



Repositorio Institucional de la Universidad Autónoma de Madrid

<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:
This is an **author produced version** of a paper published in:

Neural Networks 64 (2015): 49 – 58

DOI: <http://dx.doi.org/10.1016/j.neunet.2014.08.006>

Copyright: © 2015 Elsevier

El acceso a la versión del editor puede requerir la suscripción del recurso
Access to the published version may require subscription

Frame by Frame Language Identification in Short Utterances using Deep Neural Networks

Javier Gonzalez-Dominguez^{a,b,*}, Ignacio Lopez-Moreno^a, Pedro J. Moreno^a, Joaquin Gonzalez-Rodriguez^b

^aGoogle Inc., New York, USA

^bATVS-Biometric Recognition Group, Universidad Autonoma de Madrid, Madrid, Spain

Abstract

This work addresses the use of deep neural networks (DNNs) in automatic language identification (LID) focused on short test utterances. Motivated by their recent success in acoustic modelling for speech recognition, we adapt DNNs to the problem of identifying the language in a given utterance from the short-term acoustic features. We show how DNNs are particularly suitable to perform LID in real-time applications, due to their capacity to emit a language identification posterior at each new frame of the test utterance. We then analyse different aspects of the system, such as the amount of required training data, the number of hidden layers, the relevance of contextual information and the effect of the test utterance duration. Finally, we propose several methods to combine frame-by-frame posteriors. Experiments are conducted on two different datasets: the public NIST Language Recognition Evaluation 2009 (3 seconds task) and a much larger corpus (of 5 million utterances) known as Google 5M LID, obtained from different Google Services. Reported results show relative improvements of DNNs versus the i-vector system of 40% in LRE09 3 second task and 76% in Google 5M LID.

Keywords: DNNs, real-time LID, i-vectors.

1. Introduction

Automatic language identification (LID) refers to the process of automatically determining the language in a given speech sample [1]. The need for reliable LID is continuously growing due to several factors. Among them, the technological trend toward increased human interaction using hands-free, voice-operated devices and the need to facilitate the coexistence of a multiplicity of different languages in an increasingly globalized world.

In general, language discriminant information is spread across different structures or levels of the speech signal, ranging from low-level, short-term acoustic and spectral features to high-level, long-term features (i.e. phonotactic, prosodic). However, even though several high-level approaches are used as meaningful complementary sources of information [2] [3] [4], most LID systems still include or rely on acoustic modelling

[5][6], mainly due to their better scalability and computational efficiency.

Indeed, computational cost plays an important role, as LID systems commonly act as a pre-processing stage for either machine systems (i.e. multilingual speech processing systems) or human listeners (i.e. call routing to a proper human operator)[7]. Therefore, accurate and efficient behaviour in real-time applications is often essential, for example, when used for emergency call routing, where the response time of a fluent native operator is critical [1] [8]. In such situations, the use of high-level speech information may be prohibitive, as it often requires running one speech/phonetic recognizer per target language [9]. Lightweight LID systems are especially necessary in cases where the application requires an implementation embedded in a portable device.

Driven by recent developments in speaker verification, the current state-of-the-art in acoustic LID systems involves using i-vector front-end features followed by diverse classification mechanisms that compensate speaker and session variabilities [7] [10] [11]. The i-vector is a compact representation (typically from 400 to 600 dimensions) of a whole utterance, derived as a

*Corresponding author. Tel: +34 914977558

Email address: E-mail

address: javier.gonzalez@uam.es (Javier Gonzalez-Dominguez)

point estimate of the latent variables in a factor analysis model [12] [13]. However, while proven to be successful in a variety of scenarios, i-vector based approaches suffer from two major drawbacks when coping with real-time applications. First, i-vectors are point estimates and their robustness quickly degrades as the amount of data used to derive the i-vector decreases. Note that the smaller the amount of data, the larger the variance of the posterior probability distribution of the latent variables; and thus, the larger the i-vector *uncertainty*. Second, in real-time applications, most of the costs associated with i-vector computation occur after completion of the utterance, which introduces an undesirable latency.

Motivated by the prominence of Deep Neural Networks (DNNs), which surpass the performance of the previous dominant paradigm, Gaussian Mixture Models (GMMs), in diverse and challenging machine learning applications - including acoustic modelling [14] [15], visual object recognition [16], and many others [17] - we previously introduced a successful LID system based on DNNs in [18]. Unlike previous works on using shallow or convolutional neural networks for small LID tasks [19] [20] [21], this was, to the best of our knowledge, the first time that a DNN scheme was applied at large scale for LID, and benchmarked against alternative state-of-the-art approaches. Evaluated using two different datasets - the NIST LRE 2009 (3s task) and Google 5M LID - this scheme demonstrated significantly improved performance compared to several i-vector-based state-of-the-art systems [18].

In the current study, we explore different aspects that affect DNN performance, with a special focus on very short utterances and real-time applications. We believe the DNN-based system is a suitable candidate for this kind of application, as it could potentially generate decisions at each processed frame of the test speech segment, typically every 10ms. Through this study, we assess the influence of several factors on the performance, namely: a) the amount of required training data, b) the topology of the network, c) the importance of including the temporal context, and d) the test utterance duration. We also propose several blind techniques to combine frame by frame posteriors obtained from the DNN to get hard identification decisions.

We conduct the experiments using the following LID datasets: A dataset built from Google data, hereafter, Google 5M LID corpus and the NIST Language Recognition Evaluation 2009 (LRE'09). First, through the Google 5M LID corpus, we evaluate the performance in a real application scenario. Second, we check if the same behaviour is observed in a familiar and standard

evaluation framework for the LID community. In both cases, we focus on short test utterances (up to 3s).

The rest of this paper is organized into the following sections. Section 2 defines a reference system based on i-vectors. The proposed DNN system is presented in Section 3. The experimental protocol and datasets are described in Section 4. Next, we examine the behaviour of our scheme over a range of configuration parameters in both the task and the neural network topology. Finally, Sections 6 and 7 are devoted to presenting conclusions of the study and evaluating recommendations for future work.

2. Baseline system: i-vector

Currently, most acoustic approaches to perform LID rely on i-vector technology [22]. All such approaches, while sharing i-vectors as a feature representation, differ in the type of classifier used to perform the final language identification [23]. In the rest of this Section, we describe: a) the i-vector extraction procedure, b) the classifier used in this study, and c) the configuration details of our baseline i-vector system. Below, we describe a state-of-the-art acoustic system based on i-vectors, which will serve as our baseline i-vector system.

2.1. I-vector extraction

Based on the MAP adaptation approach in a GMM framework [24], utterances in language or speaker recognition are typically represented by the accumulated zero- and centered first-order Baum-Welch statistics, N and F , respectively, computed from a Universal Background Model (UBM) λ . For UBM mixture $m \in 1, \dots, C$, with mean, μ_m , the corresponding zero- and centered first-order statistics are aggregated over all frames of the utterance as

$$N_m = \sum_t p(m|o_t, \lambda) \quad (1)$$

$$F_m = \sum_t p(m|o_t, \lambda)(o_t - \mu_m), \quad (2)$$

where $p(m|o_t, \lambda)$ is the Gaussian occupation probability for the mixture m given the spectral feature observation $o_t \in \mathfrak{R}^D$ at time t .

