers (due to PIT) when the conversation is sparse. However, also EEND-VC exhibits some correlation, which is rather unexpected. This can be explained by the fact that EEND-VC is trained on a dataset with much less speakers' diversity than VBx and SCx (which use VoxCeleb 1 and 2). This means that it suffers more when the speech is very sparse as


Table 4

DER numerator components: missed speech (MS), false alarms (FA) and speaker confusion (SC) for each diarization system and each dataset under study. On each column we highlight in bold the best and second best values.

Method	CALLHOME			Fisher			CallCtrITA			CallCtrPOR		
	MS	FA	SC	MS	FA	SC	MS	FA	SC	MS	FA	SC
VBx	8.28	0.87	2.59	8.47	0.46	0.82	2.88	5.19	2.46	4.41	7.78	2.07
SCx	8.28	0.87	5.54	8.47	0.46	0.22	2.88	5.19	0.80	4.41	7.78	3.93
VBx+OVL	5.33	2.36	2.44	3.95	2.19	0.77	2.58	5.44	2.46	4.33	7.78	2.07
SCx+OVL	5.69	2.52	6.04	4.80	2.02	0.19	2.05	5.88	0.86	4.32	7.78	4.02
AHC-BIC	6.39	0.00	16.53	8.20	0.00	17.94	2.81	0.00	10.05	2.16	0.00	23.92
SSGD	5.92	2.84	0.94	4.86	0.72	1.56	0.41	9.99	1.79	2.34	15.62	7.20
SA-EEND	4.66	4.40	1.92	7.32	1.66	2.05	2.15	4.98	0.72	5.94	12.57	3.43
EEND-VC	3.79	3.32	1.03	6.72	0.62	0.26	2.36	4.95	0.34	6.99	6.29	0.97



(a) Pearson correlation coefficient (PCC) between DER and speech sparsity.



(b) Pearson correlation coefficient (PCC) between DER and overlap ratio.

Fig. 4. Pearson correlation coefficient (PCC) between DER and speech sparsity (a) and overlap ratio (b) for each dataset and diarization method.

there are fewer segments to cluster and thus more intra-speaker variability due to less speaker diversity in training. Moreover, when the speech is sparse EEND-VC may not be able to exploit the cannot-link mechanism in the clustering step as only one speaker can appear in a particular signal window. Also, SSGD seems to suffer from sparsity on all datasets but CallCtrPOR where it exhibits the poorest correlation among all systems indicating that the performance degradation on CallCtrPOR is due to the more challenging acoustical conditions and more pronounced babble noise compared to CallCtrITA.

Table 5

Diarization results obtained on different test sets, by training both EEND on Fisher Corpus Part I. In (a) the number of heads (H) and layers (L) is equal to the default ones considered in Table 3, whereas results in (b) are computed by swapping the (H, L) pairs of the two EEND. For each subtable, the best result on each test set is reported in **bold**, whereas the best values across subtables are underlined.

(a) Diarization results obtained on different test sets, by training both EEND on Fisher Corpus Part I; the number of heads (H) and layers (L) is equal to the one considered in Table 3.			(b) Diarization results obtained on different test sets, by training both EEND on Fisher Corpus Part I; the number of heads (H) and layers (L) for SA-EEND is the ones used for EEND-VC in Table 3, and vice versa for EEND-VC.		
Dataset	SA-EEND (H=4, L=4)	EEND-VC (H=8, L=6)	Dataset	SA-EEND (H=8, L=6)	EEND-VC (H=4, L=4)
CALLHOME	12.15	<u>11.95</u>	CALLHOME	14.96	12.54
Fisher	10.80	5.84	Fisher	13.81	5.60
CallCntrITA	11.73	<u>11.12</u>	CallCntrITA	11.42	12.47
CallCntrPOR	19.83	20.97	CallCntrPOR	29.70	<u>19.79</u>

For what concerns the overlap ratio we can observe that, in general, the systems that suffer the most are the clustering-based ones without any overlap-aware mechanism (AHC-BIC, VBx, and SCx). As expected, their overlap-aware counterparts lower the correlation value. Overall, EEND and SSGD diarization systems are the ones that exhibit the lowest correlation across the datasets as they can process overlapped speech more effectively. On CallCntrPOR it seems that the performance, for most systems except AHC-BIC and EEND-VC, is largely unaffected by the overlap ratio.

4.1.1. SA-EEND versus EEND-VC: Further comparison

In addition to the analysis done before, we compare SA-EEND and EEND-VC further in a more fair setup: when both methods are trained from scratch on the same dataset since, until this point, training data considered to obtain the EEND-VC pre-trained model differ from the ones used for SA-EEND pre-trained model e.g., for SA-EEND, SWBD-2 Phase I has not been taken into account. It makes sense to further compare the two EEND methods as their performance is highly dependent on the availability of large amounts of annotated conversations. Conversely, x-vector-based methods largely depend on the availability of datasets of pre-segmented, noisy utterances coming from a large number of different speakers to train the speaker-id feature extractor. Data for these latter do not need strict per-speaker segmentation annotation and are thus arguably easier to obtain in a more cost-effective way.

