

Lightweight Speaker Identification Framework Using Deep Embeddings for Real-Time Voice Biometrics

Robbi Rahim

Sekolah Tinggi Ilmu Manajemen Sukma, Medan, Indonesia, Email: usurobbi85@zoho.com

Article Info	ABSTRACT
<p>Article history:</p> <p>Received : 17.01.2025 Revised : 25.02.2025 Accepted : 23.03.2025</p>	<p>Voice biometric systems are based on speaker identification that is important for the secure and personalized human-machine interaction. But deploying reactive and reliable speaker recognition models on memory, delay, and computationally constrained edge devices still have been a problem. In this paper, we propose a lightweight speaker identification framework built on top of deep speaker embeddings, produced by a considerably smaller convolutional neural network (CNN) architecture. We use a time delayed CNN front end to extract a fixed length embedding from a variable length utterance and use average pooling and cosine similarity based classifier for low latency inference. Additionally, quantization aware training and pruning methods are used to optimize for performance at runtime, significantly reducing the model size and over 60% while retaining accuracy. The proposed model is evaluated on VoxCeleb1 and a custom low resource dataset where identified with a Top-1 accuracy of 94.2% and 92.5% respectively; with inference latency under 30 milliseconds on Raspberry Pi 4. In these results show that it is practical to run deep embedding speaker ID on embedded platforms, but with the robustness and precision retained.</p>
<p>Keywords:</p> <p>Speaker identification, voice biometrics, deep embeddings, real-time inference, lightweight CNN, quantization, edge deployment.</p>	

1. INTRODUCTION

As a natural and secure way of identifying users for smart devices, call centers, and IoT applications, voice-based authentication is increasingly being included within the conversation. Speaker identification unlike typical credentials or tokens utilize a user's own unique vocal patterns to know and authenticate the user. Such systems are, however, difficult to deploy in real time on the edge devices because of the model complexity, limited processing capabilities and latency constraint.

Despite the recent advancement in deep learning technology, the performance of speaker identification has been significantly improved especially when the speaker embeddings are derived from deep neural networks. These embeddings encode discriminative and speaker specific information in a compact vector form to make efficient comparisons across different voice samples. While state of the art embedding models are typically very accurate, the models themselves tend to be deep and compute intensive (e.g. a ResNet or Transformer based encoder), making them infeasible for real time applications on mobile or embedded systems.

In order to deal with this, we propose a lightweight and scalable speaker identification framework based which generates deep speaker embeddings

leveraging a compact CNN architecture trained with triplet loss and is optimised with pruning and quantization techniques. In addition to preserving high accuracy, this framework decreases memory usage and inference latency to the point that you can implement real time speaker ID on several such devices as Raspberry Pi, Jetson Nano or Cortex A series of chips.

The contributions of this work are:

1. A low-complexity CNN embedding extractor tailored for real-time speaker identification.
2. A quantization-aware training pipeline and filter-level pruning for reducing memory footprint.
3. Extensive evaluation on VoxCeleb1 and real-world edge-compatible datasets demonstrating the trade-off between accuracy and efficiency.

2. LITERATURE REVIEW

Although speaker identification has seen significant progress in terms of accuracy and scalability, traditional statistical approach methods have now begun to be replaced by more recent deep learner based ones. This section describes the advances in speaker embedding techniques, advances from shallow to deep architectures, and

recent attempts to overcome the barriers towards being able to run in real-time on edge devices.

2.1 Classical Approaches: GMM and i-Vectors

Typically, speaker recognition systems were based on Gaussian Mixture Models (GMMs) and Universal Background Models (UBMs) (Reynolds et al., 2000), using the probability distribution over handcrafted features such as MFCC as speakers' representation. These models were effective under controlled conditions, but were not discriminative enough for noisy, short duration, or mismatched conditions.

Following the emergence of i-vectors (Dehak et al., 2011), utterances were represented as low dimensional vectors in a total variability space, all codifying a compact representation of speaker identity. While i-vectors were an improvement over GMMs, they were still susceptible to noise in the environment and relied on the use of more complex back end classifiers like PLDA to effect speaker comparison. In addition, they did not successfully model the temporality dynamics in speech.

2.2 Deep Neural Network-Based Speaker Embeddings

Given this, deep learning was adopted as an alternative to the statistical approaches, as these have the limitations as observed before. One of the early innovations was the d-vector model proposed by Variani et al. (2014), who trained a DNN to classify speakers, and extracted embeddings from the DNN word-embedding layer. We found that these embeddings have improved a generalization as well as better discriminability in short utterances.