The Total Variability model, hereafter TV, can be seen as a classical FA generative model [25], with observed variables given by the *supervector* (CD x 1) of stacked statistics $F = \{F_1, F_2, \dots, F_C\}$. In the TV model, the vector of hidden variables $w \in \mathfrak{R}^L$ is known

as the utterance i -vector. Observed and hidden variables are related by the rectangular low rank matrix $T \in \mathfrak{R}^{CD \times L}$

$$N^{-1}F = Tw, \quad (3)$$

where the zero-order statistics N are represented by a block diagonal matrix $\in \mathfrak{R}^{CD \times CD}$, with C diagonal $D \times D$ blocks. The m -th component block is the matrix $N_m I_{(D \times D)}$. Given the imposed Gaussian distributions of $p(w)$ and $p(F|w)$, it can be seen that the mean of the posterior $p(w|F)$ is given by

$$w = (I + T^t \Sigma^{-1} NT)^{-1} T^t \Sigma^{-1} F, \quad (4)$$

where $\Sigma \in \mathfrak{R}^{CD \times CD}$ is the diagonal covariance matrix of F . The TV model is thus a data driven model with parameters $\{\lambda, T, \Sigma\}$. Kenny [13] provides a more detailed explanation of the derivation of these parameters, using the EM algorithm.

2.2. Classification

Since T constrains all the variability (i.e language, speaker, session), and it is shared for all the language models/excerpts, the i -vectors, w , can be seen as a new input feature to classify. Further, several classifiers - either discriminative (i.e Logistic Regression) or generative (i.e Gaussian classifier, Linear Discriminant Analysis) - can be used to perform classification [23]. In this study, we utilized LDA, followed by cosine distance (LDA_CS), as the classifier.

Even though using a more sophisticated classifier [18] would have resulted in slightly increased performance, we chose the LDA_CS considering the trade-off between performance and computational time efficiency. In this framework, the similarity measure (score) of the two given i -vectors, w_1 and w_2 , is obtained as

$$S_{w_1, w_2} = \frac{(A^t w_1)(A^t w_2)}{\sqrt{(A^t w_1)(A^t w_1)} \sqrt{(A^t w_2)(A^t w_2)}} \quad (5)$$

where A is the LDA matrix.

2.3. Feature extraction and configuration parameters

As input features for this study we used perceptual linear predictive (PLP) coefficients [26]. In particular, 13 PLP coefficients augmented with *delta* and *delta-delta* features (39 dimensions total) were extracted with a 10ms frame rate over 25ms long windows. From those features, we built a Universal Background Model of 1024 components. The Total Variability matrix was trained by using PCA and a posterior refinement of 10

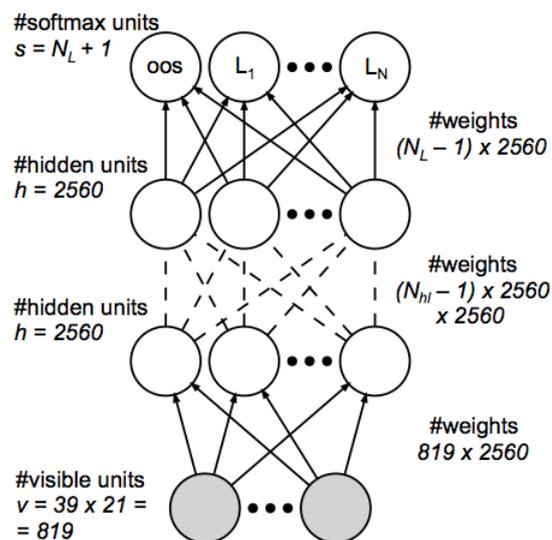


Figure 1: DNN network topology

EM iterations [22], keeping just the top 400 eigenvectors. We then derived the i -vectors using the standard methodology presented in Section 2.1. In addition, we filtered out silence frames by using an energy-based voice activity detector.

3. DNN as a language identification system

Recent findings in the field of speech recognition have shown that significant accuracy improvements over classical GMM schemes can be achieved through the use of deep neural networks, either to generate GMM features or to directly estimate acoustic model scores. Among the most important advantages of DNNs is their multilevel distributed representation of the input [15]. This fact makes the DNN an exponentially more compact model than GMMs. In addition DNNs do not require detailed assumptions about the input data distribution [27] and have proven successful in exploiting large amounts of data, reaching more robust models without lapsing into overtraining. All of these factors motivate the use of DNN in language identification. The rest of this Section describes the architecture and the practical implementation of the DNN system.

3.1. Architecture

The DNN used in this work is a fully-connected feed-forward neural network with hidden units implemented as rectified linear units (ReLU). Thus, an input at level j , x_j , is mapped to its corresponding activation y_j (input of the layer above) as

$$y_j = \text{ReLU}(x_j) = \max(0, x_j) \quad (6)$$

$$x_j = b_j + \sum_i w_{ij} y_i \quad (7)$$

where i is an index over the units of the layer below and b_j is the bias of the unit j .

The output layer is then configured as a *softmax*, where hidden units map input x_j to a class probability p_j in the form

$$p_j = \frac{\exp(x_j)}{\sum_l \exp(x_l)} \quad (8)$$

where l is an index over all the classes.

As a cost function for backpropagating gradients in the training stage, we use the cross-entropy function defined as

$$C = - \sum_j t_j \log p_j \quad (9)$$

where t_j represents the target probability of the class j for the current evaluated example, taking a value of either 1 (true class) or 0 (false class).

3.2. Implementing DNN for language identification

From the conceptual architecture explained above, we built a language identification system to work at the frame level as follows.

As the input of the net we used the same features as the i-vector baseline system (39 PLP). Specifically, the input layer was fed with 21 frames formed by stacking the current processed frame and its ± 10 left/right neighbours. Thus, the input layer comprised a total number of 819 (21×39) visible units, v .

On top of the input layer, we stacked a total number of N_{hl} (8) hidden layers, each containing h (2560) units. Then, we added the softmax layer, whose dimension (s) corresponds to the number of target languages (N_L) plus one extra output for the out-of-set (*OOS*) languages. This OOS class, devoted to non-known test languages not seen in training time, could in future allow us to use the system in open-set identification scenarios. Overall, the net was defined by a total of w free parameters (weights + bias), $w = (v + 1)h + (N_{hl} - 1)(h + 1)h + (h + 1)s$ ($\sim 48\text{M}$). The complete topology of the network is depicted in Figure 1.

Regarding the training procedure, we used asynchronous stochastic gradient descent within the Dist-Belief framework [28], a software framework that uses computing clusters with thousands of machines to train

large models. The learning rate and minibatch size were fixed to 0.001 and 200 samples¹.

Note that the presented architecture works at the frame level, meaning that each single frame (plus its corresponding context) is fed-forward through the network, obtaining a class posterior probability for all of the target languages. This fact makes the DNNs particularly suitable for real-time applications since, unlike other approaches (i.e. i-vectors), we can potentially make a decision about the language at each new frame. Indeed, at each frame, we can combine the evidence from past frames to get a single similarity score between the test utterance and the target languages. A simple way of doing this combination is to assume that frames are independent and multiply the posterior estimates of the last layer. The score s_l for the language l of a given test utterance is computed by multiplying the output probabilities p_l obtained for all its frames; or equivalently, accumulating the logs as

$$s_l = \frac{1}{N} \sum_{t=1}^N \log p(L_l | x_t, \theta) \quad (10)$$

where $p(L_l | x_t, \theta)$ represents the class probability output for the language l corresponding to the input example at time t , x_t by using the DNN defined by parameters θ .

4. Datasets and evaluation metrics

We conducted experiments on two different databases following the standard protocol provided by NIST in LRE 2009 [29]. Particularly, we used the LRE'09 corpus and a corpus generated from Google Voice services. This followed a two-fold goal: first, to evaluate the proposed methods with a large collection of real application data and second, to provide a benchmark comparable with other related works in the area by using the well-known LRE'09 framework.

4.1. Databases

4.1.1. Google 5M LID Corpus

We generated the Google 5M LID corpus dataset by randomly picking anonymized queries from several Google speech recognition services such as Voice Search or the Speech Android API. Following the user's phone Voice Search language settings, we labelled a total of ~ 5 million utterances, 150k utterances by 34

¹We define sample as the input of the DNN: the feature representation of a single frame besides those from its adjacent frames forming the context.

