# JOINT UNSUPERVISED AND SUPERVISED LEARNING FOR CONTEXT-AWARE LANGUAGE IDENTIFICATION

*Jinseok Park, Hyung Yong Kim, Jihwan Park, Byeong-Yeol Kim, Shukjae Choi, Yunkyu Lim*

42dot Inc., Seoul, Republic of Korea

{jinseok.park, hyungyong.kim, jihwan.park, byeongyeol.kim, shukjae.choi, yunkyu.lim}@42dot.ai

## ABSTRACT

Language identification (LID) recognizes the language of a spoken utterance automatically. According to recent studies, LID models trained with an automatic speech recognition (ASR) task perform better than those trained with a LID task only. However, we need additional text labels to train the model to recognize speech, and acquiring the text labels is a cost high. In order to overcome this problem, we propose context-aware language identification using a combination of unsupervised and supervised learning without any text labels. The proposed method learns the context of speech through masked language modeling (MLM) loss and simultaneously trains to determine the language of the utterance with supervised learning loss. The proposed joint learning was found to reduce the error rate by 15.6% compared to the same structure model trained by supervised-only learning on a subset of the VoxLingua107 dataset consisting of sub-three-second utterances in 11 languages.

***Index Terms***— Language identification, language recognition, joint learning, unsupervised learning.

## 1. INTRODUCTION

Language identification (LID), the process of automatically recognizing the language of a spoken utterance, is mainly used as a precursor stage of multi-lingual speech recognition and translation. Various features are available to distinguish languages, from low-level acoustic and prosodic features (tone, stress, duration, rhythm) to high-level features such as morphology and sentence syntax [1], and many LID methods have been proposed with these features [2, 3].

Over the last decade, many language identification methods with deep neural networks (DNNs) have been proposed [4, 5, 6]. Those approaches including learnable feature extractors achieved better performance than feature-based LID in challenging conditions such as utterances shorter than three seconds. Furthermore, according to recent studies [7, 8, 9], LID models trained with an automatic speech recognition (ASR) task perform better than those trained with a language identification task only. This implies that recognizing the context of utterances helps to identify the language spoken.

Although ASR helps improve LID performance, we need additional text labels to train the model for the ASR task. While language labels for each utterance are relatively easy to obtain, acquiring text labels for utterances is cost high. This follows because people roughly recognize a spoken language even if they do not know it well, whereas text labels of speech must be written by a native or expert who knows the language well.

In this paper, we propose joint unsupervised and supervised learning for context-aware language identification that does not need additional text labels. The proposed method learns the speech context through masked language modeling (MLM) loss and simultaneously trains to determine the language of the utterance with supervised learning loss. To show the effect of the proposed joint learning in wild datasets, we collected 11,000 hours of speech data in 11 languages from YouTube for training and used a subset of the VoxLingua107 [10] consisting of sub-three-second utterances. Our experimental results show that the proposed joint learning approach reduced the error rate by 15.6% compared to the same structure model trained on a supervised-only basis.

## 2. RELATED WORK

**LID with ASR:** In [7], Tang et al. remarked that phonetic information has largely been overlooked in the previous DNN-based LID and showed that the phonetic information acquired by training the ASR task improves the performance of LID. Duroselle et al. [8] conducted experiments in various combinations to measure how ASR training affects LID performance. In most cases, LID models trained with ASR generally performed better than models trained to identify languages only. In [9], Punjabi et al. showed that when their model trained jointly for ASR and LID tasks on English and Hindi datasets, performance on both two tasks improved. Although it has been experimentally proven that ASR learning helps to improve LID performance in various studies like this, a limitation is that a text label corresponding to each utterance is required to incorporate ASR learning when learning the LID model.

**BEST-RQ:** Chiu et al. [11] proposed BERT-based speech pre-training with a random-projection quantizer (BEST-RQ). They generated discrete labels from a speech signal using the random-projection quantizer and demonstrated that the network, as trained using an MLM loss to predict the discrete labels generated, showed good performance in ASR. This implies that the model learned the speech context by BEST-RQ, and we use the BEST-RQ as our unsupervised learning of LID.

