

# AM-MobileNet1D: A Portable Model for Speaker Recognition

João Antônio Chagas Nunes<sup>1</sup>, David Macêdo<sup>1,2</sup>, and Cleber Zanchettin<sup>1,3</sup>

<sup>1</sup>Centro de Informática, Universidade Federal de Pernambuco, Recife, Brasil

<sup>2</sup>Montreal Institute for Learning Algorithms, University of Montreal, Quebec, Canada

<sup>3</sup>Department of Chemical and Biological Engineering, Northwestern University, Evanston, United States of America

Emails: {jacn2, dlm, cz}@cin.ufpe.br

**Abstract**—Speaker Recognition and Speaker Identification are challenging tasks with essential applications such as automation, authentication, and security. Deep learning approaches like SincNet and AM-SincNet presented great results on these tasks. The promising performance took these models to real-world applications that becoming fundamentally end-user driven and mostly mobile. The mobile computation requires applications with reduced storage size, non-processing and memory intensive and efficient energy-consuming. The deep learning approaches, in contrast, usually are energy expensive, demanding storage, processing power, and memory. To address this demand, we propose a portable model called Additive Margin MobileNet1D (AM-MobileNet1D) to Speaker Identification on mobile devices. We evaluated the proposed approach on TIMIT and MIT datasets obtaining equivalent or better performances concerning the baseline methods. Additionally, the proposed model takes only 11.6 megabytes on disk storage against 91.2 from SincNet and AM-SincNet architectures, making the model seven times faster, with eight times fewer parameters.

**Index Terms**—Speaker Identification, SincNet, AM-SincNet, MobileNet, Portable Deep Learning.

## I. INTRODUCTION

Speaker Recognition is a very active area with applications in biometric authentication, identification, security, among others [2]. This area has two principal tasks: Speaker Identification and Speaker Verification. Figure 1 shows a representation of these two tasks. Speaker Identification (Figure 1 (a)) is the task of identifying who spoke something among a determined list of speakers. Whereas, Speaker Verification (Figure 1 (b)) aims to classify whether an audio sample belongs to a pre determined person or not. Before the Deep Learning models, most of the literature techniques to tackle this problem were based on *i*-vectors methods [3], [4], [5]. These methods try to find patterns on audio signals and classify them using techniques such as Probabilistic Linear Discriminant Analysis (PLDA) [6], heavy-tailed PLDA [7], and Gaussian PLDA [8].

The Speaker Recognition task is challenging because of the high dimensionality found on audio signals, which are complex to be modeled in low and high-level features. These features need to be good enough to distinguish different speakers. Over the years, pattern recognition, image processing, and signal processing tasks, are being successfully solved by Deep Neural approaches [9], [10], [11]. Deep learning based models learn complex patterns that may not be spotted by humans or even traditional machine learning approaches.

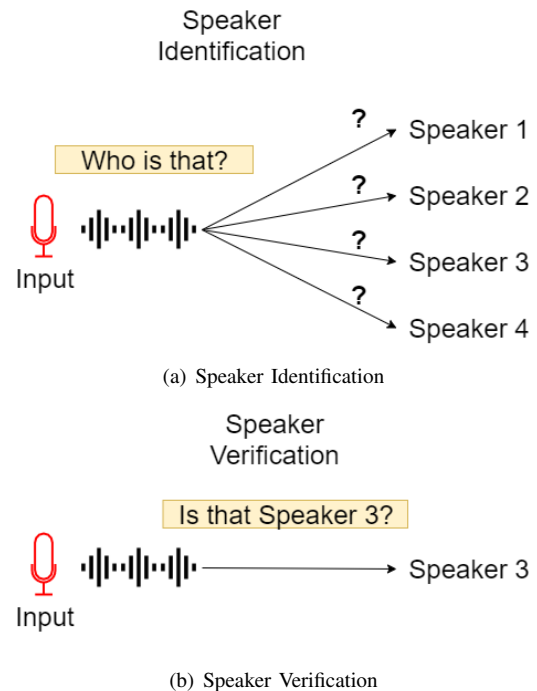


Fig. 1. Difference between Speaker Identification and Verification.

