



Baselines and Protocols for Household Speaker Recognition

Alexey Sholokhov, Xuechen Liu, Md Sahidullah, Tomi Kinnunen

► To cite this version:

Alexey Sholokhov, Xuechen Liu, Md Sahidullah, Tomi Kinnunen. Baselines and Protocols for Household Speaker Recognition. The Speaker and Language Recognition Workshop (Odyssey 2022), Jun 2022, Beijing, China. pp.185-192, 10.21437/Odyssey.2022-26 . hal-03846180

HAL Id: hal-03846180

<https://inria.hal.science/hal-03846180>

Submitted on 10 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Baselines and Protocols for Household Speaker Recognition

Alexey Sholokhov¹, Xuechen Liu^{2,3}, Md Sahidullah³, Tomi Kinnunen²

¹Independent researcher

²School of Computing, University of Eastern Finland, Finland

³Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France

sholokhovalexey@gmail.com, {tomi.kinnunen,xuechen.liu}@uef.fi, md.sahidullah@inria.fr

Abstract

Speaker recognition on household devices, such as smart speakers, features several challenges: (i) robustness across a vast number of heterogeneous domains (households), (ii) short utterances, (iii) possibly absent speaker labels of the enrollment data (passive enrollment), and (iv) presence of unknown persons (guests). While many commercial products exist, there is less published research and no publicly-available evaluation protocols or open-source baselines. Our work serves to bridge this gap by providing an accessible evaluation benchmark derived from public resources (VoxCeleb and ASVspoof 2019 data) along with a preliminary pool of open-source baselines. This includes four algorithms for active enrollment (speaker labels available) and one algorithm for passive enrollment.

1. Introduction

Speech technology has entered our living rooms. Be it a smart TV or smart home assistant connected to other devices, an increasing number of commercial products now support interaction by voice. In addition to recognizing voice commands correctly (speech recognition), it is important to be aware of the user (speaker identification) for customized experiences such as personalized playlists, recommendations and bookings.

We address such *household speaker recognition* task [1], [2] where a small group of users (household members) share a common device. The task is differentiated from generic speaker recognition tasks (addressed in standard benchmark data such as NIST SREs [3] and VoxCeleb [4], [5]) with several desiderata, intended to maximize user experience:

- **Domain robustness.** The device should be readily-usable in previously unseen domains (households);
- **User convenience.** User registration (speaker enrollment) should be done as unobtrusively as possible — ideally without any interaction with the user;
- **Adaptivity.** To maintain high performance through weeks, months, or years, the model(s) require updating (again without user interaction);
- **Short utterances.** The device should be attentive to the user with very short utterances (1-2 second range);
- **Open-set operation.** The device should be attentive only to the members of a given household, i.e. the ability to ignore commands from unknown (out of the household) users;
- **Computational efficiency.** With growing privacy concerns, it might be preferable to carry out *on-device* computations without sending the user's speech data elsewhere.

The above translate to computationally light and ideally *unsupervised* (no speaker labels required) approaches that evolve with time, as opposed to single-shot optimization on a given corpus, domain, or challenge.

In generic speaker recognition setups, the issue of speech variation across different datasets is typically addressed using *domain adaptation* (DA) of an already well-trained recognizer. Practice has shown that it is effective to first train a large speaker embedding network (that serves as a utterance-level speaker feature extractor) [6], [7], fix it across domains, followed by adaptation of back-end classifier (speaker comparator) to each domain of interest [8]–[11].

While DA methods are designed to address *small number of large domains* (many speakers), the household scenario represents the opposite — *large number of small domains* (households). As noted in [2], systems designed to address *generic* speaker recognition task on large population might be suboptimal on *specific* speaker subsets with homogeneous data. Homogeneous data can indeed be expected due to the shared device with the fixed (or infrequently changed) location. The household scenario is therefore differentiated from speaker recognition *across devices* or *on mobile devices*. In addition, DA methods are typically not designed to handle sequential (online) data arriving one utterance at a time. For these reasons, we consider lightweight back-ends with as few parameters as possible and instead focus on adapting *speaker-specific representations* (here, deep neural embeddings).

Speaker recognition for household scenarios is a relatively recent direction (at least in the published literature) [1], [2]. In [1], enrichment of speaker-specific representation using unlabeled data is achieved through semi-supervised *label propagation* [12]. Another study [2] considers a neural back-end classifier with lower-dimensional space, into which the extracted embeddings are mapped. This space is optimized for a given speaker subset. Even if we draw inspiration from these prior studies, there is an important difference in the assumptions. Specifically, we argue that in households one must prepare for the possibility of external non-target speakers, or *guests*. Therefore, the label propagation approach in [1] is not directly applicable for potential unknown (out of household) speakers.

We summarize our contributions as follows. First and foremost, given the sparse literature and lacking common methodological standards, we target our study to researchers and practitioners who look for an accessible and reproducible starting point to address household speaker recognition tasks. The methodology sections, therefore, provide a guided taxonomy of methods and suggested terminologies. Importantly, we make our experimental pipeline, methods, and protocols public¹ both for reproducibility and for promotion of further research on this topic.

Besides advancing **evaluation infrastructure** as described

¹<https://github.com/underdogliu/household-speaker-recognition>

above, we have a number of algorithmic and experimental design-related contributions too. In terms of **algorithms**, we modified several existing clustering algorithms including the one from [13] by introducing an extra “background class” and allowing for partially labeled data. Second, we consider a highly-constrained variant of *probabilistic linear discriminant analysis* (PLDA) back-end [14], [15] involving spherical covariances. This results in an equivalent scoring rule with the widely-used cosine similarity in the case of a single enrollment utterance, with an added benefit of improved score calibration for multiple enrollment utterances.

In terms of **experimental design**, while we draw inspiration from [1], [2], we suggest a number of extensions and improvements to the simulation of household speaker recognition tasks. Similar to [1], [2] we use the public VoxCeleb data [4], [5] in our experiments. In addition, we leverage another public dataset, ASVspoof 2019 physical access (PA) task [16]. Even if the primary use of this dataset is that of benchmarking spoofing countermeasures, in this work we focus on *bonafide* scenarios (no spoofing attacks). The ASVspoof 2019 PA data consists of carefully controlled simulated environments with variations in room sizes, reverberation time, and talker-to-microphone distances. As opposed to VoxCeleb (crawled from YouTube) where randomly-paired utterances may not be representative of a shared homogenous environment, the ASVspoof 2019 PA data allows a design aligned closer with the assumptions of household scenarios; speakers for a given household are selected from the same simulated environment. The experiments with both simulated environments (ASVspoof 2019) along with ‘in the wild’ style data (VoxCeleb) used in prior studies, therefore, complement one another.

2. Household Speaker Recognition

2.1. General task set-up and data assumptions

Before presenting the methodology details, we set the general stage for the household speaker recognition task. A **household** is a small set of natural persons (such as family members) whose voice data is processed by a shared device. The task of interest is *open-set speaker identification* (SID) [17] that aims at classifying an unseen utterance X either as one of the N known household members, or an additional $(N + 1)$:th ‘none of the known speakers’ class. We refer to such speakers as **guests**.

At the enrollment (registration) stage, the device collects voice data of the household members to create the speaker-specific model(s). In generic speaker recognition tasks, the speaker labels of enrollment data are typically available and trustworthy. For instance, published corpora and challenge datasets may undergo tedious scrutiny to validate the correctness of speaker labels (and other metadata) before their release. With similar quality control in mind, in the household scenario, the user might be prompted to repeat wake words or read other pre-specified text passages to create a personal profile. As such **active enrollment** could be time-consuming or otherwise inconvenient, a more attractive alternative is to collect speech data while the device is used and use speaker clustering to infer pseudo speaker labels after a sufficient amount of data has been collected. The user might then be asked for consent for profile creation. Following terminology used in some commercial products, we refer to such procedures as **passive enrollment**. In this work, we consider both active and passive enrollment strategies.

