

A MULTI PURPOSE AND LARGE SCALE SPEECH CORPUS IN PERSIAN AND ENGLISH FOR SPEAKER AND SPEECH RECOGNITION: THE DEEPMINE DATABASE

Hossein Zeinali^{1,2}, Lukáš Burget¹ and Jan “Honza” Černocký¹

¹ Brno University of Technology, Faculty of IT, IT4I Centre of Excellence, Czechia

² Sharif DeepMine Ltd., Tehran, Iran

ABSTRACT

DeepMine is a speech database in Persian and English designed to build and evaluate text-dependent, text-prompted, and text-independent speaker verification, as well as Persian speech recognition systems. It contains more than 1850 speakers and 540 thousand recordings overall, more than 480 hours of speech are transcribed. It is the first public large-scale speaker verification database in Persian, the largest public text-dependent and text-prompted speaker verification database in English, and the largest public evaluation dataset for text-independent speaker verification. It has a good coverage of age, gender, and accents. We provide several evaluation protocols for each part of the database to allow for research on different aspects of speaker verification. We also provide the results of several experiments that can be considered as baselines: HMM-based i-vectors for text-dependent speaker verification, and HMM-based as well as state-of-the-art deep neural network based ASR. We demonstrate that the database can serve for training robust ASR models.

Index Terms— speech database, text-dependent, text-independent, speaker verification, speech recognition

1. INTRODUCTION

Nowadays deep learning techniques outperform the other conventional methods in most of the speech-related tasks. Training robust deep neural networks for each task depends on the availability of powerful processing GPUs, as well as standard and large scale datasets. In text-independent speaker verification, large-scale datasets are available, thanks to the NIST SRE evaluations and other data collection projects such as VoxCeleb [1].

In text-dependent speaker recognition, experiments with end-to-end architectures conducted on large proprietary databases have demonstrated their superiority over traditional approaches [2]. Yet, contrary to text-independent speaker recognition, text-dependent speaker recognition lacks large-scale publicly available databases. The two most well-known datasets are probably RSR2015 [3] and RedDots [4]. The former contains speech data collected from 300 individuals in a controlled manner, while the latter is used primarily for evaluation rather than training, due to its small number of speakers (only 64). Motivated by this lack of large-scale dataset for text-dependent speaker verification, we chose to proceed with the collection of the DeepMine dataset, which we expect to become a standard benchmark for the task.

Apart from speaker recognition, large amounts of training data are required also for training automatic speech recognition (ASR) systems. Such datasets should not only be large in size, they should also be characterized by high variability with respect to speakers, age and dialects. While several datasets with these properties are available for languages like English, Mandarin, French, this is not

the case for several other languages, such as Persian. To this end, we proceeded with collecting a large-scale dataset, suitable for building robust ASR models in Persian.

The main goal of the DeepMine project was to collect speech from at least a few thousand speakers, enabling research and development of deep learning methods. The project started at the beginning of 2017, and after designing the database and the developing Android and server applications, the data collection began in the middle of 2017. The project finished at the end of 2018 and the cleaned-up and final version of the database was released at the beginning of 2019. In [5], the running project and its data collection scenarios were described, alongside with some preliminary results and statistics. In this paper, we announce the final and cleaned-up version of the database, describe its different parts and provide various evaluation setups for each part. Finally, since the database was designed mainly for text-dependent speaker verification purposes, some baseline results are reported for this task on the official evaluation setups. Additional baseline results are also reported for Persian speech recognition. However, due to the space limitation in this paper, the baseline results are not reported for all the database parts and conditions. They will be defined and reported in the database technical documentation and in a future journal paper.

2. DATA COLLECTION

DeepMine is publicly available for everybody with a variety of licenses for different users. It was collected using crowdsourcing [5]. The data collection was done using an Android application. Each respondent installed the application on his/her personal device and recorded several phrases in different sessions. The Android application did various checks on each utterance and if it passed all of them, the respondent was directed to the next phrase. For more information about data collection scenario, please refer to [5].

2.1. Post-Processing

In order to clean-up the database, the main post-processing step was to filter out problematic utterances. Possible problems include speaker word insertions (e.g. repeating some part of a phrase), deletions, substitutions, and involuntary disfluencies. To detect these, we implemented an alignment stage, similar to the second alignment stage in the LibriSpeech project [6]. In this method, a custom decoding graph was generated for each phrase. The decoding graph allows for word skipping and word insertion in the phrase.