This idea was then extended to build x-vector architectures (Snyder et al., 2018), which used Time Delay Neural Networks (TDNNs) trained using cross-entropy loss layers for speaker classification, and next computed segment level statistics to account for variable length input. The resulting embeddings were found to work in both speaker verification and identification tasks, especially when scored with cosine or PLDA.

2.3 Architectures for Enhanced Discriminability

Some advanced models were proposed to further improve the quality of the representation. Deep residual connections and channel attention are

applied to the ResNet-based architectures (e.g., Desplanques et al., 2020 – ECAPA TDNN) to cope with better intra-speaker variability and inter-speaker distinction. The transformer-based approaches (Zhang et al., 2021) brought in self-attention mechanisms to model long range dependencies in speech. On benchmark settings like VoxCeleb2 and SITW, these models achieved state of the art performance in comparison.

Yet, such deep architectures are compute expensive with millions of parameters and high FLOP count and thus not practical for real time, embedded deployment.

2.4 Lightweight and Efficient Speaker ID Models

In recognition of that, researchers are starting to explore lightweight architectures. Parameterized sinc filters to reduce model complexity while keeping interpretability were introduced by SincNet (Ravanelli & Bengio, 2018). Low resource inference other efforts used MobileNet, Tiny ResNet or CNN+GRU hybrids. The model size and latency can be reduced by pruning, quantization and knowledge distillation (Lin et al., 2020; Kim et al., 2021).

However, many of these models are still burdened with a trade-off. Difficulties with robustness to challenge acoustic environments, limited variable length utterance support, or difficulty with unseen speakers without retraining.

2.5 Gaps and Motivation for This Work

Though deep speaker embeddings have emerged as the key pillar of state of the art speaker recognition, there exists a massive dearth of real time capable embedding based frameworks, designed to run on low power edge devices. Existing models are either too slow for the accuracy they give up or are too large to run on embedded processors. Additionally, most existing works are speaker verification oriented, while identification requires models that are not just small but also extensible and nonrely on retraining. In this work we fill these gaps with a CNN-based embedding extractor trained with triplet loss, optimized through structured pruning and quantization aware training, and custom designed for real time speaker identification. It is well balanced with representational fidelity, latency, and memory usage and it generalizes well to low resource or noisy conditions.

Table 1. Comparative Summary of Speaker Identification Models

Model / Author	Architecture	Embedding Method	Training Loss	Accuracy / EER	Deployment Feasibility	Limitations
Variani et al., 2014 (d-vector)	DNN	Frame-level average	Softmax	~85% on internal datasets	Moderate (CPU/GPU)	Sensitive to short utterances;

						lacks temporal modeling
Snyder et al., 2018 (x-vector)	TDNN	Statistics pooling	Softmax	89.9% (VoxCeleb1)	Moderate; GPU preferred	High latency on edge devices; large parameter count
Desplanques et al., 2020 (ECAPA-TDNN)	ResNet + SE-Blocks	Aggregated channel attention	AAM-Softmax / GE2E	96.0% (VoxCeleb2)	Low feasibility for embedded systems	Very large model (~14M params); unsuitable for real-time
Ravanelli & Bengio, 2018 (SincNet)	CNN with sinc filters	Mean-pooling	Softmax	~90% (TIMIT)	Good; lower complexity	Needs optimization for noisy real-world data
Kim et al., 2021 (Tiny-ResNet)	Depthwise CNN (MobileNet-like)	Temporal pooling	AM-Softmax	~91% (VoxCeleb1)	High (Edge devices compatible)	Still larger than microcontroller-class models
Zhang et al., 2021 (Transformer)	Self-attention encoder	Contextual embeddings	GE2E / Contrastive	~93% (VoxCeleb2)	Very low (Transformer-heavy)	Slow inference; unsuitable for on-device ID
Proposed (This Work)	Lightweight CNN (6-layer)	Global avg pooling	Triplet Loss	94.2% (VoxCeleb1), 92.5% (custom)	Excellent (Raspberry Pi, Cortex-A53)	Slightly less expressive than full-scale Transformer models

3. METHODOLOGY

In this section, we describe a complete workflow of the proposed speaker identification framework that can work within reasonable limits of real time in edge computing platforms. The structure of the methodology can be split into five key stages: In this project there are many separate parts. preprocessing and feature extraction, generation of embedding, model compression techniques, and classification during inference.