Besides the test and enrollment data, we assume an additional set of unlabeled **adaptation data** to be available for

updating speaker-specific models. Such *unsupervised speaker adaptation* was introduced to the NIST speaker recognition evaluation (SRE) series in 2004 [3]. This task is *not* to be confused with techniques such as transfer learning [18], [19] or domain adaptation [8]–[11] that re-train or modify selected parameters of the recognizer itself. In this work, both the speaker embedding extractor and classifier remain fixed whereas speaker-specific models undergo frequent updates.

Since adaptation data is unlabeled, the decision as to which speaker model (if any) should be updated has to be done by the recognizer automatically. Note that an error-free (*oracle*) SID system leads to an effectively larger amount of enrollment data per user. In any practical case, however, the SID system makes errors, leading to potentially updating incorrect speaker model(s).

One of the differences between the household speaker recognition from generic speaker recognition evaluation setups is the varying amount of data available for speaker enrollment and model updating. In contrast to, *e.g.*, VoxSRC challenge setup with a single enrollment utterance per trial [20], in the household scenario we assume that multiple utterances can be used to construct the speaker model. Further, the number of utterances may vary from speaker to speaker. This corresponds to *multi-enrollment* speaker verification [21], [22].

2.2. General procedure

Moving on from the data requirements towards methodology, the core task of interest is pairwise comparison of test utterance with hypothesized model(s)². In this study, we use different *deep neural embeddings* [23] to represent speech utterances. In particular, utterance X of arbitrary duration is represented using a speaker embedding $\mathbf{x} \in \mathbb{R}^d$ of d dimensions: $\mathbf{x} = \text{dnn}(X)$. Here $\text{dnn}(\cdot)$ is either E-TDNN [7], ECAPA-TDNN [6] or a Resnet34 (details in Section 6.2) used for processing both enrollment and test data. We consider these embedding extractors fixed and focus on the system back-end.

The underlying computational task is that of creating a **speaker representation** (speaker model) for each household member using enrollment and speaker adaptation data. We consider two types of speaker representations, based on *sets* and *centroids*. In the former, a speaker is represented by a collection of embeddings extracted from the utterances of that speaker. In the latter, in turn, we maintain only a single averaged embedding (centroid) per speaker. One can always switch from the set representations to centroids, but not in the other direction. We used one of these representations depending on the speaker adaptation algorithm.

In terms of adaptation algorithms, we consider several alternatives categorized either as **online** (sequential) or **offline** (batch) methods. Whereas the former (Section 4) assumes the unlabeled adaptation data to arrive utterance-by-utterance, the latter (Section 5) assumes the adaptation data of all speakers to be available at once. While the former might be preferable from a computational point of view, the latter has the potential to benefit from the global structure of adaptation data. A common ingredient to all considered adaptation algorithms is **scoring function** (Section 3) that produces a speaker similarity score between \mathbf{x} and hypothesized speaker model(s).

Algorithm 1 outlines a generic template for active enrollment with online updates: after active speaker enrollment, data

²Closed-set speaker identification with N known speakers consists of N pairwise model-test comparisons, and the open-set SID combines this process with additional score thresholding.

Algorithm 1 Active enrollment with online updates

Require:

- Labeled enrollment data $\mathcal{X}_{\text{enroll}} = \{(\mathbf{x}_i, y_i)\}$ from one household with speakers $y_i \in \{1, 2, \dots, K\}$.
- Unlabeled adaptation data $\mathcal{X}_a = \{\mathbf{x}_t\}$ containing both household members and guests.

Ensure: One model θ_y per speaker $y \in \{1, 2, \dots, K\}$.

```

for  $y \in \{1, 2, \dots, K\}$  do ▷ Enroll each speaker
   $\theta_y \leftarrow \text{ENROLL}(\{\mathbf{x}_i \in \mathcal{X}_e : y_i = y\})$ 
end for
for  $\mathbf{x} \in \mathcal{X}_a$  do ▷ Online model update
   $k^* \leftarrow \arg \max_k \text{SCORE}(\mathbf{x}, \theta_k)$ 
  if  $\text{SCORE}(\mathbf{x}, \theta_{k^*}) > \tau$  then
     $\theta_{k^*} \leftarrow \text{UPDATE}(\mathbf{x}, \theta_{k^*})$ 
  end if
end for

```

arrive one by one. The top-scoring model is updated if the speaker similarity score exceeds an update threshold, and none of the other models are updated.

Note that the presented generic algorithm assumes discarding presumably out-of-set embeddings without any further processing. However, in principle, these embeddings could be utilized for representing the impostor class corresponding to all the possible guest speakers.

Algorithm 2 Passive enrollment

Require:

- Unlabeled adaptation data $\mathcal{X}_a = \{\mathbf{x}_t\}$ containing both household members and guests.
- $\theta = \{\theta_1, \dots, \theta_K\}, K \leftarrow \text{CLUSTER}(\{\mathbf{x}_t \in \mathcal{X}_a\})$
▷ Perform clustering
-

In the passive enrollment scenario, one needs to evaluate the ability of the speaker recognition system to identify clusters corresponding to speakers in an unlabeled adaptation set as shown in Algorithm 2. *Evaluating* the success of passive enrollment is therefore similar to the evaluation of general clustering algorithms: given a set of data points, the performance of a clustering algorithm is measured by comparing the predicted groups to the ground-truth partition. There are some minor differences, however. In particular, not all clusters are equally important. Specifically, we assume that there are several dominating clusters in the set, representing the classes of interest. These target clusters represent household members. The aim is to correctly identify only these clusters, while the predicted partitions for the rest of the data set can be arbitrary. In other words, if a pair of guest speakers were incorrectly grouped into a single cluster, it should not affect the performance metric.

3. Scoring Back-Ends

In this section, we discuss considerations related to the choice of speaker representation and scoring back-ends. Recent challenges including VoxSRC [20], [24] and VOICES 2019 [25] indicate cosine scoring to be a common choice with speaker embedding extractor networks trained with variants of *angular margin softmax losses* such as AM-Softmax and AAM-Softmax [26], [27]. Another popular choice is PLDA which is more efficient in combination with x-vector embeddings [28], [29].

As for the latter, we consider PLDA model with the full-rank speaker subspace, in the form of a *two-covariance model*

[14]. It is specified by two covariance matrices, representing between- and within-speaker covariances. We choose to use *spherical covariances* (scaled identity matrices). Despite being a highly constrained full-covariance PLDA, this low-parameter version may have an edge in the household speaker recognition scenario. In particular, a model with only two scalar parameters can be presumed to be more robust against domain mismatch compared to full-covariance PLDA which usually requires domain adaptation to the target domain [9], [30]. In our experiments, we use embedding averaging for all models apart from this spherical covariance PLDA that uses the so-called *by-the-book scoring* [21] for multiple enrollment utterances.

It can be shown that the spherical covariance PLDA is *equivalent to cosine scoring*³ in the case of a single enrollment utterance. See Appendix 10.1 for the detailed derivation.

4. Online Algorithms

We now detail the algorithm that processes data sequentially, in an online manner. In what we refer to as *centroid model with exponential smoothing*, each of the K speakers in a given household is represented by a *centroid* (average vector) of enrollment data. Call these (initial) centroids $\{c_k^{(0)}\}, k = 1, \dots, K$. At each time-step $t = 1, 2, \dots$ we obtain a new utterance represented by speaker embedding $\mathbf{x}^{(t)}$, which is compared with all the centroids to obtain K similarity scores $s_1^{(t)}, \dots, s_K^{(t)}$. Here these scores are computed either using cosine similarity or PLDA [14].

Let $k_*^{(t)} \equiv \arg \max_{1 \leq k \leq K} \{s_k^{(t)}\}$ denote the highest scoring speaker and $s_*^{(t)}$ the corresponding score. We update the centroid of the previous time step if and only if $s_*^{(t)} > \tau$, by

$$c_{k_*}^{(t)} = \alpha \mathbf{x}^{(t)} + (1 - \alpha) c_{k_*}^{(t-1)}, \quad (1)$$

where τ and α are control parameters. Update concerns only the top-scoring centroid and *no* update takes place if $s_*^{(t)} \leq \tau$.