Google 5M	
Locale/Abbrev.	Language
ar-EG	Arabic (Egypt)
ar-GULF	Arabic (Persian Gulf)
ar-LEVANT	Arabic (Levant)
bg-BG	Bulgarian
cs-CZ	Czech
de-DE	German
en-GB	English (United Kingdom)
en-IN	English (India)
en-US	English (USA)
en-ZA	English (South Africa)
es-419	Spanish (Latin America/Caribbean)
es-AR	Spanish (Argentina)
es-ES	Spanish (Spain)
fi-FI	Finish
fr-FR	French
he-IL	Hebrew (Israel)
hu-HU	Hungarian
id-ID	Indonesian
it-IT	Italian
ja-JP	Japanese
ko-KR	Korean (South Korea)
ms-MY	Malay
nl-NL	Dutch
pt-BR	Portuguese (Brazilian)
pt-PT	Portuguese (Portugal)
ro-RO	Romanian
ru-RU	Russian
sk-SK	Slovak
sr-RS	Serbian
sv-SE	Sweden
tr-TR	Turkish
zh-cmn-Hans-CN	Chinese (Mandarin)
zh-cmn-Hant-TW	Chinese (Taiwan)
zh-yue-hant-HK	Chinese (Cantonese)

LRE'09	
usen	English(USA)
span	Spanish (Latin America/Caribbean)
fars	Farsi
dari	Dari
fren	French
pash	Pashto
russ	Russian
mand	Chinese (Mandarin)

Database	$\#N_L$	Train (hours)	Test (#files)	Test length (avg. on s.)
Google 5M	34	2975	51000	4.2
LRE09_VOA_3s	8	1600	2916	3
LRE09_VOA_realtime	8	1600	11276	(0.1s to 3s)

Table 2: Data description of the Google 5M LID and LRE09 subcorpus.

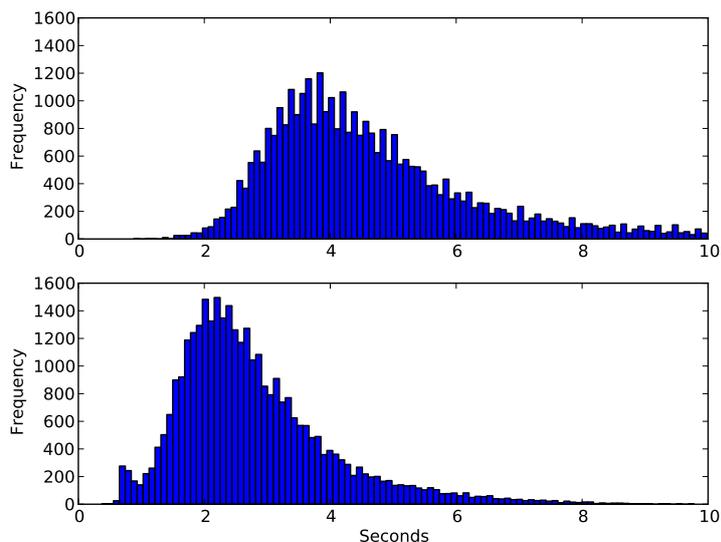


Figure 2: Histograms of durations of the Google 5M LID test utterances. Original speech signals (above) and after voice activity detection (below)

Table 1: List of the Google 5M LID (above) and LRE'09 (below) languages considered.

	Equal Error Rate (EER in %)								
	en	es	fa	fr	ps	ru	ur	zh	avg.
Iv_200h	17.22	10.92	20.03	15.3	19.98	14.87	18.74	10.09	15.89
DNN_8layers_200h	8.65	3.74	17.22	7.53	16.01	5.59	13.10	4.82	9.58

Table 3: Systems performance (ERR %) comparison per language on LRE09_VOA_3s_test . I-vector baseline system vs DNN_8layers_200h system.

different locales (25 languages + 9 dialects) yielding ~87,5h of speech per language and a total of ~2975h. A held-out test set of 1k utterances per language was created while the remainder was used for training and development. Involved languages and data description is presented in Tables 1 and 2 respectively.

An automatic speech recognition system was used to discard non-speech queries. Selected queries ranged from 1s up to 10s in duration with average speech content of 2.1s. Figure 2 shows the duration distribution before and after doing this activity detection process.

Privacy issues do not allow Google to link the user identity with the spoken utterance and therefore, determining the exact number of speakers involved in this corpus is not possible. However, it is reasonable to consider that the total number of speakers is very large.

4.1.2. Language Recognition Evaluation 2009 Dataset

The LRE evaluation in 2009 included, for the first time, data coming from two different audio sources. Besides Conversational Telephone Speech (CTS), used in the previous evaluations, telephone speech from broadcast news was used for both training and test purposes. Broadcast data were obtained via an automatic acquisition system from “Voice of America” news (VOA) where telephone and non-telephone speech is mixed.

Due to the large disparity on training material for every language (from ~10 to ~950 hours), out of the 40 initial target languages [30] we selected 8 representative languages for which up to 200 hours of audio were available: US English (en), Spanish (es), Dari (fa), French (fr), Pashto (ps), Russian (ru), Urdu (ur), Chinese Mandarin (zh) (Table 1). Further, to avoid misleading result interpretation due to the unbalanced mix of CTS and VOA, all the data considered in this dataset was part of VOA.

As test material in LRE’09, we used a subset of the NIST LRE 2009 3s condition evaluation set (as for training, we also discarded CTS test segments), yielding a total of 2916 test segments of the 8 selected languages. That makes a total of 23328 trials. We refer this test dataset as LRE09_VOA_3s_test. For evaluating performance in real-time conditions, we used the VOA test

segments for all the LRE’09 conditions (3s, 10s, 30s) with at least 3s of speech (according to our voice activity detector); that made a total of 11276 files. Then we cut these recordings to build different duration subsets ranging from 0.1s to 3s of speech. Specifically, we came up with 8 datasets of 11276 files with durations: 0.1s, 0.2s, 0.5s, 1s, 1.5s, 2.0s, 2.5s, 3.0s. We refer those test datasets as the LRE09_VOA_realtime_test benchmark.

4.2. Evaluation metrics

In order to assess the performance we used Accuracy and Equal Error Rate (EER)² metrics. Language identification rates are measured in terms of Accuracy, understanding this as the % of correctly identified trials when making hard decisions (by selecting the top scored language) Language detection rates are measured in terms of per-language EER and for the sake of clarity we do not deal with the problem of setting optimal thresholds (calibration) as we previously did in [18].

5. Experimental results

5.1. Global performance

We start our study by comparing the performance of the proposed DNN scheme with the baseline i-vector system on the LRE09_VOA_3s_test corpus. Table 3 summarizes this comparison in terms of EER. Results show how the DNN approach largely outperforms the i-vector system, obtaining up to a ~40% relative improvement. An even larger improvement is obtained on the Google 5M corpus, where we found an average relative gain of ~76% (see Figure 3). Those results are especially remarkable since they are found on short test utterances and demonstrate the ability of the DNN to exploit discriminative information in large datasets.

It is also worth analyzing the errors made by the DNN system as a function of the *similarity* of the different languages. We present in Figure A.7 (Appendix A) the confusion matrix obtained using the DNN system on the

²EER is the point on ROC or DET curve where false acceptance and true reject rates are equal.

