



HAL
open science

Cross-corpora spoken language identification with domain diversification and generalization

Spandan Dey, Md Sahidullah, Goutam Saha

► **To cite this version:**

Spandan Dey, Md Sahidullah, Goutam Saha. Cross-corpora spoken language identification with domain diversification and generalization. *Computer Speech and Language*, 2023, 81 (June 2023), pp.101489. <10.1016/j.csl.2023.101489>. <hal-03984643>

HAL Id: hal-03984643

<https://inria.hal.science/hal-03984643v1>

Submitted on 13 Feb 2023

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.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Cross-Corpora Spoken Language Identification with Domain Diversification and Generalization

Spandan Dey^{a,*}, Md Sahidullah^{b,**}, Goutam Saha^a

^a*Department of Electronics & Electrical Communication Engineering, Indian Institute of Technology Kharagpur, Kharagpur, India-721302*

^b*Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France*

Abstract

This work addresses the cross-corpora generalization issue for the low-resourced spoken language identification (LID) problem. We have conducted the experiments in the context of Indian LID and identified strikingly poor cross-corpora generalization due to corpora-dependent non-lingual biases. Our contribution to this work is twofold. First, we propose *domain diversification*, which diversifies the limited training data using different audio data augmentation methods. We then propose the concept of maximally diversity-aware cascaded augmentations and optimize the augmentation fold-factor for effective diversification of the training data. Second, we introduce the idea of *domain generalization* considering the augmentation methods as pseudo-domains. Towards this, we investigate both domain-invariant and domain-aware approaches. Our LID system is based on the state-of-the-art *emphasized channel attention, propagation, and aggregation based time delay neural network* (ECAPA-TDNN) architecture. We have conducted extensive experiments with three widely used corpora for Indian LID research. In addition, we conduct a final blind evaluation of our proposed methods on the Indian subset of VoxLingua107 corpus collected in the wild. Our experiments demonstrate that the proposed domain diversification is more promising over commonly used simple augmentation methods. The study also reveals that domain generalization is a more effective solution than domain diversification. We also notice that domain-aware learning performs better for same-corpora LID, whereas domain-invariant learning is more suitable for cross-corpora generalization. Compared to basic ECAPA-TDNN, its proposed domain-invariant extensions improve the cross-corpora EER up to 5.23%. In contrast, the proposed domain-aware extensions also improve performance for same-corpora test scenarios.

Keywords: Language identification (LID), cross-corpora evaluation, audio augmentation, domain generalization, domain adversarial training, multitask learning.

*Corresponding author.

**Md Sahidullah is presently an Independent Researcher and Consultant.

Email addresses: sd21@iitkgp.ac.in (Spandan Dey), md.sahidullah@inria.fr (Md Sahidullah), gsaha@ece.iitkgp.ac.in (Goutam Saha)

1. Introduction

In the era of voice assistants, smart homes, and numerous other innovative gadgets, the speech-based *human-computer interaction* (HCI) is becoming a part of our daily life. Across the globe, the number of users relying upon modern-day speech applications, such as *automatic speech recognition* (ASR), *speaker recognition and verification*, *speech synthesis*, *speech emotion recognition* (SER), will keep on increasing [14]. For ease of use, the speech applications should operate efficiently on multiple input languages. A front-end *language identification* (LID) block is generally used for multilingual applicability [69]. The LID system predicts the language from the input, and the following speech application then adapts the mode of operation accordingly [42].

For efficient real-world deployment of the speech applications, improving the generalization of the front-end LID module is important. For LID systems, the generalized classifier should be robust against several non-lingual sources, such as speaker identity, gender, age, dialects, and accents, mismatches due to channel and background environments [17]. We can assume that diversity in non-lingual effects is expected to increase in larger speech corpora with greater diversity in data collection settings. Such corpora may improve the generalization by reducing the mismatch with unknown test utterances [83]. Therefore, the availability of a large and diversified corpus is important for developing robust LID systems.

Speech corpora collected through the initiatives, such as *NIST LRE* [31, 74] and *OLR challenges* [44, 84], consist of large-scale multilingual speech corpora with carefully verified ground truths. These corpora are extensively used to develop and evaluate LID algorithms using deep neural network (DNN) architectures [30, 49, 57], which substantially advanced LID research in the recent decade. However, many of the widely spoken languages in the world are still low-resourced due to the lack of standard large-scale corpora. The LID research with those low-resource languages is commonly conducted with independently created small-scale in-house corpora, which have limited non-lingual variations [18]. As a consequence, training state-of-the-art DNN-based LID systems using such small corpus is prone to overfitting. To address this issue, researchers have experimented with different approaches, such as applying data augmentation, feature selection, regularizers, cross-validation, or model ensembling [17, 70, 78]. Still, achieving generalization for the LID systems with low-resource languages remains the key challenge [17].

In this work, we aim to improve the generalization of the low-resourced LID systems under the *cross-corpora* framework. In LID literature, most work considers the performance by training system on the train set and evaluating it on the dedicated evaluation set from the same corpus. However, for the low-resourced corpora, the data diversity within the dataset can be limited. As a consequence, even a decent performance on its own evaluation set may question its generalization on evaluations with different data. Hence, we explore cross-corpora evaluations to assess the generalization ability of the LID system. For each corpus, we independently train LID models and evaluate them with the test subsets for all the available corpora. Our preliminary investigation [17] with different low-resourced LID corpora revealed that the cross-corpora LID performances are strikingly inferior compared to the same-corpora performances. In that study, we used vanilla x-vector time delay neural network (TDNN) architecture [82] using *mel frequency cepstral coefficients* (MFCC) features. The current work utilizes state-of-the-art *emphasized channel attention, propagation, and aggregation in TDNN* (ECAPA-TDNN) architecture [16] and advances this further for better generalization.

In this study, our key focus is to address the cross-corpora generalization by reducing the

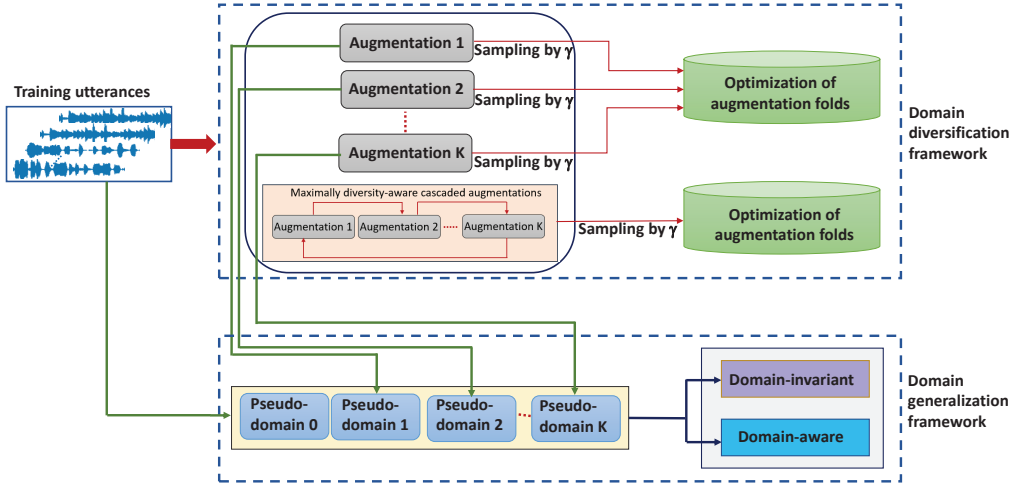


Fig. 1. Overview of the proposed frameworks for improved cross-corpora generalization in LID. Here, γ denotes the augmentation fold-factor (explained in Section 5.3 and K denotes the total number of signal-level augmentation categories.

performance mismatch with the same-corpora evaluation. We propose novel domain diversification strategies and further frame the cross-corpora LID problem by introducing the *domain generalization* (DG) concepts [10] in this field. One possible way to address cross-corpora performance issues is by following the *domain adaptation* (DA) [63, 100] approaches. In the DA framework, the training corpus is used as the source domain, and we have access to limited labeled or unlabeled data from the evaluation condition, which we refer to as the target domain. In the cross-corpora scenario, similarly, the training-set of one corpus can be considered as the source domain, and the training-set utterances from the remaining corpora can be considered as individual target domains. However, applying DA methods in a straightforward way is not possible in some practical scenarios for two reasons. Firstly, we may not have any specific target domain audio data at all while training the LID systems. Secondly, even if we have some target domain data (which can be the other corpora except the training corpus in our problem), applying DA can improve the generalization on that specific target corpus, and the generalization to unseen conditions will remain a key concern. Therefore, we apply the concepts of DG in this work. This is a more recent and practical solution, and it does not require specific target domain data [103]. Rather, in the DG framework, multiple pseudo-target domains are created from the source domain, and the model is trained to achieve better generalization across multiple pseudo-domains. We create different pseudo-domains following diverse domain diversification methods. These pseudo-domains are then utilized to develop multiple DG approaches. We propose domain-invariant and domain-aware DG approaches for learning language-discriminating knowledge by extending the state-of-the-art ECAPA-TDNN architecture.

In Fig. 1, we have presented the overview of our proposed solutions towards better generalization of LID systems. Fig. 1 presents the main two approaches that we have explored: domain diversification and domain generalization.

The major contributions of this work are summarized as follows:

- We explore several audio data augmentations, including commonly used methods as well as several new ways of increasing the data diversity. The comprehensive study of different

augmentation techniques could also be useful for other speech-based classification tasks.

- To maintain a trade-off between the amount of original data and augmented data, we conduct a systematic study to analyze the impact of augmentation *fold-factor* for the generalization of TDNN-based LID classifiers. The analysis also provides an optimal augmentation fold-factor in terms of computation complexity and generalization performance.
- We further increase the diversity of the low-resourced corpora by proposing the *maximally diversity-aware cascaded augmentation*. It generates more realistic augmented audio data based on maximum mutual domain diversity measures while preserving the perception and intelligibility of the spoken language.
- Considering the different domain-diversified audio categories as multiple pseudo-domains, we introduce the concept of DG for improving the cross-corpora performance. We extend the state-of-the-art ECAPA-TDNN architecture by introducing two opposite but functionally similar DG approaches suitable for the unknown target domains. The first one is domain-invariant approach, where the LID model is trained to get confused in the task of identifying different pseudo-domain. In the domain-invariant approach, we explore adversarial training and *maximum mean discrepancy* (MMD) approaches. In the second domain-aware approach, along with learning the language discriminating cues, we train the LID model to recognize the different pseudo-domains. We explore multitask learning with pseudo-domain classification as an additional task in the domain-aware approach.
- We further extend the conventional MMD-based DG approach by proposing the *multi-domain* MMD (MD-MMD) algorithm. Instead of considering the pseudo-domains altogether as a single target domain, the proposed method computes diversity-weighted multi-domain MMD losses and jointly optimizes the combined MD-MMD losses.
- We conduct systematic benchmarking and analysis of the same-corpora and cross-corpora LID performances with three independently collected LID corpora in the Indian languages. Finally, we perform a blind evaluation of the algorithms in a fourth unseen corpus.

The rest of the paper is organized as follows: Section 2 formally introduces the cross-corpora evaluation paradigm. Section 3 describes the corpora and the experimental methodology used for this work. The proposed domain diversification techniques for improving the cross-corpora performance are discussed in Section 4. The cross-corpora problem is then addressed from domain generalization approaches in Section 5. In Section 6, we present the experimental results. We provide additional analysis of the cross-corpora problem, including blind evaluations on an unseen in-the-wild corpus in Section 7. Finally, we conclude the paper in Section 8 by providing a summary and recommendations.

2. Problem definition & related works

In this section, we have formally introduced our cross-corpora evaluation protocol and discussed its applications in various speech-based classification tasks.

2.1. Formulation of cross-corpora evaluation

In the cross-corpora evaluation paradigm, the model is trained on one corpus and tested on multiple corpora. If the test samples are taken from the hold-out part of the training corpus, then it is called *same-corpora* evaluation. When the test samples are taken from other corpora, it is called *cross-corpora* evaluation. Next, we present a formal definition of the cross-corpora evaluation followed by the perspectives and constraints.

Definition: Let us consider a cross-corpora experimental setting with C number of corpora. Let the data of the i -th corpus be denoted as $\{(\mathbf{X}_k^i, y_k^i)\}_{k=1}^{N^i}$, where $i \in [1, C]$ and N^i denoting the number of samples in corpus i . \mathbf{X}_k is the features of the k -th utterance of the i -th corpus, and y_k is the corresponding label.

For each corpus i , the feature space $\mathcal{X}^i = [\mathbf{X}_1^i, \mathbf{X}_2^i, \dots, \mathbf{X}_{N^i}^i]$ is partitioned into disjoint training and evaluation parts \mathcal{X}_T^i and \mathcal{X}_E^i . The corresponding label space $\mathcal{Y}^i = [y_1^i, y_2^i, \dots, y_{N^i}^i] \in \{1, 2, \dots, L\}$ (L denotes the total number of target classes) is also partitioned as \mathcal{Y}_T^i and \mathcal{Y}_E^i respectively.

Let us consider a mapping function $f(\mathbf{X}; \mathbf{W}) : \mathbf{X} \rightarrow \mathbf{y}$ with \mathbf{W} denoting the learnable parameter of the mapping function. Associated with f , the corresponding multi-class classification loss function is denoted as $\mathcal{L}(\mathbf{y}, f(\mathbf{X}; \mathbf{W}))$, which needs to be minimized during training.

Consider, a model is trained with the training set from the i -th corpus, $(\mathcal{X}_T^i, \mathcal{Y}_T^i)$ and is tested with the evaluation set from the j -th corpus, \mathcal{X}_E^j . For the test utterances, the posterior prediction probability vector for this train-test pair is denoted as $p(\mathcal{Y}_E^j | \mathcal{X}_E^j; \mathbf{W}^i)$, where \mathbf{W}^i denotes the learned parameters trained on \mathcal{X}_T^i . For $j = i$, the evaluation follows the *same-corpora* configuration. Our key interest is in the *cross-corpora* conditions, i.e., when $j \neq i$. We then consider a performance metric computation function g , that takes the true class labels (\mathcal{Y}) and predicted labels ($p(\mathcal{Y} | \mathcal{X}; \mathbf{W})$) as inputs and computes a scalar value indicating the classification performance.