These methods are more complex but usually produce better results than conventional approaches. Indeed, Convolutional Neural Networks (CNN) [12] have already proved to be the best choice for image classification, detection, and recognition tasks [13], [14], [15].

An exciting Deep Learning-based model to Speaker Recognition is the SincNet model [16]. It is a hybrid model, like [17], [18], [19], [20], [21], that uses Sinc convolutional layers in its first layer. The Sinc convolutional layer uses parameterized Sinc functions to convolve the waveform of the audio signal and uses this information to Speaker Identification. A more efficient approach to the problem is the AM-SincNet [22], which is a combination of the SincNet and the AM-Softmax loss function [1]. The AM-SincNet solves the problem of the traditional SincNet by replacing the Softmax layer with an improved AM-Softmax layer. In fact, this change was enough to show an improvement of approximately 40% in the Frame




Fig. 2. Comparison between Softmax (A) and AM-Softmax (B). Adapted from [1].

Error Rate when compared to the traditional SincNet.

These models are being applied to real-world purposes to aggregate value, performance, and usability to end-user applications. However, end-user applications are becoming more mobile. The mobile devices revolutionized many aspects of how we live our lives. From text messaging, to map apps for localization, and always having a camera to document what we see and experience, the smartphones, and other portable computations are incredibly useful devices. As the massive use of these dispositive, more server applications are migrating to the client device for better usability, performance, and quality.

The problem with Deep Learning models in this context is that they are not mobile; they are often large size models that require substantial computational effort, memory, storage size, and energy to be executed. Most of them need dedicated hardware, like Graphics Processing Units (GPUs), to have a reasonable inference time.

Some deep learning approaches, as the MobileNet model [23], are designed to work on mobile devices. The MobileNet architecture introduces the concept of depthwise separable convolutions that enables it to be lighter and faster than the traditional deep learning models. The depthwise separable convolutions break the convolution process into two parts: the depthwise and the pointwise operation. It impacts on reducing the number of parameters on the convolutional layer and improving the time to proceed a convolution.

In this paper, we proposed two new deep learning models based on the MobileNetV2 [24] and the AM-Softmax [1] loss function to deal directly with raw audio signals for Speaker Identification task. The proposed models are MobileNet1D, and AM-MobileNet1D, both of them evaluated against the SincNet and AM-SincNet models on the TIMIT [25] and MIT [26] datasets. The experiments consider the Frame Error Rate (FER), Classification Error Rate (CER), the inference time, the number of parameters, and the model size. On the tested scenarios, the proposed methods show to be competitive, getting equally good, or even better results than the SincNet and AM-SincNet. In addition, the proposed methods are about 7 times faster, have 8 times fewer parameters, and take up to 8 times less space on disk than the SincNet and the AM-SincNet.

The following sections are organized as: Section II, we present the background of this paper, the proposed methods are introduced in Section III, Section IV explains how the experiments were built, the results are discussed at Section V, and finally at Section VI we made our conclusions.

## II. BACKGROUND

Deep learning models usually use the Softmax loss function on its final layer to map the extracted features into different classes. It works by delimiting a linear surface, which is used as a decision boundary to split the samples into their respective categories. The problem with the Softmax decision boundary is that it does not maximize the distance of samples from different classes, which may cause misclassification of samples that are placed too close to the linear separation surface. This characteristic may spoil the Softmax efficiency on tasks that require to measure the distance between samples to make a decision. A new loss function based on the Softmax was proposed by [1] to solve this problem, it is called Additive Margin Softmax (AM-Softmax) and introduces an additive margin to the Softmax decision boundary. The additive margin forces samples from the same class to lay closer to each other, while maximizes the distance of samples from different classes.

The Additive Margin Softmax (AM-Softmax), proposed by [1] to deal with face recognition problems, is a based Softmax loss function that introduces an additive margin of separation to the boundaries of each class. It means that each linear surface of separation has an additional region (the additive




Fig. 3. MobileNetV2 Inverted Residual Block. Adapted from [24].