**JUST:** Bai et al. [12] proposed joint unsupervised and supervised training (JUST) for multilingual ASR. Since JUST trains the model by jointly calculating unsupervised loss and supervised loss, the model will be optimized to a single loss, in contrast to the conventional self-supervised models which are optimized for each loss in different steps. JUST also showed good performance when training the existing multilingual ASR data only without additional data. We use this joint learning concept in this paper for that the LID models recognize the speech context by end-to-end learning without additional training data.

## 3. TWO LEARNING METHODS

### 3.1. Supervised learning

In order to train the LID model by supervised learning, we first prepare pair data of input speech and language label $\mathbf{y}$. We then extract the features $\{\mathbf{x}_t\}_{t=1}^T$ for $T$ time steps from the input speech signal. The features are masked along the time axis with the masking probability $\rho_m$ for masking size $m$. We skip the masking step if the masking size is zero or in the inference process.

After calculating the represented vectors by passing features to the encoder layers, the language probability vector $\mathbf{p} \in \mathbb{R}^N$ is calculated by a pooling layer, a linear classifier, and a softmax layer. Linear classifiers have output size $N$ signifying the number of language classes. Finally, cross-entropy loss for supervised learning is calculated with language probability $\mathbf{p}$ and language label $\mathbf{y}$, as

$$\mathcal{L}_s = -\sum_{i=1}^N y_i \log p_i. \qquad (1)$$

### 3.2. Joint learning

We apply BEST-RQ [11] as unsupervised learning to help the model catch the context from the input speech signal. First, we generate codebook vectors $\{\mathbf{c}_j\}_{j=1}^M$ and a random projection matrix $A \in \mathbb{R}^{D \times F}$ with Xavier uniform initialization [13] before starting training; here, $D$ is the dimension of the codebook vector, $F$ is the frequency dimension of the input feature, and $M$ is the number of codebook vectors. We then normalize the codebook vectors and fix $\{\mathbf{c}_j\}_{j=1}^M$ and $A$.
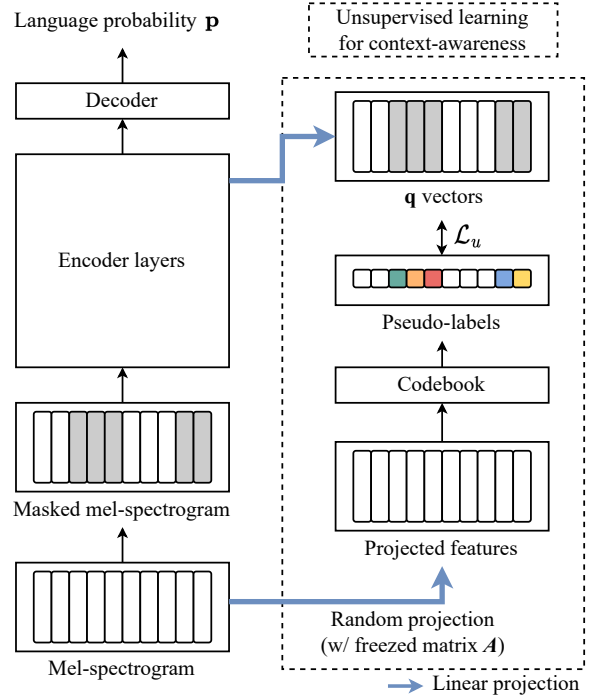


**Fig. 1**: The overall learning process. The left side represents the normal language identification model (baseline) and the right side represents the unsupervised learning process (the sub-sampling concept is omitted).

During training, we generate pseudo-labels $\mathbf{z}_t$ for each time step $t$ by comparing distances among the $\{\mathbf{c}_j\}_{j=1}^M$ to the normalized projected feature calculated by normalizing and multiplying $A$ and $\mathbf{x}_t$:

$$\mathbf{z}_t \leftarrow \arg\min_j |Norm(A\mathbf{x}_t) - \mathbf{c}_j|. \qquad (2)$$

We generate vectors $\{\mathbf{q}_t\}_{t=1}^T$ from the intermediate vectors of an encoder layer by linear projection and softmax. The model will be trained to classify the pseudo-label of the masked position $\mathcal{T}$, and the unsupervised loss is as follows:

$$\mathcal{L}_u = -\frac{1}{|\mathcal{T}|}\sum_{t \in \mathcal{T}}\sum_{j=1}^M z_{t,j} \log q_{t,j}. \qquad (3)$$

In the above process, if the encoder includes a subsampling layer, the input features are stacked to the same degree as the sub-sampling factor. The random projection matrix is also generated to have the same dimension as the stacked input feature.