Consider, $(\mathcal{X}_T^i, \mathcal{Y}_T^i)$ is used to train a standalone LID model (f^i) trained with corpus i . The key objective of cross-corpora evaluation is to assess and minimize the performance mismatch by f^i when it is evaluated on \mathcal{X}_E^i and \mathcal{X}_E^j , respectively. The cross-corpora performance mismatch (for corpora i and j) is formulated as:

$$D_g(i, j) = \left| g(\mathcal{Y}_E^i, p(\mathcal{Y}_E^i | \mathcal{X}_E^i; \mathbf{W}^i)) - g(\mathcal{Y}_E^j, p(\mathcal{Y}_E^j | \mathcal{X}_E^j; \mathbf{W}^i)) \right| \quad (1)$$

The aim of this work is to minimize the performance mismatch ($D_g(i, j), \forall \{i, j\} \in [1, C]$) without degrading the same-corpora performance. We can achieve this by estimating the parameters of the LID model (\mathbf{W}^i) such that the following condition is finally satisfied.

$$\min_{\mathbf{W}^i | \mathcal{X}_T^i} \mathcal{L}(\mathcal{Y}_E^i, f(\mathcal{X}_E^i; \mathbf{W}^i)) \text{ s.t } D_g(i, j) \rightarrow 0 \quad (2)$$

Eq. 2 is needed to be ensured $\forall \{i, j\} \ni j \neq i$.

Perspectives & constraints: While deploying the trained standalone models in the real-world, test samples can be from any unseen domains. Therefore, generalizing over specific target domains may not be the most efficient solution for such cases. Our goal is to improve generalization over unseen domains with the evaluation performance on multiple corpora as a reference. Hence, apart from the training corpus, utterances from the remaining corpora is not utilized as target domains. In our evaluation framework, we apply the following constraints considering the challenges in the real-world application of the LID systems.

- Independent models (f^i) are trained using only the training data from a single corpus ($\mathcal{X}_T^i, \mathcal{Y}_T^i$), where $i \in [1, C]$. Utterances from the training samples of other corpora $\mathcal{X}_T^j \ni j \neq i$ are not used during the system development.
- For improving cross-corpora performance mismatch between train corpus i and evaluation corpus j , we prohibit using the feedback from the same-corpora evaluation of the LID model trained with corpus j . This is to ensure the generalization are not biased towards corpus j .
- For a fair performance comparison, we follow a closed-set LID problem with a fixed set of languages. All the corpora share the common label space: $\mathcal{Y} = \{\cap \mathcal{Y}^i\}_{i=1}^C$.

One may consider the cross-corpora study as an application of the commonly practiced domain adaptation (DA) methods used in the machine learning literature [63, 100]. Following Eq. 2, under the DA formulation of the cross-corpora study, corpora i and j can be considered as source and target domains, respectively. However, there is a fundamental difference between the objective of our cross-corpora generalization and the goal of the DA approaches. DA is effective for scenarios when the target data is known prior. However, we are interested in improving LID generalization during real-world deployments where the test data is unknown and can come from diverse non-lingual domains. For such *out-of-domain* generalization, the formulation of our cross-corpora evaluation task is more appropriate. Unlike DA, we do not use subset of utterances from corpus j as target domain data to compute \mathbf{W}^i . Rather, we follow domain generalization (DG) [103] (which is fundamentally similar to DA), where instead of using specific target domains, we create multiple pseudo-domains from the source corpus i using several domain diversification methods. The pseudo-domain data are then used to compute \mathbf{W}^i .

2.2. Related works on cross-corpora evaluations

The majority of the research in cross-corpora evaluation is so far conducted for computer vision applications, including object detection [39, 15], facial expression recognition [104], deepfake detector [61], person re-identification [94]. In speech processing applications, cross-corpora evaluations are often adopted to assess the robustness, especially for tasks, such as emotion recognition and speaker recognition anti-spoofing, where dataset variability is limited due to the constraints in data collection.

In the speech emotion recognition (SER) field, the study reported in [77] performed cross-corpora SER with speaker-wise feature normalization using six different corpora. Cross-corpora SER performance is severely degraded except when the test data has background similarities with the train data. In [102], the authors used six emotional speech corpora and improved cross-corpora performance by adding unlabeled emotional utterances to the agglomerated training data. Several other works in SER also investigated cross-corpora evaluations [91, 90, 86, 79, 26].

In speaker recognition anti-spoofing, the work in [67] conducted cross-corpora experiments with two corpora, ASVspoof2015 and BTAS2016, which covered common spoofing attacks, speech synthesis, and voice conversion. This work demonstrated poor generalization in different data conditions. In [40], the authors used two anti-spoofing corpora, ASVspoof and AVspoof, and improved cross-corpora generalization by combining multiple systems with score fusion. Another recent study with seven seven anti-spoofing corpora demonstrated different levels of generalization with various classifiers [13]. The study further correlated cross-corpora performance with different signal-level descriptors.

Cross-corpora evaluations were also explored to detect the presence of dysarthria in speech [29], music mood prediction [35], multi-modal voice activity detection (VAD) [5], speech enhancement [64], and speech recognition [88]. In all cases, the authors demonstrated poor generalization due to large mismatch in data conditions.

Studies with cross-corpora evaluations are limited in language recognition research compared to other speech-related tasks discussed above. A few studies included cross-corpora experimental results even though it was not explicitly mentioned in the papers [2, 23, 45, 46, 47, 89]. The work in [89] trained the LID systems on subsets of VoxLingua107 and evaluated them on KALAKA-3 and LRE07. They evaluated the robustness of different statistical and neural network models for cross-corpora LID. However, they did not conduct additional modifications to improve the cross-corpora generalization. Another recent work [23] proposed a discriminative hierarchical model for language recognition and used a combination of NIST LRE, RATS, NIST SRE, VoxLingua107, Panarabic, CALLHOME and CALLFRIEND databases for training and a combination of LRE15, LRE17, LASRS, KALAKA and Crowdscore (CROWD) databases for evaluation. The aim of this work was to develop a discriminative probabilistic linear discriminant analysis (PLDA) framework as a better alternative to the commonly used generative modeling. The study in [45] incorporated hierarchical phonotactical information using multitask learning in LID. They evaluated their proposed algorithms on the NIST LRE17 and AP17-OLR databases independently. The LRE17 evaluation utterances are sampled from narrow-band MLS14 and wide-band VAST databases, which were not used in training and development in LRE17. In [46], the authors separately used the NIST LRE17 and AP17-OLR corpora and removed the redundancies in the pre-trained multilingual wav2vec2.0 XLSR model based speech representations by using squeeze-excitation (SE) and linear bottleneck modules. In [2], the authors used a custom training set for the unconstrained LID task of the OLR21 challenge. They combined training utterances of OLR21 with the Mozilla Common Voice [59] data. The authors then used the pre-trained wav2vec 2.0 model and fine-tuned it using the custom training set.

Most of the above works violated the cross-corpora evaluation standards by mixing data from multiple corpora, and their key objectives were not focused on explicitly reducing the cross-corpora performance mismatch. Moreover, they mainly focus on major languages with rich linguistic resources. Our current work focuses on improving cross-corpora language recognition using databases that are developed independently and contain limited resources.

3. Experimental setup

This section describes our experimental setup, including the details of speech corpora, feature extraction, and neural network architecture.

3.1. Description of corpora

We consider three corpora that are most widely used in the Indian LID literature. They are: *IIITH-ILSC* (IIITH) [92], *LDC 2017S14* (LDC) [38], and *IITKGP-MLILSC* (KGP) [51]. For our experiments, we have selected the five languages, *Bengali, Hindi, Punjabi, Tamil, and Urdu*, which are common to all three corpora. In Table 1, we have summarized the meta-data of the three corpora.

The meta-data comparison in Table 1 lists the differences and similarities among the three corpora. The duration of speech indicates that these are low-resource corpora compared to the LID corpora in major languages. Out of these three, the KGP is the smallest one, with just four

Table 1: Meta-data comparison of the three speech corpora used in our experiments. We have shown the collective information for all the common subsets of the five languages considered for this study. (BN: Broadcast news, CTS: Conversational telephone speech.)

Meta-data	IITH	LDC	KGP
Duration	15 hours	118 hours	4 hours
Speakers	250	378	67
Utterances	12,663	577	155
Mode of speech	BN and CTS	CTS	BN
Accent & dialect	Mostly standard	Variations exist	Standard
Environment	Studio, real-world	Real-world	Studio
Gender ratio	Balanced	Not maintained	Balanced
Audio channel	Mono	Stereo	Mono
Audio format	16 kHz (.wav)	8 kHz (.flac)	8 kHz (.wav)

hours of speech data, whereas the LDC is the largest one, with about 24 hours of audio data per language. IITH and KGP data are mostly collected from broadcast news (BN). The LDC data contains conversational telephone speech (CTS) recordings from the NIST LRE 2011 dataset. The accents and dialects are also varied across the corpora. For example, the Bengali utterances present in the IITH and KGP corpora contain dialects spoken mostly in India, whereas the LDC corpus contains dialects spoken in different parts of Bangladesh.

The IITH and KGP corpora have dedicated evaluation protocols including data-split for training and test set. For LDC, there was no prior partitioning of the data. We randomly sample 80% of the audio recordings as training data and the remaining are used for test. LDC utterances are stereo audio with each channel corresponding to one of the two speakers in the telephone call. For each language, one speaker is common in all the telephone call recordings. To ensure speaker disjoint partitions, we only consider the audio channel excluding the common speaker in the LDC test set. For all corpora, the training utterances are further partitioned with an 80 : 20 ratio into speaker-disjoint training and validation sets.

3.2. Data pre-processing & feature extraction

Firstly, we process the corpora to have a consistent sampling rate of 8 kHz, and an audio format of 16-bit PCM (.wav). Then we remove the non-speech regions using an energy-based voice activity detector (VAD) [3]. We finally compute 20-dimensional *mel frequency cepstral coefficients* (MFCCs) with 20 filters, and this is followed by utterance-level cepstral mean subtraction (CMS). For MFCC extraction, we use Hamming window of 25 ms and hop-length of 10 ms. While the *mel filter banks* (MFB) feature is frequently used with the neural network based classifiers, several works [37, 41, 81] have also applied MFCCs as input to the TDNN-based classifiers as used in our work. Note that the existing work with the three used corpora [92, 53, 52, 11] also reported their performances mostly using MFCCs. We segment the training speech utterances with an active speech of 3 s duration to create the training examples. For utterances with a duration of more than 3 s, we segment them into multiple chunks and discard the remaining portion. If the utterance duration is less than 3 s (which is very rare), we discard the utterance. The test utterances are also similarly segmented with a duration of 3 s.

3.3. Architecture description

Baseline TDNN: For the baseline reference, we use the *x-vector TDNN* [82] architecture trained on MFCC (with CMS post-processing) feature following our preliminary investigation in [17]. The model consists of five TDNN layers followed by a statistics pooling (Stat pooling) layer and two fully connected (FC) layers [81]. Table 2 shows the architectural details.

Table 2: Architecture details of the baseline x-vector TDNN with input feature dimension F , temporal dimension T , and the number of classes L .

Layer	Context	Input dimension	Output dimension
TDNN-ReLU	$[t - 2 : t + 2]$	$F * 5 * T$	$512 * T$
TDNN-ReLU	$\{t - 2, t, t + 2\}$	$512 * 3 * T$	$512 * T$
TDNN-ReLU	$\{t - 3, t, t + 3\}$	$512 * 3 * T$	$512 * T$
TDNN-ReLU	$\{t\}$	$512 * T$	$512 * T$
TDNN-ReLU	$\{t\}$	$512 * T$	$1500 * T$
Stat pooling	T	$1500 * T * 1$	$3000 * 1$
FC-ReLU	-	$3000 * 1$	$512 * 1$
FC-ReLU	-	$512 * 1$	$512 * 1$
FC-Softmax	-	$512 * 1$	$L * 1$

ECAPA-TDNN: We have utilized the ECAPA-TDNN architecture [16], which outperformed other basic TDNN architectures for speaker verification tasks in [16]. This is also very recently adopted for language recognition tasks [37, 85]. The ECAPA-TDNN model modifies the basic x-vector architecture in different ways. Instead of the TDNN layers, three squeeze-excitation (SE) based residual blocks (SE-Res2) are used. Multi-layer feature aggregation (MFA) is performed with the outputs of these three blocks. The aggregated output is then fed to a channel and context attentive pooling layer (Attn Pooling) followed by a fully connected layer. We use batch normalization (BN) for 1d-convolutional (Conv1D) layers and used a drop-out of 0.25 for the linear layers. Table 3 presents a detailed description of this architecture.

The TDNN architectures are implemented using the PyTorch library [66] of Python. We use a batch size of 32, AdamW optimizer [50], *reduce on plateau* based learning rate scheduler with an initial value of 0.001. We set the maximum epoch as 30 and apply early stopping criteria with five epochs on the validation loss. For the loss function, we have used the *angular margin* (AM) softmax [93] with scale parameter as 30 and margin as 0.2.

Table 3: Architecture details of the ECAPA-TDNN with input feature dimension F , temporal dimension T , and number of classes L . (Here, BN denotes batch normalization).

Layer	Skip connect	Input dimension	Output dimension
Conv1D+ReLU+BN	-	$F * T$	$512 * T$
SE-Res2-Block	-	$512 * T$	$512 * T$
SE-Res2-Block	-	$512 * T$	$512 * T$
SE-Res2-Block	-	$512 * T$	$512 * T$
Conv1D+ReLU	layer 2,3,4	$512 * 3 * T$	$1536 * T$
Attn Pooling+BN	-	$1536 * T * 1$	$3072 * 1$
FC-ReLU	-	$3072 * 1$	$192 * 1$
FC-Softmax	-	$192 * 1$	$L * 1$

The neural LID models can be employed in two different ways. The first is an end-to-end architecture that builds one single module with speech as input and language labels as output [30]. The other type trains an embedding extractor first, which is followed by a backend classifier based on PLDA or cosine similarity [81]. The end-to-end system is relatively simple and straightforward without the burden of training additional back-end for scoring. On the other hand, the additional backend-based approach is more flexible, especially when tackling out-of-set languages [54]. Our present work focuses on closed-set LID problems with fixed five languages, and for brevity, we consider end-to-end architecture.

3.4. Performance evaluation metrics

Following the literature, we use *equal error rate* (EER) [8] and *cost average* (C_{avg}) as performance evaluation metrics. EER is defined as the error rate at the decision threshold where the false acceptance and false rejection rates are equal. For multi-class scenarios, we compute one-vs-all EERs for each target language and report the average value. The C_{avg} is used as the primary evaluation metric in the NIST language recognition evaluations (LRE) [74] and OLR challenges [44]. It is defined as follows [74]:

$$C_{\text{avg}} = \frac{1}{N} \sum_{L_t} \left[p_{\text{Target}} \cdot p_{\text{Miss}}(L_t) + \sum_{L_n} p_{\text{Non-target}} \cdot p_{\text{FA}}(L_t, L_n) \right] \quad (3)$$

where L_t and L_n are the targets and non-target languages. p_{Miss} and p_{FA} are the probability of miss and false alarm, respectively. p_{Target} is the prior probability of the target languages, and it is set at 0.5. $p_{\text{Non-target}} = (1 - p_{\text{Target}})/(L - 1)$, where, L is the total number of languages. The lower value of EER and C_{avg} indicates better language recognition performance.