The update rule (1) corresponds to exponential smoothing [31] where the updated centroid is a weighted average of the current embedding and the current centroid. Thus, $0 \leq \alpha \leq 1$ serves as *smoothing factor* that controls the rate at which the centroids evolve. Larger values of α correspond to a stronger emphasis on the most recent observation. Note that simple averaging of embeddings fits this update rule too if one relaxes the assumption of the smoothing factor to be a constant and sets it instead as $\alpha_t = 1/(t + 1)$. In summary, the method has two control parameters α and τ . We address their role in the experimental part. For practical purposes, they are tuned with the aid of a development set.

One can see that the update rule (1) does not track the number of performed updates. However, the information about the number of averaged embeddings is required for the PLDA by-the-book scoring. At the same time, the number of updates may not be an adequate measure of the accumulated information since it depends on the value of the smoothing factor (*e.g.* consider $\alpha \approx 1$). We address this issue by proposing a heuristic rule to define the zero-order statistics (counts) in the case of weighted averaging of embeddings. See Appendix 10.2 for the detailed discussion.

5. Offline Algorithms

We now turn attention to offline algorithms that operate under the assumption that the whole adaptation set is available at once.

³This could be the reason why such model, to the authors knowledge, was not mentioned in the existing literature.

Recall that in the active enrollment scenario a small labeled enrollment set is augmented by an unlabeled adaptation set. This naturally fits the formulation of *semi-supervised clustering* [32], where additional knowledge about the true labels of a subset of data points is assumed. This observation led us to extend two clustering algorithms to their semi-supervised versions, detailed in the next two subsections.

5.1. Semi-supervised k-means clustering

We propose straightforward modifications to the widely known k-means algorithm [33] for the household adaptation setup. First, we modified the traditional (unsupervised) k-means algorithm to fix the cluster assignments for a subset of embeddings with known speaker labels. Second, we introduced an additional “background class” to allow data points with no cluster assignments, further referred to as outliers. These data points correspond to the guest speakers. Specifically, similar to the centroid model described above, we use a pre-defined threshold τ to determine whether a given data point should be assigned to the closest cluster or labeled as an outlier (a guest). This restricts the clusters to have a bounded diameter. Here, the threshold is a control parameter which should be estimated from data.

Similar to the centroid model, the cluster centers (centroids) are initialized by computing speaker-specific average embeddings from the enrollment data.

5.2. Semi-supervised variational Bayesian clustering

Our second proposed offline algorithm can be viewed as variational Bayesian inference [34] applied to a Gaussian mixture model [35]. In fact, it is a slightly modified version of the clustering algorithm described in [13], in which the speaker embeddings are assumed to have been generated by the two-covariance PLDA model [14] discussed in Section 3. That is, speakers are modeled by Gaussians with shared covariance but speaker-specific means. By assuming that K speakers are present in a given collection of embeddings, this results in a mixture of K Gaussian components. More details are given in the Appendix 10.3.

We modified the mixture model described in [13] where each component corresponds to a speaker. In particular, we add an extra label to present ‘background class’. In practice, we add one $(K + 1)$ th component to the mixture model to represent the marginal distribution of embeddings. This component is assumed to model all of the unknown (guest) speakers in the household scenario. Also, similar to the semi-supervised k-means clustering described above, we fixed the component assignments for a subset of embeddings with known speaker labels.

This model requires specifying the prior probabilities corresponding to the mixture components $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_{K+1}\}$. To define the prior probability for the ‘background class’, we apply a threshold τ as in the other algorithms. In detail, we define this probability as $\pi_{K+1} = \sigma(\tau)$, where $\sigma(\cdot)$ denotes the sigmoid function. Accordingly, the remaining probabilities π_k are initialized such that $\sum_k \pi_k = 1$ for $k = 1, \dots, K + 1$.

5.3. Label propagation

Inspired by recent related work [1], we also consider a label propagation approach [36]. Since the original algorithm is not directly applicable to our *open-set* scenario with potentially unknown speakers, we introduced an additional pre-filtering step. First, all the embeddings that are not sufficiently close to the initial speaker models are labeled as outliers and removed. This

Dataset	Target	Non-target (Known/Unknown)
VoxCeleb-dev	27874	80126 / 108000
VoxCeleb-eval	27706	80294 / 108000
ASVspoof 2019	75600	216000 / 291600

Table 1: Trial statistics of the datasets used for our experiments. (Known/Unknown refers to the non-target trials from the same household speakers/guests.)

outlier filtering is again controlled by a pre-defined threshold τ . After filtering, we assume the remaining set to consist of household members only. Thus, the conventional label propagation algorithm is applied to this set.

6. Experimental Setup

6.1. Household speaker recognition protocols

We have chosen two public datasets for our experiments: ASVspoof 2019 and VoxCeleb. By drawing inspiration from prior studies [1], [2], we design new evaluation protocols for the household speaker recognition task.

ASVspoof 2019: We consider the bonafide audio data of the *physical access* (PA) subset of ASVspoof 2019 representing different room conditions [16]. The dataset consists of 27 different room conditions with various room sizes, reverberation times, and talker-to-mic distances. In our experiments, we use the evaluation set consisting of 48 speakers (21 male and 27 female). We set the household sizes as 4, 6, 8, and 10, where each household consists of an equal number of male and female speakers. We create the enrollment list, adaptation list, and trial list considering audio files from the same room condition. The speakers are randomly selected. We also have the same number of guest speakers corresponding to a room but chosen randomly from the remaining speakers. We use three utterances for speaker enrollment, whereas the remaining utterances are used for adaptation. We additionally consider the audio data from the guest speakers for adaptation. The adaptation list contains utterances from the household and guest speakers in a random order. The trial lists consist of only same-gender trials as cross-gender non-target trials are relatively easier than the other.

VoxCeleb: We consider VoxCeleb1 data consisting of 1251 speakers, where audio data was originally collected from YouTube videos [4]. Following the work reported in [2], we have chosen four utterances for enrollment, 10 for evaluation, and 13 for adaptation. Each of the utterances was taken from different videos and we have considered speakers with at least 27 video recordings. Similar to ASVspoof 2019 PA, we have gender-balanced households. The households with sizes 4, 6, 8, and 10 were randomly sampled to create 400 simulated households. Each household has the same number of guests as the number of household members. Similar to ASVspoof 2019, there are no cross-gender trials. We have created two sets of ASV protocols from VoxCeleb1 with disjoint set of speakers. One of them is used as development set whereas the other set is used as test set.

To sum up, we designed three different protocols VoxCeleb-dev, VoxCeleb-eval, ASVspoof with trial statistics summarized in Table 1. The VoxCeleb-dev protocol serves as the development set and the remaining two are used for final evaluation.

6.2. Speaker embedding extractors and back-end scoring

X-vector. This speaker embedding extractor is based on *Extended TDNN* (E-TDNN) [7] x-vector. It is trained on *dev* set of VoxCeleb2 [5], with 5994 speakers. It acquires 80-dimensional

mel filterbank outputs as acoustic features. Having the same architecture from [7], we introduce two modifications: 1) We replace statistics pooling with attentive statistics pooling [37]; 2) For loss function, we substitute the conventional cross-entropy one with additive angular margin softmax [38].

SpeechBrain. This pre-trained model is recently proposed by SpeechBrain [39] with pre-trained model available⁴, based on *ECAPA-TDNN* [6], another x-vector variant with a hybrid of convolutional and residual building blocks at layers before the pooling layer. It also imports the aforementioned two modifications and the same acoustic features. The pre-trained model is trained on *train* set of VoxCeleb1 [4] and *dev* set of VoxCeleb2, which results in 7205 speakers. Both speaker embedding extractors reached promising performance in our pilot experiments⁵.