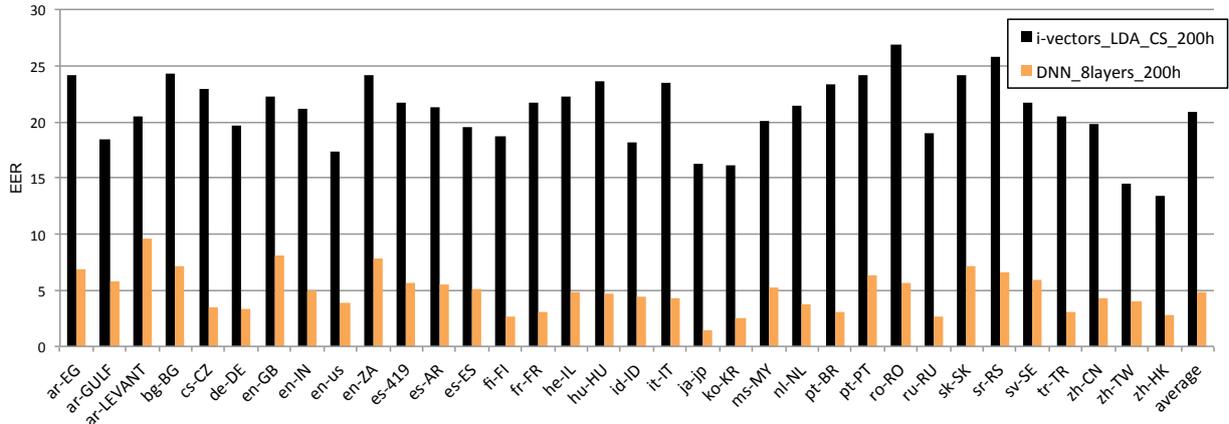


Figure 3: Systems performance (ERR %) comparison per language on Google 5M LID corpus. I-vector baseline system vs DNN_8layers_200h system.

	Equal Error Rate (EER in %)								
	en	es	fa	fr	ps	ru	ur	zh	avg.
DNN_2layers_200h	12.66	5.04	19.67	8.60	17.84	8.75	14.78	5.54	11.61
DNN_4layers_200h	8.53	3.58	16.19	5.82	15.42	6.38	11.24	3.16	8.79
DNN_8layers_200h	8.65	3.74	17.22	7.53	16.01	5.59	13.10	4.82	9.58

Table 4: Effect of using different number of hidden layers. Systems performance (ERR %) per language on LRE09_VOA_test_3s

Google 5M LID corpus. Confusion submatrices around dialects (i.e ar-EG/ar-GULF/ar-LEVANT) illustrate the difficulty of dialect identification from spectral features in short utterances [31]. These results suggest that exploiting just acoustic information might be not enough to reach accurate identification when dealing with dialects [32][33][34].

5.2. Number of hidden layers and training material

In this section, we evaluate two related aspects when training a DNN: the number of hidden layers and the amount of training material used. On one hand, we want to exploit the ability showed by DNNs to improve the recognition performance while increasing the training, avoiding overfitting. On the other hand, we aim to get the lightest architecture possible without losing accuracy.

We started by fixing the available training material to its maximum in LRE'09 (200h per language) and then reducing the number of hidden layers from 8 (DNN_8layers_200h) to 4 (DNN_4layers_200h) and 2 (DNN_2layers_200h). Table 4 summarizes those results. The net with 4-hidden layers seems to be more discriminative than the 2-hidden layers, and more interestingly,

than the one with 8-hidden layers. In particular, on average, the DNN_4layers_200h outperforms by $\sim 8\%$ in terms of EER the DNN_8layers_200h system, using half as many parameters.

As a further step, we swept the number of hours used per language from 1h to 200h for the three nets. Figure 4 shows the % Accuracy as a function of the training hours per language. As expected, the bigger the amount of training data, the better the performance. However, the slope of this gain degrades when reaching 100h per language. Indeed, from the 2 layer system, increasing the training material incurs in a minor degradation mostly due to underfitting. Again, it is clear from the results the need for a convenient tradeoff between the training data and number of parameters to optimize. In particular, our best configuration contains $\sim 21M$ parameters for $\sim 648M$ training samples.

5.3. Real-time identification

Taking now as reference the net with best performance so far (DNN_4_200h) we explored the performance degradation when limiting the test duration. The goal is to gain some insight about how long a test utterance must be to consider the identification *accurate*, a main concern in real-time applications.

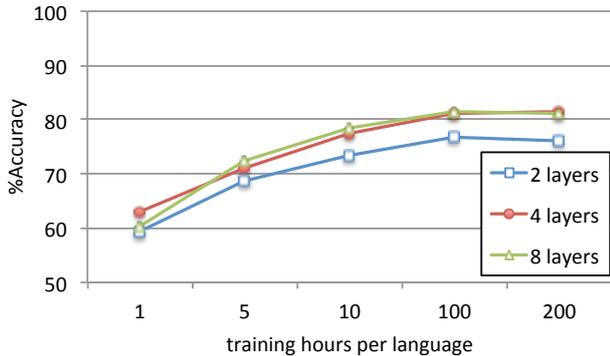


Figure 4: DNN system performance (%Accuracy) in function of the training time per language and the number of hidden layers. Results on LRE09_VOA_3s_test

Figure 5 shows the average accuracy as a function of the test durations, for both test corpus LRE09_VOA_realtime_test and Google 5M LID. We highlight here two main points. Notice first that up to 0.5s of speech (according to our voice activity detection) the identification accuracy is very poor (rates under 50% accuracy). Very quick decisions can lead systems to a bad user experience in real-time applications. Second, as expected, the larger the test duration, the better the performance. However, this practically saturates after 2s. This suggests that a decision can be taken at this point without significant loss of accuracy even when we increase the number of target languages from 9 to 34.

A more detailed analysis per language can be seen in Table A.7 (Appendix A) for all the 34 languages involved in the Google 5M LID corpus, where we show that the previous conclusion holds true also for each individual language. Confusion matrices on LRE09_VOA_realtime_test are also collected in Appendix A.

5.4. Temporal context

So far we have been using a fixed right/left context of ± 10 frames respectively. That is, the input of our network, as mentioned in Section 3, is formed by stacking the features of every frame with its corresponding 10 to the left and 10 to the right neighbours. We explore in this section the effect of including a shorter/wider context for language identification.

The motivation behind using temporal information from a large number of frames lies in the idea of incorporating additional high-level information (i.e. phonetic, phonotactic and prosodic information). This idea has been largely and successfully exploited in language identification by using long-term phonotactic/prosodic

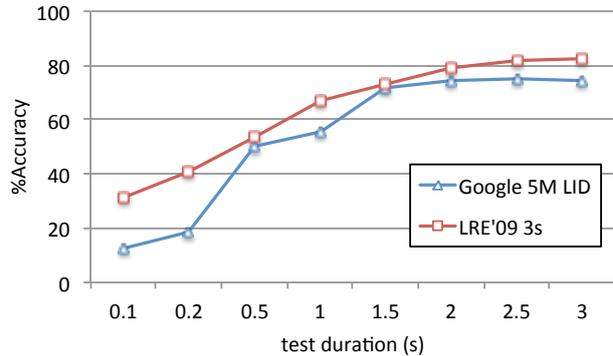


Figure 5: DNN_4layers_200h system performance (% Accuracy) in function of the test utterance duration. Results on LRE09_VOA_realtime_3s.

tokenizations [35][3] or, in acoustic approaches, by using shifted-delta-cepstral features [36].

We modify the input of the network by stacking each frame with a symmetric context that ranges from 0 to 50 left and right neighbour frames; that is, we sweep from a context-free scheme to a maximum context that spans 0.5s to the left and 0.5s to the right (a total of 1s context).

Table 5 summarizes the obtained results on the LRE09_VOA_realtime_test (3s subcorpus) using the DNN_4_200h network. The importance of the context is apparent from first two rows. We observe a relative improvement of $\sim 49\%$ from the ± 10 context scheme with respect to the context-free one. We find the lowest EER when using a ± 20 frames of context. After this value the EER increases. This behaviour can be explained by understanding that as we demand our net to learn more 'high-level' rich features, we are also increasing the size of the input; therefore forcing the net to learn more complex features from the same amount of data. Figure 6 collects the top 10 filters for a given minibatch (those which produce highest activations in the minibatch) extracted from the first hidden layer for the DNN_4_200h network. The distribution of those weights evidences how the DNN is using the context information.