We train these two models, with the same hyper-parameters as explained in Section 3.2.6 on Fisher Corpus Part I using the same split as described in Section 3.1.2. We also change the maximum number of local speakers (and thus the output layer size) from three to two, because in this work we consider only two-speaker scenarios, whereas the pre-trained model from Kinoshita et al. (2021a) was trained by assuming a maximum of three local speakers.

Both models are trained for 300 epochs, and we report results obtained by performing inference with the model checkpoint that achieved the lowest validation loss during the training phase. As done for the two pre-trained models, we tune the threshold and, for post-processing smoothing, the median filtering length on the validation set of each of the four datasets separately.

All results are reported in Table 5a. For EEND-VC we suppose that the total number of speakers is known a-priori and we exploit this information during the clustering phase. As expected, EEND-VC outperforms SA-EEND on most datasets, whereas SA-EEND slightly outperforms EEND-VC on CallCntrPOR. EEND-VC is significantly better on Fisher, again, as its reliance on clustering for speaker tracking allows for better generalization to longer sequences even if it has been trained on shorter ones.

Comparing Table 5a with Table 3, which both consider (H = 4, L = 4) for SA-EEND and (H = 8, L = 6) for EEND-VC, we can notice that the performance of EEND-VC decreases for all datasets except Fisher, where instead it improves due to the perfectly matched training and test. For SA-EEND we observe the same except that now the performance is also slightly better for CallCntrPOR.

In general, the performance loss by training on Fisher Corpus Part I, which is a much smaller dataset than the ones used for both EEND systems in Table 3, is quite significant. Overall this result suggests that both these approaches are especially data hungry, as the dataset used here is not that small, and that state-of-the-art results are possible only when a large amount of data are available (e.g., several thousands of hours, as used in EEND-VC and SA-EEND original works).

Finally, we can analyze how these two EEND-based approaches trained on Fisher compare with the SSGD-based pipeline which was also trained using Fisher and whose results were reported in Table 3. SSGD is now more competitive with respect to EEND-VC but overall this latter still affords lower DER on most datasets, especially on CallCntrPOR and Fisher. Both EEND-based approaches also appear more robust than SSGD on the two CallCntr datasets which have mismatched acoustic characteristics and different speech statistics compared to Fisher and CALLHOME. SSGD is more sensitive to such a mismatch.

In order to also carry out a comparison in which their number of parameters is the same, we trained both EEND models on Fisher Corpus Part I, setting for SA-EEND the same number of heads (H) and layers (L) used to obtain the results of EEND-VC in Table 5a, and vice versa for EEND-VC. Then, SA-EEND was trained with H = 8 and L = 6, whereas EEND-VC with H = 4 and L = 4; we report these results in Table 5b.

Therefore, by comparing Tables 5a and 5b, we can analyze two situations in which the number of parameters of the two EEND is on par; in particular, instances with (H = 4, L = 4) and (H = 8, L = 6).

In the latter case, we can see that EEND-VC performs better than SA-EEND on all test sets. For the former case, SA-EEND outperforms EEND-VC only on CALLHOME and CallCntrITA test sets, whereas is worse on Fisher and CallCntrPOR.




Fig. 5. RTF versus DER for the algorithms in exam. For each, we report average values and standard deviations (using bars) across the four datasets in exam.

SA-EEND is the one that suffers the most from the usage of an (H, L) pair different from the one considered in their pre-trained models.

Overall, it can be seen that the best EEND on all test sets still remains the pre-trained EEND-VC model of Table 3, thanks to the larger amount of data of the training set. The same states even when these EEND models are trained on the same data (i.e., Fisher Corpus Part I) when their number of layers is not the same as they have, respectively, in Table 3. Concerning experiments in which both models are trained on the same data and the number of layers on par, EEND-VC outperforms SA-EEND on all test sets when (H = 8, L = 6), and even on Fisher and CallCntrPOR when (H = 4, L = 4).

4.2. Computational cost

In order to visualize the accuracy of each algorithm together with its performance from a computational perspective, we use Figs. 5, 6, and 7; in order to have a more detailed overview, the precise values are also reported in Table 6.

In Fig. 5, we plot the RTF versus DER for each system under study. Both DER and RTF values have been averaged across the four datasets. The DER and RTF standard deviations are reported using bars. We can observe that both EEND methods in general offer a very good trade-off between performance and processing time, with EEND-VC being the best overall. Clustering-based techniques have very similar RTFs overall and overlap-aware variants (VBx+OVL and SCx+OVL) only add a small computational overhead as the most time-consuming operation consists in the embedding extraction.

SSGD has the highest RTF due to the fact that neural speech separation is quite expensive, especially as most current state-of-the-art models are largely based on learnt over-complete filterbanks (Luo et al., 2020; Luo and Mesgarani, 2019) with small kernels. On the other hand, as said, SSGD can be more convenient in all applications when it is necessary to transcribe both speakers, as diarization comes with a negligible computational overhead after separation. As expected, the AHC-BIC system is the fastest but its performance in terms of DER is also the worst amongst all.