CLOVA. This pre-trained model is proposed by and described in [40]. It is based on ResNet34 architecture [41] and trained on *dev* set of VoxCeleb2 [5], with 5994 speakers. The input acoustic feature for training this model is 64-dimensional mel filterbank, with online data augmentation [40]. This pre-trained model is provided by CLOVA⁶. Pilot experiments confirm the promising performance of this model⁷.

Moreover, when extracting the speaker embeddings, before feeding the input audio into the pre-trained models, we perform trimming with a threshold of -20 dB, followed by random chunking with a length of 2 seconds, elaborating the short-duration property of the audios under the household scenario.

We used different scoring back-ends depending on the type of embedding. For the x-vector embeddings, we used the conventional back-end including linear discriminant analysis, length normalization, and the full-rank PLDA. For other embeddings cosine similarity or the proposed spherical covariance, PLDA were used.

6.3. Training details

Active enrollment. We used a subset of speakers from the VoxCeleb1 *dev* set for setting all the tunable parameters. To this end, we created two trial lists: one for selecting parameters of the algorithms and scoring back-end via grid-search and the other is for training a global linear calibration [42]. Naturally, speakers in this development set do not overlap with the speakers used to create the evaluation protocol. We also used the VoxCeleb2 *dev* set with 5994 speakers for training the full covariance PLDA back-end for the x-vector embeddings. Apart from the PLDA model, it includes the whitening transform with linear dimensionality reduction, followed by length-normalization.

Passive enrollment. We use *agglomerative hierarchical clustering* (AHC) [43] with cosine scoring for passive enrollment. AHC is an iterative method that merges the existing clusters if similarities between the clusters are higher than a predefined threshold. As a primary study on this subject, we employ cosine similarity as the similarity measure.

6.4. Performance evaluation

Active enrollment. We have a common test set composed of three types of trials. A *target* trial consists of matched enrollment and test identities, *known non-target* trial of two unmatched identities from the household, and *unknown non-target* of unmatched enrollment identity with a guest. We report two different kinds of *equal error rates* (EERs, %). Both use targets as the positive class while the negative class is either one of the two types of non-targets. EERs are computed from scores pooled across all households.

Passive enrollment. We first perform scoring on labeled test segments against all detected clusters which are supposed to represent individual speakers. Similar to adaptation data, the test data includes both household members and guest speakers. Each test segment is labeled with the closest cluster if the corresponding score is above a given threshold. Otherwise, the test segment is labeled as belonging to an ‘unknown’ speaker, that is, not a household member. Finally, we pool all the predictions across the households and employ micro-averaged *Jaccard error rate* (JER) [44] designed to measure the quality of clustering results. The numerical values of JER range from 0% to 100%. Lower values of JER indicate better clustering. We adopt JER as we are only interested in finding the clusters of household speakers. For finding the clustering threshold, we first perform a grid search within a certain range and directly perform the evaluation; we then find the optimal value on the subset from VoxCeleb-dev (same as the aforementioned one for the active enrollment) and evaluate with the found threshold. Thresholds are individually optimized for the type of embeddings.

7. Results for Active Enrollment

7.1. Online adaptation with the centroid model

We begin by examining the basic properties of the online algorithm described in Section 4. Specifically, we consider a special case of the update rule (1) that corresponds to simple averaging — summing embeddings with equal weights. We consider three different scoring methods, all of which are equivalent to cosine scoring in the case of a single enrollment utterance: cosine scoring with embedding averaging (CSEA), cosine scoring with score averaging (CSSA), and the proposed spherical covariance PLDA (sph PLDA).

Figure 1 displays the two different kinds of EERs on ASVspoof data as a function of the model update threshold, τ (we observed similar trends on VoxCeleb). As can be seen, with a high threshold the metrics approach the case of *no adaptation*. Indeed, if τ is set too high, no embeddings will be accepted for updating speaker models. On the opposite end, too low τ may lead to speaker models updated by embeddings of wrong speakers. There is (scoring method dependent) ‘sweet spot’ that leads to substantial improvement over no adaptation. In terms of the scoring approaches, the proposed spherical covariance PLDA achieves lower EERs than either of the cosine scoring variants on the full range of thresholds.

7.2. Comparison of online and offline methods

We now compare all the adaptation algorithms — online and offline — for different speaker embeddings extractors. The parameters of each algorithm were tuned on the development set (to give the lowest EERs) and kept fixed at the evaluation step. Besides the algorithms described above, we consider two other useful points of reference. The first one, **no adaptation**, which corresponds to a conventional speaker verification evaluation setup without additional adaptation data. The second

⁴<https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

⁵EER is 2.21% for E-TDNN and 0.90% for SpeechBrain ECAPA-TDNN on the VoxCeleb1 test set [4], scored via PLDA back-end from Kaldi and cosine similarity, respectively.

⁶https://github.com/clovaai/voxceleb_trainer

⁷Reported EER is 1.18% on the VoxCeleb1 test set [4].

Algorithm	Scoring	ASVspoof			VoxCeleb	
		CLOVA	SpeechBrain	x-vector	CLOVA	x-vector
No adaptation	CSEA	3.75 / 3.82	4.42 / 4.44	–	1.87 / 1.74	–
No adaptation	PLDA	3.73 / 3.83	4.12 / 4.13	9.55 / 9.64	1.74 / 1.74	3.85 / 3.87
Centroid, $\alpha = 1/(t+1)$	CSEA	3.03 / 4.06	3.54 / 3.60	–	1.39 / 1.40	–
Centroid, $\alpha = 1/(t+1)$	PLDA	3.04 / 3.52	3.06 / 3.08	11.88 / 17.57	1.38 / 1.38	3.04 / 3.32
Centroid, $\alpha = 0.1$	CSEA	3.76 / 5.11	3.53 / 3.57	–	1.51 / 1.52	–
Centroid, $\alpha = 0.1$	PLDA	3.41 / 3.98	3.10 / 3.15	20.31 / 26.94	1.36 / 1.35	3.05 / 3.38
k-means	CSEA	3.07 / 4.17	3.50 / 3.51	–	1.46 / 1.46	–
k-means	PLDA	3.34 / 5.59	3.10 / 3.14	12.52 / 18.74	1.32 / 1.41	3.02 / 3.19
VB	CSEA	3.21 / 4.79	3.59 / 3.61	–	1.48 / 1.61	–
VB	PLDA	3.31 / 4.96	3.14 / 3.17	13.48 / 19.31	1.34 / 1.45	3.21 / 4.03
LP	CSEA	3.00 / 3.65	3.52 / 3.58	–	1.48 / 1.51	–
Oracle (error-free adaptation)	CSEA	2.84 / 2.91	3.29 / 3.29	–	1.25 / 1.13	–
Oracle (error-free adaptation)	PLDA	2.91 / 3.01	2.79 / 2.77	7.94 / 8.03	1.25 / 1.20	2.69 / 2.74

Table 2: EERs (%) with known / unknown non-targets for different embedding extractors and algorithms evaluated on the ASVspoof and the VoxCeleb protocols. From top to bottom, the first pair of rows represent the case where the adaptation set is empty. The next block shows results for the online centroid model with two different options to define the smoothing factor. The third block demonstrates different offline algorithms described in section 5, where VB stands for variational Bayesian clustering and LP is the label propagation. Finally, the last two rows show the Oracle adaptation corresponds to including the adaptation set to the enrollment set. Two scoring back-ends were used: cosine scoring with embedding averaging (CSEA) and PLDA. Since the training set for the SpeechBrain embedding extractor includes the Voxceleb1 data, we did not use it for the second protocol.

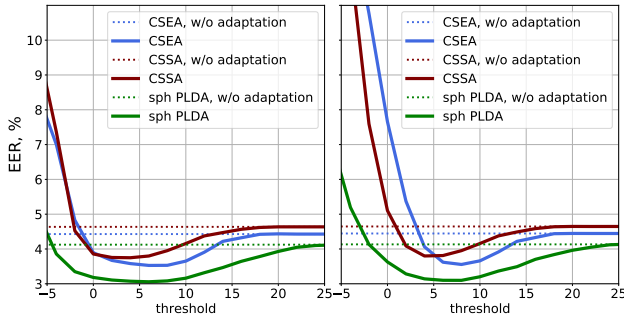


Figure 1: Comparison of scoring back-ends on the ASVspoof protocol. All the scores were calibrated by linear logistic regression. The presented results were obtained for the SpeechBrain embeddings. Left: EERs with known non-targets. Right: EERs with unknown non-targets.

one, **oracle**, including the adaptation data as part of the enrollment set; this corresponds to an error-free adaptation algorithm that predicts the speaker labels correctly.