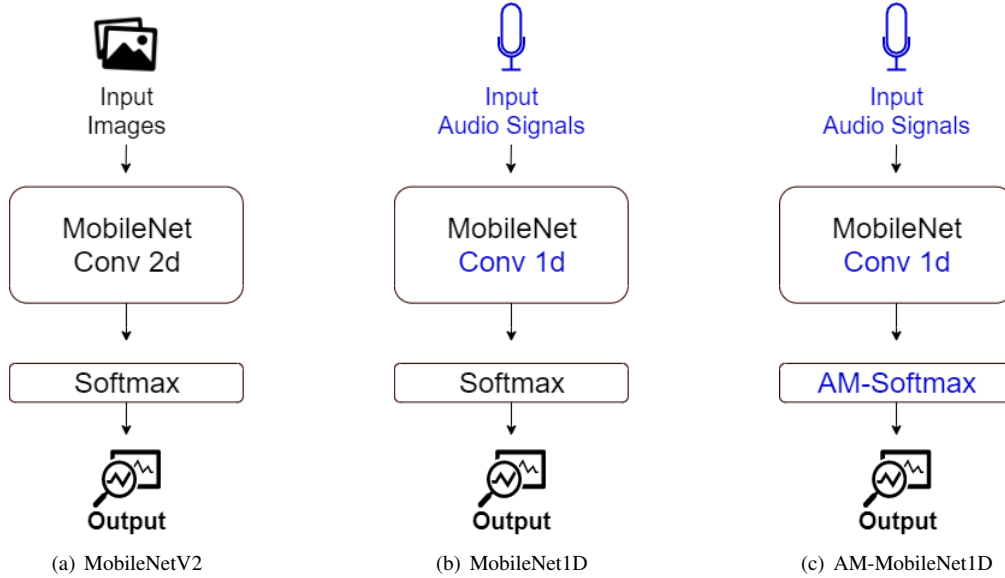


Fig. 4. MobileNet and the proposed architectures.

margin) that may not be used for any sample, a safe zone. An implication of the additive margin used on the AM-Softmax is that it maximizes the distance between samples from different classes, while at the same time, it forces the samples from the same class to be closer to each other. Figure 2 shows the behavior of the samples when the AM-Softmax is used comparing traditional Softmax (A) and the AM-Softmax (B). The AM-Softmax added two more parameters in comparison to traditional Softmax. The new parameter  $s$  scale the samples and parameter  $m$  controls the size of the additive margin. The AM-Softmax equation is written as:

$$Loss = -\frac{1}{n} \sum_{i=1}^n \log \frac{\phi_i}{\phi_i + \sum_{j=1, j \neq y_i}^c \exp(s(W_j^T f_i))} \quad (1)$$

$$\phi_i = \exp(s(W_{y_i}^T f_i - m)) \quad (2)$$

where  $W$  is the weight matrix, and  $f_i$  is the input from the  $i$ -th sample for the last fully connected layer. The  $W_{y_i}^T f_i$  is also known as the target logit for the  $i$ -th sample. The  $s$  and  $m$  are the parameters responsible for scaling and additive margin, respectively.

Besides being applied to face recognition, the AM-Softmax was also applied to speaker recognition by [22]. The AM-SincNet improves the results of the SincNet [16]. The SincNet overcomes traditional convolutional networks on the mentioned task. However, it uses the Softmax loss function on its last layer. According to [22], the Softmax loss function was spoiling the SincNet results, so the authors proposed the AM-SincNet model. The AM-SincNet model is based both in the SincNet architecture and the AM-Softmax loss function. The AM-SincNet improves the SincNet results by replacing a traditional Softmax layer to an improved AM-Softmax layer. The method was evaluated on TIMIT [25] dataset and reduced the Frame Error Rate in about 40% compared to the

traditional SincNet. Despite the good results on the speaker identification task, if we consider the mobile application of the model, the SincNet is not appropriated. The model needs too much computational effort to make predictions in a reasonable inference time.