Joint learning uses the losses of both supervised and unsupervised learning at the same time, and the final loss $\mathcal{L}$ is calculated as

$$\mathcal{L} = (1-\lambda)\mathcal{L}_s + \lambda\mathcal{L}_u, \qquad (4)$$

where $\lambda$ has a weighting value between $[0, 1]$ and is set to 0.5 in our experiments. Fig. 1 shows the overall learning process.

## 4. EXPERIMENTS

### 4.1. Training setup

**Features:** All speech signals were converted into the log-mel spectrogram. The features were extracted with window size 400, hop size 160, fast Fourier transform size 512, and 80 filters. The sample rate was 16,000.

**Architecture:** We used the Conformer(S) [14] encoder including the sub-sampling layer. All parameters are the same as for the original Conformer(S): the number of encoder layers is 16, the encoder dimension is 144, and the number of attention heads is 4. For the decoder, we used an average pooling layer and a linear classifier.

**Unsupervised learning parameters:** We set the codebook vector dimension $D$ as 16 and the number of codebook vectors $M$ as 64, 256, and 1024. The **q** vectors were generated from the output of the 15th encoder layer. The masking probability $\rho_m$ was set differently depending on the masking size to make the average of the total masking region 35%. Masking regions can be overlapped. We conducted the experiment by dividing the masking size into six steps from 80 ms to 480 ms.

**Datasets and training parameters:** Training and evaluation were conducted for 11 languages[1]. To train with a large amount of data for each language class, we collected data from YouTube manually. We listed over a hundred channels for each language, including news, documentary, talk show, beauty channels, etc., and downloaded audio data from each channel. We used our voice activity detection to extract speech segments from the downloaded audio data, acquiring 1,000 hours of utterances for each language for a total of 11,000 hours. The total number of segments is 4.8 million, with an average segment length of 8.16 seconds. For training, the Adam optimizer [15] was used with the peak learning rate set to 0.001 and warm-up set to 5,000 steps. All input signals were cropped to three seconds at a random position during training, and the batch size was 256.

### 4.2. Performance on VoxLingua107-T11

We used the VoxLingua107 [10] dataset for the validation. VoxLingua107 contains 107 languages in the training dataset, while the dev set has only 33 languages. In particular, there are no dev sets for hi, id, ko, and vi among the languages we are targeting. The number of utterances for each language in the dev set is also small, with an average of 48 utterances.

---

[1] English (en), French (fr), German (de), Hindi (hi), Indonesian (id), Italian (it), Japanese (ja), Korean (ko), Mandarin Chinese (zh), Spanish (es), and Vietnamese (vi).
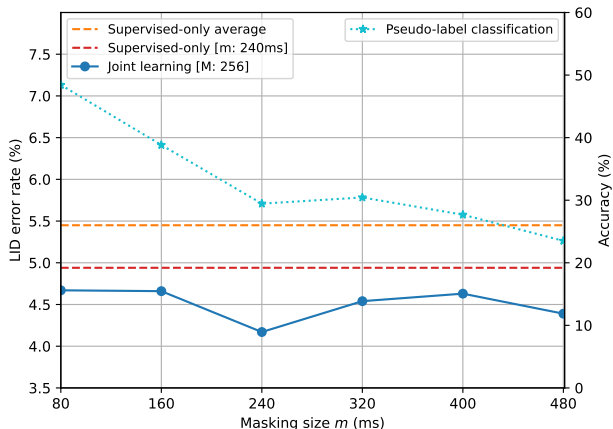


**Fig. 2**: The results for masking size $m$. The left axis represents the error rates (%) of LID; the right axis represents the pseudo-label classification accuracy (%).