Table 2 displays the results on both the ASVspoof and VoxCeleb evaluation protocols. The results are presented for two different scoring methods. In general, scoring is independent of the adaptation algorithm, however, in some cases (*e.g.* k-means) the algorithm also uses a similarity function as a part of its decision making process. While that function may be different from the one used for scoring verification trials, we did not consider such options in our experiments.

First, one can see that adaptation indeed helps to reduce the EER in all cases, apart from x-vector embeddings on the ASVspoof data. This can be explained by the sensitivity of the full-covariance PLDA model to domain mismatch as we did not apply any domain adaptation⁸ techniques such as [9], [30]. This is reinforced by noting that x-vector embeddings also benefit from speaker adaptation when evaluated with the matched (VoxCeleb) domain.

Note that the metrics reported in Table 2 may substantially

⁸Not to be confused with speaker adaptation studied in this work.

deviate from the values estimated using in-domain development data. Due to domain mismatch, the threshold tuned on the development set may be suboptimal in other domains. Nonetheless, with motivations noted in Section 1, the authors have purposefully refrained from extensive domain-specific scoring method optimization — we wanted to expose the methods to handle the unknown.

Another important and perhaps surprising observation from Table 2 is that the online algorithm achieved comparable (or sometimes even better) performance to offline algorithms. This finding might be explained by the relatively large amount of data available for adaptation. We hypothesize that such an amount is sufficient for saturation in performance while lower amounts of adaptation data might benefit from offline methods.

Finally, one can see that all the offline algorithms have a comparable performance by achieving up to 20% relative reduction in EER with known non-targets and up to 15% relative reduction in EER with unknown non-targets. With this observation, a possible choice could be using the arguably simplest algorithm – semi-supervised k-means with thresholding.

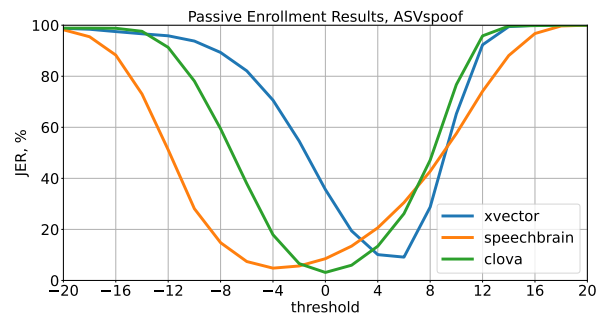


Figure 2: Results (in terms of JER, %) for passive enrollment with three types of speaker embeddings.

8. Results for Passive Enrollment

The results of passive enrollment with varying thresholds are illustrated in Fig. 2 for the ASVspoof dataset for all three types of speaker embeddings. SpeechBrain yields lower JER than the

Embedding	ASVspoof	VoxCeleb
CLOVA	27.42	4.55
SpeechBrain	7.41	-
x-vector	70.64	8.86

Table 3: JERs (%) of different embeddings on the two evaluation sets. The threshold value is optimized individually. Since the training set for the SpeechBrain embedding extractor includes the VoxCeleb1 data, we do not evaluate it here.

other two embeddings for a wide range of thresholds. Performance becomes worse till 100% JER is beyond a certain range of threshold values applied on the similarity measures. For a certain lower value of the threshold, the number of predicted clusters is less, hence resulting in a high false alarm. Conversely, for a certain higher value of the threshold, the number of predicted clusters is very high and that leads to a high miss rate.

In our final experiment, we use the VoxCeleb *dev* data to set the thresholds (one per each embedding type). These thresholds then remain fixed on the *evaluation* part of ASVspoof and VoxCeleb data. The results in Table 3 indicate that SpeechBrain performs best on ASVspoof with a JER of 7.41%, while the performance is worst for the x-vector embedding extractor. This might be due to dataset mismatch leading to a threshold that fails to generalize from VoxCeleb on ASVspoof. Concerning the relative performance of CLOVA and x-vector, the former outperforms the latter by demonstrating better generalization ability on both datasets. Nevertheless, the best performances via grid search over different clustering thresholds (Fig. 2) indicate scope for improvement in passive enrollment.

All in all, the above findings are observations that serve as a baseline and future work on passive enrollment may start from here.

9. Conclusion

The aim of this work has been the creation of an accessible and reproducible entry point to address household speaker recognition tasks. Besides detailing the task and assumptions, we designed a public evaluation benchmark for both active and passive enrollment; and online and offline algorithms. From the various methods considered, the simple online update rule based on the weighted average of embeddings achieved competitive performance to offline algorithms. The simplicity and computationally cheap updates make it a viable solution for household speaker recognition.

Our planned future work includes (1) improving the baselines to improve domain robustness. For instance, it would be interesting to generalize existing domain adaptation methods to handle online data. (2) We also plan to analyze the online methods in terms of their sensitivity to the *order* of data. (3) Despite the promising preliminary results, it is also important to analyze the impact of the critical control parameters on the convergence properties of the algorithms. An important difference with generic speaker recognition tasks is that the model(s) are allowed to evolve with time which complicates analysis. Nonetheless, it is preferable to find principled alternatives over grid search. (4) Finally, we plan to include more methods, especially to the passive enrollment scenarios which we only slightly touched upon. We may also need new performance metrics for a more aligned comparison of active and passive scenarios.

10. Appendix

10.1. PLDA with spherical covariance \equiv cosine scoring

Let $\mathbf{x}_{i,j} \in \mathbb{R}^d$ denote j th speaker embedding (observation) of speaker i . The *two-covariance* PLDA [14] assumes these observations to have been generated by a linear Gaussian model of the form

$$\mathbf{x}_{i,j} = \boldsymbol{\mu} + \mathbf{y}_i + \boldsymbol{\varepsilon}_{i,j},$$

where $\boldsymbol{\mu}$ is global mean of the embedding space, \mathbf{y}_i is latent speaker identity variable with prior distribution $\mathbf{y}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{B})$ and $\boldsymbol{\varepsilon}_{i,j}$ a residual with prior $\boldsymbol{\varepsilon}_{i,j} \sim \mathcal{N}(\mathbf{0}, \mathbf{W})$. The model is equivalently specified by the following distributions:

$$\begin{aligned} p(\mathbf{y}_i) &= \mathcal{N}(\mathbf{y}_i | \mathbf{0}, \mathbf{B}) \\ p(\mathbf{x}_{i,j} | \mathbf{y}_i) &= \mathcal{N}(\mathbf{x}_{i,j} | \boldsymbol{\mu} + \mathbf{y}_i, \mathbf{W}). \end{aligned} \quad (2)$$

The two matrices \mathbf{B} and \mathbf{W} model between- and within-speaker covariances, respectively. By imposing subspace structure on these matrices (rank $< d$) one may obtain other flavors of PLDA [15] as special cases. Here, we consider the highly-constrained case of *spherical covariances*,

$$\begin{aligned} \mathbf{B} &= \sigma_B^2 \mathbf{I} \\ \mathbf{W} &= \sigma_W^2 \mathbf{I}, \end{aligned} \quad (3)$$

where σ_B^2 and σ_W^2 are between-speaker and within-speaker variances and \mathbf{I} denotes an identity matrix.