A lightweight and fast deep learning model is the MobileNetV2 [24], which is a new version of the MobileNet [23], both of them were built to work with images in classification, detection, and segmentation tasks. MobileNetV2 is a reasonable choice when applications need to run on mobile devices that have limited computational resources and cannot support large size models. In comparison to the MobileNet, the MobileNetV2 introduces the Inverted Residual Block, shown in Figure 3, where the thickness of each block indicates how much channels it has. The Inverted Residual Block connects the bottlenecks layers reducing the amount of memory needed in comparison to the MobileNet Residual Block that connects layers with higher number of channels. What makes the MobileNet architecture unique is the depthwise separable convolutions, which is also part of the MobileNetV2 model. The depthwise separable convolution is a new way to make time-efficient convolutions, and it is built by splitting the traditional process of convolution into two other operations: depthwise convolution and pointwise convolution. They work by sharing the convolution effort between them. The depthwise convolution applies a single convolution filter on each channel of the input sample, while the pointwise convolution uses a  $1 \times 1$  convolution filter on the output of the depthwise operation. Thus, the pointwise work is to combine the outputs from the depthwise operation into the  $n$  required depthwise separable convolutional filters. This modification drastically reduces the computational effort needed to process the data and the model size, while speeds up the entire process when compared to traditional convolutions. According to [23], the

use of the depthwise separable convolutions of size  $3 \times 3$  helps to reduce the computational effort of a model in 8 to 9 times when compared with traditional convolutional layers.

Another advantage of the MobileNet and MobileNetV2 is that their principal operation, the depthwise separable convolutions, can be built in any modern deep learning framework because the depthwise and the pointwise convolutions are made from traditional convolution operations. According to [24], the MobileNetV2 overcomes, in both accuracy and model complexity, state-of-the-art real-time detectors on COCO dataset [27] for the object detection task. Combined with SSDLite detection module [24], which is a modified version of the Single Shot Detector (SSD) [28], it also consumes 20 times less computational effort and has 10 times fewer parameters than the YOLOv2 [29].

Yet, the MobileNets are designed to computer vision tasks, so their operations, including the depthwise separable convolutions, are built to work with images. Thus, to work directly with audio signals in the speaker recognition problem, these models need some adaptation.

### III. PROPOSED MODELS

In this section, we propose the lightweight deep learning models MobileNet1D and AM-MobileNet1D to tackle the speaker recognition problem, especially the speaker identification tasks. First, we present the adaptation of the MobileNetV2 for speaker recognition (called MobileNet1D). Then we describe how the modified MobileNet1D built the AM-MobileNet1D for the same task.

#### A. MobileNet1D

The MobileNet [23] is a convolutional neural network built to be fast and light, which makes it appropriate to work with high dimensional data like audio signals. Although, the MobileNet architecture is modeled to work with two-dimensional structs like images. However, unlike images, the audio signals are modeled in one-dimensional. Indeed, it is necessary to adapt the MobileNet architecture to fit audio signals. Thus, to tackle the speaker recognition problem, we made some adaptation on the MobileNetV2 [24] architecture.

Nevertheless, to adapt the MobileNetV2 architecture, we first go to its core, the depthwise separable convolution. Every element on the depthwise separable convolution is designed to work with two-dimensional structs, the convolutions, the batch normalization, and the pooling layers. For this reason, it is necessary to adapt these basic operations on the depthwise and pointwise layers. Thus, we first change the convolutions on the depthwise struct to proceed 1d convolutions instead of 2d. Batch normalization and pooling operations were also replaced from 2d to 1d implementations. These modifications not only prepared MobileNet to process audio signals, but it also speeds up the network by reducing its complexity and model size. Figure 4 shows an illustration of the MobileNet with the proposed MobileNet1D and AM-MobileNet1D methods. In Figure 4 (b), the audio signals are fed into the MobileNet architecture adapted with 1d depthwise separable convolutions,

then the extracted features feed a Softmax layer which will output the probabilities of the sample belong to each class.

#### B. AM-MobileNet1D

As presented in Figure 4 (b), the proposed method MobileNet1D uses a Softmax loss function to map the extracted features into their respective classes. However, the Softmax loss function may not be the best choice for this type of problem. Indeed, the Additive Margin Softmax (AM-Softmax) loss function has shown to be a more sophisticated approach for classification problems. It let the models achieve better results not only in the face recognition problem where it was first proposed but also in the speaker recognition problem, as shown in [22]. The AM-Softmax has improved significantly the Frame Error Rate measured in the AM-SincNet in comparison to the calculated in the traditional SincNet.