From Fig. 5 (and even better from Table 6) it is possible to see that the standard deviation of the RTF (computed by averaging it across the four datasets) of DNN-based clustering systems is about 0.02, for SSGD is less than half, whereas for EEND systems and AHC-BIC is two orders of magnitude lower. The low standard deviation of RTF means that inference time is approximately proportional to the audio file duration. This aspect, in addition to indicating the low dependency of inference time on the different characteristics of each audio file (e.g., duration, speech sparsity, etc.), also reflects the reliability of the setup (no background tasks, etc.) used to conduct the experiments. Thus, the slightly greater standard deviation of clustering-based systems could be due to their greater susceptibility to speech sparsity variability, since they rely on x-vectors, which are computed only on speech segments, after the VAD.

In Fig. 6, we plot the maximum RAM occupation in inference versus DER for each system under study. Again for both we plot the macro-average across all datasets. We can see that most systems have very similar maximum RAM requirements: VBx, SCx, their OVL variants, and the two EEND models are under 1 GB of RAM. As expected, the SSGD model instead is the most memory-demanding due to the aforementioned fact about the filterbanks, which increases the amount of information the model needs to process.

Since it could be argued that, depending on the actual implementation, it is usually possible to trade-off speed (and thus here RTF) for memory occupation, in Fig. 7 we plot the RTF versus the maximum RAM occupation, again by macro averaging across the datasets. It is evident that, for the time versus space complexity trade-off, the systems that lie closer to the origin offer intrinsically better computational parsimony. We can see that the AHC-BIC and the two EEND approaches are by far the most efficient, whereas




Fig. 6. Maximum RAM occupation in Megabytes (MB) versus DER for the algorithms in exam. For each, we report average values and standard deviations (using bars) across the four datasets in exam.




Fig. 7. RTF versus RAM for the algorithms in exam. For each, we report average values and standard deviations (using bars) across the four datasets in exam.

Table 6

DER, RTF, and RAM obtained by performing inference with the different diarization algorithms under study. For each of them, we report the mean value, averaged on the four test sets, along with the associated standard deviation.

Metric	Unit	Diarization systems							
		VBx	SCx	VBx+OVL	SCx+OVL	AHC-BIC	SSGD	SA-EEND	EEND-VC
DER		11.57	12.21	10.42	11.54	22.00	13.54	12.95	9.41
		± 1.71	± 3.24	± 2.57	± 3.75	± 5.43	± 6.94	± 5.35	± 2.80
RTF		$1.31e-1$	$1.28e-1$	$1.52e-1$	$1.49e-1$	$6.45e-4$	$4.25e-1$	$1.64e-3$	$1.54e-3$
		$\pm 2.20e-2$	$\pm 2.18e-2$	$\pm 2.20e-2$	$\pm 2.18e-2$	$\pm 1.11e-4$	$\pm 9.36e-3$	$\pm 2.06e-4$	$\pm 3.67e-4$
RAM	MB	627	627	633	633	15	1814	508	577
		± 18	± 18	± 15	± 15	± 0.02	± 419	± 31	± 11

SSGD is the least efficient. The clustering-based approaches are largely equivalent despite their differences, as, again, the bulk of the processing is due to the extraction of the x-vectors.

From Fig. 7, it could be seen that SSGD is the system that suffers the most the variability between the different datasets, because its RAM's standard deviation is about 400 MB; in particular, the longer audio duration the greater the RAM occupation. In contrast, DNN-based clustering systems and EEND models experience a RAM's standard deviation of only a few tens of MB, and even lower for the AHC-BIC. EEND-VC shows stronger robustness to file duration with respect to SA-EEND, as its RAM occupation is about three times lower.

In general, we can say that EEND-VC affords the most favorable trade-off between performance and computational efficiency. SSGD achieves decent performance but is the least efficient. The clustering-based methods, excluding AHC-BIC, offer similar computational trade-offs and thus, VBx+OVL, which performs slightly better overall, should be preferred.

5. Conclusion

In this paper, we present an experimental review of speaker diarization in the CTS domain and investigate clustering-based, end-to-end, and separation-guided diarization methods. In total, we considered eight different methods and analyzed their performance in terms of diarization accuracy and computational requirements using four different datasets.

Our goal is to give quantitative guidance concerning which diarization approach might be the most appropriate depending on the specific use-case scenario, e.g., the characteristics of the audio recordings and the computational capabilities of the machine that can be used for inference.

Among the diarization methods, we considered VBx and a clustering-diarization pipeline based on spectral clustering (SCx), as well as their overlap-aware modifications (VBx+OVL, SCx+OVL), plus an approach based on AHC, BIC, and hand-crafted features, denoted as AHC-BIC. Regarding end-to-end methods, we tested SA-EEND and EEND-VC. The SSGD pipeline instead was taken from a recent work focused on low-latency diarization in the CTS domain. The datasets used are from multiple different languages and have different characteristics in terms of speech sparsity and presence of overlapped speech. We included both publicly available, widely used CTS datasets such as CALLHOME as well as two proprietary datasets from industry.