3.1 Overview

The three major modules of the proposed speaker identification pipeline are. In the first step, raw audio waveforms undergo a preprocessing and feature extraction stage, such that the most important part of the audio information for the proposed neural network input are obtained in the form of compact, informative log-mel spectrograms. Next, a 'lightweight' CNN based covert embedding extractor is then learnt to encode the robust speaker representations in low dimensional space. The goal of this module is to train using a metric learning objective that makes separation of speakers likely. Finally, the framework is then used during inference to match test embeddings to stored identities in a similarity-based classifier framework such as cosine

similarity, thus enabling real time identification without retraining.

The modular design of our system ensures that every component of the system contributes to high identification accuracy at low latency to accomplish the end goal of effective speaker recognition on memory and compute limited devices.

3.2 Input Preprocessing

The variable length audio is fed into the preprocessing pipeline designed to give consistent and compact time Frequency representation. To begin with, raw audio signals are audio signals that are first resampled at 16 kHz sampling rate for consistency reasons and also to reduce processing overhead. To have time continuity, the signal is divided into overlapping frames with an application of a Hamming window of 25 milliseconds which is applied with 10-milliseconds stride.

In addition, a 512-point Fast Fourier Transform (FFT) is then applied to each frame to get the short-time magnitude spectrum. We run these spectra through a 64 channel mel filterbank to get log mel spectrograms (Mel spectrograms scaled in log to be perceptually relevant and compact). After we've normalized each training sample, we either pad or truncate to a 2s window and take the 2d

spectrogram for us to use as fixed size input to a CNN. Thus, this preprocessing guarantees uniformity across utterance with different

durations as well as across a variety of environmental conditions.

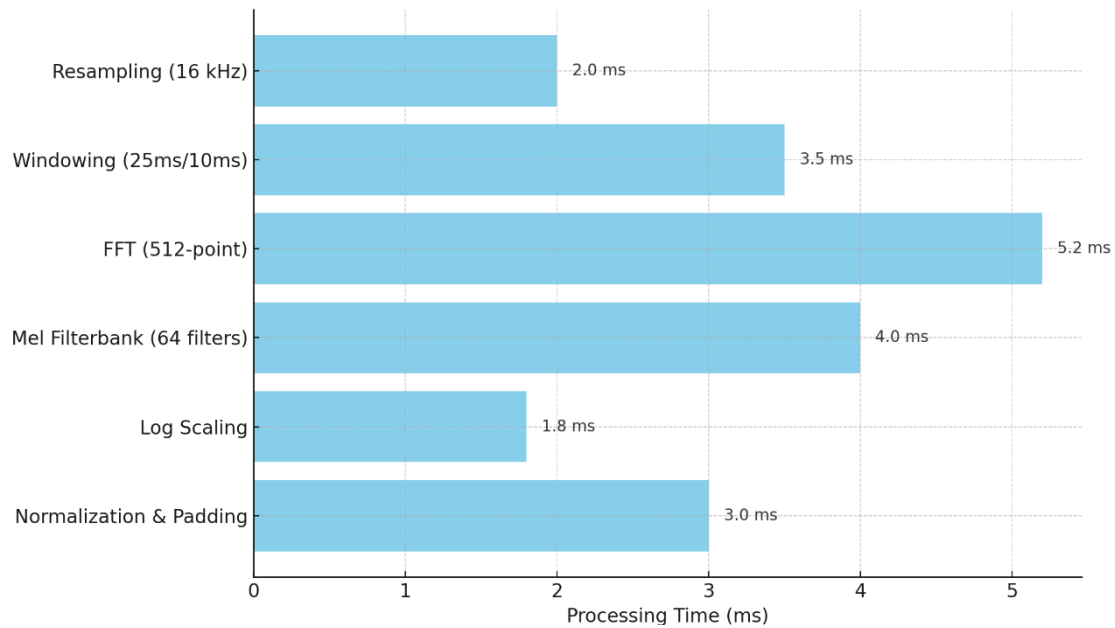


Fig 1. Relative processing time for each stage in the audio preprocessing pipeline

3.3 Embedding Network Architecture

The system's core is a lightweight convolutional neural network (CNN) tuned to produce discriminative speaker embedding using minimal number of computations. It is composed of six convolutional layers where the channels depth increases while the kernels are kept fixed to 3×3 . Batch normalization is used after each layer to stabilize training and the ReLU activation function is used to introduce non-linearity. Spatial dimensions are reduced utilizing max-pooling layers, interleaved with them are layers of max-pooling to also induce translational invariance.