4. Cross-corpora LID: Domain diversification

Low-resourced LID corpora contain a substantially lesser number of utterances. Mostly the data collection source and configurations are very limited. So, the amount of non-lingual diversities in terms of varying backgrounds, recording instruments, room environments, and speaker populations are limited. These create corpora-specific non-lingual biases across the databases leading to inferior cross-corpora performances [17]. Therefore, this section explores extensive audio augmentation methods to diversify the low-resourced corpora. In the speech processing literature, audio augmentation techniques, such as noise and non-speech addition, volume, pitch, and speed perturbations, are commonly used to artificially increase the amount and variations in the training data [60, 75, 96]. For addressing the cross-corpora generalization problem, we focus on the extensive application of augmentation-based approaches for two reasons: (i) increasing the training data helps avoiding overfitting issues [81], and (ii) diversifying the training data can help reducing the corpora mismatch. To incorporate diverse sources of randomness in the training data, we explore several existing and new audio augmentation techniques. We further propose a new concept called *maximally diversity-aware cascaded augmentation* approach and optimize the *augmentation fold-factor*.

4.1. Description of the explored augmentation methods

Based on the manner of perturbing the audio data, we have classified the explored augmentation techniques into eight categories. We primarily focus on signal-level augmentation categories

Table 4: Summary of the different signal-level and feature-level augmentation categories and their respective sub-categories applied in the LID framework.

Additive non-speech augmentation (A^1)	Signal parameter perturbations (A^2)	Bandwidth augmentation (A^3)	Environmental variations (A^4)	Speech enhancement (A^5)	Audio encodings (A^6)	Restoration from lossy codecs (A^7)	Feature-level augmentation (A^8)
Babble Music Noise	Pitch Shift Speed Volume	Lower cutoff Upper cutoff Telephone- -channel BPF	Smart phone Live hall Radio broadcast Vinyl recording	Spectral subtraction MMSE-based De-reverberation	a-law ima-adpcm oki-adpcm u-law	.aac .gsm .mp3 .ogg .opus .wma	SpecAug Mixup

that operate directly on the raw waveform. Further, we have also explored feature-level augmentations. For each category, there are different sub-categories. The different augmentation categories are chosen based on the corpora-analysis conducted in our preliminary investigation [17], and they introduce diversities that help in reducing the corpora mismatch. We first studied the impact of each category on the LID performances. Then, we have pooled and sampled utterances from the different categories to accumulate the different ways of domain diversification they possess. We summarize the applied augmentation categories in Table 4. A detailed description of these augmentation categories is presented as follows.

4.1.1. Additive non-speech augmentation

The *long-term average spectrum* (LTAS) and signal-to-noise ratio (SNR) histogram analysis in our preliminary investigation indicated that the channel effects and backgrounds may vary across different low-resourced LID corpora [17]. Following this, non-speech effects, such as additive noise, babble noise, and music clips, are added to the training utterances. This augmentation category is followed from [82] and implemented with the Kaldi tool’s [68] VoxCeleb recipe¹. This uses non-speech samples from the MUSAN corpus [80].

4.1.2. Signal parameter perturbations

The utterances across the corpora can also vary based on the speaking style. For example, pitch, speaking rate, and loudness can be different across the speakers. To reduce the corpora mismatch in such perspectives, we randomly alter different signal parameters using the audio degradation toolbox [55].

- Pitch: Random shifted by $[-4, 4]$ semitones without changing the utterance lengths.
- Shift: Training utterances are randomly segmented into two parts and swapped.
- Speed: The speed of the playback is changed by randomly choosing $\Gamma \in [-15, +15]$. The length of the augmented audio (l) becomes $l' = ((100 - \Gamma)/100)l$. $\Gamma > 0$ increases playback speed and $\Gamma < 0$ reduces it [55].
- Volume: Random gains in the range of $[-30, 40]$ dB.

¹<https://github.com/kaldi-asr/kaldi/tree/master/egs/voxceleb/v2>

4.1.3. Bandwidth augmentation

Following the LTAS comparison in [17], we observe that the spectrums of different corpora may vary prominently in the different frequency ranges. Hence, we process the utterances with filters of random cutoff frequencies.

- Upper cutoff augmentation: Filtered with a fixed lower cutoff frequency (f_L) of 20 Hz. The upper cutoff frequencies (f_U) are randomly selected $f_U \in [2500, 3500]$ Hz.
- Lower cutoff augmentation: Filtered with a fixed $f_U = f_s/2$ (f_s denotes sampling frequency) and random $f_L \in [50, 200]$ Hz.
- Telephone channel band-pass filtering (BPF): To replicate the filtering effects that occurred in the old analog telephone channels (following the LDC meta-data of Table 1), we have filtered the audios using a band-pass filter with fixed $f_L = 300$ Hz and variable $f_U \in [3000, 4000]$ Hz.

4.1.4. Convolution with different room environment responses

Following Table 1, the IIITH and KGP utterances are mostly collected in recording studios. The LDC data is recorded outdoors and in smaller rooms. For simulating the different recording environments, utterances are convolved with randomly chosen room impulse responses from the RIR corpus following the Kaldi tool [68]. Further, using the audio degradation toolbox [55], we transform the utterance environment into vinyl recordings, live hall recordings, smartphone playback, and radio broadcasting.

4.1.5. Speech enhancement

The SNR histogram analysis in [17] shows that the utterances across the corpora can be corrupted with different levels of noise. To reduce the mismatch due to the varying noise levels, we propose speech enhancements as an augmentation category. The Voicebox tool [7] is used to randomly apply one of the three speech enhancement algorithms to the utterances: spectral subtraction [4], log MMSE noise estimation based enhancement [27], and de-reverberation [20].

4.1.6. Altering audio encoders

The intermediate data recording, processing, and storage across the corpora may not follow the same encoding format. Following the augmentations in [25], we explore several encoding formats for data augmentation. We have randomly converted the encoding format of the utterances among *linear pulse code modulation* (L-PCM), a-law, u-law, Ima, and Oki adaptive delta PCM (ADPCM).²

4.1.7. Restoration from lossy audio codecs

Lossy codecs allow degrading the audio fidelity up to some range for data compression. During intermediate processings, the utterances across the corpora could be stored in different lossy codec formats and then converted into lossless `.wav` (for IIITH and KGP) or `.flac` (for LDC). Augmentation by lossy codecs diversifies utterances by incorporating codec-dependent partial information losses. Using the SoX library,³ we randomly convert the utterances in different lossy codecs, such as `.aac`, `.mp3`, `.gsm`, `.opus` and then restored them to `.wav`.

²<https://github.com/FFmpeg/FFmpeg>

³<https://github.com/rabitt/pysox>

4.1.8. Feature-level augmentation

The augmentation methods discussed so far are applied at the signal level. We also explore the suitability of two widely used augmentation methods which are applied at the feature level.

SpecAug [65]: This feature-level augmentation method randomly modifies the filter-bank features by (i) warping along the time axis, (ii) time masking, and (iii) frequency masking along random time-frequency locations. It increases the robustness against deformation along the time axis and partial information loss in specific frequency channels and time frames [65]. SpecAug is primarily designed for spectrograms or MFBs. Due to masking, the standard SpecAug zeros out some of the MFB values, and the remaining values remain the same. However, in this work, we utilize the SpecAug on MFCCs. We do this by applying SpecAug directly on the MFBs that are generated as an intermediate stage during MFCC extraction. During the decorrelation by the final DCT stage, SpecAug perturbs the entire feature matrix. In this work, we utilize SpecAug with MFCCs mainly due to a fair comparison with the baseline and the existing LID literature with the three corpora which used MFCCs. Nevertheless, several earlier works also applied SpecAug on MFCCs [21, 72, 87, 97].

Mixup [99]: This augmentation method generates new training example $(\mathbf{X}_{\text{mix}}, y_{\text{mix}})$ by linear interpolation of randomly selected training examples (\mathbf{X}_i, y_i) and (\mathbf{X}_j, y_j) . This is defined as:

$$\begin{aligned}\mathbf{X}_{\text{mix}} &= \theta \mathbf{X}_i + (1 - \theta) \mathbf{X}_j, \\ y_{\text{mix}} &= \theta y_i + (1 - \theta) y_j,\end{aligned}\tag{4}$$

where \mathbf{X}_i and \mathbf{X}_j denote the feature inputs of two independent examples, y_i and y_j are the corresponding labels presented in one-hot encoding. θ is the mixing parameter, and it follows beta-distribution where two shape parameters are equal and lie between 0 and ∞ . Mixup creates soft labels y_{mix} , which increase the robustness against corrupt labels and adversarial examples [99].

4.2. Optimizing augmentation fold

The data augmentation methods help in improving the generalization of neural network architectures such as TDNN [82]. Studies have demonstrated that data augmentation with several methods helps in training larger models [25]. However, we suspect that to solve the issue of insufficient training data for the low-resourced corpora, if we use excessively large numbers of augmented samples, it can increase the computation. Further, with too many augmented utterances compared to the original samples, the LID model can start learning the acoustic variations and diverge from learning the true language-discriminating cues present in the original training data [83]. In spite of notable advancements for exploring different audio augmentation techniques, we observe that the existing literature [25, 36, 96] does not provide comprehensive analysis on the amount of audio data to be augmented for training. In this work, we extensively explore finding the optimum (with respect to the same-corpora and cross-corpora LID performance) fraction of the augmented utterances to the original utterances in the combined (augmented + original) training set.

Suppose, for any i -th corpus, the total number of non-augmented original training utterances be N_O^i . For each signal-level augmentation category A^k , where $k \in \{1, 2, \dots, K\}$, we have created one augmented instance for all of its S^k sub-categories. Hence, the total number of augmented utterances corresponding to A^k is $M^k = S^k N_O^i$. If we randomly sample η^k fractions from these M^k utterances, then the total augmented utterances (N_A^i) are $N_A^i = \sum_{k=1}^K M^k \eta^k$. If $\eta^k = \gamma / (KS^k)$ (where γ is a non-zero positive real number), the total utterances in the combined augmented

training set (N_{Tot}^i) is:

$$N_{\text{Tot}}^i = N_O^i + N_A^i = (1 + \gamma)N_O^i \quad (5)$$

In this work, γ is termed as the *augmentation fold-factor*. It denotes the ratio of augmented utterances to the original non-augmented training utterances in the combined training set. We create different combined training sets for $\gamma = 0.5, 1, 2, 3$, and 4. Apart from evaluating the LID performance for each augmentation category, we also train independent LID models for different augmentation fold-factors and analyze their performances.

4.3. Cascading augmentation categories

The domain diversification approach we follow aims to replicate the non-lingual diversities present in the unseen cross-corpora evaluation utterances to improve generalization. For this, we create more realistic augmented audios that simulate the different real-world acoustic scenarios. Therefore, beyond the conventional ways of independently exploring the different audio augmentation categories, we propose an algorithm called *maximally diversity-aware cascaded augmentations* that applies multiple data augmentation categories together with the categories being selected based on the introduced diversity.

In the literature, several augmentation approaches include multiple signal processing stages to perturb the data. For example, the room environment-based augmentations (A^5) includes re-verberation effect followed by noise additions. However, such examples are limited, and their potentialities are very little explored. In this work, we present one of the very first theoretical and experimental studies to propose efficient ways of cascading multiple augmentation categories. In the cascaded augmentation, the signal processing steps of categories A^i and A^j can be cascaded to form a new augmentation category $A^{(i,j)}$. Considering the diverse signal processing operations of the explored augmentation categories, the order of cascading is assumed to be non-commutative, $A^{(i,j)} \neq A^{(j,i)}$. We can pipeline two or more different augmentation categories to generate diverse augmented audio data. Here, we have to deal with a trade-off situation, the amount of diversification of the utterances versus preserving the semantic language labels of the augmented utterances. We address this trade-off by limiting the cascading to only two signal-level augmentation categories.

Our method relies on the distances between the utterance embeddings for different augmentation categories. We train the LID model and extract the embeddings Z^k of the training utterances belonging to the signal-level augmentation categories $A^k \forall k \in [1, K]$. We use the output from the linear layer before the final softmax layer to extract the 192-dimensional utterance-level embeddings (Z^k). Then, we calculate the *diversity* among the embedding distributions $\mathbf{D}^{(i,j)} = \mathbf{D} \|p(Z^i), p(Z^j)\|$. Here, $i, j \in [1, K]$ denote two different signal-level augmentation categories and $i \neq j$. To measure the diversity, we use three metrics [106]: *cosine distance*, *KL divergence*, and *symmetric KL divergence*. For all the three distance metrics, we compute the embedding distances ($\mathbf{D}^{(0,k)}$) between the original utterances (denoted as A^0) and augmented utterances from A^k . We then perform correlation analysis between the $\mathbf{D}^{(0,k)}$ values and respective c_{avg}^k scores. Here, c_{avg}^k denotes the same-corpora LID performance (on validation data) when the training data is formed by combining augmented utterances of category A^k with the original training utterances. Based on experimental results (see Section 6.5), we select the KL divergence as diversity measure metric because it shows the most negative correlation values. The negative correlation implies that the performance becomes better (i.e., lower C_{avg}) with higher domain diversity. Further, unlike the other two distance metrics, KL divergence is not symmet-

Algorithm 1 Maximally diversity-aware cascaded augmentation.

```
1: procedure (Cascade  $A^i$  and  $A^j$  to  $A^{(i,j)}$ ) ▷  $i, j \in [1, K]$  and  $i \neq j$ 
2:   for  $i = 1 : K$  do
3:      $Z^i \leftarrow A^i$  ▷  $Z$  : TDNN embedding
4:     for  $m = 1 : K$  do
5:       if  $i \neq m$  then
6:          $\mathbf{D}_{\text{KL}}^{(i,m)} \leftarrow \mathbf{D}_{\text{KL}} \|p(Z^i), p(Z^m)\|$  ▷  $\mathbf{D}_{\text{KL}}$  : KL divergence
7:       if  $\mathbf{D}_{\text{KL}}^{(i,j)} > \mathbf{D}_{\text{KL}}^{(i,m)}, \forall m \in [1, K] \neq j$  then
8:         Cascade  $A^i$  followed by  $A^j$ .
```

ric, $\mathbf{D}_{\text{KL}}^{(i,j)} \neq \mathbf{D}_{\text{KL}}^{(j,i)}$. This is also concurrent with the non-commutative property of the cascaded augmentation approach.

For K augmentation categories, there could be a total of $\binom{K}{2}$ cascaded pairs. We avoid repeating the same augmentation category. Therefore, the effective number of ways is $\binom{K}{2} - K$. Correspondingly, we have one pair-wise KL divergence score $\mathbf{D}_{\text{KL}}^{(i,j)}$, where, $i, j \in [1, K]$ and $i \neq j$. The KL divergence scores between the embedding of two different augmentation categories are an indication of the variations in diversity they contribute to the training data. To maximize the diversity gain for the cascaded augmentations, for each augmentation category $i \in [1, K]$, we cascade the augmented category $j \in [1, K] \neq i$ such that: $\mathbf{D}_{\text{KL}}^{(i,j)} > \mathbf{D}_{\text{KL}}^{(i,m)}, \forall m \in [1, K] \neq j$. The generated cascaded augmented utterances are denoted as $A^{(i,j)}$. The summary of our proposed approach for generating the maximally diversity-aware cascaded augmented utterances is presented in Algorithm 1.