Regarding the clustering-based approaches, we found that the computational burden of overlap-aware VBx and SCx are comparable with their non overlap-aware versions, as most resources are due to the x-vector extraction step. In low overlapped speech scenarios (e.g., call-center customer care), their performance is comparable but, as expected, overlap-aware versions bring noticeable improvement on the datasets with more overlapped speech featuring more colloquial speech. VBx and its overlap-aware variant (VBx+OVL) were found to perform the best overall among clustering-based techniques. The AHC-BIC approach has the worst performance overall among all the methods but is also the most computationally efficient by far and thus may be the only viable option for edge applications with severe hardware constraints.

End-to-end approaches on the other hand offer better computational efficiency than all clustering approaches except the AHC-BIC. Overall EEND-VC performs the best among all methods and thus has the best trade-off between computing requirements and DER. SA-EEND instead suffers from training versus inference sequence length mismatch, leading to performance degradation in recordings with longer duration. Both SA-EEND and EEND-VC are highly susceptible to the characteristics of the training data and in particular to its size, and state-of-the-art results are only possible when trained on several thousands of hours of data.

Finally, SSGD can achieve good performance but its performance degrades in presence of acoustic mismatches between training and test data. EEND-based approaches are instead more robust to this. This is due to the fact that SSGD needs to be trained with datasets for which the oracle single speaker sources are available, a condition which is not usually satisfied on real-world data, leading to a possible domain mismatch. End-to-end fine-tuning could be a potential solution, and it could be explored in future works. Another drawback of SSGD is that it is the most demanding diarization technique computationally-wise among the ones considered here. This is because it relies on speech separation which is inherently a more challenging task than diarization. This aspect can only be tolerated if speech separation has another purpose aside from diarization (e.g., ASR).

Our contribution cannot be considered exhaustive, like any experimental review, and our intention is that it can be found useful as a reference for future research directions and industrial applications. Possible future extensions could focus on different scenarios such as meeting scenarios where the overlap conditions and the number of speakers are different as well as include other diarization systems such as ones based on TS-VAD.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Luca Serafini, reports financial support was provided by Fondazione CARITRO. Giovanni Morrone reports financial support was provided by Fondazione CARITRO.

Data availability

Two datasets we use are available under the LDC and are widely used in the literature. Other two are proprietary and confidential and cannot be disclosed.

Acknowledgments

We would like to thank Dr. Keisuke Kinoshita at NTT Communication Science Laboratories for providing the pre-trained EEND-VC model.

This work has been supported by the AGEVOLA project (SIME code 2019.0227), funded by the Fondazione CARITRO.