Nevertheless, the AM-Softmax introduces two new parameters when compared to the traditional Softmax. They are the  $s$  parameter which is responsible for scaling the features and the  $m$  parameter that controls the size of the additive margin. In [1], where the AM-Softmax loss function was proposed, it was established that the scaling factor  $s$  needs to be high enough to work well. It also can be learned by the training process of the network. However, it would make the network convergence too slow. Thus, [1] states that a fixed value of  $s$  works just fine and it proposes to set  $s = 30$ . On the other hand, the margin size of  $m$  may vary in the range  $0 \leq m \leq 1$ . In fact, [22] made several experiments on the dataset TIMIT to see the influence of the margin parameter on the Frame Error Rate (FER), it concludes that at the end of training every  $m$  value tested in the range  $0, 35 \leq m \leq 0, 80$  converges the FER to the same place. Also, none of the  $m$  values tested on AM-SincNet got higher (worse) error rates than the SincNet using the Softmax loss function.

So, the proposed Additive Margin MobileNet1D (AM-MobileNet1D) is a combination of the MobileNet1D described in the last subsection and the AM-Softmax loss function. It replaces the last Softmax layer of the MobileNet1D for an improved Additive Margin Softmax layer. With this adaptation, the AM-MobileNet1D gets two new parameters from the AM-Softmax function, the scaling parameter  $s$ , and the margin size  $m$ . Following [1], we set  $s = 30$  for the proposed model. In addition, for the margin size parameter, as [22] stated that every tested value of  $m$  converged the FER to the same place, we fixed  $m = 0.5$  in our experiments. Figure 4 (c) shows an illustration of the proposed AM-MobileNet1D where the input audio signals fed a MobileNet1D and then the extracted features are mapped into their classes by an AM-Softmax loss function.

### IV. EXPERIMENTS

In this section, we describe the experimental protocol to evaluate the proposed models, MobileNet1D and AM-MobileNet1D, considering the baseline methods SincNet and AM-SincNet on the Speaker Identification task.

TABLE I  
COMPARISON OF THE SPEAKER RECOGNITION MODELS ON THE TIMIT AND MIT DATASET USING THE METRICS FRAME ERROR RATE (FER), CLASSIFICATION ERROR RATE (CER), INFERENCE TIME, NUMBER OF PARAMETERS, AND MODEL SIZE.

	TIMIT		MIT		Inference Time (ms)[↓]	Parameters (total)[↓]	Model Size (mb)[↓]
	FER (%) [↓]	CER (%) [↓]	FER (%) [↓]	CER (%) [↓]			
SincNet	44.64	1.15	34.65	1.27	$42.60 \pm 1.5766$	22,787,134	91.2
AM-SincNet	27.86	0.50	<b>30.30</b>	<b>0.52</b>	$41.17 \pm 0.7529$	22,787,134	91.2
MobileNet1D	26.50	0.57	38.75	2.54	<b><math>5.88 \pm 0.1584</math></b>	<b>2,825,942</b>	<b>11.6</b>
AM-MobileNet1D	<b>21.30</b>	<b>0.43</b>	35.55	2.19	<b><math>5.85 \pm 0.1673</math></b>	<b>2,825,942</b>	<b>11.6</b>

### A. Datasets

We use two different speaker datasets. The first one is the well-know TIMIT [25] dataset, which contains audio samples from speakers of the eight main American dialects. It contains 630 different speakers, and each one of them was instructed to read 10 phonetically rich sentences. The dataset is actually divided into two distinct sets, the training set and the testing set. To a fear comparison, we follow the same protocol as [16] and [22], where the experiments use only the training set of the TIMIT dataset, which contains 462 speakers and 3,969 audio samples.

The second dataset used to evaluate the models is the MIT [26] dataset, which is built from audio samples collected from mobile devices in a variety of environments in order to induce the presence of noise. Like the TIMIT dataset, the MIT dataset is also divided into two distinct sets, the set of authentic users and the set of imposters. The data were collected in two separate sessions for the authentic users and one more session for the imposters. Each session lasted for 20 minutes and was recorded on different days. In order to simulate normal daily conditions, the sessions were recorded in three different places: a quiet office, a mildly noisy lobby, and a busy street intersection. The MIT dataset contains audio samples from 88 different speakers, 39 women, and 49 men. Analyzing the imposters' set of the MIT dataset, we saw that not every authentic user has its own imposter. So, to avoid the imbalance problem, we only use the authentic users set, as the imposters set is not complete. The authentic users set contains 5,184 audio samples from 48 different speakers, where each speaker recorded exactly 108 samples.