Although the number of parameters of the input layer is affected by the size of the contextual window, the input layer represent less than the 25% of the model size. Thus, it seems that DNNs can lead to better modelling of the contextual information than competing approaches, such as GMM-based systems, which are traditionally more affected by the *curse of dimensionality*. Note that the relative gains reported in this analysis ($\sim 50\%$) surpass previous attempts reported in the literature in including contextual information using the

Equal Error Rate (EER in %)									
	en	es	fa	fr	ps	ru	ur	zh	avg.
No Context	19.07	9.65	24.82	13.17	21.64	14.28	19.39	12.38	16.8
± 10	8.42	3.62	15.89	5.46	14.54	6.31	10.05	3.47	8.47
± 20	7.71	3.88	15.49	6.11	12.9	6.09	10.5	4.00	8.33
± 30	9.44	4.53	16.24	7.95	14.4	7.96	12.07	5.23	9.72
± 40	12.05	5.08	17.41	9.71	15.47	9.14	13.1	6.27	11.03
± 50	9.85	5.71	19.26	8.8	14.54	7.76	13.37	6.51	10.72

Table 5: Effect of using different left/right input contexts for the DNN_4layers.200h system. System performance (ERR %) on LRE09_VOA_realtime_test (3s).

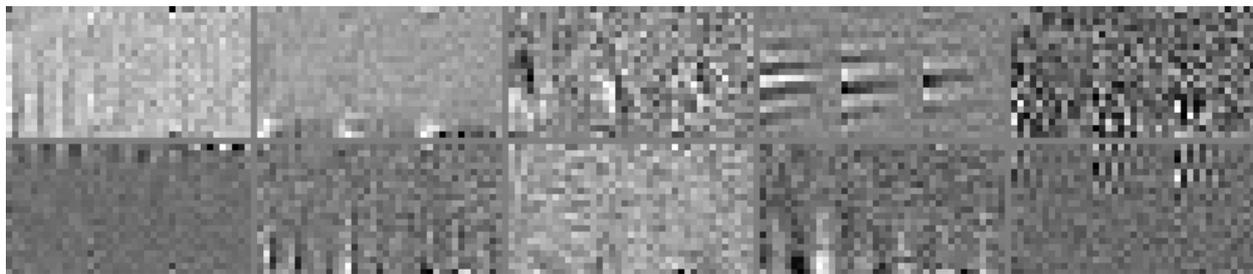


Figure 6: Visualization of top 10 filters (those which produce highest activations in the given minibatch) of the first hidden layer using a ± 10 context. Each filter is composed by 21 rows (number of frames stacked as input) and 39 columns (feature dimension).

GMM paradigm [36]. We refer also to [37] for an extensive comparison of different features in language identification over an i-vector based framework.

5.5. Frame-by-frame posteriors combination

One of the features that make DNNs particularly suitable for real-time applications is their ability to generate frame-by-frame posteriors. Indeed we can derive decisions about the language identification at each frame. Here we aim to study how can we combine frame posteriors into a single utterance-level score.

Probably the most standard way to perform this combination is assuming frame independence and using the product rule (see Section 3). That is, simply compute the product of the posteriors frame-by-frame as the new single score vector. Another common and simple approach used in the literature is plurality voting, where, at each frame, the language associated with the highest posterior receives a single *vote* while the rest receive none. The voting scheme aims to control the negative effect of outlier scores. The score for a given language l , s_l , is then computed by counting the received votes over all the frames as

$$s_l = \sum_{t=1}^N \delta(p(L_l|x_t, \theta)), \quad (11)$$

with δ function defined as

$$\delta(p(L_l|x_t, \theta)) \begin{cases} 1, & \text{if } l == \underset{l}{\arg \max}(p(L_l|x_t, \theta)) \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

A more interesting approach, among *blind* techniques (no need for training), is to weight the posteriors of every frame as a function of the entropy of its posterior distribution. The idea here is to penalize those frames whose distribution of posteriors across the set of languages tends to be uniform (high entropy). This approach was successfully applied in [38], resulting in a performance improvement when working with mismatched training and test datasets. The resulting score for language l , s_l , is computed as

$$s_l = \sum_{t=1}^N \log\left(\frac{1}{h_t} p(L_l|x_t, \theta)\right) \quad (13)$$

where the weight for frame t is the inverse of its entropy

$$h_t = - \sum_{l=1}^N p(L_l|x_t, \theta) \log_2 p(L_l|x_t, \theta) \quad (14)$$

Table 6 compares these three different combination schemes; product, voting and entropy on the

	Equal Error Rate (EER in %)								
	en	es	fa	fr	ps	ru	ur	zh	avg.
Product	8.53	3.58	16.19	5.82	15.42	6.38	11.24	3.16	8.79
Voting	12.66	5.04	19.67	8.60	17.84	8.75	14.78	5.54	11.61
Entropy	8.65	3.74	17.22	7.53	16.01	5.59	13.10	4.82	9.58

Table 6: Comparison of different frame combination schemes for the DNN_4layers_200h. Systems performance (ERR %) per language. Results on LRE09_VOA_3s_test

LRE09_VOA_3s_test corpus. Results show a better performance of the simple product rule compared to the other approaches, with voting the worst choice. This result suggests that making binary decisions at a frame level leads to a performance degradation. Although the entropy scheme does not help in this scenario, it should be considered when working with more noisy environments.

6. Conclusion

In this work, we present a detailed analysis of the use of deep neural networks (DNNs) for automatic language identification (LID) of short utterances. Guided by the success of DNNs for acoustic modelling in speech recognition, we explore the capacity of DNNs to learn language information embedded in speech signals.

Through this study, we also explore the limits of the proposed scheme for real-time applications, evaluating the accuracy of the system when using very short test utterances ($\leq 3s$). We find, for our proposed DNN scheme, that while more than 0.5s is needed to obtain over 50% accuracy rates, 2s are enough to reach accuracy rates of over 90%. Further, we experiment with the amount of training material, the number of hidden layers and the combination of frame posteriors. We also analyze the relevance of including the temporal context, which is critical to achieving high performance in LID.

Results using NIST LRE 2009 (8 languages selected) and Google 5M LID datasets (25 languages + 9 dialects) demonstrate that DNNs outperform current state-of-art i-vector-based approaches when dealing with short test durations. Finally, we demonstrated that using a frame-by-frame approach, DNNs can be successful applied for real-time applications.

7. Future Work

We identified several areas where further investigation is needed. Among them, establishing a more appropriate combination of frame posteriors obtained in

DNNs; exploring different fusions among DNNs and i-vector systems [39]; and dealing with unbalanced training data. Note that even though we proposed different ways of combining posteriors, all of them were blind (no need for training), as we focused on real-time applications and simple approaches. However, other, data-driven methods could be useful at the time of combining those posteriors.

Further neural network architectures should also be explored. For instance, recurrent neural networks might be an elegant solution to incorporating contextual information. Also, convolutional neural networks might help reduce the number of parameters of our model.

Another promising approach is that of using activations of the last hidden layer as bottleneck features. Then, i-vector-based systems or another classification architecture could be trained over those bottleneck features, rather than over classical features, such as PLP or MFCC.