**Table 1**: LID classification error rate (%) on VoxLingua-T11. The top three rows show the results according to the model structure and the bottom four rows show the results according to learning methods.

| Method | VoxLingua107-T11(3s) |
| --- | --- |
| w/o masking (#parameters) | |
| ECAPA-TDNN (15.7M) [16] | 6.54% |
| ATP-Conformer (9.5M) [6] | 6.08% |
| Supervised-only (10.3M) | 5.80% |
| w/ masking [$M/m$] | |
| Supervised-only [- / 240 ms] | 4.94% |
| Joint learning [64 / 240 ms] | 4.23% |
| Joint learning [256 / 240 ms] | **4.17%** |
| Joint learning [1024 / 240 ms] | 4.33% |

For this reason, we composed 11 language evaluation datasets (VoxLingua107-T11), collecting sub-three-second utterances from VoxLingua107 training data. The total number of utterances is 3,255, and their average length is 2.54 seconds.

Table 1 shows the classification error rate (%) according to model structure (top three rows) and learning methods (bottom four rows). First, we checked that the attentive temporal pooling Conformer-based LID (ATP-Conformer) [6] showed better performance than the ECAPA-TDNN-based LID [16, 17, 18] (even though the ATP-Conformer has a lower number of parameters). As our base model is slightly larger than the ATP-Conformer, there was a slight performance improvement.

In order to compare the performance of models trained by supervised-only versus joint learning for masking size,

**Fig. 3**: The $t$-SNE plots of two models, trained by supervised-only (left) and joint learning (right). Spanish (es) and Italian (it) were separated more clearly in the joint learning model than in the supervised-only model.

we trained models with variable masking sizes from 0 ms to 480 ms (Table 1 and Fig. 2). Both methods showed the best performance when training with a masking size of 240 ms. Comparing the results of the third and fourth lines of Table 1, it can be seen that even the models trained with the same supervised-only learning show performance differences depending on the presence of masking, which seems to be the effect of SpecAugment [19]. Joint learning once again significantly improved performance. As a result, the best error rates for supervised-only and joint learning are 4.94% and 4.17%, yielding a relative 15.6% error rate reduction by the proposed joint learning method.

Fig. 2 shows the LID error rates and pseudo-label classification accuracy for masking size. Pseudo-label accuracy was 48.43% when the masking size was 80 ms and decreased to 23.51% according to the masking size increase to 480 ms. We found that pseudo-label classification accuracy is not directly related to the LID error rate and that the proposed joint learning method showed a lower LID classification error rate than the best model trained as supervised-only for all masking cases. The supervised-only models recorded an average error rate of 5.45% while the joint learning models recorded an average error rate of 4.50%.

### 4.3. Performance for each language

Table 2 shows the performance of two models that were trained by different methods for each language. Both models performed relatively poorly in Vietnamese, Indonesian, and Spanish. There were many cases where Vietnamese was incorrectly recognized as Indonesian, and Spanish was incorrectly recognized as Italian (or the reverse). Comparing the average $F_1$ scores, the joint learning model performed better than the supervised-only model, and there was a significant performance improvement in Indonesian (0.909 → 0.931) and Spanish (0.874 → 0.912). Fig. 3 shows the $t$-SNE plots

**Table 2**: $F_1$ score, precision ($P$), and recall ($R$) of the models as trained by two methods for each language.

| Lang | Supervised-only best | | | Joint learning best | | |
|------|-------|-------|-------|-------|-------|-------|
|      | $F_1$ | $P$ | $R$ | $F_1$ | $P$ | $R$ |
| en | 0.948 | 0.988 | 0.912 | **0.954** | 0.988 | 0.923 |
| fr | **0.969** | 0.965 | 0.974 | 0.968 | 0.973 | 0.962 |
| de | **0.967** | 0.967 | 0.967 | 0.961 | 0.937 | 0.985 |
| hi | 0.961 | 0.945 | 0.977 | **0.976** | 0.957 | 0.995 |
| id | 0.909 | 0.887 | 0.933 | **0.931** | 0.902 | 0.962 |
| it | 0.944 | 0.947 | 0.941 | **0.953** | 0.940 | 0.967 |
| ja | **0.966** | 0.965 | 0.968 | 0.962 | 0.974 | 0.950 |
| ko | 0.981 | 0.994 | 0.968 | **0.983** | 0.992 | 0.974 |
| zh | **0.953** | 0.920 | 0.987 | 0.952 | 0.911 | 0.996 |
| es | 0.874 | 0.876 | 0.872 | **0.912** | 0.956 | 0.872 |
| vi | 0.923 | 0.962 | 0.887 | **0.936** | 0.985 | 0.892 |
| Avg | 0.945 | 0.947 | 0.944 | **0.953** | 0.956 | 0.953 |

of the two models trained by the different methods, and we can see that Spanish and Italian were separated more clearly in the joint learning model.