Each dataset was divided into two different sets, the training set, and the testing set. For the TIMIT dataset, five utterances of each speaker were used to train the network, while the remaining three were used to test. On the other hand, for the MIT dataset, 70% of the samples from each speakers were selected for the training process, while the remaining samples were used for testing. In both datasets, the samples that compose the training and the testing set were chosen randomly. Yet, in each dataset, the number of samples was balanced between the classes.

### B. Training

The pre-processing step on the audio samples are the same used in [16] and [22], where a normalization was applied to the audio signals. Also, for the TIMIT dataset, the non-speech interval was removed from the beginning and the end of the sentences. Moreover, the waveform of each audio sample was divided into 200ms frames with 10ms overlap; these frames were used to feed the network.

For training purposes, every network model was configured to use the RMSprop optimizer. Also, learning rate was set to  $lr=0.001$ , with  $\alpha=0.95$ , and  $\epsilon=10^{-7}$ . We set mini-batches to be at size 128. The AM-Softmax comes with two more parameters compared to the traditional Softmax, the parameter  $s$  for scaling the features and the parameter  $m$  to control the size of the additive margin. According to [1], the scaling factor could be learned by the network, but it would take too long to optimize its value. Also, [1] states that the scaling factor  $s$  needs to be high enough in order to work well. Thus the author suggests setting  $s$  to a fixed value of  $s=30$  to speed up the training process. On the other hand, for the  $m$  parameter, [22] made several experiments with the AM-SincNet model in the TIMIT dataset to see the influence of the margin parameter on the FER, according to the authors every tested value of  $m$  converges the FER to the same place, thus in our experiments with the AM-MobileNet1D we set  $m=0.5$ . We also have added an *epsilon* constant of value  $10^{-11}$  to the AM-Softmax equation in order to avoid a division by zero in the required places.

Each model was trained for exactly 360, which shows to be enough to exploit the model's efficiency. We also set a *seed* = 1234 to maintain the consistency of the experiments. The experiments run in a Core i5-9400F machine with a NVIDIA RTX 2060 GPU. We use CUDA 10.1 and Pytorch 1.0.

### C. Metrics

To evaluate the proposed models and compare them to the baseline approaches, we used the Frame Error Rate (FER) and Classification Error Rate (CER), two widely used metrics on the literature for the Speaker Recognition problem. We also compare the models by their inference time, which is




Fig. 5. Frame Error Rate (FER) over the training epochs.

calculated using the average time in milliseconds that the machine takes to process 6,561 batches of size 128, their number of parameters, and the model size in megabytes.

The performed experiments may be reproduced by using the code that we made available online at the GitHub <sup>1 2</sup>.

## V. RESULTS

In this section, we explore the results of the experiments. The proposed methods show to be competitive in terms of FER and CER. The methods obtained an equivalent or even better results than the SincNet and AM-SincNet models on the TIMIT and MIT dataset. Additionally, the proposed methods, MobileNet1D and AM-MobileNet1D, are about 7 times faster, have 8 times fewer parameters, and take up to 8 times less space on disk than the SincNet and AM-SincNet.

Table I shows the results of the experiments on TIMIT and MIT dataset using the metrics Frame Error Rate (FER) and Classification Error Rate (CER). The best result in each dataset is highlighted in bold. Based on these results, it is possible to see that in every scenario, the AM-Softmax loss function improved the performance of the models. Indeed, the benefit of using the AM-Softmax loss function on the Speaker Recognition models is quite clear, mainly on the SincNet model for the TIMIT dataset. The FER and CER calculated on the TIMIT dataset for the proposed method MobileNet1D are about 40% and 50% smaller (better) than the ones measured from SincNet, respectively. Also, when comparing the MobileNet1D with the AM-SincNet in the same dataset, they are practically equivalent, the MobileNet1D got a slightly better FER, while the AM-SincNet got a slightly better CER. For this dataset, the model with the best results was the proposed AM-MobileNet1D, it got the lowest (best) FER and CER over the models. According to Table I, the