References

- [1] Y. Muthusamy, E. Barnard, R. Cole, Reviewing automatic language identification, *Signal Processing Magazine*, IEEE 11 (4) (1994) 33–41. doi:10.1109/79.317925.
- [2] M. Zissman, Comparison of Four Approaches to Automatic Language Identification of Telephone Speech, *IEEE Trans. Acoust., Speech, Signal Processing* 4 (1) (1996) 31–44.
- [3] L. Ferrer, N. Scheffer, E. Shriberg, A Comparison of Approaches for Modeling Prosodic Features in Speaker Recognition, in: *International Conference on Acoustics, Speech, and Signal Processing*, 2010, pp. 4414–4417. doi:10.1109/ICASSP.2010.5495632.
- [4] D. Martinez, E. Lleida, A. Ortega, A. Miguel, Prosodic features and formant modeling for an i-vector-based language recognition system, in: *Acoustics, Speech and Signal Processing (ICASSP)*, 2013 IEEE International Conference on, 2013, pp. 6847–6851. doi:10.1109/ICASSP.2013.6638988.
- [5] P. Torres-Carrasquillo, E. Singer, T. Gleason, A. McCree, D. Reynolds, F. Richardson, D. Sturim, The MITLL NIST LRE 2009 Language Recognition System, in: *Acoustics Speech and Signal Processing (ICASSP)*, 2010 IEEE International Conference on, 2010, pp. 4994–4997. doi:10.1109/ICASSP.2010.5495080.
- [6] J. Gonzalez-Dominguez, I. Lopez-Moreno, J. Franco-Pedroso, D. Ramos, D. Toledano, J. Gonzalez-Rodriguez, Multilevel and Session Variability Compensated Language Recognition:

- ATVS-UAM Systems at NIST LRE 2009, Selected Topics in Signal Processing, *IEEE Journal of* 4 (6) (2010) 1084–1093. doi:10.1109/JSTSP.2010.2076071.
- [7] H. Li, B. Ma, K. A. Lee, Spoken language recognition: From fundamentals to practice, *Proceedings of the IEEE* 101 (5) (2013) 1136–1159. doi:10.1109/JPROC.2012.2237151.
- [8] E. Ambikairajah, H. Li, L. Wang, B. Yin, V. Sethu, Language identification: A tutorial, *Circuits and Systems Magazine, IEEE* 11 (2) (2011) 82–108. doi:10.1109/MCAS.2011.941081.
- [9] M. A. Zissman, K. Berkling, Automatic Language Identification, *Speech Communication* 35 (1-2) (2001) 115–124.
- [10] N. Brummer, S. Cumani, O. Glembek, M. Karafiát, P. Matejka, J. Pesan, O. Plchot, M. M. Soufifar, E. V. de, J. Černocký, Description and Analysis of the Brno276 System for LRE2011, in: *Proceedings of Odyssey 2012: The Speaker and Language Recognition Workshop, International Speech Communication Association, 2012*, pp. 216–223.
- [11] D. Sturim, W. Campbell, N. Dehak, Z. Karam, A. McCree, D. Reynolds, F. Richardson, P. Torres-Carrasquillo, S. Shum, The MIT LL 2010 Speaker Recognition Evaluation System: Scalable Language-Independent Speaker Recognition, in: *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on, 2011*, pp. 5272–5275. doi:10.1109/ICASSP.2011.5947547.
- [12] N. Dehak, P. A. Torres-Carrasquillo, D. A. Reynolds, R. Dehak, Language Recognition via i-vectors and Dimensionality Reduction., in: *INTERSPEECH, ISCA, 2011*, pp. 857–860.
- [13] P. Kenny, P. Oullet, V. Dehak, N. Gupta, P. Dumouchel, A Study of Interspeaker Variability in Speaker Verification, *IEEE Trans. on Audio, Speech and Language Processing* 16 (5) (2008) 980–988.
- [14] A. Mohamed, G. Dahl, G. Hinton, Acoustic Modeling using Deep Belief Networks, *Audio, Speech, and Language Processing, IEEE Transactions on* 20 (1) (2012) 14–22. doi:10.1109/TASL.2011.2109382.
- [15] G. Hinton, L. Deng, D. Yu, G. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, B. Kingsbury, Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups, *Signal Processing Magazine, IEEE* 29 (6) (2012) 82–97. doi:10.1109/MSP.2012.2205597.
- [16] D. Ciresan, U. Meier, L. Gambardella, J. Schmidhuber, Deep Big Simple Neural Nets Excel on Handwritten Digit Recognition, *CoRR abs/1003.0358*.
- [17] D. Yu, L. Deng, Deep Learning and its Applications to Signal and Information Processing [Exploratory DSP], *Signal Processing Magazine, IEEE* 28 (1) (2011) 145–154. doi:10.1109/MSP.2010.939038.
- [18] I. Lopez-Moreno, J. Gonzalez-Dominguez, O. Plchot, D. Martinez, J. Gonzalez-Rodriguez, P. Moreno, Automatic Language Identification using Deep Neural Networks, *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, to appear.
- [19] R. Cole, J. Inouye, Y. Muthusamy, M. Gopalakrishnan, Language Identification with Neural Networks: A Feasibility Study, in: *Communications, Computers and Signal Processing, 1989. Conference Proceeding., IEEE Pacific Rim Conference on, 1989*, pp. 525–529. doi:10.1109/PACRIM.1989.48417.
- [20] M. Leena, K. Srinivasa Rao, B. Yegnanarayana, Neural Network Classifiers for Language Identification using Phonotactic and Prosodic Features, in: *Intelligent Sensing and Information Processing, 2005. Proceedings of 2005 International Conference on, 2005*, pp. 404–408. doi:10.1109/ICISIP.2005.1529486.
- [21] G. Montavon, Deep Learning for Spoken Language Identification, in: *NIPS workshop on Deep Learning for Speech Recognition and Related Applications, 2009*.
- [22] N. Dehak, P. Kenny, R. Dehak, P. Dumouchel, P. Ouellet, Front-End Factor Analysis for Speaker Verification, *Audio, Speech, and Language Processing, IEEE Transactions on* 19 (4) (2011) 788 – 798.
- [23] D. Martinez, O. Plchot, L. Burget, O. Glembek, P. Matejka, Language Recognition in iVectors Space., in: *INTERSPEECH, ISCA, 2011*, pp. 861–864.
- [24] D. Reynolds, Speaker Identification and Verification Using Gaussian Mixture Speaker Models, *Speech Communication* 17 (1-2) (1995) 91–108.
- [25] C. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st Edition, Springer, 2007.
- [26] H. Hermansky, Perceptual Linear Predictive (PLP) Analysis of Speech, *The Journal of the Acoustical Society of America* 87 (4) (1990) 1738–1752.
- [27] A. rahman Mohamed, G. E. Hinton, G. Penn, Understanding how Deep Belief Networks perform Acoustic Modelling., in: *ICASSP, IEEE, 2012*, pp. 4273–4276.
- [28] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, Q. Le, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, A. Ng, Large Scale Distributed Deep Networks, in: P. Bartlett, F. Pereira, C. Burges, L. Bottou, K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25, 2012*, pp. 1232–1240.
- [29] NIST, The 2009 NIST SLR Evaluation Plan, www.itl.nist.gov/iad/mig/tests/lre/2009/LRE09_EvalPlan_v6.pdf (2009).
- [30] G. Liu, C. Zhang, J. H. L. Hansen, A Linguistic Data Acquisition Front-End for Language Recognition Evaluation, in: *Proc. Odyssey, Singapore, 2012*.
- [31] P. A. Torres-Carrasquillo, D. E. Sturim, D. A. Reynolds, A. McCree, Eigen-channel Compensation and Discriminatively Trained Gaussian Mixture Models for Dialect and Accent Recognition., in: *INTERSPEECH, 2008*, pp. 723–726.
- [32] F. Biadsy, Automatic Dialect and Accent Recognition and its Application to Speech Recognition, Ph.D. thesis, Columbia University (2011).
- [33] W. Baker, D. Eddington, L. Nay, Dialect Identification: The Effects of Region of Origin and Amount of Experience, *American Speech* 84 (1) (2009) 48–71.
- [34] G. Liu, Y. Lei, J. H. Hansen, Dialect Identification: Impact of Difference between Read versus Spontaneous Speech, in: *EUSIPCO-2010, 2003-2006*.
- [35] D. Reynolds, W. Andrews, J. Campbell, J. Navratil, B. Peskin, A. Adami, Q. Jin, D. Klusacek, J. Abramson, R. Mihaescu, J. Godfrey, D. Jones, B. Xiang, The SuperSID Project: Exploiting High-Level Information for High-Accuracy Speaker Recognition, in: *IEEE International Conference on Acoustics, Speech, and Signal Processing, Vol. 4, 2003*, pp. 784–787.
- [36] P. A. Torres-Carrasquillo, E. Singer, M. A. Kohler, J. R. Deller, Approaches to Language Identification Using Gaussian Mixture Models and Shifted Delta Cepstral Features, in: *ICSLP, Vol. 1, 2002*, pp. 89–92.
- [37] M. Li, S. Narayanan, Simplified Supervised i-vector Modeling with Application to Robust and Efficient Language Identification and Speaker Verification, *Computer, Speech, and Language*.
- [38] H. Misra, H. Bourlard, V. Tyagi, New Entropy based Combination rules in HMM/ANN Multi-Stream ASR, in: *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03). 2003 IEEE International Conference on, Vol. 2, 2003*, pp. II-741–4 vol.2. doi:10.1109/ICASSP.2003.1202473.
- [39] G. Saon, H. Soltau, D. Nahamoo, M. Picheny, Speaker Adaptation of Neural Network Acoustic Models using i-