## 5. CONCLUSION

We proposed context-aware language identification by jointly unsupervised and supervised learning. In order to show the effect of our proposed joint learning approach on the wild datasets, we collected 11,000 hours of speech data for 11 languages from YouTube for training and used a subset of the VoxLingua107 with sub-three-second utterances. In experiments, the proposed joint learning reduced the error rate by 15.6% compared to the same structure model trained by supervised-only learning with the same training data.

# 6. REFERENCES

[1] Haizhou Li, Bin Ma, and Kong Aik Lee, "Spoken language recognition: from fundamentals to practice," *Proceedings of the IEEE*, vol. 101, no. 5, pp. 1136–1159, 2013.

[2] Marc A Zissman, "Comparison of four approaches to automatic language identification of telephone speech," *IEEE Transactions on Speech and Audio Processing*, vol. 4, no. 1, pp. 31, 1996.

[3] Marc A Zissman and Kay M Berkling, "Automatic language identification," *Speech Communication*, vol. 35, no. 1-2, pp. 115–124, 2001.

[4] Ignacio Lopez-Moreno, Javier Gonzalez-Dominguez, Oldrich Plchot, David Martinez, Joaquin Gonzalez-Rodriguez, and Pedro Moreno, "Automatic language identification using deep neural networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 5337–5341.

[5] Javier Gonzalez-Dominguez, Ignacio Lopez-Moreno, Pedro J Moreno, and Joaquin Gonzalez-Rodriguez, "Frame-by-frame language identification in short utterances using deep neural networks," *Neural Networks*, vol. 64, pp. 49–58, 2015.

[6] Quan Wang, Yang Yu, Jason Pelecanos, Yiling Huang, and Ignacio Lopez Moreno, "Attentive temporal pooling for conformer-based streaming language identification in long-form speech," *arXiv preprint arXiv:2202.12163*, 2022.

[7] Zhiyuan Tang, Dong Wang, Yixiang Chen, Lantian Li, and Andrew Abel, "Phonetic temporal neural model for language identification," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 1, pp. 134–144, 2017.

[8] Raphaël Duroselle, Md Sahidullah, Denis Jouvet, and Irina Illina, "Modeling and training strategies for language recognition systems," in *INTERSPEECH*, 2021.

[9] Surabhi Punjabi, Harish Arsikere, Zeynab Raeesy, Chander Chandak, Nikhil Bhave, Ankish Bansal, Markus Müller, Sergio Murillo, Ariya Rastrow, Andreas Stolcke, et al., "Joint ASR and language identification using RNN-T: An efficient approach to dynamic language switching," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2021, pp. 7218–7222.

[10] Jörgen Valk and Tanel Alumäe, "VoxLingua107: a dataset for spoken language recognition," in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 652–658.

[11] Chung-Cheng Chiu, James Qin, Yu Zhang, Jiahui Yu, and Yonghui Wu, "Self-supervised learning with random-projection quantizer for speech recognition," in *International Conference on Machine Learning*. PMLR, 2022, pp. 3915–3924.

[12] Junwen Bai, Bo Li, Yu Zhang, Ankur Bapna, Nikhil Siddhartha, Khe Chai Sim, and Tara N Sainath, "Joint unsupervised and supervised training for multilingual ASR," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6402–6406.

[13] Xavier Glorot and Yoshua Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.

[14] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al., "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*, 2020.

[15] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.

[16] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *INTERSPEECH*, 2020.

[17] Rohan Kumar Das, Ruijie Tao, and Haizhou Li, "HLT-NUS submission for 2020 NIST conversational telephone speech SRE," *arXiv preprint arXiv:2111.06671*, 2021.

[18] Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, Jianyuan Zhong, Ju-Chieh Chou, Sung-Lin Yeh, Szu-Wei Fu, Chien-Feng Liao, Elena Rastorgueva, François Grondin, William Aris, Hwidong Na, Yan Gao, Renato De Mori, and Yoshua Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, arXiv:2106.04624.

[19] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin Dogus Cubuk, and Quoc V. Le, "SpecAugment: A simple augmentation method for automatic speech recognition," in *INTERSPEECH*, 2019.