For speaker comparison, the PLDA model (2) is used for computing log-likelihood ratio (LLR) score of same-speaker vs. different-speaker hypotheses for a pair of enrollment (\mathbf{x}_e) and test (\mathbf{x}_t) embeddings. For the null (same-speaker) hypothesis one assumes \mathbf{x}_e and \mathbf{x}_t to share common latent identity variable while for the alternative (different-speaker) hypothesis these latent factors are assumed to be different. While the interested reader may refer to details in [45], [46], the resulting LLR score can be written in closed form as

$$\begin{aligned} \ell &= \mathbf{x}_e^T \mathbf{P} \mathbf{x}_t + \mathbf{x}_t^T \mathbf{P} \mathbf{x}_e + \mathbf{x}_e^T \mathbf{Q} \mathbf{x}_e + \mathbf{x}_t^T \mathbf{Q} \mathbf{x}_t + \mathbf{c}^T (\mathbf{x}_t + \mathbf{x}_e) + k, \end{aligned}$$

where the matrices \mathbf{P} , \mathbf{Q} , the vector \mathbf{c} and scalar k depend only on the PLDA model parameters $(\boldsymbol{\mu}, \mathbf{B}, \mathbf{W})$. With the spherical covariance assumption (3), \mathbf{P} and \mathbf{Q} turn out to be scaled identity matrices too, say $\mathbf{P} = \alpha \mathbf{I}$, $\mathbf{Q} = \beta \mathbf{I}$. Further, the vector $\mathbf{c} = \gamma \boldsymbol{\mu}$ is scaled version of the mean.

We can now rewrite the LLR score as

$$\begin{aligned} \ell &= \mathbf{x}_e^T (\alpha \mathbf{I}) \mathbf{x}_t + \mathbf{x}_t^T (\alpha \mathbf{I}) \mathbf{x}_e + \mathbf{x}_e^T (\beta \mathbf{I}) \mathbf{x}_t + \mathbf{x}_t^T (\beta \mathbf{I}) \mathbf{x}_e \\ &\quad + \gamma \boldsymbol{\mu}^T (\mathbf{x}_e + \mathbf{x}_t) + k \\ &= 2\alpha \mathbf{x}_e^T \mathbf{x}_t + \beta (\|\mathbf{x}_e\|^2 + \|\mathbf{x}_t\|^2) + \gamma \boldsymbol{\mu}^T (\mathbf{x}_e + \mathbf{x}_t) + k. \end{aligned} \quad (4)$$

We proceed by noting that subtracting an arbitrary global offset $\mathbf{b} \in \mathbb{R}^d$ from all speaker embeddings, $\mathbf{x}_{i,j} \leftarrow \mathbf{x}_{i,j} - \mathbf{b}$, preserves pairwise distances and angles between vectors and therefore retains the geometric structure of the speaker embedding space. This leads to an otherwise identical PLDA model but with global mean $\boldsymbol{\mu} - \mathbf{b}$. Therefore, by *choosing* $\mathbf{b} = \boldsymbol{\mu}$, the new model has a global mean $\mathbf{0}$ and (4) becomes

$$\ell = 2\alpha \mathbf{x}_e^T \mathbf{x}_t + \beta (\|\mathbf{x}_e\|^2 + \|\mathbf{x}_t\|^2) + k.$$

If both the enrollment and test embeddings are length-normalized ($\|\mathbf{x}_e\| = \|\mathbf{x}_t\| = 1$), the score further simplifies to

$$\ell = 2\alpha \mathbf{x}_e^T \mathbf{x}_t + 2\beta + k,$$

which is an affine transform of the inner product $\mathbf{x}_e^T \mathbf{x}_t$ (cosine score) between the enrollment and test embeddings. Thus, PLDA scoring with spherical covariance matrices and cosine scoring are mappable to one another through order-preserving transform and therefore present equivalent scoring rules.

10.2. Note on PLDA and many-to-many verification

Consider *set-to-set* speaker verification where a trial is represented by a pair of sets with arbitrary cardinalities: $\mathbf{X}_e = \{\mathbf{x}_{e,1}, \dots, \mathbf{x}_{e,N}\}$ and $\mathbf{X}_t = \{\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,M}\}$, referred to as the *enrollment* and the *test* set, respectively.

Assuming that the observations are generated from the PLDA model with the latent identity variable, denoted by \mathbf{y} , the verification likelihood ratio can be written as a function of the posterior distributions of \mathbf{y} [47]

$$\text{LR}(\mathbf{X}_e, \mathbf{X}_t) = \frac{p(\mathbf{X}_e, \mathbf{X}_t | \text{same})}{p(\mathbf{X}_e, \mathbf{X}_t | \text{diff})} = \int \frac{p(\mathbf{y} | \mathbf{X}_e) p(\mathbf{y} | \mathbf{X}_t)}{p(\mathbf{y})} d\mathbf{y}, \quad (5)$$

where the integral admits the closed-form expression. These posteriors can be seen to be Gaussian with the mean parameters depending only on the average of feature vectors:

$$\begin{aligned} p(\mathbf{y} | \mathbf{x}_1, \dots, \mathbf{x}_N) &= \mathcal{N}(\mathbf{y} | \mathbf{m}, \Sigma), \\ \Sigma &= (\mathbf{B}^{-1} + N \cdot \mathbf{W}^{-1})^{-1}, \\ \mathbf{m} &= N \cdot \Sigma \mathbf{W}^{-1} \mathbf{c}, \end{aligned}$$

where $\mathbf{c} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i$ is the centroid (average vector) of the set. One can see that the log-likelihood ratio can be computed even if both sets are represented *implicitly* by pairs of (centroid, cardinality):

$$\text{LR}(\mathbf{X}_e, \mathbf{X}_t) \equiv \text{LR}((\mathbf{c}_e, N), (\mathbf{c}_t, M)), \quad (6)$$

where \mathbf{c}_e and \mathbf{c}_t are the centroids of \mathbf{X}_e and \mathbf{X}_t , respectively. The counts, N and M , are known as the *zero-order statistics*.

It worth noting that embeddings averaging is a commonly used strategy for speaker verification in the case of multiple enrollment utterances [21]. And PLDA can be also seen as an instance of ‘centroid based’ scoring, with a difference that the information about the set cardinality is preserved. Intuitively, the information about set cardinality allows to make judgments about uncertainty: the more vectors are in the set, the higher confidence of the mean estimate. The PLDA model has advantage of taking this information into account.

Let us now recall the centroid algorithm where the speakers are represented by centroids, and the centroids are updated by the exponential smoothing rule:

$$\mathbf{c}^{(t)} = \alpha \mathbf{x}^{(t)} + (1 - \alpha) \mathbf{c}^{(t-1)}.$$

By direct substitution of the update rule back into itself, one can see that it corresponds to the weighted averaging:

$$\begin{aligned} \mathbf{c}^{(N)} &= \sum_{i=1}^N p_i \mathbf{x}^{(i)} \\ p_i &= \begin{cases} (1 - \alpha)^{N-1} & \text{if } i = 1 \\ \alpha(1 - \alpha)^{(N-i)} & \text{if } i > 1 \end{cases} \end{aligned}$$

with $\mathbf{c}^{(0)} = \mathbf{x}_1$.

Now, if one wants to use the updated centroid for verification against the test utterance \mathbf{x}_{test} , it is not immediately obvious how to define the zero-order statistic n in $\text{LR}((\mathbf{c}^{(N)}, n), \mathbf{x}_{\text{test}})$, given by (6). For example, updating a centroid with a very small $\alpha \approx 0$ does not benefit from the information available after collecting several new data points. Therefore, setting n with the number of collected embeddings would not be adequate.

Here we propose a heuristic to set the count parameter in the case of unequal weights. First, let us notice that if $\alpha = 0$ or $\alpha = 1$, this boils down to the single enrollment verification, that is, the input count is 1. This corresponds to either $p_1 = 1$ or $p_N = 1$ with other weights being zero. Also, if $p_i = 1/N$ for all $i = 1 \dots N$, then the cardinality is N .