References

- Anguera, X., Bozonnet, S., Evans, N., Fredouille, C., Friedland, G., Vinyals, O., 2012. Speaker diarization: A review of recent research. *IEEE/ACM Trans. Audio Speech Lang. Process.* 20 (2), 356–370.
- Bai, Z., Zhang, X.-L., 2021. Speaker recognition based on deep learning: An overview. *Neural Netw.* 140, 65–99.
- Boeddeker, C., Heitkaemper, J., Schmalenstroer, J., Drude, L., Heymann, J., Haeb-Umbach, R., 2018. Front-end processing for the CHiME-5 dinner party scenario. In: *Proc. of CHiME-5 Workshop on Speech Processing in Everyday Environments*. pp. 35–40.
- Bredin, H., 2017. TristoutNet: triplet loss for speaker turn embedding. In: *Proc. of ICASSP*. IEEE, pp. 5430–5434.
- Bredin, H., Laurent, A., 2021. End-to-end speaker segmentation for overlap-aware resegmentation. In: *Proc. of Interspeech*. pp. 3111–3115.
- Bredin, H., Yin, R., Coria, J.M., Gelly, G., Korshunov, P., Lavechin, M., Fustes, D., Titeux, H., Bouaziz, W., Gill, M.-P., 2020. Pyannote.Audio: neural building blocks for speaker diarization. In: *Proc. of ICASSP*. IEEE, pp. 7124–7128.
- Brown, A., Huh, J., Chung, J.S., Nagrani, A., Garcia-Romero, D., Zisserman, A., 2022. Voxsrc 2021: The third voxceleb speaker recognition challenge. *arXiv preprint arXiv:2201.04583*.
- Bullock, L., Bredin, H., Garcia-Perera, L.P., 2020. Overlap-aware diarization: Resegmentation using neural end-to-end overlapped speech detection. In: *Proc. of ICASSP*. IEEE, pp. 7114–7118.
- 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: *Proc. of ICASSP*. IEEE, pp. 4133–4136.
- Chen, S., Gopalakrishnan, P., et al., 1998. Speaker, environment and channel change detection and clustering via the bayesian information criterion. In: *Proc. of DARPA Broadcast News Transcription and Understanding Workshop*. pp. 127–132.
- Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., Li, J., Kanda, N., Yoshioka, T., Xiao, X., et al., 2022. WavLM: Large-scale self-supervised pre-training for full stack speech processing. *IEEE J. Sel. Top. Sign. Process.* 16 (6), 1505–1518.
- 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: *Proc. of ICASSP*. IEEE, pp. 7284–7288.
- Cheng, M., Wang, W., Zhang, Y., Qin, X., Li, M., 2022. Target-speaker voice activity detection via sequence-to-sequence prediction. *arXiv preprint arXiv:2210.16127*.
- Chung, J.S., Huh, J., Mun, S., Lee, M., Heo, H.S., Choe, S., Ham, C., Jung, S., Lee, B.-J., Han, I., 2020. In defence of metric learning for speaker recognition. In: *Proc. of Interspeech*. pp. 2977–2981.
- Chung, J.S., Nagrani, A., Zisserman, A., 2018. VoxCeleb2: Deep speaker recognition. In: *Proc. of Interspeech*. pp. 1086–1090.
- Cieri, C., Miller, D., Walker, K., 2004. The Fisher corpus: A resource for the next generations of speech-to-text. In: *LREC*, Vol. 4. pp. 69–71.
- Coria, J.M., Bredin, H., Ghannay, S., Rosset, S., 2021. Overlap-aware low-latency online speaker diarization based on end-to-end local segmentation. In: *Proc. of ASRU*. IEEE, pp. 1139–1146.
- Cornell, S., Omologo, M., Squartini, S., Vincent, E., 2022. Overlapped speech detection and speaker counting using distant microphone arrays. *Comput. Speech Lang.* 72, 101306.
- Dehak, N., Kenny, P.J., Dehak, R., Dumouchel, P., Ouellet, P., 2010. Front-end factor analysis for speaker verification. *IEEE/ACM Trans. Audio Speech Lang. Process.* 19 (4), 788–798.
- Desplanques, B., Thienpondt, J., Demuynck, K., 2020. ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification. In: *Proc. of Interspeech*. pp. 3830–3834.
- Diez, M., Burget, L., Landini, F., Černocký, J., 2019. Analysis of speaker diarization based on Bayesian HMM with eigenvoice priors. *IEEE/ACM Trans. Audio Speech Lang. Process.* 28, 355–368.
- Diez, M., Burget, L., Matejka, P., 2018. Speaker diarization based on Bayesian HMM with eigenvoice priors. In: *Proc. of Odyssey*. pp. 147–154.
- Fan, Y., Kang, J., Li, L., Li, K., Chen, H., Cheng, S., Zhang, P., Zhou, Z., Cai, Y., Wang, D., 2020. CN-Celeb: a challenging chinese speaker recognition dataset. In: *Proc. of ICASSP*. IEEE, pp. 7604–7608.
- Fang, X., Ling, Z.-H., Sun, L., Niu, S.-T., Du, J., Liu, C., Sheng, Z.-C., 2021. A deep analysis of speech separation guided diarization under realistic conditions. In: *Proc. of APSIPA ASC*. IEEE, pp. 667–671.
- Fujita, Y., Kanda, N., Horiguchi, S., Nagamatsu, K., Watanabe, S., 2019a. End-to-end neural speaker diarization with permutation-free objectives. In: *Proc. of Interspeech*. 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: *Proc. of ASRU*. IEEE, pp. 296–303.
- Fujita, Y., Watanabe, S., Horiguchi, S., Xue, Y., Nagamatsu, K., 2020. End-to-end neural diarization: Reformulating speaker diarization as simple multi-label classification. *arXiv preprint arXiv:2003.02966*.
- Gish, H., Siu, M.-H., Rohlicek, J.R., 1991. Segregation of speakers for speech recognition and speaker identification. In: *Proc. of ICASSP*. pp. 873–876.
- Han, E., Lee, C., Stolcke, A., 2021. BW-EDA-EEND: Streaming end-to-end neural speaker diarization for a variable number of speakers. In: *Proc. of ICASSP*. IEEE, pp. 7193–7197.
- Hatch, A.O., Kajarekar, S., Stolcke, A., 2006. Within-class covariance normalization for SVM-based speaker recognition. In: *Proc. of Interspeech*. pp. 1471–1474.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proc. of CVPR*. pp. 770–778.
- Hermansky, H., 1990. Perceptual linear predictive (PLP) analysis of speech. *J. Acoust. Soc. Am.* 87 (4), 1738–1752.
- Horiguchi, S., Fujita, Y., Watanabe, S., Xue, Y., Garcia, P., 2022a. Encoder-decoder based attractors for end-to-end neural diarization. *IEEE/ACM Trans. Audio Speech Lang. Process.* 30, 1493–1507.
- 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: *Proc. of Interspeech*. pp. 269–273.
- Horiguchi, S., Watanabe, S., Garcia, P., Takashima, Y., Kawaguchi, Y., 2022b. Online neural diarization of unlimited numbers of speakers. *arXiv preprint arXiv:2206.02432*.
- Horiguchi, S., Yalta, N., Garcia, P., Takashima, Y., Xue, Y., Raj, D., Huang, Z., Fujita, Y., Watanabe, S., Khudanpur, S., 2021. The Hitachi-JHU DIHARD III system: Competitive end-to-end neural diarization and x-vector clustering systems combined by DOVER-Lap. In: *3rd DIHARD Speech Diarization Challenge Workshop*.