vectors, in: Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on, 2013, pp. 55–59. doi:10.1109/ASRU.2013.6707705.

Appendix A. Extended Results

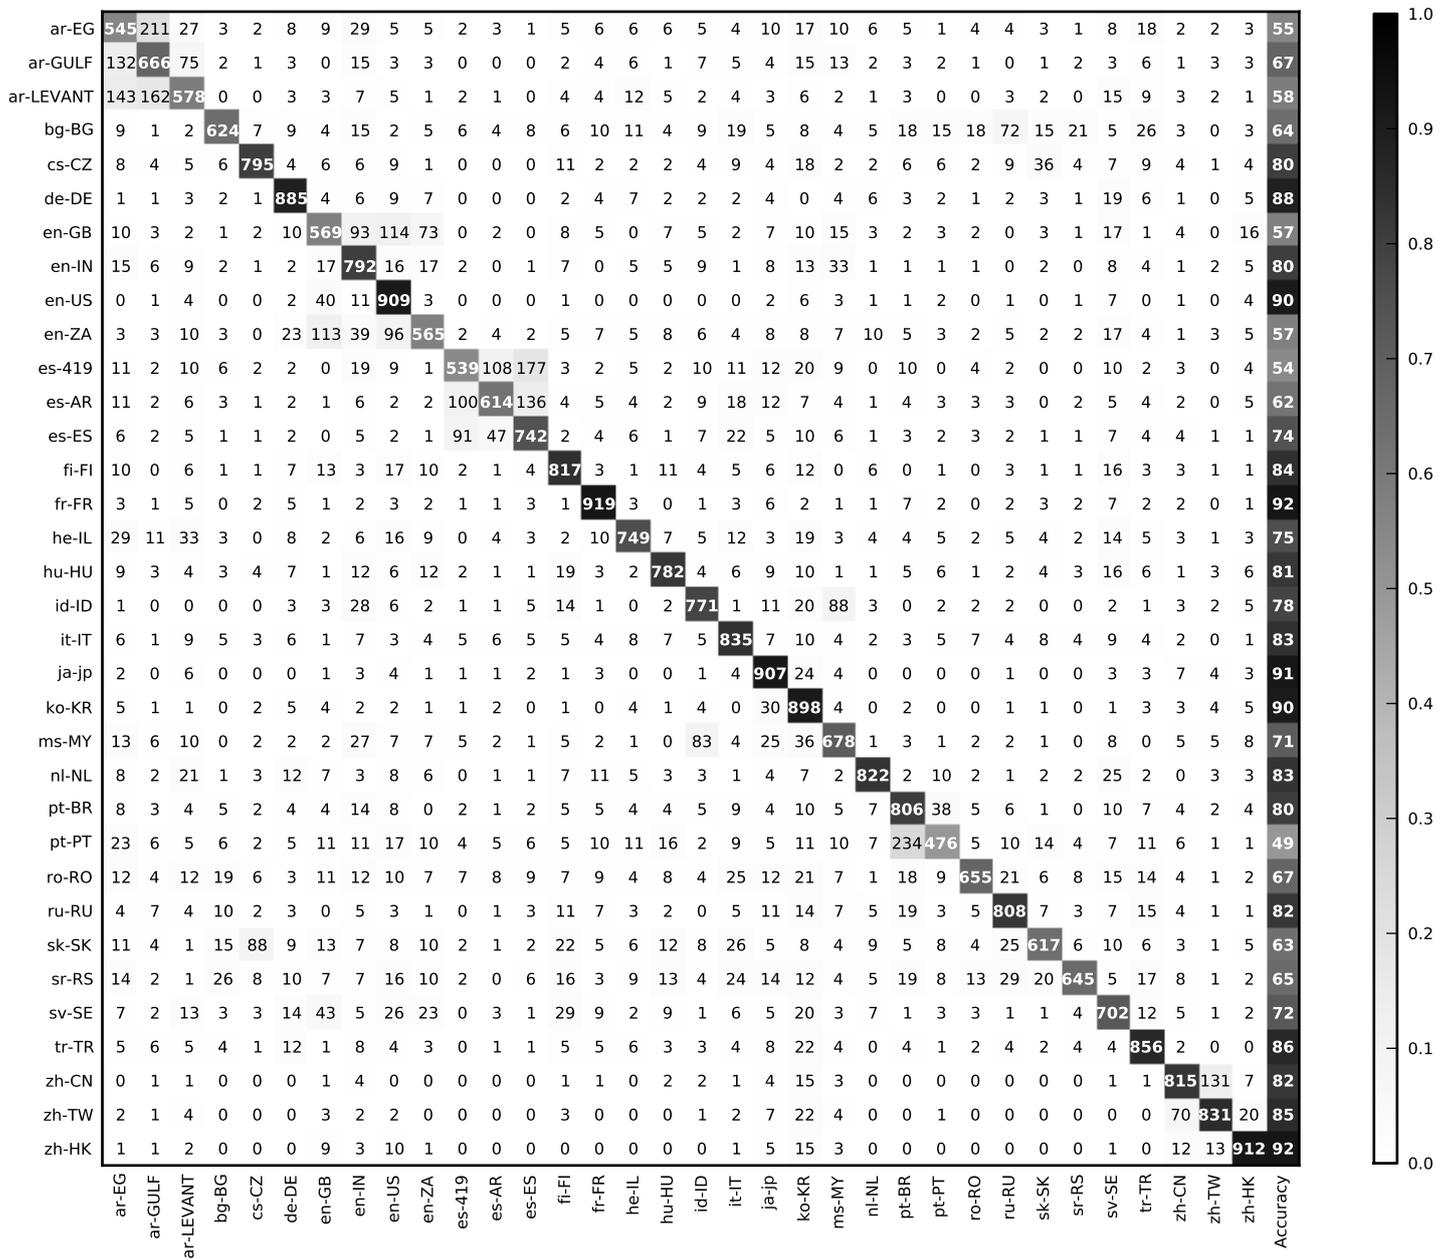


Figure A.7: Confusion matrix obtained by evaluating the DNN.8_200h system on the Google 5M LID corpus.