FER measured on the TIMIT dataset for the AM-MobileNet1D was 19% better than the MobileNet1D, 23% better than the AM-SincNet, and 52% better than the traditional SincNet. In the same way, in terms of CER it was 24% better than the MobileNet1D, 14% better than the AM-SincNet, and 62% better than the SincNet. On the other hand, the SincNet and the AM-SincNet overcome the proposed models in terms of FER and CER for the MIT dataset. Still, the proposed models are quite good on this dataset showing an accuracy of about 98%.

Table I also shows the results of inference time in milliseconds, the number of parameters, and the model size (in megabytes) for each model. The inference time results are in the format *time ± standard deviation*. The results show that the proposed models are about 7 times faster, have 8 times fewer parameters, and take up to 8 times less space on disk than the SincNet and the AM-SincNet. The proposed MobileNet1D and AM-MobileNet1D take only 11.6 megabytes on disk, against 91.2mb from the SincNet and AM-SincNet. In fact, these results show that the proposed models are suitable for being embedded on Mobile devices, showing high performance and low inference latency for the speaker recognition problem.

Figures 5 (a) and (b) show the evolution of the Frame Error Rate (FER) over the training epochs on the TIMIT and MIT dataset, respectively. There, the SincNet and AM-SincNet are plotted in red, while the MobileNet1D and AM-MobileNet1D are in blue. Also, the methods using the AM-Softmax loss function are dashed. On both figures, the methods using the AM-Softmax loss function perform better than the ones using the traditional Softmax. From Figure 5 (a) it is clear the difference on FER from both proposed models, MobileNet1D and AM-MobileNet1D, to the traditional SincNet. Also, the MobileNet1D and the AM-SincNet are equivalent in terms of FER on the TIMIT dataset. Figure 5 (b) shows the results for

<sup>1</sup><https://github.com/joaoantoniocn/AM-MobileNet1D>

<sup>2</sup><https://github.com/joaoantoniocn/AM-SincNet>



the MIT dataset. There we can see that unlike Figure 5 (a), the results are closer to each other, showing a modest difference between them. Yet, the AM-SincNet has performed better on this dataset, while the MobileNet1D and the SincNet presented almost the same result.

In general, the proposed AM-MobileNet1D show high performance in the TIMIT dataset (the best of all models), while got less than 1% increase in the CER on the MIT dataset compared to the SincNet. Indeed, the great advantages of the proposed MobileNet1D and AM-MobileNet1D are its inference time, which is about 7 times faster than the SincNet and the AM-SincNet, and its model size which takes only 11.6 megabytes on disk. Thus, making it an excellent choice to tackle the speaker recognition problem, especially in environments where the system needs a quick or embedded answer.

## VI. CONCLUSION

This paper has proposed new methods for directly processing waveform audio on mobile devices for the speaker recognition task. The proposed methods MobileNet1D and AM-MobileNet1D got about 7 times faster results with a low error rate compared to the hybrid methods SincNet and AM-SincNet. The AM-MobileNet1D also got the best Frame Error Rate (FER) and Classification Error Rate (CER) for the TIMIT dataset. Another advantage of the proposed methods lies in their model size, which takes only 11.6 megabytes on the disk.

For future work, we should consider larger datasets like VoxCeleb1 [30] and VoxCeleb2 [31]. VoxCeleb1 contains over 100k samples distributed in about 1.2k celebrities, and the samples were extracted from videos uploaded on YouTube. In the same way, VoxCeleb2 [31] has over a million samples from over 6k speakers. Thus, it would be good to see how the MobileNet1D and the AM-MobileNet1D behaves on larger datasets.