We also cascade two feature-level augmentation techniques, SpecAug, and mixup, independently on top of different signal-level augmentation methods. We perform this because the sources of audio perturbations covered by the signal and feature-level augmentations are different. Hence, cascading them can further enhance the diversity in the training data.

5. Cross-corpora LID: Domain generalization

Cross-corpora mismatches in non-lingual factors can be considered as a *domain shift* problem [10]. Domain shift caused by co-variance shifts, concept shifts, changes the feature-label joint distribution between the training (source) and testing (target) data samples, $p(\mathcal{X}^{\text{src}} \mathcal{Y}^{\text{src}}) \neq p(\mathcal{X}^{\text{tar}} \mathcal{Y}^{\text{tar}})$ [58]. In this section, we address the cross-corpora domain shift problem from domain generalization [34] perspectives.

5.1. Domain adaptation vs. domain generalization

A common approach followed in the machine learning literature is to incorporate information from the target domain samples during training to mitigate the domain shift problem. The model, trained on the source domain, is usually adapted to that particular target domain [76]. This approach is known as the *domain adaptation* [100, 95]. In our context, we can consider the training corpus as the source domain and the remaining two corpora as individual target domains. However, following the constraints of the cross-corpora evaluations presented in Section 2, we refrain from using any data of the target domains. Moreover, adapting to any specific target domain does not guarantee generalization for any unseen test domains during real-world deployments [19]. Considering these factors, we have selected an alternative approach of leveraging information from the source domain and improving the generalization of the LID models. This

approach is called *domain generalization* [103]. The concept of DG was first formally introduced in the computer vision field [6]. It is gradually becoming an attraction in the research community for a realistic solution to the domain shift problem [10] when, rather than any specific target domain, we want to achieve robustness in real-world unseen conditions.

5.2. Pseudo-domains: Solving for multiple source domains

In speech processing literature, the DG approaches are less explored due to the difficulties in collecting multi-source labeled data [28]. Hence, we first propose a solution for availing data from multiple source domains into our framework. In Section 4.1, we have discussed the different signal-level augmentation categories, applied for domain diversification. Here, we propose to utilize those augmentation categories as well as the original condition as multiple source domains for solving cross-corpora LID from a DG perspective. We have termed those synthetic domains as *pseudo-domains* since they are artificially created from the source domain data. This approach of utilizing audio augmentations in DG is more appropriate for our current study with limited training data from low-resourced languages.

Let us denote each pseudo-domain as $\mathcal{D}^{\text{dom}} = \{(\mathbf{X}_i^{\text{dom}}, \mathbf{y}_i^{\text{dom}})\}_{i=1}^{N_{\text{dom}}}$, where $\text{dom} \in [0, K]$ with K denoting the total number of signal-level augmentation categories, N_{dom} denoting the number of training samples for the corresponding pseudo-domain. The utterances $\{(\mathbf{X}_i^{\text{dom}}, \mathbf{y}_i^{\text{dom}})\}$ are assumed to be sampled from the domain-specific joint distribution $p(\mathcal{X}^{\text{dom}}, \mathcal{Y}^{\text{dom}})$. Without adapting to any specific target domain data $\mathcal{D}^{\text{tar}} = \{(\mathbf{X}_i^{\text{tar}}, \mathbf{y}_i^{\text{tar}})\}_{i=1}^{N_{\text{tar}}}$ (where N_{tar} denotes the number of samples of the target domain), the goal in DG is to generalize over unseen target domain distributions $p(\mathcal{X}^{\text{tar}}, \mathcal{Y}^{\text{tar}})$, where $\text{tar} \in [K + 1, \infty)$ [34]. In the following sub-sections, we have discussed the different DG approaches to address the cross-corpora LID problem.

5.3. Domain adversarial ECAPA-TDNN architecture

This work further extends the state-of-the-art ECAPA-TDNN architecture with an additional adversarial domain classifier branch for learning domain-invariant language information. As shown in Fig. 2, the proposed architecture has an encoder (\mathcal{E}) which is composed of the frame-level SE-Res2 blocks followed by the attentive pooling layer. For each utterance, the encoder maps the input feature vector \mathbf{X} to Z . After the encoder, we have used two segment-level classification branches; LID branch ($\mathcal{D}_{\text{lang}}$) and domain branch (\mathcal{D}_{dom}). $\mathcal{D}_{\text{lang}}$ is composed of two fully connected layers. For each i -th utterance, it outputs the language posterior probabilities $p(\mathbf{y}|\mathbf{X}_i; \theta_E, \theta_L)$. Here, θ_E and θ_L denote the learnable parameters of the encoder and LID branch, respectively. The additional domain classifier branch \mathcal{D}_{dom} (with learnable parameters denoted as θ_D) predicts the pseudo-domain posterior probabilities $p(\mathbf{y}^{\text{dom}}|\mathbf{X}_i; \theta_E, \theta_D)$ as a $K + 1$ class softmax vector with K signal-level augmentation categories and the additional class of original non-augmented utterances. Here, we do not include the sub-category classes in the domain classifier because all of them perturb the audio in a similar manner. Further, too many similar output classes due to the sub-categories can make the domain classification task challenging which can hinder the main LID task. Here, $\mathcal{Y}^{\text{dom}} \in \{0, 1, \dots, K\}$ denotes the pseudo-domain label space. The \mathcal{D}_{dom} branch consists of a Conv1D layer followed by two fully connected layers. It is connected to the \mathcal{E} module by a *gradient reversal layer* (GradRev) [24] that flips the sign of the gradient flows into the encoder during the backpropagation. During training, the pseudo-domain data $\{(\mathbf{X}^{\text{dom}}, \mathbf{y}, \mathbf{y}^{\text{dom}})\} \forall \text{dom} \in [0, K]$ is fed as input. We train the model by optimizing the combined adversarial loss function (\mathcal{L}_{adv}) defined as




Fig. 2. Adversarial ECAPA-TDNN architecture for domain generalization. Here, $K = 7$ (with “+1” denoting the original non-augmented class) denotes the total number of pseudo-domain corresponding to each signal-level category. L denotes the number of target languages which is five in the present closed-set LID study.

$$\begin{aligned} \mathcal{L}_{adv}(\theta_E, \theta_L, \theta_D) = \sum_i [(1 - \lambda_{adv}) \mathcal{L}_{lang}(y_i, \mathcal{D}_{lang}(\mathcal{E}(\mathbf{X}_i)); \theta_E, \theta_L) \\ - \lambda_{adv} \mathcal{L}_{dom}(\mathcal{D}_{dom}(y_i^{dom}, \mathcal{E}(\mathbf{X}_i)); \theta_E, \theta_D)]. \end{aligned} \quad (6)$$

Here, i denotes index of individual utterances, $\mathbf{X}_i \in \mathcal{X}$. The total loss (\mathcal{L}_{adv}) is composed by the language classification loss \mathcal{L}_{lang} and domain classification loss \mathcal{L}_{dom} . The parameter $\lambda_{adv} \in [0, 1]$ decides the weights of the two losses in the total loss. By optimizing the min-max problem, \mathcal{L}_{lang} trains θ_L and θ_E to improve the LID accuracy. Simultaneously, θ_D is estimated to improve the pseudo-domain recognition performance. The gradient reversal layer flips the sign of $\delta \mathcal{L}_{dom} / \delta \theta_D$, resulting in θ_E to learn domain-invariant language discriminative information.

5.4. Multitask learning for cross-corpora generalization

We have explored multitask learning (MTL) [101, 73, 9] in the cross-corpora LID framework with the aim of learning pseudo-domain aware representations that can improve generalization. Joint learning of multiple features or multiple tasks has been shown to improve the robustness for cross-channel LID test samples [43]. Using MTL, our LID model jointly learns an additional eight-class domain classification task [103]. Thus, the LID is expected to capture the domain-specific language discriminative information generalized across multiple pseudo-domains. In this experiment, the extension to the ECAPA-TDNN architecture is similar to that of the adversarial scenario (see Section 5.3) except for the removal of the gradient reversal layer [1]. We can express the total loss to be optimized as

$$\begin{aligned} \mathcal{L}_{MTL}(\theta_E, \theta_L, \theta_D) = \sum_i [\mathcal{L}_{lang}(y_i, \mathcal{D}_{lang}(\mathcal{E}(\mathbf{X}_i)); \theta_E, \theta_L) \\ + \lambda_{MTL} \mathcal{L}_{dom}(y_i^{dom}, \mathcal{D}_{dom}(\mathcal{E}(\mathbf{X}_i)); \theta_E, \theta_D)]. \end{aligned} \quad (7)$$

In Eq. 6 and Eq. 7, the major difference \mathcal{L} lies in the sign of the \mathcal{L}_{dom} component. So, basically,

these two experiments attempt to solve the domain-mismatch problem following two opposite strategies: domain-invariant and domain-aware [1]. In domain generalization, we do not have information about any specific target domain. Therefore, both of these opposite but complementary DG approaches can be suitable across the different unseen domains that may be encountered in real-world applications.

5.5. MMD-based domain generalization approaches

The most fundamental assumption in machine learning is that the training and evaluation data are independent and identically distributed (i.i.d.). However, this assumption is often violated due to the domain mismatch. For improving generalization, utilizing the pseudo-domains, our idea is to train the LID model in such a way that the distribution mismatch among the training conditions is reduced. In this work, we explore *maximum mean discrepancy* (MMD) [32] based techniques for domain generalization. We investigate MMD for two main reasons. First, this efficiently minimizes the distribution shift between source and target data [32]. Second, the language recognition literature successfully used MMD for reducing domain mismatch [22].

MMD is the metric for measuring the distance between two distributions [32]. In DA literature, MMD-based loss (\mathcal{L}_{MMD}) is optimized to minimize the statistical moments differences between the source and target domains [105]. In this work, we first utilize the pseudo-domains as a target domain so that the MMD-based DA technique is applied under the DG configurations. Thereafter, we extend the MMD-based DG model by proposing multi-domain extension. For the conventional MMD-based DG experiments, we have treated the non-augmented training utterances as source domain (\mathcal{S}^{src}) and all the pseudo-domains (\mathcal{S}^{dom}) as a target domain.

Let the feature space distribution of the source and target domain be denoted as $p^{\text{src}} = p(\mathcal{X}^{\text{src}})$ and $p^{\text{dom}} = p(\mathcal{X}^{\text{dom}})$, respectively. We first compute the MMD-loss as the squared difference between the means of the source and target representations mapped to a *reproducing kernel Hilbert space* (RKHS) [32] and this is defined as

$$d_{\mathcal{H}} = \left\| \sum_{\mathbf{X}_i^{\text{src}} \in \mathcal{S}^{\text{src}}} \phi(\mathbf{X}_i^{\text{src}}) - \sum_{\mathbf{X}_i^{\text{dom}} \in \mathcal{S}^{\text{dom}}} \phi(\mathbf{X}_i^{\text{dom}}) \right\|_{\mathcal{H}}^2. \quad (8)$$

In Eq. 8, i is the utterance indices, $\mathbf{X}^{\text{src}} \in \mathcal{X}^{\text{src}}$, $\mathbf{X}^{\text{dom}} \in \mathcal{X}^{\text{dom}}$. \mathcal{H} denotes the RKHS space corresponding to a positive semi-definite kernel \mathcal{K} . The MMD-loss is jointly optimized with the LID loss to make the LID model robust against the domain shifts. For practical purposes, an empirical estimate of the MMD-loss is computed by Eq. 9 assuming a positive definite kernel $\mathcal{K}(\mathbf{X}^{\text{src}}, \mathbf{X}^{\text{dom}}) = \langle \phi(\mathbf{X}^{\text{src}}), \phi(\mathbf{X}^{\text{dom}}) \rangle$. Here, ϕ maps the inputs from feature space to the RKHS, and $\langle \cdot, \cdot \rangle$ operator denotes the inner product.

$$\hat{d}_{\mathcal{H}} = \left\| \frac{1}{N_{\text{src}}} \sum_{i=1}^{N_{\text{src}}} \phi(\mathbf{X}_i^{\text{src}}) - \frac{1}{N_{\text{dom}}} \sum_{i=1}^{N_{\text{dom}}} \phi(\mathbf{X}_i^{\text{dom}}) \right\|_{\mathcal{H}}^2 \quad (9)$$

Here, N_{src} and N_{dom} denote the number of source and target domain utterances, respectively. Further simplification of Eq. 9 leads to:

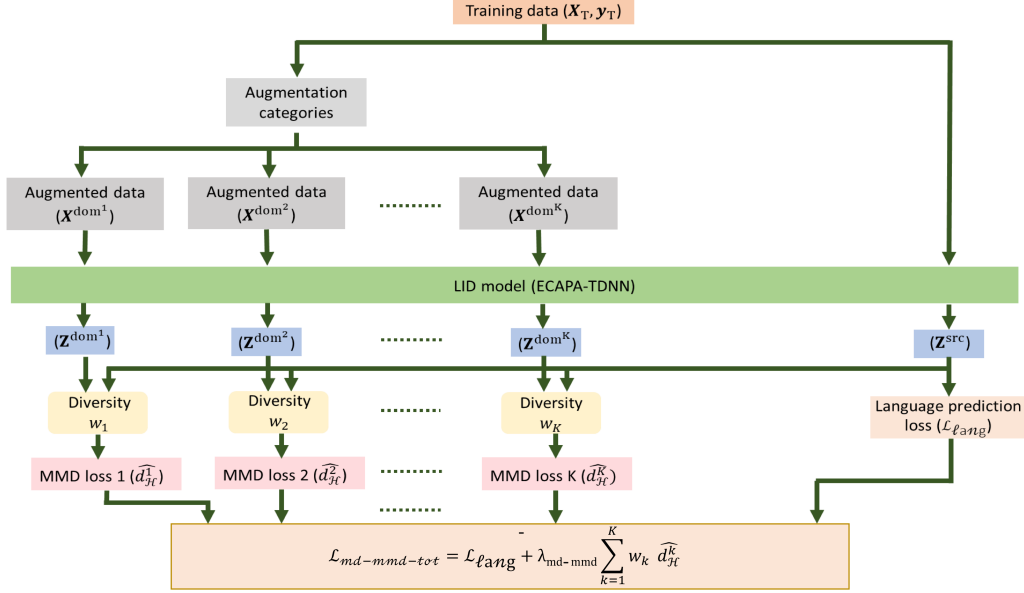


Fig. 3. Working principle of the MD-MMD based ECAPA-TDNN model with $\mathcal{L}_{\text{lang}}$ denoting the AM-softmax language prediction loss and $\hat{d}_{\mathcal{H}}^k$ denoting individual multi-domain MMD loss components.