- Huang, Z., Watanabe, S., Fujita, Y., Garcia, P., Shao, Y., Povey, D., Khudanpur, S., 2020. Speaker diarization with region proposal network. In: Proc. of ICASSP. IEEE, pp. 6514–6518.
- Huh, J., Brown, A., Jung, J.-w., Chung, J.S., Nagrani, A., Garcia-Romero, D., Zisserman, A., 2023. Voxsrc 2022: The fourth VoxCeleb speaker recognition challenge. arXiv preprint arXiv:2302.10248.
- 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: Proc. of DARPA Speech Recognition Workshop. pp. 61–66.
- Junqua, J.-C., Mak, B., Reaves, B., 1994. A robust algorithm for word boundary detection in the presence of noise. IEEE Trans. Speech Audio Process. 2 (3), 406–412.
- Kanda, N., Boeddeker, C., Heitkaemper, J., Fujita, Y., Horiguchi, S., Haeb-Umbach, R., 2019. Guided source separation meets a strong ASR backend: Hitachi/Paderborn University joint investigation for dinner party ASR. In: Proc. of Interspeech. pp. 1248–1252.
- Kenny, P., Boulianne, G., Ouellet, P., Dumouchel, P., 2007. Speaker and session variability in GMM-based speaker verification. IEEE/ACM Trans. Audio Speech Lang. Process. 15 (4), 1448–1460.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization. In: Proc. of ICLR.
- Kinoshita, K., Delcroix, M., Iwata, T., 2022a. Tight integration of neural-and clustering-based diarization through deep unfolding of infinite gaussian mixture model. In: Proc. of ICASSP. IEEE, pp. 8382–8386.
- Kinoshita, K., Delcroix, M., Tawara, N., 2021a. Advances in integration of end-to-end neural and clustering-based diarization for real conversational speech. In: Proc. of Interspeech. pp. 3565–3569.
- Kinoshita, K., Delcroix, M., Tawara, N., 2021b. Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds. In: Proc. of ICASSP. IEEE, pp. 7198–7202.
- Kinoshita, K., von Neumann, T., Delcroix, M., Boeddeker, C., Haeb-Umbach, R., 2022b. Utterance-by-utterance overlap-aware neural diarization with Graph-PIT. In: Proc. of Interspeech. pp. 1486–1490.
- Kolbæk, M., Yu, D., Tan, Z.-H., Jensen, J., 2017. Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks. IEEE/ACM Trans. Audio Speech Lang. Process. 25 (10), 1901–1913.
- Koluguri, N.R., Park, T., Ginsburg, B., 2022. TitaNet: Neural model for speaker representation with 1D depth-wise separable convolutions and global context. In: Proc. of ICASSP. IEEE, pp. 8102–8106.
- Landini, F., Glembek, O., Matějka, P., Rohdin, J., Burget, L., Diez, M., Silnova, A., 2021a. Analysis of the BUT diarization system for VoxConverse challenge. In: Proc. of ICASSP. IEEE, pp. 5819–5823.
- Landini, F., Lozano-Diez, A., Burget, L., Diez, M., Silnova, A., Zmoliková, K., Glembek, O., Matejka, P., Stafylakis, T., Brümmer, N., 2021b. BUT system description for the third DIHARD speech diarization challenge. In: Proc. of 3rd DIHARD Speech Diarization Challenge Workshop.
- Landini, F., Lozano-Diez, A., Diez, M., Burget, L., 2022. From simulated mixtures to simulated conversations as training data for end-to-end neural diarization. pp. 5095–5099.
- 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. Comput. Speech Lang. 71, 101254.
- Li, Y., Gao, F., Ou, Z., Sun, J., 2018. Angular softmax loss for end-to-end speaker verification. In: Proc. of ISCSLP. IEEE, pp. 190–194.
- Li, Q., Kreyssig, F.L., Zhang, C., Woodland, P.C., 2021. Discriminative neural clustering for speaker diarisation. In: 2021 IEEE Spoken Language Technology Workshop. SLT, IEEE, pp. 574–581.
- Liu, D., Kubala, F., 1999. Fast speaker change detection for broadcast news transcription and indexing. In: Proc. of Eurospeech. pp. 1031–1034.
- Lloyd, S., 1982. Least squares quantization in PCM. IEEE/ACM Trans. Audio Speech Lang. Process. 28 (2), 129–137.