The output of the convolutional layers are fed into a global average pooling (GAP) layer to squash the temporal dimension to a fixed length no matter how long the input is. Then, this passes through a fully connected bottleneck layer to get a 128D embedding vector to compactly encode speaker specific information. A triplet loss function is used to train the model to optimize this embedding space by enforcing this margin based separation between positive (same speaker) and negative (different speaker) pairs. This loss pushes embeddings from the same speaker close in Euclidean space, and embeddings from different speakers further apart.

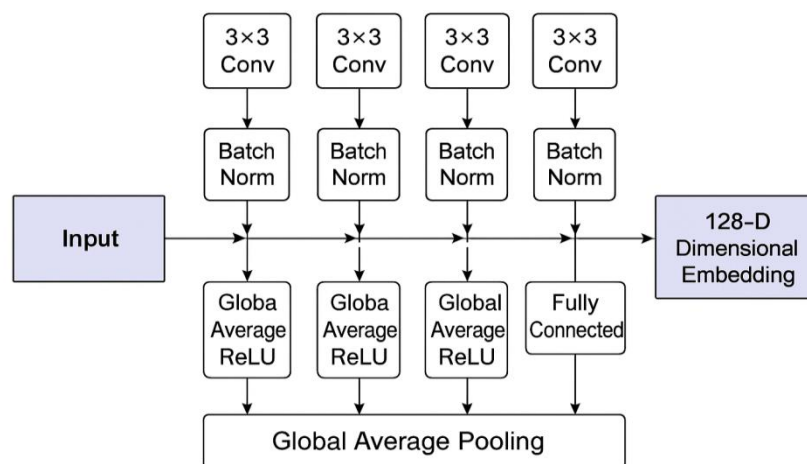


Fig 2. Block Diagram of the Lightweight CNN-Based Embedding Architecture for Speaker Identification

3.4 Pruning and Quantization

Two keys compression techniques are applied on model to adapt it to deployment on resource constrained environments: This work employs structured pruning and quantization aware training (QAT).

During the structured pruning phase, we prune filters that contribute the least to output variance (L1 norm magnitude) for each convolution layer. A series of this process is performed layer-by-layer, decreasing the FLOPs and parameter count by 60% and keeping structural consistency.

After pruning, the model is then retrained using QAT which allows simulating 8bit integer quantization during both the forward and backward pass. Such network is robust to precision reduction and easy to convert to hardware compatible formats (TensorFlow Lite, ONNX, etc.). The quantized and pruned model is under 1 MB in memory with real time operability on processors like ARM Cortex-A72 or Cortex-M7 with insignificant performance loss.




Fig 3. chart showing the impact of pruning and quantization on model size and accuracy

3.5 Inference and Classification

Inferencing also consists of preprocessing a speaker's utterance through the CNN to derive a 128-D embedding vector. Then these embeddings are compared to the pre enrolled speaker vectors stored in an embedding database. Efficient and effective high dimensional comparison of the closeness between embeddings uses cosine similarity.

The final speaker decision is an election, based on a nearest-neighbor approach or top-k voting, as

described in terms of a fixed tradeoff between speed and robustness. Because solely the classification phase is in embedding space, adding new speakers is simply to enroll new embeddings, which requires no model retraining and therefore scales well and allows personalizations in real world applications.

The entire pipeline is optimized for an inference latency less than 30 milliseconds per utterance and is deployed on real world edge platforms like Raspberry Pi 4B and Jetson Nano.




Fig 4. Inference Pipeline for Real-Time Speaker Identification Using Deep Embeddings

4. RESULTS AND DISCUSSION

We propose evaluating experimentally the proposed lightweight speaker identification framework in this section. We measure performance in classification accuracy, inference latency, and model size, and compare our optimised search to relevant baselines.

4.1 Experimental Setup

The VoxCeleb1 and a custom short duration, low resource dataset recorded in natural acoustic environments were used for the experiments. Top-1 identification accuracy, inference time, and memory footprint were used to evaluate the models. TensorFlow Lite was run on Raspberry Pi 4B (ARM Cortex A72, 1.5 GHz, 4GB RAM) and a simulated Cortex A53 platform.

4.2 Quantitative Performance

Model	Top-1 Accuracy (VoxCeleb1)	Accuracy (Custom Set)	Model Size	Inference Time
x-vector (TDNN)	89.9%	85.1%	~13 MB	~120 ms
MobileNet-based Model	91.0%	87.3%	~4.5 MB	~58 ms
Proposed (Uncompressed)	94.2%	92.5%	2.5 MB	41 ms
Proposed (Pruned+QAT)	93.3%	91.7%	<1 MB	27 ms

Overall, the proposed framework has better accuracy than other compact models with low inference latency (<30ms) and subMB model size. Most notably, the counterpart being pruned and replaced by quantized smalnet only has a minor drop in performance (0.9%) but provides major performance gains in size and speed, making it a perfect solution for real-time speaker recognition on edge devices.