$$\hat{d}_{\mathcal{H}} = \frac{1}{N_{\text{src}}^2} \sum_{i=1}^{N_{\text{src}}} \sum_{j=1}^{N_{\text{src}}} \phi(\mathbf{X}_i^{\text{src}}) \phi(\mathbf{X}_j^{\text{src}}) + \frac{1}{N_{\text{dom}}^2} \sum_{i=1}^{N_{\text{dom}}} \sum_{j=1}^{N_{\text{dom}}} \phi(\mathbf{X}_i^{\text{dom}}) \phi(\mathbf{X}_j^{\text{dom}}) - \frac{2}{N_{\text{src}} N_{\text{dom}}} \sum_{i=1}^{N_{\text{src}}} \sum_{j=1}^{N_{\text{dom}}} \phi(\mathbf{X}_i^{\text{src}}) \phi(\mathbf{X}_j^{\text{dom}}) \quad (10)$$

Using multiple-kernel extensions in MMD-loss enhances the classification performance [48]. In this work, we use kernel \mathcal{K} as the combination of M radial basis function (RBF) with different variances:

$$\mathcal{K}(\mathbf{X}_i, \mathbf{X}_j) = \sum_{m=1}^M \exp\left(-\frac{(\mathbf{X}_i - \mathbf{X}_j)^T (\mathbf{X}_i - \mathbf{X}_j)}{\sigma_m^2}\right) \quad (11)$$

Where,

$$\sigma_m^2 = \frac{2^m}{N_{\text{src}} N_{\text{dom}}} \sum_{i=1}^{N_{\text{src}}} \sum_{j=1}^{N_{\text{dom}}} (\mathbf{X}_i - \mathbf{X}_j)^T (\mathbf{X}_i - \mathbf{X}_j) \quad (12)$$

The combined loss is defined as

$$\mathcal{L}_{\text{mmd-tot}} = \mathcal{L}_{\text{lang}} + \lambda_{\text{mmd}} \hat{d}_{\mathcal{H}}. \quad (13)$$

In Eq. 13, λ_{mmd} is the scalar to determine the relative weights of the LID loss and MMD loss. The joint optimization learns the language-discriminating representations while producing a more similar representation of the source and target domain.

In Eq. 13, all the pseudo-domains are collectively treated as a single target domain \mathcal{S}^{dom} with the single MMD-loss component $\hat{d}_{\mathcal{H}}$. To utilize the finer domain differences among the pseudo-domains more efficiently, we propose to extend the conventional MMD approach into *multi-domain MMD* (MD-MMD) loss. The proposed MMD method efficiently incorporates the

different sources of diversity for each pseudo-domain. For each pseudo-domain $\mathcal{S}^{\text{dom}^k}$, where $k = [1 : K]$, we compute separate MMD loss component $\hat{d}_{\mathcal{H}}^k$:

$$\hat{d}_{\mathcal{H}}^k = \left\| \frac{1}{N_{\text{src}}} \sum_{i=1}^{N_{\text{src}}} \phi(\mathbf{X}_i^{\text{src}}) - \frac{1}{N_{\text{dom}}^k} \sum_{i=1}^{N_{\text{dom}}^k} \phi(\mathbf{X}_i^{\text{dom}^k}) \right\|_{\mathcal{H}}^2 \quad (14)$$

Here, N_{dom^k} denotes the number of training utterances for $\mathcal{S}^{\text{dom}^k}$. Each component is then optimized simultaneously with $\mathcal{L}_{\text{lang}}$. The total loss function for the MD-MMD is:

$$\mathcal{L}_{\text{md-mmd-tot}} = \mathcal{L}_{\text{lang}} + \lambda_{\text{md-mmd}} \sum_{k=1}^K w_k \hat{d}_{\mathcal{H}}^k \quad (15)$$

In Eq. 15, the w_k are the relative weights for each pseudo-domains in the MD-MMD loss. One simple way is to set $w_k = 1/K \forall k \in [1, K]$. However, different pseudo-domains offer different levels of domain diversity in the source domain. Therefore, the domain with the most diversity measure from the source domain can be weighted the most. In the proposed approach, w_k are set according to the corresponding normalized domain diversity measures (see Section 5.3) from the source domain:

$$w_k = \left[\mathbf{D}_{KL}^{\{0,k\}} / \sum_{k=1}^K \mathbf{D}_{KL}^{\{0,k\}} \right]. \quad (16)$$

The detailed working principle of the proposed MD-MMD model is illustrated in Fig. 3.

6. Results & analysis

In this section, we present and analyze the LID performance results of the different experiments that we conduct in this study. Following the NIST LRE [74] and OLR challenges [44], the LID performances are reported using $EER (\%) / C_{\text{avg}} * 100$ (in the Tables denoted as EER/C_{avg}). We primarily focus on the mismatch between the cross-corpora and same-corpora performance metrics during the evaluations. For the proposed approaches, compared to the baseline, the relative improvements in the same-corpora LID performances are also studied. We first compare the LID performances of the x-vector baseline TDNN and the ECAPA-TDNN architecture. Then, we individually assess the LID performance of each augmentation category. Following that, we train and evaluate the different domain diversification approaches discussed earlier. Finally, we utilize the diversified data as pseudo-domains and evaluate the LID performances of the DG extensions.

6.1. Baseline LID performances

We first train LID models using the baseline x-vector TDNN and ECAPA-TDNN architectures. The corresponding same-corpora and cross-corpora evaluation results are presented in Table 5. The results demonstrate that ECAPA-TDNN outperforms the x-vector TDNN for most of the evaluation cases. We observe that only for some cases during the evaluation with KGP-test data, x-vector shows slightly improved performances. The additional modifications that are modeled in the ECAPA-TDNN compared to the x-vector potentially degrade the robustness towards the non-lingual mismatches contained in the KGP-test data. KGP-test subset contains

Table 5: LID performance evaluation of the TDNN architectures in $EER(\%) / C_{avg} * 100$.

Training corpus	Baseline TDNN			ECAPA-TDNN		
	IIITH-test	LDC-test	KGP-test	IIITH-test	LDC-test	KGP-test
	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}
IIITH	12.40 / 13.82	43.10 / 46.23	30.30 / 29.91	9.74 / 11.51	42.43 / 44.82	34.83 / 32.62
LDC	50.08 / 47.34	22.95 / 26.40	45.60 / 45.92	46.35 / 43.21	21.86 / 25.70	42.95 / 39.59
KGP	37.55 / 34.91	51.80 / 46.72	11.18 / 10.82	36.52 / 33.59	47.25 / 45.49	12.36 / 8.75

Table 6: LID performances (in $EER(\%) / C_{avg} * 100$) due to individual signal-level augmentation category in the IIITH training data.

Testing corpus	Without augmentation	Non-speech addition (A^1)	Signal parameter perturbation (A^2)	Bandwidth augmentation (A^3)	Environmental variations (A^4)	Speech enhancements (A^5)	Audio encoding (A^6)	Lossy codec restoration (A^7)
IIITH	9.74 / 11.51	8.43 / 8.73	10.12 / 11.76	9.76 / 11.56	9.48 / 9.46	9.21 / 9.06	8.93 / 9.01	9.63 / 11.42
LDC	42.43 / 44.82	40.33 / 36.78	41.15 / 43.64	43.77 / 46.11	39.62 / 37.80	39.84 / 38.32	39.32 / 36.44	42.71 / 44.70
KGP	34.83 / 32.62	33.23 / 39.01	39.87 / 38.70	42.06 / 37.32	37.08 / 33.72	35.06 / 31.91	32.36 / 29.98	38.26 / 35.59

only 60 chunks (each of 3 s duration) from each language. Whereas IIITH-test contains a total of 4878 numbers of 3 s chunks, and the LDC-test contains a total of 24617 3 s chunks. Due to the extremely small size of the KGP-test, there can be chances of sample biases that lead to inferior ECAPA-TDNN performance. Nevertheless, since ECAPA-TDNN achieves better LID performance for the majority of the LID evaluations, we only consider the ECAPA-TDNN based LID models for the subsequent experiments. For both architectures, we observe a prominent performance mismatch between the same-corpora and cross-corpora evaluations. The cross-corpora LID experiments achieve EER of more than 30%, indicating unreliable language predictions. The higher EER also reflects that achieving decent generalization over different LID corpora is a challenging task.

6.2. LID performances: Individual augmentation categories

For each of the seven signal-level augmentation categories A^k , for each corresponding sub-category, we create an augmented copy of the training utterances. Thereafter, for each augmentation category, we randomly sample (with fold-factor $\gamma = 1$) the augmented utterances and combine them with the original non-augmented training data (A^0). We thus create seven sets of combined augmented audio data and train individual LID models. Table 6 shows the LID performance comparison for each of the signal-level augmentation categories. For comparative assessments, the LID performance of the ECAPA-TDNN models, trained with the non-augmented original training data, are also presented in Table 6.

For all three training corpora, we observe LID performance improvements in similar ranges due to the explored augmentation categories. Hence, in Table 6 and in the following domain diversification experiments, we have presented the LID performances with only the IIITH-trained scenario. From Table 6, we observe that for same-corpora evaluation, the non-speech addition category performs the best. For both LDC and KGP cross-corpora evaluations, codec-encoding augmentation shows the best performance improvements. The proposed speech-enhancement based augmentation and environmental augmentations also exhibit consistently promising performance improvements for both the cross-corpora evaluations. We then try to associate these cross-corpora evaluation observations with the available meta information. Contrary to the broadcast news studio recordings of the IIITH and KGP, the LDC data contains telephonic conversations, recorded in both indoor and outdoor environments. So, the diverse background noise

Table 7: Impact of augmentation fold-factors on LID performance (in EER (%) / $C_{\text{avg}} * 100$) using the IIITH as training corpus.

Fold-factor (γ)	IIITH-test	LDC-test	KGP-test
	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}
0.5	9.18 / 10.92	39.43 / 43.79	37.56 / 34.89
1	8.75 / 10.01	39.61 / 43.96	38.65 / 36.70
2	8.00 / 9.14	38.19 / 42.71	38.88 / 35.52
3	7.68 / 9.28	38.64 / 42.40	36.80 / 32.92
4	8.03 / 9.68	38.89 / 43.20	35.73 / 34.30

related mismatches can be suppressed using speech enhancement based augmentations. The environmental augmentations helps in reducing corpora-mismatch in terms of recording room environments. LID performance improvements due to the codec encoding augmentation can be justified by the fact that the intermediate data processing and storage protocols of the different corpora may be different.

6.3. LID performance of different augmentation fold-factors

We experimentally optimize the augmentation fold-factor (γ) that controls the ratio of the augmented and original training utterances in the combined augmented training set. We have trained independent LID models using $\gamma = \{0.5, 1, 2, 3, 4\}$ and compared their LID performances in Table 7. The results indicate that as γ increases from 0.5 to 3, in general, both same-corpora and cross-corpora LID performance improve. For IIITH-trained LID model, best same-corpora EER of 7.68% is achieved for $\gamma = 3$. For the same train-test pair, the lowest C_{avg} of 9.14 is attained for $\gamma = 2$. For the cross-corpora evaluations with LDC-test, we observe that the EER attains the lowest value at $\gamma = 2$. Whereas, the lowest C_{avg} for KGP-test is achieved for $\gamma = 3$. This echoes the fact that the performances of the TDNN-based models improve with more training data incurred due to augmentation [82]. But for almost all the cases, the LID performances start to degrade beyond $\gamma = 3$. This is an interesting observation that can be justified by saying that beyond $\gamma = 3$, there are too many acoustically diversified training samples compared to the original training data. So, the LID models may start learning the non-lingual diversities instead of the actual language-discriminating cues present in the original data [83]. Moreover, a smaller value of the fold-factor, $\gamma = 2, 3$, also reduces computation and the model training time. For the subsequent experiments, we use $\gamma = 2$ to sample utterances from different signal-level augmentation categories. This sampled augmentation set is denoted as “signal-level ($\gamma = 2$)” in Table 10 and “Augmentation ($\gamma = 2$)” in Table 11. Sampled augmentation sets with $\gamma = 2$ or 3 is optimum in terms LID performances in our work. Further, for $\gamma = 2$, we also need a lesser computational cost. It is an interesting future work to verify the optimum augmentation fold-factors across the different speech-based classification tasks.

6.4. LID performance of signal-feature level cascaded augmentations

We apply cascaded data augmentation by performing feature-level augmentations during the feature extraction steps of the signal-level augmented utterances. We choose all seven signal-level augmentation categories and sample them with $\gamma = 2$. For feature-level augmentation, we evaluate SpecAug and mixup separately. The corresponding LID performances (for IIITH as training data) are presented in Table 8.

Table 8: LID performances (in EER (%) / $C_{\text{avg}} * 100$) by cascading the signal-level and feature-level augmentations using the IITH as training corpus.

Augmentation	IITH-test	LDC-test	KGP-test
	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}
Signal-level ($\gamma = 2$)	8.00 / 9.14	38.19 / 42.71	38.88 / 35.52
Signal-level followed by SpecAug	9.82 / 11.20	48.67 / 45.16	39.74 / 40.24
Signal-level followed by Mixup	9.26 / 10.93	39.70 / 40.61	45.01 / 38.94
Diversity-aware cascading	9.13 / 10.74	39.70 / 43.16	35.17 / 29.32
Signal-level followed by cascading	8.61 / 10.22	39.16 / 43.04	37.75 / 35.84

Analyzing Table 8, we observe that cascading the different signal-level and feature-level augmentation categories does not help in improving the LID generalization. It may be possible that the signal-feature augmented utterances are over-diversified, and the language classification task becomes challenging. As a consequence, in Section 6.5, we have reported the experiments of the diversity-aware cascaded augmentations by considering the multiple signal-level augmentation categories.

6.5. LID performance: Diversity-aware cascaded augmentations

To augment the utterances in a more realistic manner, we propose cascading different signal-level augmentation categories to create further diversified training utterances. For effective diversification, we find out the pair-wise maximally diversified augmentation categories for cascading.

6.5.1. Metrics for measuring the domain diversity

For each signal-level augmentation category A^k , embeddings (Z^k) are extracted for the corresponding utterances. For embedding extraction, we use the LID model trained on the combined augmentation IITH training set (with $\gamma = 2$). We also extract embeddings for the original non-augmented IITH utterances and denote them as Z^0 . Next, we measure the embedding distances between Z^0 and $Z^k \forall k \in [1, 7]$ using three different metrics, cosine distance (\mathbf{D}_{cos}), KL-divergence (\mathbf{D}_{KL}), and symmetric KL divergence ($\mathbf{D}_{\text{KL-s}}$). For each metric, we have correlated the embedding distance with the corresponding validation C_{avg} (primary evaluation metric) scores and presented the results in Table 9. As the embedding distances are not guaranteed to follow Gaussian distribution, we use the *Spearman correlation coefficients* for the correlation computation. We have finally selected the KL divergence for measuring the domain diversity distance because it shows the most negative correlation with the C_{avg} scores. The more the KL divergence between the original data and an augmentation category, the more the diversity expected, and eventually the lesser the corresponding C_{avg} . Further, the non-commutative characteristics (not present in $\mathbf{D}_{\text{KL-s}}$) of the KL divergence also support the basic assumption of the cascaded augmentations: $A^{i,j} \neq A^{j,i}$.