- Luo, Y., Chen, Z., Yoshioka, T., 2020. Dual-Path RNN: efficient long sequence modeling for time-domain single-channel speech separation. In: Proc. of ICASSP. IEEE, pp. 46–50.
- Luo, Y., Mesgarani, N., 2019. Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation. IEEE/ACM Trans. Audio Speech Lang. Process. 27 (8), 1256–1266.
- 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: Proc. of ICASSP. IEEE, pp. 4828–4831.
- Medennikov, I., Korenevsky, M., Prisyach, T., Khokhlov, Y., Korenevskaya, M., Sorokin, I., Timofeeva, T., Mitrofanov, A., Andrusenko, A., Podluzhny, I., et al., 2020. Target-speaker voice activity detection: A novel approach for multi-speaker diarization in a dinner party scenario. In: Proc. of Interspeech. pp. 274–278.
- Miao, Y., Zhang, H., Metze, F., 2015. Speaker adaptive training of deep neural network acoustic models using i-vectors. IEEE/ACM Trans. Audio Speech Lang. Process. 23 (11), 1938–1949.
- Morrone, G., Cornell, S., Raj, D., Serafini, L., Zovato, E., Brutti, A., Squartini, S., 2022. Low-latency speech separation guided diarization for telephone conversations. In: Proc. of SLT. IEEE.
- Nagrani, A., Chung, J.S., Zisserman, A., 2017. VoxCeleb: A large-scale speaker identification dataset. In: Proc. of Interspeech. pp. 2616–2620.
- Ng, A., Jordan, M., Weiss, Y., 2001. On spectral clustering: Analysis and an algorithm. Adv. Neural Inf. Process. Syst. 14.
- Ning, H., Liu, M., Tang, H., Huang, T.S., 2006. A spectral clustering approach to speaker diarization. In: Ninth International Conference on Spoken Language Processing.
- O’Shaughnessy, D., 1988. Linear predictive coding. IEEE Potentials 7 (1), 29–32.
- Park, T.J., Han, K.J., Kumar, M., Narayanan, S., 2019. Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap. IEEE Signal Process. Lett. 27, 381–385.
- Park, T.J., Kanda, N., Dimitriadis, D., Han, K.J., Watanabe, S., Narayanan, S., 2022. A review of speaker diarization: Recent advances with deep learning. Comput. Speech Lang. 72, 101317.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. Adv. NIPS 32, 8026–8037.
- Peddinti, V., Chen, G., Manohar, V., Ko, T., Povey, D., Khudanpur, S., 2015. JHU ASPIRE system: Robust LVCSR with TDNNs, iVector adaptation and RNN-LMS. In: Proc. of ASRU. IEEE, pp. 539–546.
- Przybocki, M., Alvin, M., 2001. 2000 NIST Speaker Recognition Evaluation LDC2001S9. URL <https://catalog.ldc.upenn.edu/LDC2001S97>.
- Raj, D., Garcia-Perera, L.P., Huang, Z., Watanabe, S., Povey, D., Stolcke, A., Khudanpur, S., 2021a. DOVER-Lap: A method for combining overlap-aware diarization outputs. In: Proc. of SLT. IEEE, pp. 881–888.
- Raj, D., Huang, Z., Khudanpur, S., 2021b. Multi-class spectral clustering with overlaps for speaker diarization. In: Proc. of SLT. IEEE, pp. 582–589.
- Raj, D., Khudanpur, S., 2021. Reformulating DOVER-Lap label mapping as a graph partitioning problem. In: Proc. of Interspeech. pp. 2351–2355.
- Reddy, C.K., Gopal, V., Cutler, R., 2021. DNSMOS: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors. In: Proc. of ICASSP. IEEE, pp. 6493–6497.
- 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.
- Ryant, N., Church, K., Cieri, C., Cristia, A., Du, J., Ganapathy, S., Liberman, M., 2019. Second DIHARD challenge evaluation plan. Tech. Rep, Linguistic Data Consortium.
- Ryant, N., Church, K., Cieri, C., Du, J., Ganapathy, S., Liberman, M., 2021. Third DIHARD challenge evaluation plan. In: Proc. of Interspeech. pp. 3570–3574.
- Saon, G., Soltau, H., Nahamoo, D., Picheny, M., 2013. Speaker adaptation of neural network acoustic models using i-vectors. In: Proc. of ASRU. IEEE, pp. 55–59.