For text-dependent and text-prompted parts of the database, such errors are not allowed. Hence, any utterances with errors were removed from the enrollment and test lists. For the speech recognition part, a sub-part of the utterance which is correctly aligned to the corresponding transcription is kept. After the cleaning step, around

Table 1. Statistics of the DeepMine database.

	Count
Finished sessions	22,741
Finished sessions containing English	18,916
Recorded utterances	544,533
Unique speakers	1,858
Respondents with at least 5 utterances	1,969
Respondents with at least 1 session	1,796
Respondents with at least 4 sessions	1,104
Respondents with at least 8 sessions	813
Respondents with at least 16 sessions	607
Respondents with at least 60 sessions	98

190 thousand utterances with full transcription and 10 thousand with sub-part alignment have remained in the database.

2.2. Statistics

After processing the database and removing problematic respondents and utterances, 1969 respondents¹ remained in the database, with 1149 of them being male and 820 female. 297 of the respondents could not read English and have therefore read only the Persian prompts. About 13200 sessions were recorded by females and similarly, about 9500 sessions by males, i.e. women are over-represented in terms of sessions, even though their number is 17 % smaller than that of males. Other useful statistics related to the database are shown in Table 1.

The last status of the database, as well as other related and useful information about its availability can be found on its website, together with a limited number of samples².

3. DEEPMINE DATABASE PARTS

The DeepMine database consists of three parts. The first one contains fixed common phrases to perform text-dependent speaker verification. The second part consists of random sequences of words useful for text-prompted speaker verification, and the last part includes phrases with word- and phoneme-level transcription, useful for text-independent speaker verification using a random phrase (similar to Part4 of RedDots). This part can also serve for Persian ASR training. Each part is described in more details below. Table 2 shows the number of unique phrases in each part of the database. For the English text-dependent part, the following phrases were selected from part1 of the RedDots database, hence the RedDots can be used as an additional training set for this part:

1. “My voice is my password.”
2. “OK Google.”
3. “Artificial intelligence is for real.”
4. “Actions speak louder than words.”
5. “There is no such thing as a free lunch.”

¹By “respondent”, we mean a unique combination of speaker and mobile phone. Hence, one speaker with several different mobile devices is considered as several respondents.

²<http://data.deepmine.ir/en/>

Table 2. Numbers of unique phrases in each phrase type. “3-months/4-digits” means a random subset of months/digits used as test phrase. “Persian name and family” means a sequence of several Persian full names.

	Count	Part
Persian text-dependent	5	1
English text-dependent	5	1
Persian months-name, 3-months	1,320	2
Persian months-name, 12-months	11,620	2
English digits, 4-digits	5,040	2
English digits, 10-digits	6,448	2
Persian name and family	8,305	3
Persian transcribed phrases	166,838	3

3.1. Part1 - Text-dependent (TD)

This part contains a set of fixed phrases which are used to verify speakers in text-dependent mode. Each speaker utters 5 Persian phrases, and if the speaker can read English, 5 phrases selected from Part1 of the RedDots database are also recorded.

We have created three experimental setups with different numbers of speakers in the evaluation set. For each setup, speakers with more recording sessions are included in the evaluation set and the rest of the speakers are used for training in the background set (in the database, all background sets are basically training data). The rows in Table 3 corresponds to the different experimental setups and shows the numbers of speakers in each set. Note that, for English, we have filtered the (Persian native) speakers by the ability to read English. Therefore, there are fewer speakers in each set for English than for Persian. There is a small “dev” set in each setup which can be used for parameter tuning to prevent over-tuning on the evaluation set.

For each experimental setup, we have defined several official trial lists with different numbers of enrollment utterances per trial in order to investigate the effects of having different amounts of enrollment data. All trials in one trial list have the same number of enrollment utterances (3 to 6) and only one test utterance. All enrollment utterances in a trial are taken from different consecutive sessions and the test utterance is taken from yet another session. From all the setups and conditions, the 100-spkr with 3-session enrollment (3-sess) is considered as the main evaluation condition. In

Table 3. Number of speakers in different evaluation sets of text-dependent part.