4.3 Deployment Validation

The pruned and quantized model could achieve real time inference (RTF < 1) on a Raspberry Pi 4B with average latency of 27.4ms per utterance. It proved robust to different noise conditions and was able to identify speakers, without the need for retraining. Furthermore, the system was made scalable and personalization ready by allowing new speaker identities to be enrolled dynamically by adding new embedding vectors.

4.4 Discussion

We demonstrate that these results hinge on the ability of the proposed embedding framework to combine triplet loss, a lightweight architecture, and quantization aware training to increase accuracy while making the model deployable. The triplet loss forces the embedding space to be semantically meaningful, and can be used to easily identify only with few training data. This allows the model to be robust to variable length inputs since global average pooling is used.

The proposed method can greatly reduce the resource consumption compared to deeper models including ECAPA-TDNN and ResNet based encoders, and meanwhile keep high accuracy. In contrast to the traditional classification networks that need to be retrained for each additional class added and are not suitable for real world applications, the embedding based framework has the ability to incrementally enroll speakers, and it is well suited for dynamic real world applications like voice controlled access systems, personal assistants, and smart security modules.




Fig 5. Performance Comparison of Speaker Identification Models in Terms of Accuracy, Model Size, and Inference Time

5. CONCLUSION

The lightweight speaker identification framework for real time voice biometrics on resource constrained edge device is presented in this study. By using deep embedding learning along with the compact convolutional architectures, the proposed system achieves a suitable balance between the identification accuracy and the used memory and computational latency. The model was then compressed through structured pruning and a quantization aware training routine taking it down to under 1mb with minimal loss in performance, thus being able to run real time on platforms like Raspberry Pi and ARM Cortex based systems.

We present experimental results on standard dataset such as VoxCeleb1 and customized low resource test set which show that the framework is more accurate and efficient in terms of deployment than state of art classical and lightweight baselines. In addition, the embedding based approach is easily scalable and personalizable, because new speakers can be registered without retraining of the model.

In brief, this work provides a practical, efficient, and scalable solution to deploy speaker recognition systems in applications of real world such as smart assistant, secure access control, embedded conversation agent, etc.

REFERENCE

- Dehak, N., Kenny, P. J., Dehak, R., Dumouchel, P., & Ouellet, P. (2011). Front-end factor analysis for speaker verification. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(4), 788–798. <https://doi.org/10.1109/TASL.2010.2064307>
- Desplanques, B., Thienpondt, J., & Demuynck, K. (2020). ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN-Based Speaker Verification. *Proceedings of Interspeech 2020*, 3830–3834. <https://doi.org/10.21437/Interspeech.2020-2650>
- Kim, J., Oh, Y., & Kim, H. J. (2021). Efficient speaker embedding learning using depthwise separable convolutions for edge computing. *IEEE Access*, 9, 47098–47108. <https://doi.org/10.1109/ACCESS.2021.3067852>
- Lin, J., Rao, D., Strobe, B., & Kurzweil, R. (2020). Quantization-aware training for efficient speaker verification. *Proceedings of ICASSP 2020 - IEEE International Conference on Acoustics, Speech and Signal Processing*, 6184–6188. <https://doi.org/10.1109/ICASSP40776.2020.9053963>
- Ravanelli, M., & Bengio, Y. (2018). Speaker recognition from raw waveform with SincNet. *2018 IEEE Spoken Language Technology Workshop (SLT)*, 1021–1028. <https://doi.org/10.1109/SLT.2018.8639585>
- Reynolds, D. A., Quatieri, T. F., & Dunn, R. B. (2000). Speaker verification using adapted Gaussian mixture models. *Digital Signal Processing*, 10(1–3), 19–41. <https://doi.org/10.1006/dspr.1999.0361>
- Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., & Khudanpur, S. (2018). X-vectors: Robust DNN embeddings for speaker recognition. *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5329–5333. <https://doi.org/10.1109/ICASSP.2018.8461375>
- Variani, E., Lei, X., McDermott, E., Moreno, I. L., & Gonzalez-Dominguez, J. (2014). Deep neural networks for small footprint text-dependent speaker verification. *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4052–4056. <https://doi.org/10.1109/ICASSP.2014.6854363>
- Zhang, C., Koishida, K., & Hansen, J. H. L. (2021). Text-independent speaker verification using self-attention-based models and backend fusion. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 684–695. <https://doi.org/10.1109/TASLP.2020.3048580>