Table 9: Different domain diversity measures and their correlations with C_{avg} on IITH validation data.

	A^1	A^2	A^3	A^4	A^5	A^6	A^7	Correlation
\mathbf{D}_{cos}	0.127	0.106	0.221	0.123	0.122	0.009	0.008	-0.107
\mathbf{D}_{KL}	0.353	0.133	0.257	0.188	0.205	0.171	0.065	-0.428
$\mathbf{D}_{\text{KL-s}}$	0.303	0.177	0.252	0.211	0.232	0.204	0.128	-0.428
C_{avg}	7.11	11.30	11.07	7.82	7.98	7.20	9.95	-

6.5.2. LID performance: Maximally diversity-aware cascaded augmentation

Next, we present the LID performance for the proposed cascaded augmentation method. The aim of this experiment is to increase the diversity in the training data more realistically by cascading signal-level augmentation categories, which are mutually the most diverse. Considering all the $\mathbf{D}_{KL}^{(i,j)}$ values, for each A^i , we cascade its utterances with another augmentation category A^j such that $\mathbf{D}_{KL}^{(i,j)}$ maximum. Table 10 shows the pairwise divergence of all augmentation pairs for the IITH corpus. For each augmentation category, we create one cascaded augmented set. We then sample and combine these cascaded sets with the original non-augmented utterances by $\gamma = 2$ and train the LID model. Further, we train another LID model with the combined utterances of original data, signal-level augmented data ($\gamma = 1$), and cascaded augmented data ($\gamma = 1$). The last two rows of Table 8 show the LID performance using the maximally diversity-aware cascaded augmented utterances.

Table 10: Pairwise KL divergences between the embeddings of the explored signal-level augmentation categories for the IITH corpus.

	A^1	A^2	A^3	A^4	A^5	A^6	A^7
A^1	0.00	0.311	0.385	0.458	0.713	0.169	0.312
A^2	0.313	0.00	0.271	0.151	0.178	0.228	0.101
A^3	0.518	0.332	0.00	0.210	0.609	0.391	0.309
A^4	0.341	0.155	0.191	0.00	0.287	0.278	0.194
A^5	0.529	0.173	0.532	0.329	0.00	0.340	0.251
A^6	0.187	0.191	0.296	0.311	0.393	0.00	0.180
A^7	0.257	0.104	0.280	0.174	0.210	0.195	0.00

The proposed maximally diversity-aware cascaded augmentation shows a promising LID performance. Interestingly, for this cascaded augmentation, the relative improvements of the cross-corpora performance are even more than corresponding same-corpora performance improvements. One can keep on cascading more than two augmentation methods to incorporate more diversity in the training data. However, this will increase computational complexity. Furthermore, incorporating more diversity in the data may distort the language-specific cues in the audio data.

We have also measured the KL divergence of the cascaded-augmented embeddings with the non-augmented embeddings; $\mathbf{D}_{KL}^{(0,\text{cascaded})} = 0.286$. Comparing with the $\mathbf{D}_{KL}^{(0,k)}$ values for $k \in [1, 7]$ (from Table 9), we find $\mathbf{D}_{KL}^{(0,\text{cascaded})}$ to be higher for six out of the seven augmented categories. In this work, for each A^i row of Table 10, if the column A^p shows the highest value, we cascade the two augmentation categories as $A^{(i,p)}$. We could also cascade two augmented categories with the highest overall $\mathbf{D}_{KL}^{(i,j)}$ values. However, by doing so, it is not guaranteed that the diversities for all the explored augmentation categories, $A^k \forall k = [1, 7]$, is preserved by cascading. The objective of this study is to improve generalization on multiple completely independent corpora. The proposed cascading algorithm, which incorporates diversity by ensuring audio from all the explored augmented categories, could be more beneficial in such unseen evaluation scenarios.

6.6. LID performance: Domain generalization perspectives

We have performed experiments with three DG techniques: adversarial, multitasking, and MMD as discussed in Section 5. All these approaches extend the state-of-the-art ECAPA-TDNN

Table 11: LID performance (in EER (%) / $C_{\text{avg}} * 100$) of the proposed domain generalization approaches. For comparison, LID performance with the sampled (with $\gamma = 2$) signal-level augmented utterances are also presented.

Approaches	Training corpus: IITH			Training corpus: LDC			Training corpus: KGP		
	IITH-test	LDC-test	KGP-test	IITH-test	LDC-test	KGP-test	IITH-test	LDC-test	KGP-test
	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}	EER / C_{avg}
Augmentation ($\gamma = 2$)	8.00 / 9.14	38.19 / 42.71	38.88 / 35.52	32.91 / 32.53	20.10 / 24.22	48.25 / 41.53	33.67 / 33.79	47.25 / 45.49	8.94 / 7.26
Adversarial	9.64 / 10.88	35.40 / 40.68	28.28 / 22.92	34.98 / 36.77	20.80 / 25.14	44.94 / 43.59	36.76 / 34.35	48.26 / 42.73	9.24 / 8.98
Multitasking	7.97 / 9.03	36.29 / 41.89	29.47 / 25.69	36.32 / 36.40	21.28 / 25.47	44.27 / 44.50	38.97 / 36.12	48.12 / 42.59	8.79 / 8.70
MMD	9.06 / 9.78	39.58 / 45.30	33.15 / 26.28	38.89 / 35.81	22.10 / 26.02	46.32 / 45.88	40.58 / 38.15	48.44 / 45.05	11.98 / 13.70
MD-MMD	9.61 / 10.99	36.45 / 41.68	32.56 / 32.68	38.61 / 32.41	21.07 / 24.46	39.32 / 39.94	40.63 / 34.73	48.47 / 44.46	11.71 / 12.00

architecture and jointly optimize multiple cost functions with the trade-off parameter λ . Following [1], we set λ experimentally for better convergence of training. For the adversarial method, we have kept $\lambda = 0.001$ for the initial 15 epochs. For the other two methods, the initial value of λ is 0.1 for the first five epochs. After keeping the initial λ value, we keep on increasing the λ by 0.01 per epoch for all the DG approaches. We have shown the comparative LID performances of different DG methods in Table. 11. In this table, for comparative assessments, we have also reported the LID performance of augmented data generated by sampling utterances (with $\gamma = 2$) from different signal-level augmentation categories.

We have observed a trend that the multitasking method yields the best performance for the same-corpora condition. For the cross-corpora evaluations, the adversarial model is found to be the most effective. The evaluation utterances of the same-corpora should have higher degrees of domain similarity with the training data. Therefore, the additional insights about the pseudo-domains learned from domain-aware language learning, improve the LID performance for such cases. Whereas the cross-corpora domains can be even more diverse, and in such cases incorporating domain invariance is more useful. For the IITH-trained model, the performance improvements for KGP test data are more compared to the LDC test data. It can be justified as the meta information reveals the IITH and KGP data are relatively similar in various aspects compared to the LDC data. The proposed MD-MMD shows prominent performance improvements as compared to the conventional MMD, especially for the cross-corpora scenarios. The MD-MMD method treats each pseudo-domain separately and weighs them as per the diversity distances. It enhances both same-corpora and cross-corpora LID performance. MMD-based methods aim to reduce the distribution mismatch between the source (here, non-augmented original training data) and the target (here, augmented data with $\gamma = 2$). It is reported in the existing literature that, while focusing on reducing the source-target distribution mismatch, MMD-based methods reduce the inter-class separability [98, 12]. This can be the reason why both the two MMD-based methods do not outperform the adversarial or multitask learning based DG approaches in most cases.

The explored DG approaches clearly outperform the baseline system for both same-corpora and cross-corpora evaluations. Considering the IITH-trained system, from the baseline (in Table 5) C_{avg} improves by 5.55% for the LDC-test utterances. The same for the KGP test utterances is 6.99%. Compared to the ECAPA-TDNN model, the corresponding relative improvements in C_{avg} are 4.14% and 9.70%, respectively, for the LDC and KGP test corpora. The DG approaches further outperform the explored domain diversification approaches in most cases. The DG approaches utilize the augmented training data as pseudo-domains. So, we compare them with the ECAPA-TDNN model trained on augmented training set ($\gamma = 2$). The comparison reveals a relative C_{avg} improvement of 2.03% and 12.60%, respectively, for the LDC and KGP test corpora. If we study the LID performances for the different experiments we conduct, we can observe a clear trend that whenever LDC/KGP trained LID models are evaluated on KGP/LDC test utter-

ances, the performance remains inferior compared to other cross-corpora train-test pairs. The meta information comparison, presented in Table 1, reveals a prominently higher corpora mismatch between these two corpora. For example, KGP is the smallest database, and its utterances are collected from broadcast news recordings. Whereas LDC is the largest database containing conversational telephone speech (CTS) data. The LDC utterances have a strong influence on regional dialects, and they are recorded in diverse backgrounds. Therefore, achieving cross-corpora generalization between such diverse corpora-pair is challenging. The language-specific cross-corpora study can provide further insights for solving this challenging situation. Still, compared to the baseline, we get the best-case EER improvements of 4.55% for the LDC-test and 6.28% for KGP-test data. On the other hand, for both same-corpora and cross-corpora scenarios, IITH utterances show consistent performance improvements. IITH contains both BN and CTS data, and it has a moderately large size. We believe the meta characteristics of this corpus have some commonalities with both corpora. So, achieving cross-corpora generalization over the IITH corpus is relatively easier. Overall, the LID performance improvements we achieve are promising, considering the fact that the corpora we use are low-resourced and developed independently with distinct differences in the meta-data characteristics.

7. Further analysis and discussion

In this section, we provide a blind cross-corpora evaluation of the proposed approaches on another unseen corpus, additional analysis, and ablation studies. We also discuss several prospects of the current study.

7.1. Evaluation on unseen corpus

We have conducted extensive experiments so far on the three corpora and have demonstrated substantial improvements in cross-corpora experiments. This is indeed a valid argument that a machine learning algorithm should ideally refrain from using any kind of feedback from the evaluation set to optimize its configurations. Rather, the optimization should be made solely on the validation set. However, one may argue that the configuration and parameters of the proposed methods are optimized for the selected three corpora. Therefore, to further ensure the generalization ability of the proposed methods, we perform experiments on a fourth corpus *VoxLingua107* (VoxLingua) [89], which is completely unseen during the system development process. VoxLingua is a vast and diverse corpus, with utterances automatically collected from wild sources using online video streaming platforms. We have only conducted the final evaluations with this corpus by randomly selecting 500 utterances of 3 s duration from each of the five languages. This likely includes speech utterances of diverse audio quality with various background noises that represent evaluation in the wild.

For this blind evaluation, we have considered the LID models trained on the IITH corpus. Fig. 4 shows the *detection error trade-off* (DET) curve as well as EERs for the same-corpora and cross-corpora evaluations following the different experiments, discussed earlier. We observe that all the explored approaches considerably outperform the baseline. Compared to the blind cross-corpora performance of the ECAPA-TDNN model, we achieve relative EER improvements of 4.40% and 5.23% for domain diversification and domain generalization methods, respectively. Among the domain diversification experiments, the proposed diversity-aware cascaded augmentation performs the best in the blind evaluation. Whereas the adversarial learning method outperforms all other DG methods for this evaluation.




Fig. 4. DET plots of the baselines and proposed approaches for the same-corpora (IITH-test) and unseen cross-corpora (VoxLingua-test) conditions.

7.2. Assessing domain diversification

We have proposed the maximally diversity-aware cascaded augmentations for effective diversification of the training data that can improve the cross-corpora generalization. We further analyze here whether the domain diversity is improved due to such augmentation methods. The domain diversity measures adopted in this work are data-driven, where we use embeddings from the LID model trained on augmented data. We now conduct this diversity analysis with a LID-independent model trained on different audio data. We select a pre-trained model trained for acoustic scene classification tasks. In particular, we choose *OpenL3* embedding to explore the background and environmental differences across the three corpora [33].

Fig. 5 visualizes the non-augmented training set embeddings for the three databases excluding and including IITH augmented training set. We observe that three corpora consisting of non-augmented original training data (Fig. 5 (a)) are somehow separated by the *OpenL3* embeddings, and this indicates the presence of strong non-lingual differences. We observe that the *OpenL3* embeddings of the IITH-augmented utterances (Fig. 5 (b)) are well distributed across all three corpora demonstrating effective domain diversification with the proposed method.

We have also analyzed the overall LTAS of the three databases along with the augmented IIITH LTAS in Fig. 6. The augmented spectrum shows closer proximity with the KGP-LTAS spectrum, indicating improved cross-corpora performance. We also present a language-specific LTAS comparison using two languages: Bengali and Punjabi. These two plots indicate that the cross-corpora mismatch is non-uniform across the languages. It promises a potential future extension of this work with language-specific cross-corpora performance analysis and generalization.

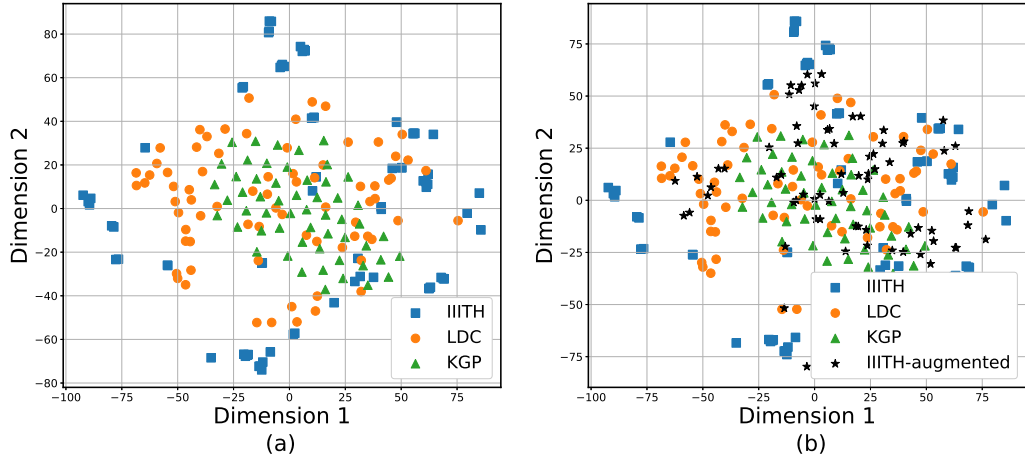


Fig. 5. T-SNE visualization of OpenL3 embeddings for (a) the training sets of the IIITH, LDC, and KGP corpora, (b) with further inclusion of the IIITH-augmented ($\gamma = 2$) set.