% Accuracy		Test utterance duration							
Locale	Language	0.1s	0.2s	0.5s	1.0s	1.5s	2.0s	2.5s	3.0s
ar-EG	Arabic (Egypt)	12	14	38	41	52	56	58	57
ar-GULF	Arabic (Persian)	19	22	46	53	61	64	69	68
ar-LEVANT	Arabic (Levant)	43	51	54	48	65	64	61	62
bg-BG	Bulgarian	7	11	38	51	65	64	68	72
cs-CZ	Czech	2	6	45	69	71	75	79	81
de-DE	German	18	27	62	72	84	88	87	89
en-GB	English (United)	4	12	31	39	49	54	54	54
en-IN	English (India)	27	30	56	63	73	74	76	78
en-US	English (USA)	22	29	62	70	85	87	89	91
en-ZA	English (South)	4	7	34	45	46	51	56	57
es-419	Spanish (Latin)	6	8	25	41	47	50	52	55
es-AR	Spanish (Argentina)	6	8	35	50	53	56	58	61
es-ES	Spanish (Spain)	5	9	48	54	67	70	72	73
fi-FI	Finish	14	23	55	75	76	80	82	82
fr-FR	French	14	25	69	83	90	93	94	95
he-IL	Hebrew (Israel)	4	10	46	60	60	67	68	70
hu-HU	Hungarian	8	16	48	71	72	80	82	82
id-ID	Indonesian	13	21	45	62	69	72	75	76
it-IT	Italian	8	13	42	58	75	78	80	81
ja-JP	Japanese	18	25	68	87	89	91	94	95
ko-KR	Korean (South)	16	25	68	91	89	91	92	92
ms-MY	Malay	17	25	44	59	63	70	72	72
nl-NL	Dutch	6	12	56	68	76	80	80	81
pt-BR	Portuguese (Brazilian)	11	18	47	74	74	78	80	81
pt-PT	Portuguese (Portugal)	6	8	28	53	42	43	48	49
ro-RO	Romanian	7	12	34	43	56	61	64	66
ru-RU	Russian	5	11	52	70	83	85	85	85
sk-SK	Slovak	10	13	30	40	48	51	55	58
sr-RS	Serbian	6	9	35	54	55	59	60	62
sv-SE	Sweden	10	16	42	62	65	70	73	71
tr-TR	Turkish	5	10	55	78	79	82	83	85
zh-cmn-Hans-CN	Chinese (Mandarin)	12	16	54	76	75	76	80	82
zh-cmn-Hant-TW	Chinese (Taiwan)	12	22	63	78	80	83	83	85
zh-yue-hant-HK	Chinese (Cantonese)	15	25	68	81	88	91	90	91

Table A.7: Systems performance (Accuracy %) by language and test utterance duration on Google 5M Database.

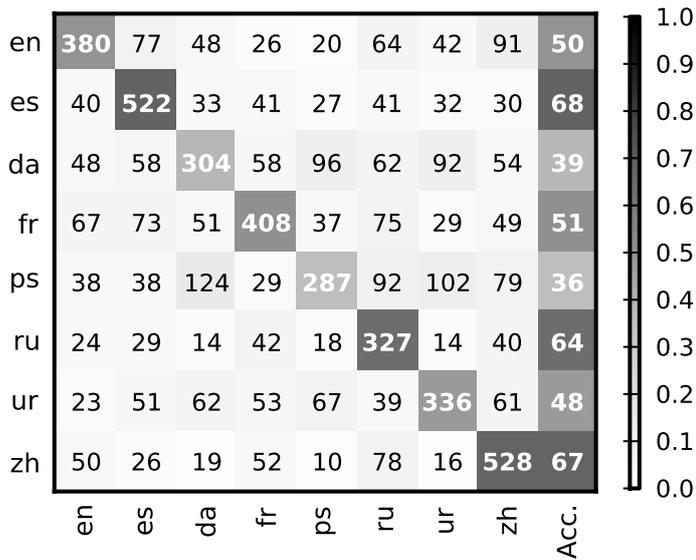


Figure A.8: DNN_4layers_200h confusion matrix on LRE'09 (0.5s test)

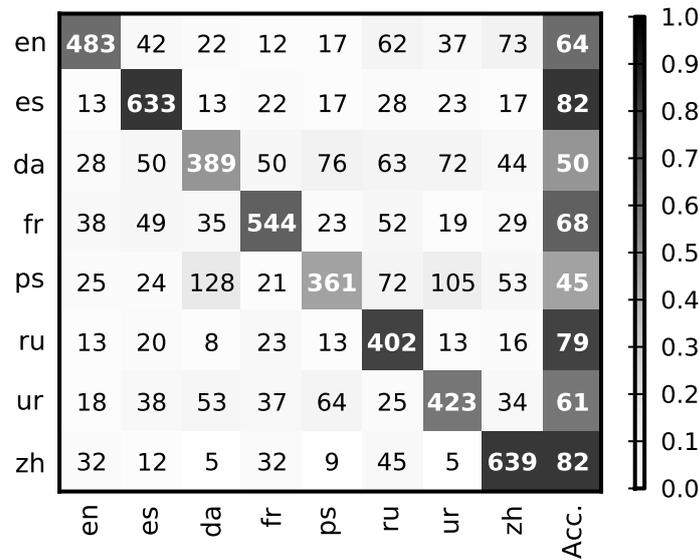


Figure A.9: DNN_4layers_200h confusion matrix on LRE'09 (1s test)

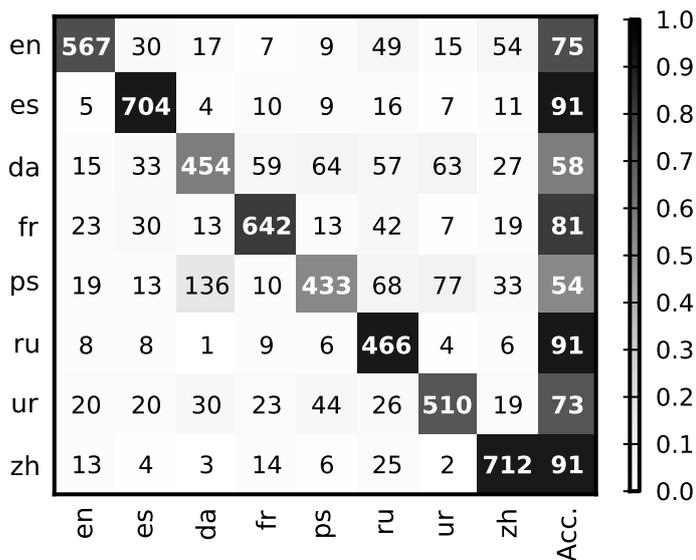


Figure A.10: DNN_4layers_200h confusion matrix on LRE'09 (2s test)

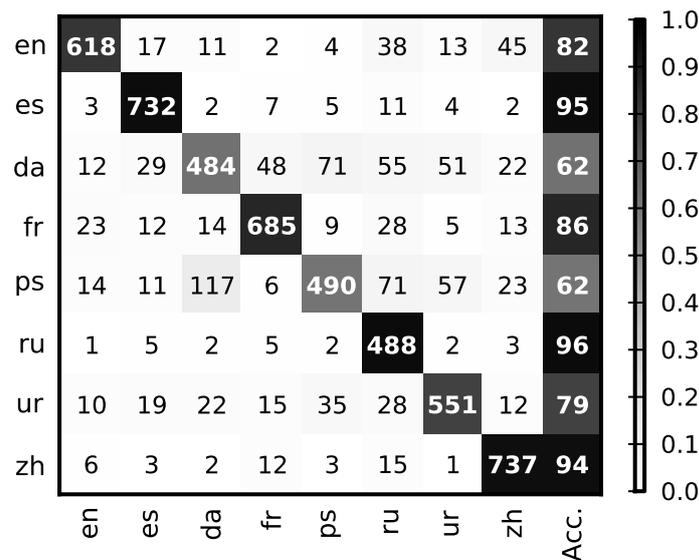


Figure A.11: DNN_4layers_200h confusion matrix on LRE'09 (3s test)