Noting that the entropy of the discrete uniform distribution $\mathbf{p} = [p_1, p_2, \dots, p_N]$ where $p_i = 1/N$ equals to $\log N$, we calculate the zero-order statistic as

$$n = e^{\mathcal{H}(\mathbf{p})}.$$

where $\mathcal{H}(\mathbf{p}) = -\sum_i p_i \log p_i$ is the entropy. Then, for the maximum entropy, achieved for the uniform distribution, $n = N$. At the same time, if one of the weights equals to 1, then $n = 1$, which is also consistent with the special case.

Another issue related to setting the zero-order statistic is discussed in studies [48]–[50] in the context of compensating apparently inadequate assumption of feature independence. The proposed solutions suggest to scale down zero-order statistics $n_r = r \cdot n$ with the recommended factor of $r = 0.2 - 0.3$ [13], [48] or simply setting $n = 1$ independently of the set cardinality [21]. The latter heuristic is nothing else than a commonly adopted recipe of embedding averaging.

10.3. VB clustering details

We assume that the extracted utterance embeddings were generated by the PLDA model with the full-rank speaker subspace in the form of two-covariance model [14]. This model, specified by (2), is parameterized by two matrices \mathbf{B} and \mathbf{W} interpreted as between- and within-speaker covariances. We assume the data have zero-mean.

We construct a finite mixture model for partitioning a set of embeddings $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ into K clusters corresponding to individual speakers. The speakers are represented by Gaussians with a shared covariance matrix \mathbf{W} and speaker-specific latent mean vectors $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_K\}$.

Further, each embedding \mathbf{x}_i is associated with a binary one-hot vector \mathbf{z}_i , whose components indicate whether this observation belongs to the corresponding mixture component. For instance, if the embedding \mathbf{x}_i is generated from component k then $z_{i,k} = 1$ with other vector components equal to zero.

The latent variables $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ are drawn independently from the categorical distributions governed by the mixing coefficients π_k :

$$p(\mathbf{Z}) = \prod_{i=1}^N \prod_{k=1}^K \pi_k^{z_{i,k}}$$

Conditioned on \mathbf{Z} , the embeddings are assumed to be independently drawn from a speaker-specific Gaussian distribution so that

$$p(\mathbf{X} | \mathbf{Y}, \mathbf{Z}) = \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(\mathbf{x}_i | \mathbf{y}_k, \mathbf{W})^{z_{i,k}}$$

We further modified the standard mixture model described above by introducing one extra label to present ‘background class’. To this end, we added one component to the mixture model that represents the marginal distribution of embeddings:

$$p(\mathbf{x}) = \int p(\mathbf{x} | \mathbf{y}) p(\mathbf{y}) d\mathbf{y} = \mathcal{N}(\mathbf{x} | \mathbf{0}, \mathbf{B} + \mathbf{W})$$

This component, indexed as component $K + 1$, corresponds to the embeddings which does not belong to any of K clusters. In the context of household speaker recognition this component is intended to model all the unknown speakers (guests) at once, since discriminating those speakers is irrelevant to the task.

$$p(\mathbf{x}_i | \mathbf{Y}, \mathbf{z}_i) = \begin{cases} \mathcal{N}(\mathbf{x}_i | \mathbf{0}, \mathbf{B} + \mathbf{W}) & \text{if } z_{i,K+1} = 1 \\ \mathcal{N}(\mathbf{x}_i | \mathbf{y}_k, \mathbf{W})^{z_{i,k}} & \text{otherwise} \end{cases}$$

The clustering problem requires finding the most likely partition of the data, represented by \mathbf{Z} . Following [13], we use

variational Bayesian inference [34] to approximate the posterior distribution $p(\mathbf{Z}|\mathbf{X}) \approx \prod_i q(\mathbf{z}_i)$. Similar to the semi-supervised k-means clustering described above, posterior distribution $q(\mathbf{z})$ for a subset of labeled data points are initialized by one-hot representations of their known labels and kept fixed during inference. After the convergence, the inferred cluster assignments can be found as $k_i^* = \arg \max_k q(\mathbf{z}_{i,k})$. Detailed derivations and the update rules can be found in [13].

Finally, the joint distribution of all of the random variables is given by

$$p(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = p(\mathbf{X}|\mathbf{Y}, \mathbf{Z})p(\mathbf{Y})p(\mathbf{Z}).$$

The clustering problem requires finding the most likely partition of the data, represented by \mathbf{Z} . This can be formulated as finding the most likely partition $\mathbf{Z}^* = \arg \max_{\mathbf{Z}} p(\mathbf{Z}|\mathbf{X})$ where $p(\mathbf{Z}|\mathbf{X}) = \int p(\mathbf{Z}|\mathbf{Y}, \mathbf{X}) d\mathbf{Y}$ is the clustering (partition) posterior distribution. Since direct optimization of this objective is intractable due to vast combinatorial search space, one has to resort to available approximation techniques. We use variational Bayesian inference [34] to approximate the joint posterior $q(\mathbf{Z}, \mathbf{Y}) \approx p(\mathbf{Z}, \mathbf{Y}|\mathbf{X})$ assuming that the approximate posterior is factorizes as $q(\mathbf{Z}, \mathbf{Y}) = q(\mathbf{Z})q(\mathbf{Y})$. We use variational inference because is often much faster than Monte-Carlo based [51] or greedy algorithms [52].

Following [13], we search for such $q(\mathbf{Z})$ and $q(\mathbf{Y})$ that maximize the weighted *evidence lower bound* (ELBO) augmented with a pair of scaling factors F_A and F_B :

$$\begin{aligned} \mathcal{L}(q(\mathbf{Z}), q(\mathbf{Y})) = & F_A \cdot \mathbb{E}_{q(\mathbf{Z})q(\mathbf{Y})} [\log p(\mathbf{X}|\mathbf{Y}, \mathbf{Z})] \\ & - F_B \cdot \text{KL}(q(\mathbf{Y})||p(\mathbf{Y})) - \text{KL}(q(\mathbf{Z})||p(\mathbf{Z})). \end{aligned}$$

Here, $\mathbb{E}_q[\cdot]$ denotes expectation with respect to q and $\text{KL}(\cdot||\cdot)$ is the Kullback-Leibler divergence [34]. Despite that the theoretically correct variational Bayesian inference assumes $F_A = F_B = 1$, tuning the values of these scaling factors helps to compensate for wrong model assumptions of statistical independence and to improve the clustering performance. For a more in-depth discussion of their interpretations, please refer to [13].

The assumption of factorized posterior leads to the inference algorithm consisting of iterative updates of factors $q(\mathbf{Z})$ and $q(\mathbf{Y})$, with each update having a closed-form solution. Similar to the semi-supervised k-means clustering described above, posterior distribution $q(\mathbf{z})$ for a subset of labeled data points are initialized by one-hot representations of their known labels and kept fixed during inference. After the convergence, the inferred cluster assignments can be found as $k_i^* = \arg \max_k q(\mathbf{z}_{i,k})$. Detailed derivations and the update rules can be found in [13].