- Sari, L., Moritz, N., Hori, T., Le Roux, J., 2020. Unsupervised speaker adaptation using attention-based speaker memory for end-to-end ASR. In: Proc. of ICASSP. IEEE, pp. 7384–7388.
- Schwartz, R., Dodge, J., Smith, N.A., Etzioni, O., 2020. Green AI. *Commun. ACM* 63 (12), 54–63.
- Shafey, L.E., Soltau, H., Shafraan, I., 2019. Joint speech recognition and speaker diarization via sequence transduction. In: Proc. Interspeech 2019. pp. 396–400.
- Siu, M.-H., Yu, G., Gish, H., 1992. An unsupervised, sequential learning algorithm for the segmentation of speech waveforms with multiple speakers. In: Proc. of ICASSP. IEEE, pp. 189–192.
- Snyder, D., Chen, G., Povey, D., 2015. MUSAN: A music, speech, and noise corpus. arXiv preprint [arXiv:1510.08484](https://arxiv.org/abs/1510.08484).
- Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., Khudanpur, S., 2018. X-vectors: Robust DNN embeddings for speaker recognition. In: Proc. of ICASSP. IEEE, pp. 5329–5333.
- Stolcke, A., Yoshioka, T., 2019. DOVER: A method for combining diarization outputs. In: Proc. of ASRU. IEEE, pp. 757–763.
- Tranter, S.E., Reynolds, D.A., 2006. An overview of automatic speaker diarization systems. *IEEE/ACM Trans. Audio Speech Lang. Process.* 14 (5), 1557–1565.
- Tritschler, A., Gopinath, R.A., 1999. Improved speaker segmentation and segments clustering using the bayesian information criterion. In: Proc. of Eurospeech. pp. 679–682.
- Variani, E., Lei, X., McDermott, E., Moreno, I.L., Gonzalez-Dominguez, J., 2014. Deep neural networks for small footprint text-dependent speaker verification. In: Proc. of ICASSP. IEEE, pp. 4052–4056.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Adv. NIPS* 30, 6000–6010.
- Wagstaff, K., Cardie, C., Rogers, S., Schrödl, S., et al., 2001. Constrained K-means clustering with background knowledge. In: Proc. of ICML. pp. 577–584.
- Wang, W., Cai, D., Lin, Q., Yang, L., Wang, J., Wang, J., Li, M., 2021a. The dku-dukeece-lenovo system for the diarization task of the 2021 voxceleb speaker recognition challenge. arXiv preprint [arXiv:2109.02002](https://arxiv.org/abs/2109.02002).
- Wang, Q., Downey, C., Wan, L., Mansfield, P.A., Moreno, I.L., 2018. Speaker diarization with LSTM. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, IEEE, pp. 5239–5243.
- Wang, Y., He, M., Niu, S., Sun, L., Gao, T., Fang, X., Pan, J., Du, J., Lee, C.-H., 2021b. USTC-NELSLIP system description for DIHARD-III challenge. In: Proc. of 3rd DIHARD Speech Diarization Challenge Workshop.
- Wang, W., Li, M., 2022. Incorporating end-to-end framework into target-speaker voice activity detection. In: Proc. of ICASSP. IEEE, pp. 8362–8366.
- Wang, W., Lin, Q., Cai, D., Li, M., 2022a. Similarity measurement of segment-level speaker embeddings in speaker diarization. *IEEE/ACM Trans. Audio Speech Lang. Process.* 30, 2645–2658.
- Wang, W., Lin, Q., Li, M., 2022b. Online target speaker voice activity detection for speaker diarization. arXiv preprint [arXiv:2207.05920](https://arxiv.org/abs/2207.05920).
- Wang, W., Qin, X., Cheng, M., Zhang, Y., Wang, K., Li, M., 2022c. The DKU-DukeECE diarization system for the VoxCeleb speaker recognition challenge 2022. arXiv preprint [arXiv:2210.01677](https://arxiv.org/abs/2210.01677).
- Wang, W., Qin, X., Li, M., 2022d. Cross-channel attention-based target speaker voice activity detection: Experimental results for the m2met challenge. In: Proc. of ICASSP. IEEE, pp. 9171–9175.
- Wang, Z.-Q., Wang, D., 2017. Unsupervised speaker adaptation of batch normalized acoustic models for robust ASR. In: Proc. of ICASSP. IEEE, pp. 4890–4894.
- Watanabe, S., Mandel, M., Barker, J., Vincent, E., Arora, A., Chang, X., Khudanpur, S., Manohar, V., Povey, D., Raj, D., Snyder, D., Subramanian, A.S., Trmal, J., Yair, B.B., Boeddeker, C., Ni, Z., Fujita, Y., Horiguchi, S., Kanda, N., Yoshioka, T., Ryant, N., 2020. CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings. In: Proc. of CHiME-6 Workshop on Speech Processing in Every-Day Environments. pp. 1–7.
- Xia, W., Lu, H., Wang, Q., Tripathi, A., Huang, Y., Moreno, I.L., Sak, H., 2022. Turn-to-diarize: Online speaker diarization constrained by transformer transducer speaker turn detection. In: ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, IEEE, pp. 8077–8081.
- Xue, Y., Horiguchi, S., Fujita, Y., Watanabe, S., Garcia, P., Nagamatsu, K., 2021. Online end-to-end neural diarization with speaker-tracing buffer. In: Proc. of SLT. IEEE, pp. 841–848.
- Yamaguchi, M., Yamashita, M., Matsunaga, S., 2005. Spectral cross-correlation features for audio indexing of broadcast news and meetings. In: Proc. of Interspeech. pp. 613–616.
- Yu, F., Zhang, S., Fu, Y., Xie, L., Zheng, S., Du, Z., Huang, W., Guo, P., Yan, Z., Ma, B., Xu, X., Bu, H., 2022. M2met: The icassp 2022 multi-channel multi-party meeting transcription challenge. In: ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, pp. 6167–6171. <http://dx.doi.org/10.1109/ICASSP43922.2022.9746465>.
- Zeghidour, N., Grangier, D., 2021. Wavesplit: End-to-end speech separation by speaker clustering. *IEEE/ACM Trans. Audio Speech Lang. Process.* 29, 2840–2849.
- Zeghidour, N., Teboul, O., Grangier, D., 2021. DIVE: End-to-end speech diarization via iterative speaker embedding. In: Proc. of ASRU. IEEE, pp. 702–709.
- Zeinali, H., Sameti, H., Stafylakis, T., 2018. DeepMine speech processing database: Text-dependent and independent speaker verification and speech recognition in Persian and English. In: Proc. of Odyssey. pp. 386–392.
- Zhang, A., Wang, Q., Zhu, Z., Paisley, J., Wang, C., 2019. Fully supervised speaker diarization. In: Proc. of ICASSP. IEEE, pp. 6301–6305.