Fig. 6. LTAS comparison of the three corpora and the IIITH-augmented data (with $\gamma = 2$) for (a) Bengali, (b) Punjabi, and (c) all training utterances.

7.3. Linguistic diversity v/s non-lingual diversity in domain diversification

In Section 4.2, we discuss the impact of the fold-factor, i.e. the ratio of the number of augmented utterances to the number of non-augmented original training utterances in the training set. In this case, we follow the common practice for augmented file selection as discussed in the literature [81, 25]. This involves the generation of augmented copies for each method, followed

Table 12: Impact of same-corpora and cross-corpora LID performance (in EER (%) / $C_{\text{avg}} * 100$) for the content vs. non-lingual diversities in the augmented training sets. Here, three augmentation sets are created using the IITH corpus, each having different degrees of content and non-lingual diversities.

Augmentation	IITH	LDC	KGP
	EER/ C_{avg}	EER/ C_{avg}	EER/ C_{avg}
First scenario	8.75 / 10.01	39.61 / 43.96	38.65 / 36.70
Second scenario	7.40 / 9.41	42.31 / 45.99	28.00 / 29.75
Third scenario	7.00 / 8.98	46.15 / 41.55	31.00 / 30.37

by a random selection of augmented files from the pooled data to meet the desired fold-factor. Due to the random sampling from the pooled data, one utterance can have multiple augmented copies, i.e., content is repeated multiple times with varying non-lingual diversities. Whereas some of the utterances may not have any augmentation copy at all. In this section, we further explore the impact of content vs. non-lingual diversity within the augmented data. Such studies are not widely explored in the existing audio augmentation literature, especially for speaker and language recognition problems. Hence we find the inclusion of such discussion important for the augmentation research community.

We create three augmentation sets (with $\gamma = 1$) using three scenarios. In the *first scenario*, we follow the usual convention of sampling the augmented data, which is discussed at the beginning of this section and is followed in Section 6.3. As mentioned above, in this scenario, there is a possibility of repetition of some content, whereas, for some utterances, there is a possibility of having no augmented copy. In the *second scenario*, we apply a constraint that only one augmented version for each audio file is randomly selected from all augmented versions of the same file. Note that both scenarios consist of an equal number of audio files (twice the number of files in the original data due to $\gamma = 1$), but the former has lesser content variability than the latter. We also consider a *third scenario* which further restricts the content variability by randomly selecting 1/2 of the original utterances. Thereafter, for each of the selected utterances, we randomly select any three augmented copies (from seven different signal-level augmentation categories). The third augmentation set also contains an equal number of total audio files as the other two conditions, but it has lesser content variability and more non-lingual variability.

Table 12 shows the LID results for three augmentation scenarios with IITH as the training corpus. Due to the randomness in audio file selection, we have conducted each experiment five times and we have reported the average performance. The results in Table 12 show that applying constraints on the content repetition helps to improve the LID performance for the two test conditions: IITH and KGP. The meta information says that IITH and KGP both are mainly broadcast news data. So, the manner of articulation, accent, and dialect remains standard for these databases. Whereas LDC being conversational data, should contain diverse variations in accents and dialects. For such a database in evaluation, the constraint on content repetition during data augmentation does not help. We consider further exploration of content vs. non-lingual diversity as an interesting future work in the context of domain diversification.

7.4. Domain-invariant vs. domain-aware approaches

The LID performance of the explored domain generalization methods reveals a trend that the adversarial models are suitable for cross-corpora generalization while the multitask learning models help to improve same-corpora LID performance. Therefore, we have two apparently opposite approaches to address the domain information during language training: *domain-invariant*

and *domain-aware* [1]. The same-corpora data is more likely to have relatively higher domain similarity with the created pseudo-domains compared to the cross-corpora utterances. So, the additional task of pseudo-domain classification in the multitask learning architecture potentially makes the LID model aware of the possible diversities in the same-corpora test utterances. On the contrary, the domains of the cross-corpora evaluation utterances could be less similar to the pseudo-domains. In such scenarios, the domain-invariant approach helps improve the cross-corpora LID performance by making the model less sensitive to domain diversity. Future work could consider exploring an algorithm for dynamically selecting either domain-aware or domain-invariant models for a given test audio [1].

7.5. Cross-corpora performance gap: Remaining challenges and solutions

The proposed methods in this work show substantial improvements in cross-corpora evaluations. However, the performance gap between the same-corpora and cross-corpora evaluations is still considerably high. This could be due to several reasons: (i) The data augmentation methods used for domain diversification do not accurately simulate the real-world domains, and the audio data simulation is itself a challenging task [71]. The domain generalization problem formulated in this study also relies on augmented data. We speculate that augmented data with more realistic conditions would help to reduce the performance gap. (ii) Given the constraints in our problem formulation, we do not use any target domain data, and this limits the scope of optimization of the models for one specific domain. (iii) The domain-invariant approach, which works well for the cross-corpora condition, does not entirely remove the domain-related characteristics from the latent representation. However, the main motivation of the current study is to assess the limits of domain diversification and domain generalization in challenging cross-corpora conditions with a focus on relative improvement over the existing baseline. Following the short-duration NIST LRE evaluation conditions [74], we also use a very short test speech of duration 3 s, which makes the LID problem more difficult [30]. The state-of-the-art LID systems show about 15% EER on 3 s duration condition of LRE17 where the systems are both trained and tested on conversational telephone speech [62]. Even though the systems with such higher EERs are unreliable, we believe that the relative improvements we achieve provide a promising direction for solving this challenging task of cross-corpora generalization with short-duration test speech.

In the current work, the application of the proposed domain diversification and domain generalization is limited to the ECAPA-TDNN, a specific type of model for language classification which trains architecture in a discriminative way. The study can be extended with other discriminative and generative approaches. Recently, self-supervised learning (SSL) [56] are being widely explored for different tasks, including language recognition [46], which exploits a large amount of unlabeled audio data from other sources. We consider cross-corpora evaluation with SSL frameworks and assessing the performance gaps with same-corpora evaluation as interesting future research avenues.

8. Conclusions

This study has explored the suitability of domain diversification and domain generalization for improving cross-corpora generalization in low-resourced LID tasks. We have explored existing and new data audio data augmentation methods for domain diversification. For domain generalization, we investigate adversarial learning, multitask learning as well as MMD loss, and its new variants with pseudo-domains. In addition to extensive experiments on three widely used

Indian LID corpora, we also perform a blind evaluation on an unseen corpus to further validate the effectiveness of the proposed approaches. We summarize the main findings of this work as follows:

- Different evaluation corpora can contain diverse non-lingual characteristics. Hence, along with the conventional audio augmentations, exploring new sources of perturbing the audio data further diversify the training data, which is effective for unseen cross-corpora generalization. For example, the proposed speech enhancement or codec-based augmentations are most promising for improving robustness against LDC and KGP test data for the IITH-trained LID models.
- The fundamental motivation for applying domain diversification is to reduce the cross-corpora mismatch by simulating the non-lingual characteristics related to the potential unseen domain. Hence, augmenting audio data in a more realistic manner is important. For this reason, we propose the maximally diversity-aware cascaded augmentation with an optimized fold-factor, which is a more effective solution in blind evaluation than the conventional augmentation methods.
- We aim to achieve better LID generalization across diverse unseen domains with limited training data. Hence, instead of attaining robustness for a specific target domain using the conventional domain adaptation methods, we opt for domain generalization. It utilizes the different augmentation categories as pseudo-domains to further improve the LID performance compared to the domain diversification methods.
- We explore domain generalization in two fundamentally opposite ways using the domain-invariant and domain-aware methods. Domain-invariant LID models are trained using gradient reversal or MMD-based methods in such a way that it becomes less sensitive to non-lingual domain-specific mismatches. In the domain-aware cases, we rather update the LID model to additionally learn domain-specific cues using multitask learning. Our study reveals that domain-invariant learning is more effective for cross-corpora generalization, whereas domain-awareness improves same-corpora LID performance.

Although the proposed domain diversification and domain generalization largely improve the cross-corpora LID performance, investigations are still required to reduce the same-corpora and cross-corpora performance gap further. Exploring other advanced data augmentation methods using voice conversion, text-to-speech synthesis, and style transfer could be promising. We address the cross-corpora problem from the perspectives of the entire database. Investigating the cross-corpora LID in a language-specific manner can reveal further insights into this problem. Our further investigation of the spoken content v/s non-lingual variability indicates another promising future direction.

Acknowledgements

The first author would also like to thank the Linguistic Data Consortium (LDC), University of Pennsylvania, USA, for the LDC Data Scholarship. We also thank Prof. KS Rao (IIT Kharagpur) and Prof. AK Vuppala (IIIT Hyderabad) for their help with the KGP and IITH datasets. Finally, the authors would like to thank the associate editor and the reviewers for their careful reading, detailed comments, and constructive suggestions, which substantially enhanced the content of the manuscript.

References

- [1] Adi, Y., Zeghidour, N., Collobert, R., Usunier, N., Liptchinsky, V., Synnaeve, G., 2019. To reverse the gradient or not: an empirical comparison of adversarial and multi-task learning in speech recognition, in: ICASSP, IEEE. pp. 3742–3746.
- [2] Alumäe, T., Kukk, K., 2022. Pretraining approaches for spoken language recognition: TalTech submission to the OLR 2021 challenge, in: Odyssey: The Speaker and Language Recognition Workshop, ISCA. pp. 240–247.
- [3] Benyassine, A., Shlomot, E., Su, H., Massaloux, D., Lamblin, C., Petit, J., 1997. A silence compression scheme for use with g. 729 optimized for v. 70 digital simultaneous voice and data applications (recommendation g. 729 annex b). *IEEE Commun. Mag* 35, 64–73.
- [4] Berouti, M., Schwartz, R., Makhoul, J., 1979. Enhancement of speech corrupted by acoustic noise, in: ICASSP, IEEE. pp. 208–211.
- [5] Beyan, C., Shahid, M., Murino, V., 2021. RealVAD: A real-world dataset and a method for voice activity detection by body motion analysis. *IEEE Transactions on Multimedia* 23, 2071–2085.
- [6] Blanchard, G., Lee, G., Scott, C., 2011. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in Neural Information Processing Systems* 24, 2178–2186.
- [7] Brookes, M., et al., 1997. Voicebox: Speech processing toolbox for MATLAB. Software, available [Mar. 2011] from www.ee.ic.ac.uk/hp/staff/dmb/voicebox/voicebox.html 47.
- [8] Brümmer, N., Du Preez, J., 2006. Application-independent evaluation of speaker detection. *Computer Speech & Language* 20, 230–275.
- [9] Caruana, R., 1997. Multitask learning. *Machine learning* 28, 41–75.
- [10] Cha, J., Chun, S., Lee, K., Cho, Han-Cheol and Park, S., Lee, Y., Park, S., 2021. SWAD: Domain generalization by seeking flat minima, in: *Advances in Neural Information Processing Systems*, pp. 22405–22418.
- [11] Chakraborty, J., Chakraborty, B., Bhattacharya, U., 2021. DenseRecognition of spoken languages, in: *International Conference on Pattern Recognition (ICPR)*, IEEE. pp. 9674–9681.
- [12] Chen, Y., Song, S., Li, S., Wu, C., 2019. A graph embedding framework for maximum mean discrepancy-based domain adaptation algorithms. *IEEE Transactions on Image Processing* 29, 199–213.
- [13] Chettri, B., Hautamäki, R.G., Sahidullah, M., Kinnunen, T., 2021. Data quality as predictor of voice anti-spoofing generalization, in: *INTERSPEECH, ISCA*. pp. 1659–1663.
- [14] Clark, L., Doyle, P., Garaialde, D., Gilmartin, E., Schlögl, S., Edlund, J., Aylett, M., Cabral, J., Munteanu, C., Edwards, J., et al., 2019. The state of speech in HCI: Trends, themes and challenges. *Interacting with Computers* 31, 349–371.
- [15] Deng, W., Zheng, L., Ye, Q., Kang, G., Yang, Y., Jiao, J., 2018. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification, in: *CVPR*, pp. 994–1003.
- [16] Desplanques, B., Thienpondt, J., Demuyne, K., 2020. ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification, in: *INTERSPEECH, ISCA*. pp. 1–5.
- [17] Dey, S., Saha, G., Sahidullah, M., 2021. Cross-corpora language recognition: A preliminary investigation with Indian languages, in: *EUSIPCO, IEEE*. pp. 546–550.
- [18] Dey, S., Sahidullah, M., Saha, G., 2022. An overview on Indian spoken language recognition from machine learning perspective. *ACM Transactions on Asian and Low-Resource Language Information Processing* 21, 1–45.
- [19] Ding, Z., Fu, Y., 2017. Deep domain generalization with structured low-rank constraint. *IEEE Transactions on Image Processing* 27, 304–313.
- [20] Doire, C.S., et al., 2016. Single-channel online enhancement of speech corrupted by reverberation and noise. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 25, 572–587.
- [21] Du, C., Li, H., Lu, Y., Wang, L., Qian, Y., 2021. Data augmentation for end-to-end code-switching speech recognition, in: *Spoken Language Technology Workshop (SLT)*, IEEE. pp. 194–200.
- [22] Duroselle, R., Jouvét, D., Illina, I., 2020. Metric learning loss functions to reduce domain mismatch in the x-vector space for language recognition, in: *INTERSPEECH*, pp. 447–451.
- [23] Ferrer, L., Castan, D., McLaren, M., Lawson, A., 2022. A discriminative hierarchical PLDA-based model for spoken language recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 30, 2396–2410.
- [24] Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V., 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17, 2096–2030.
- [25] Garcia-Romero, D., Sell, G., Mccree, A., 2020. MagNetO: X-vector magnitude estimation network plus offset for improved speaker recognition, in: *Odyssey: The Speaker and Language Recognition Workshop, ISCA*. pp. 1–8.
- [26] Gerczuk, M., Amiriparian, S., Ottl, S., Schuller, B., 2021. EmoNet: A transfer learning framework for multi-corpus speech emotion recognition. *IEEE Transactions on Affective Computing* , 1–1.