11. References

- [1] L. Chen, V. Ravichandran, and A. Stolcke, "Graph-based label propagation for semi-supervised speaker identification," *CoRR*, vol. abs/2106.08207, 2021.
- [2] Z. Tan, Y. Yang, E. Han, *et al.*, "Improving speaker identification for shared devices by adapting embeddings to speaker subsets," in *ASRU*, 2021, pp. 1124–1131.
- [3] C. S. Greenberg, L. P. Mason, S. O. Sadjadi, *et al.*, "Two decades of speaker recognition evaluation at the national institute of standards and technology," *Computer Speech and Language*, vol. 60, p. 101 032, 2020.
- [4] A. Nagrani, J. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Interspeech*, 2017, pp. 2616–2620.
- [5] J. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [6] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Interspeech*, 2020, pp. 3830–3834.
- [7] D. Snyder, D. Garcia-Romero, G. Sell, *et al.*, "Speaker recognition for multi-speaker conversations using X-vectors," in *ICASSP*, 2019, pp. 5796–5800.
- [8] B. Sun, J. Feng, and K. Saenko, "Return of frustratingly easy domain adaptation," in *AAAI*, 2016, pp. 2058–2065.
- [9] K. A. Lee, Q. Wang, and T. Koshinaka, "The CORAL+ algorithm for unsupervised domain adaptation of PLDA," in *ICASSP*, 2019, pp. 5821–5825.
- [10] M. J. Alam, G. Bhattacharya, and P. Kenny, "Speaker Verification in Mismatched Conditions with Frustratingly Easy Domain Adaptation," in *Odyssey*, 2018, pp. 176–180.
- [11] P.-M. Bousquet and M. Rouvier, "On Robustness of Unsupervised Domain Adaptation for Speaker Recognition," in *Interspeech*, 2019, pp. 2958–2962.
- [12] D. Zhou, O. Bousquet, T. N. Lal, *et al.*, "Learning with local and global consistency," in *NIPS*, 2003, pp. 321–328.
- [13] M. Díez, L. Burget, F. Landini, *et al.*, "Analysis of speaker diarization based on bayesian HMM with eigenvoice priors," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 355–368, 2020.
- [14] N. Brümmer and E. de Villiers, "The speaker partitioning problem," in *Odyssey*, 2010, p. 34.
- [15] A. Sizov, K.-A. Lee, and T. Kinnunen, "Unifying probabilistic linear discriminant analysis variants in biometric authentication," in *S+SSPR, Joint IAPR International Workshop*, ser. Lecture Notes in Computer Science, vol. 8621, Springer, 2014, pp. 464–475.
- [16] X. Wang, J. Yamagishi, M. Todisco, *et al.*, "ASVspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Computer Speech and Language*, vol. 64, p. 101 114, 2020.
- [17] P. S. A. M. Ariyaeeinia J. Fortuna and A. Malagaonkar, "Verification effectiveness in open-set speaker identification," in *ICVIP*, vol. 153, 2006, pp. 618–624.
- [18] R. W. M. Ng, X. Liu, and P. Swietojanski, "Teacher-student training for text-independent speaker recognition," in *IEEE SLT Workshop*, 2018, pp. 1044–1051.
- [19] S. Wang, Y. Yang, T. Wang, *et al.*, "Knowledge distillation for small foot-print deep speaker embedding," in *ICASSP*, 2019, pp. 6021–6025.
- [20] J. S. Chung, A. Nagrani, E. Coto, *et al.*, "VoxSRC 2019: The first VoxCeleb speaker recognition challenge," *CoRR*, vol. abs/1912.02522, 2019.
- [21] P. Rajan, A. Afanasyev, V. Hautamaki, *et al.*, "From single to multiple enrollment i-vectors: Practical PLDA scoring variants for speaker verification," *Digital Signal Processing*, vol. 31, pp. 93–101, 2014.
- [22] C. Zeng, X. Wang, E. Cooper, *et al.*, "Attention back-end for automatic speaker verification with multiple enrollment utterances," *CoRR*, vol. abs/2104.01541, 2021.
- [23] Z. Bai and X. Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [24] A. Nagrani, J. S. Chung, J. Huh, *et al.*, "VoxSRC 2020: The second VoxCeleb speaker recognition challenge," *CoRR*, vol. abs/2012.06867, 2020.
- [25] M. K. Nandwana, M. Lomnitz, C. Richey, *et al.*, "The VOICES from a Distance Challenge 2019: Analysis of Speaker Verification Results and Remaining Challenges," in *Odyssey*, 2020, pp. 165–170.

- [26] J. Deng, J. Guo, X. Niannan, *et al.*, “ArcFace: Additive angular margin loss for deep face recognition,” in *CVPR*, 2019.
- [27] J. S. Chung, J. Huh, S. Mun, *et al.*, “In defence of metric learning for speaker recognition,” in *Interspeech*, 2020, pp. 2977–2981.
- [28] D. Snyder, D. Garcia-Romero, G. Sell, *et al.*, “X-vectors: Robust DNN embeddings for speaker recognition,” in *ICASSP*, 2018, pp. 5329–5333.
- [29] H. Zeinali, S. Wang, A. Silnova, *et al.*, “BUT system description to VoxCeleb speaker recognition challenge 2019,” *CoRR*, vol. abs/1910.12592, 2019.
- [30] H. Aronowitz, “Compensating inter-dataset variability in PLDA hyper-parameters for robust speaker recognition,” in *Odyssey*, 2014.
- [31] R. G. Brown, *Smoothing, Forecasting and Prediction of Discrete Time Series*. Englewood-Cliffs, NJ: Prentice-Hall, 1963.
- [32] X. Zhu, “Semi-supervised learning literature survey,” Computer Sciences, University of Wisconsin-Madison, Tech. Rep. 1530, 2005.
- [33] A. K. Jain, M. N. Murty, and P. J. Flynn, *Data clustering: a review*, 3. New York, NY, USA: ACM, 1999, vol. 31, pp. 264–323.
- [34] C. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer Science+Business Media, LLC, 2006.
- [35] A. Corduneanu and C. Bishop, “Variational bayesian model selection for mixture distributions,” in *AISTATS*, 2001, pp. 27–34.
- [36] D. Zhou, O. Bousquet, T. N. Lal, *et al.*, “Learning with local and global consistency,” in *NIPS*, 2003, pp. 321–328.
- [37] K. Okabe, T. Koshinaka, and K. Shinoda, “Attentive statistics pooling for deep speaker embedding,” in *Interspeech*, 2018, pp. 2252–2256.
- [38] F. Wang, J. Cheng, W. Liu, *et al.*, “Additive margin softmax for face verification,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [39] M. Ravanelli, T. Parcollet, P. Plantinga, *et al.*, *SpeechBrain: A general-purpose speech toolkit*, 2021.
- [40] H. S. Heo, B.-J. Lee, J. Huh, *et al.*, “Clova baseline system for the VoxCeleb speaker recognition challenge 2020,” *arXiv preprint arXiv:2009.14153*, 2020.
- [41] K. He, X. Zhang, S. Ren, *et al.*, “Deep residual learning for image recognition,” in *CVPR*, 2016, pp. 770–778.
- [42] D. A. van Leeuwen, N. Brummer, and A. Swart, “A comparison of linear and non-linear calibrations for speaker recognition,” in *Odyssey*, 2014.
- [43] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, ser. Springer Series in Statistics. New York, NY, USA: Springer New York Inc., 2001.
- [44] N. Ryant, K. Church, C. Cieri, *et al.*, “The Second DIHARD Diarization Challenge: Dataset, Task, and Baselines,” in *Interspeech*, 2019, pp. 978–982.
- [45] L. Burget, O. Plchot, S. Cumani, *et al.*, “Discriminatively trained probabilistic linear discriminant analysis for speaker verification,” in *ICASSP*, 2011, pp. 4832–4835.
- [46] J. Rohdin, S. Biswas, and K. Shinoda, “Discriminative PLDA training with application-specific loss functions for speaker verification,” in *Odyssey*, 2014, pp. 26–32.
- [47] J. Villalba, N. Brümmer, and N. Dehak, “Tied variational autoencoder backends for i-vector speaker recognition,” in *Interspeech*, 2017, pp. 1004–1008.
- [48] T. Stafylakis, P. Kenny, V. Gupta, *et al.*, “Compensation for inter-frame correlations in speaker diarization and recognition,” in *ICASSP*, 2013, pp. 7731–7735.
- [49] A. McCree, G. Sell, and D. Garcia-Romero, “Extended variability modeling and unsupervised adaptation for PLDA speaker recognition,” in *Interspeech*, 2017, pp. 1552–1556.
- [50] —, “Speaker diarization using leave-one-out gaussian PLDA clustering of DNN embeddings,” in *Interspeech*, 2019, pp. 381–385.
- [51] C. Andrieu, N. de Freitas, A. Doucet, *et al.*, “An introduction to MCMC for machine learning,” *Mach. Learn.*, vol. 50, no. 1-2, pp. 5–43, 2003.
- [52] A. Silnova, N. Brummer, J. Rohdin, *et al.*, “Probabilistic embeddings for speaker diarization,” in *Odyssey*, 2020, pp. 24–31.