## REFERENCES

- [1] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, July 2018.
- [2] H. Beigi, *Fundamentals of Speaker Recognition*. Springer Publishing Company, Incorporated, 2011.
- [3] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech and Language Processing*, 2010.
- [4] A. Kanagasundaram, R. Vogt, D. Dean, S. Sridharan, and M. Mason, "i-vector based speaker recognition on short utterances," in *Interspeech 2011*, 2011, pp. 2341–2344.
- [5] R. Travadi, M. Van Segbroeck, and S. Narayanan, "Modified-prior i-Vector Estimation for Language Identification of Short Duration Utterances," in *Proceedings of the Fifteenth Annual Conference of the International Speech Communication Association*, Singapore, Sep. 2014, pp. 3037–3041.
- [6] S. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," *2007 IEEE 11th International Conference on Computer Vision*, pp. 1–8, 2007.
- [7] P. Matejka, O. Glembek, F. Castaldo, M. J. Alam, O. Plchot, P. Kenny, L. Burget, and J. Cernocky, "Full-covariance ubm and heavy-tailed plda in i-vector speaker verification," *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 4828–4831, 05 2011.
- [8] S. Cumani, O. Plchot, and P. Laface, "Probabilistic linear discriminant analysis of i-vector posterior distributions," in *ICASSP. IEEE*, 2013, pp. 7644–7648. [Online]. Available: <http://dblp.uni-trier.de/db/conf/icassp/icassp2013.htmlCumaniPL13>
- [9] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *ArXiv*, vol. abs/1705.02304, 2017.
- [10] J.-W. Jung, H.-S. Heo, I.-H. Yang, H.-J. Shim, and H.-J. Yu, "Avoiding speaker overfitting in end-to-end dnn using raw waveform for text-independent speaker verification," in *Annual Conference of the International Speech Communication Association*, 09 2018, pp. 3583–3587.
- [11] M. S. Hussain and M. A. Haque, "Swishnet: A fast convolutional neural network for speech, music and noise classification and segmentation," *CoRR*, vol. abs/1812.00149, 2018. [Online]. Available: <http://arxiv.org/abs/1812.00149>
- [12] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," in *Proceedings of the IEEE*, 1998, pp. 2278–2324.
- [13] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017. [Online]. Available: <http://dx.doi.org/10.1109/CVPR.2017.683>
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.
- [15] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, 06 2015.
- [16] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*, Dec 2018, pp. 1021–1028.
- [17] E. Variani, X. Lei, E. McDermott, I. Lopez-Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *ICASSP. IEEE*, 2014, pp. 4052–4056.
- [18] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A novel scheme for speaker recognition using a phonetically-aware deep neural network," *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 1695–1699, 05 2014.
- [19] S. H. Ghahlehjeh and R. C. Rose, "Deep bottleneck features for i-vector based text-independent speaker verification," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 555–560.
- [20] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *INTERSPEECH*, F. Lacerda, Ed. ISCA, 2017, pp. 999–1003. [Online]. Available: <http://dblp.uni-trier.de/db/conf/interspeech/interspeech2017.htmlSnyderGPK17>
- [21] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333, 2018.
- [22] J. A. Chagas Nunes, D. Macêdo, and C. Zanchettin, "Additive margin sincnet for speaker recognition," in *2019 International Joint Conference on Neural Networks (IJCNN)*, July 2019, pp. 1–5.
- [23] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *CoRR*, vol. abs/1704.04861, 2017. [Online]. Available: <http://arxiv.org/abs/1704.04861>
- [24] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.
- [25] J. Garofolo, L. Lamel, W. Fisher, J. Fiscus, and D. Pallett, "Darpa timit acoustic-phonetic continous speech corpus cd-rom. nist speech disc 1-1.1," *NASA STI/Recon Technical Report N*, vol. 93, p. 27403, 01 1993.
- [26] R. H. Woo, A. Park, and T. J. Hazen, "The mit mobile device speaker verification corpus: Data collection and preliminary experiments," in *2006 IEEE Odyssey - The Speaker and Language Recognition Workshop*, June 2006, pp. 1–6.
- [27] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer International Publishing, 2014, pp. 740–755.

- [28] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," *Lecture Notes in Computer Science*, p. 21–37, 2016.
- [29] J. Redmon and A. Farhadi, "Yolo9000: Better, faster, stronger," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 6517–6525.
- [30] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," in *INTERSPEECH*, 2017.
- [31] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *Interspeech 2018*, Sep 2018.