- [27] Gerkmann, T., Hendriks, R.C., 2011. Unbiased MMSE-based noise power estimation with low complexity and low tracking delay. *IEEE Transactions on Audio, Speech, and Language Processing* 20, 1383–1393.
- [28] Gideon, J., McInnis, M.G., Provost, E.M., 2021. Improving cross-corpus speech emotion recognition with adversarial discriminative domain generalization (ADDoG). *IEEE Transactions on Affective Computing* 12, 1055–1068.
- [29] Gillespie, S., Logan, Y.Y., Moore, E., Laures-Gore, J., Russell, S., Patel, R., 2017. Cross-database models for the classification of dysarthria presence., in: *INTERSPEECH, ISCA*. pp. 3127–3131.
- [30] Gonzalez-Dominguez, J., et al., 2015. Frame-by-frame language identification in short utterances using deep neural networks. *Neural Networks* 64, 49–58.
- [31] Greenberg, C.S., Martin, A.F., Przybocki, M.A., 2012. The 2011 NIST language recognition evaluation, in: *INTERSPEECH, ISCA*. pp. 34–37.
- [32] Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., Smola, A., 2006. A kernel method for the two-sample-problem. *Advances in Neural Information Processing Systems* 19.
- [33] Grollmisch, S., Cano, E., Kehling, C., Taenzer, M., 2021. Analyzing the potential of pre-trained embeddings for audio classification tasks, in: *EUSIPCO, IEEE*. pp. 790–794.
- [34] Gulrajani, I., Lopez-Paz, D., 2021. In search of lost domain generalization, in: *International Conference on Learning Representations*.
- [35] Hu, X., Yang, Y.H., 2017. Cross-dataset and cross-cultural music mood prediction: A case on Western and Chinese pop songs. *IEEE Transactions on Affective Computing* 8, 228–240.
- [36] Iqbal, T., Helwani, K., Krishnaswamy, A., Wang, W., 2021. Enhancing audio augmentation methods with consistency learning, in: *ICASSP, IEEE*. pp. 646–650.
- [37] Kang, W., Alam, M.J., Fathan, A., 2022. Deep learning-based end-to-end spoken language identification system for domain-mismatched scenario, in: *Language Resources and Evaluation Conference*, pp. 7339–7343.
- [38] Karen, J., et al., 2017. Multi-language conversational telephone speech 2011 – South Asian LDC2017S14. web download. Philadelphia: Linguistic Data Consortium. - .
- [39] Khosla, A., Zhou, T., Malisiewicz, T., Efros, A.A., Torralba, A., 2012. Undoing the damage of dataset bias, in: *European Conference on Computer Vision, Springer*. pp. 158–171.
- [40] Korshunov, P., Marcel, S., 2019. A cross-database study of voice presentation attack detection, in: *Handbook of Biometric Anti-Spoofing, Springer*, pp. 363–389.
- [41] Kumawat, P., Routray, A., 2021. Applying TDNN architectures for analyzing duration dependencies on speech emotion recognition, in: *INTERSPEECH, ISCA*. pp. 3410–3414.
- [42] Li, H., Ma, B., Lee, K.A., 2013. Spoken language recognition: from fundamentals to practice. *Proceedings of the IEEE* 101, 1136–1159.
- [43] Li, L., Li, Z., Liu, Y., Hong, Q., 2021. Deep joint learning for language recognition. *Neural Networks* 141, 72–86.
- [44] Li, Z., Zhao, M., Hong, Q., Li, L., Tang, Z., Wang, D., Song, L., Yang, C., 2020. AP20-OLR challenge: Three tasks and their baselines, in: *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC), IEEE*. pp. 550–555.
- [45] Liu, H., Garcia Perera, L.P., Khong, A., Styles, S., Khudanpur, S., 2022a. PHO-LID: A unified model incorporating acoustic-phonetic and phonotactic information for language identification, in: *INTERSPEECH, ISCA*. pp. 2233–2237.
- [46] Liu, H., Perera, L.P.G., Khong, A.W., Chng, E.S., Styles, S.J., Khudanpur, S., 2022b. Efficient self-supervised learning representations for spoken language identification. *IEEE Journal of Selected Topics in Signal Processing* 16, 1296–1307.
- [47] Liu, H., Perera, L.P.G., Khong, A.W., Dauwels, J., Styles, S.J., Khudanpur, S., 2022c. Enhancing language identification using dual-mode model with knowledge distillation, in: *Odyssey: The Speaker and Language Recognition Workshop, ISCA*. pp. 248–254.
- [48] Long, M., Cao, Y., Wang, J., Jordan, M., 2015. Learning transferable features with deep adaptation networks, in: *International conference on machine learning, PMLR*. pp. 97–105.
- [49] Lopez-Moreno, I., Gonzalez-Dominguez, J., Martinez, D., Plchot, O., Gonzalez-Rodriguez, J., Moreno, P.J., 2016. On the use of deep feedforward neural networks for automatic language identification. *Computer Speech & Language* 40, 46–59.
- [50] Loshchilov, I., Hutter, F., 2018. Decoupled weight decay regularization, in: *ICLR*.
- [51] Maity, S., A.K., V., Rao, K., Nandi, D., 2012. IITKGP-MLLSC speech database for language identification, in: *National Conference on Communication (NCC), IEEE*. pp. 1–5.
- [52] Mandava, T., Vuddagiri, R.K., Vydana, H.K., Vuppala, A.K., 2019. An investigation of LSTM-CTC based joint acoustic model for Indian language identification, in: *Automatic Speech Recognition and Understanding Workshop (ASRU), IEEE*. pp. 389–396.
- [53] Mandava, T., Vuppala, A.K., 2019. Attention based residual-time delay neural network for Indian language identification, in: *International Conference on Contemporary Computing (IC3), IEEE*. pp. 1–5.

- [54] Martinez, D., Plchot, O., Burget, L., Glembek, O., Matějka, P., 2011. Language recognition in ivectors space, in: INTERSPEECH, ISCA.
- [55] Mauch, M., Ewert, S., 2013. The audio degradation toolbox and its application to robustness evaluation, in: International Society for Music Information Retrieval Conference (ISMIR), Curitiba, Brazil.
- [56] Mohamed, A., Lee, H.y., Borgholt, L., Havtorn, J.D., Edin, J., Igel, C., Kirchhoff, K., Li, S.W., Livescu, K., Maaløe, L., Sainath, T.N., Watanabe, S., 2022. Self-supervised speech representation learning: A review. *IEEE Journal of Selected Topics in Signal Processing* 16, 1179–1210.
- [57] Monteiro, J., Alam, J., Falk, T.H., 2019. Residual convolutional neural network with attentive feature pooling for end-to-end language identification from short-duration speech. *Computer Speech & Language* 58, 364–376.
- [58] Moreno-Torres, J.G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N.V., Herrera, F., 2012. A unifying view on dataset shift in classification. *Pattern Recognition* 45, 521–530.
- [59] Mozilla, . Common Voice. <https://commonvoice.mozilla.org/en>. Accessed: 2020-01-12.
- [60] Mushtaq, Z., Su, S.F., 2020. Environmental sound classification using a regularized deep convolutional neural network with data augmentation. *Applied Acoustics* 167, 107389.
- [61] Nadimpalli, A.V., Rattani, A., 2022. On improving cross-dataset generalization of deepfake detectors, in: CVPR, IEEE/CVF. pp. 91–99.
- [62] Padi, B., Mohan, A., Ganapathy, S., 2020. Towards relevance and sequence modeling in language recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28, 1223–1232.
- [63] Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q., 2010. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* 22, 199–210.
- [64] Pandey, A., Wang, D., 2022. Self-attending RNN for speech enhancement to improve cross-corpus generalization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 30, 1374–1385.
- [65] Park, D.S., et al., 2019. SpecAugment: A simple data augmentation method for automatic speech recognition, in: INTERSPEECH, ISCA. pp. 2613–2617.
- [66] Paszke, A., Gross, S., Massa, F., Lerer, A., 2019. PyTorch: An imperative style, high-performance deep learning library., in: *NeurIPS*, pp. 8024–8035.
- [67] Paul, D., Sahidullah, M., Saha, G., 2017. Generalization of spoofing countermeasures: A case study with ASVspoof 2015 and BTAS 2016 corpora, in: ICASSP, IEEE. pp. 2047–2051.
- [68] Povey, D., et al., 2011. The Kaldi speech recognition toolkit, in: *Automatic Speech Recognition and Understanding Workshop (ASRU)*, IEEE.
- [69] Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I., 2022. Robust speech recognition via large-scale weak supervision. *OpenAI Blog* .
- [70] Reddy, V.R., Maity, S., Rao, K.S., 2013. Identification of Indian languages using multi-level spectral and prosodic features. *International Journal of Speech Technology* 16, 489–511.
- [71] Ribas, D., Vincent, E., Calvo, J.R., 2016. A study of speech distortion conditions in real scenarios for speech processing applications, in: *Spoken Language Technology Workshop (SLT)*, pp. 13–20.
- [72] Rossenbach, N., Zeyer, A., Schlüter, R., Ney, H., 2020. Generating synthetic audio data for attention-based speech recognition systems, in: ICASSP, IEEE. pp. 7069–7073.
- [73] Ruder, S., 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* .
- [74] Sadjadi, S.O., Kheyrkhan, T., Tong, A., Greenberg, C.S., Reynolds, D.A., Singer, E., Mason, L.P., Hernandez-Cordero, J., 2018. The 2017 NIST language recognition evaluation., in: *Odyssey: The Speaker and Language Recognition Workshop*, ISCA.
- [75] Salamon, J., Bello, J.P., 2017. Deep convolutional neural networks and data augmentation for environmental sound classification. *IEEE Signal Processing Letters* 24, 279–283.
- [76] Sarfjoo, S., Madikeri, S., Motlicek, P., Marcel, S., 2020. Supervised domain adaptation for text-independent speaker verification using limited data, in: INTERSPEECH, pp. 3815–3819.
- [77] Schuller, B., Vlasenko, B., Eyben, F., Wöllmer, M., Stuhlsatz, A., Wendemuth, A., Rigoll, G., 2010. Cross-corpus acoustic emotion recognition: Variances and strategies. *IEEE Transactions on Affective Computing* 1, 119–131.
- [78] Shen, P., Lu, X., Li, S., Kawai, H., 2017. Conditional generative adversarial nets classifier for spoken language identification., in: INTERSPEECH, ISCA. pp. 2814–2818.
- [79] Singh, P., Saha, G., Sahidullah, M., 2021. Non-linear frequency warping using constant-Q transformation for speech emotion recognition, in: *International Conference on Computer Communication and Informatics (ICCCI)*, IEEE. pp. 1–6.
- [80] Snyder, D., Chen, G., Povey, D., 2015. MUSAN: A music, speech, and noise corpus. *arXiv preprint arXiv:1510.08484* .
- [81] Snyder, D., et al., 2018a. Spoken language recognition using x-vectors., in: *Odyssey: The Speaker and Language Recognition Workshop*, pp. 105–111.
- [82] Snyder, D., et al., 2018b. X-vectors: Robust DNN embeddings for speaker recognition, in: ICASSP, IEEE. pp. 5329–5333.

- [83] Sturm, B.L., 2014. A simple method to determine if a music information retrieval system is a “horse”. *IEEE Transactions on Multimedia* 16, 1636–1644.
- [84] Tang, Z., Wang, D., Song, L., 2019. AP19-OLR challenge: Three tasks and their baselines, in: *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, IEEE. pp. 1917–1921.
- [85] Thienpondt, J., Desplanques, B., Demuynck, K., 2022. Tackling the score shift in cross-lingual speaker verification by exploiting language information, in: *ICASSP*, IEEE. pp. 7187–7191.
- [86] Toledo-Ronen, O., Sorin, A., 2013. Voice-based sadness and anger recognition with cross-corpora evaluation, in: *ICASSP*, IEEE. pp. 7517–7521.
- [87] Tong, F., Zhao, M., Zhou, J., Lu, H., Li, Z., Li, L., Hong, Q., 2021. ASV-Subtools: Open source toolkit for automatic speaker verification, in: *ICASSP*, IEEE. pp. 6184–6188.
- [88] Tsakalidis, S., Byrne, W., 2005. Acoustic training from heterogeneous data sources: Experiments in Mandarin conversational telephone speech transcription, in: *ICASSP*, IEEE. pp. I–461.
- [89] Valk, J., Alumäe, T., 2021. VoxLingua107: a dataset for spoken language recognition, in: *Spoken Language Technology Workshop (SLT)*, IEEE. pp. 652–658.
- [90] Vlasenko, B., Philippou-Hübner, D., Wendemuth, A., 2013. Parameter optimization issues for cross-corpora emotion classification, in: *Humaine Association Conference on Affective Computing and Intelligent Interaction*, IEEE. pp. 454–459.
- [91] Vlasenko, B., Prylipko, D., Böck, R., Wendemuth, A., 2014. Modeling phonetic pattern variability in favor of the creation of robust emotion classifiers for real-life applications. *Computer Speech & Language* 28, 483–500.
- [92] Vuddagiri, et al., 2018. IIITH-ILSC speech database for Indian language identification., in: *Spoken Language Technologies for Under-Resourced Languages*, IEEE. pp. 56–60.
- [93] Wang, F., Cheng, J., Liu, W., Liu, H., 2018a. Additive margin softmax for face verification. *IEEE Signal Processing Letters* 25, 926–930.
- [94] Wang, J., Zhu, X., Gong, S., Li, W., 2018b. Transferable joint attribute-identity deep learning for unsupervised person re-identification, in: *CVPR*, IEEE. pp. 2275–2284.
- [95] Wang, M., Deng, W., 2018. Deep visual domain adaptation: A survey. *Neurocomputing* 312, 135–153.
- [96] Wei, S., Zou, S., Liao, F., et al., 2020. A comparison on data augmentation methods based on deep learning for audio classification, in: *Journal of Physics: Conference Series*, IOP Publishing.
- [97] Xia, W., Zhang, C., Weng, C., Yu, M., Yu, D., 2021. Self-supervised text-independent speaker verification using prototypical momentum contrastive learning, in: *ICASSP*, IEEE. pp. 6723–6727.
- [98] Yan, H., Ding, Y., Li, P., Wang, Q., Xu, Y., Zuo, W., 2017. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation, in: *CVPR*, IEEE. pp. 2272–2281.
- [99] Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D., 2018. mixup: Beyond empirical risk minimization, in: *ICLR*.
- [100] Zhang, Y., Deng, B., Tang, H., Zhang, L., Jia, K., 2020. Unsupervised multi-class domain adaptation: Theory, algorithms, and practice. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- [101] Zhang, Y., Yang, Q., 2021. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 1–1.
- [102] Zhang, Z., Weninger, F., Wöllmer, M., Schuller, B., 2011. Unsupervised learning in cross-corpus acoustic emotion recognition, in: *Automatic Speech Recognition and Understanding Workshop (ASRU)*, IEEE. pp. 523–528.
- [103] Zhou, K., Liu, Z., Qiao, Y., Xiang, T., Loy, C.C., 2022. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- [104] Zhu, R., Zhang, T., Zhao, Q., Wu, Z., 2015. A transfer learning approach to cross-database facial expression recognition, in: *International Conference on Biometrics (ICB)*, IEEE. pp. 293–298.
- [105] Zhu, Y., Zhuang, F., Wang, J., Ke, G., Chen, J., Bian, J., Xiong, H., He, Q., 2020. Deep subdomain adaptation network for image classification. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1713–1722.
- [106] Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q., 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE* 109, 43–76.