Lang	Set name	Male			Female		
		back	dev	eval	back	dev	eval
Persian	100-spkr	851	10	100	607	10	100
	200-spkr	741	20	200	497	20	200
	300-spkr	631	30	300	387	30	300
English	100-spkr	722	8	85	510	8	88
	200-spkr	620	18	175	415	17	174
	300-spkr	538	29	245	318	27	261

Table 4. Numbers of trials in each different evaluation set for all possible trial types in text-dependent SV for 3-sess enrollment. TC: Target-Correct, TW: Target-Wrong, IC: Imposter-Correct, IW: Imposter-Wrong. K stands for thousand and M stands for million trials.

Set name	Male				Female			
	TC	TW	IC	IW	TC	TW	IC	IW
100-spkr	409K	1.63M	30.7M	30.7M	667K	2.67M	60.0M	60.0M
200-spkr	452K	1.81M	58.0M	58.0M	792K	3.17M	126M	126M
300-spkr	480K	1.92M	85.6M	85.6M	835K	3.34M	176M	176M

Table 4, the number of trials for Persian 3-sess are shown for the different types of trial in the text-dependent speaker verification (SV). Note that for Imposter-Wrong (IW) trials (i.e. imposter speaker pronouncing wrong phrase), we merely create one wrong trial for each Imposter-Correct (IC) trial to limit the huge number of possible trials for this case. So, the number of trials for IC and IW cases are the same.

3.2. Part2 - Text-prompted (TP)

For this part, in each session, 3 random sequences of Persian month names are shown to the respondent in two modes: In the first mode, the sequence consists of all 12 months, which will be used for speaker enrollment. The second mode contains a sequence of 3 month names that will be used as a test utterance. In each 8 sessions received by a respondent from the server, there are 3 enrollment phrases of all 12 months (all in just one session), and 7×3 other test phrases, containing fewer words. For a respondent who can read English, 3 random sequences of English digits are also recorded in each session. In one of the sessions, these sequences contain all digits and the remaining ones contain only 4 digits.

Similar to the text-dependent case, three experimental setups with different number of speaker in the evaluation set are defined (corresponding to the rows in Table 5). However, different strategy is used for defining trials: Depending on the enrollment condition (1- to 3-sess), trials are enrolled on utterances of all words from 1 to 3 different sessions (i.e. 3 to 9 utterances). Further, we consider two conditions for test utterances: **seq** test utterance with only 3 or 4 words and **full** test utterances with all words (i.e. same words as in enrollment but in different order). From all setups an all conditions, the 100-spkr with 1-session enrolment (1-sess) is considered as the main evaluation condition for the text-prompted case. In Table 5, the numbers of trials (sum for both seq and full conditions) for Persian 1-sess are shown for the different types of trials in the text-prompted SV. Again, we just create one IW trial for each IC trial.

3.3. Part3 - Text-independent (TI)

In this part, 8 Persian phrases that have already been transcribed on the phone level are displayed to the respondent. These phrases are chosen mostly from news and Persian Wikipedia. If the respondent is unable to read English, instead of 5 fixed phrases and 3 random digit strings, 8 other Persian phrases are also prompted to the respondent to have exactly 24 phrases in each recording session.

This part can be useful at least for three potential applications. First, it can be used for text-independent speaker verification. The second application of this part (same as Part4 of RedDots) is text-

prompted speaker verification using random text (instead of a random sequence of words). Finally, the third application is large vocabulary speech recognition in Persian (explained in the next subsection).

Based on the recording sessions, we created two experimental setups for speaker verification. In the first one, respondents with at least 17 recording sessions are included to the evaluation set, respondents with 16 sessions to the development and the rest of respondents to the background set (can be used as training data). In the second setup, respondents with at least 8 sessions are included to the evaluation set, respondents with 6 or 7 sessions to the development and the rest of respondents to the background set. Table 6 shows numbers of speakers in each set of the database for text-independent SV case.

For text-independent SV, we have considered 4 scenarios for enrollment and 4 scenarios for test. The speaker can be enrolled using utterances from 1, 2 or 3 consecutive sessions (1sess to 3sess) or using 8 utterances from 8 different sessions. The test speech can be one utterance (1utt) for short duration scenario or all utterances in one session (1sess) for long duration case. In addition, test speech can be selected from 5 English phrases for cross-language testing (enrollment using Persian utterances and test using English utterances). From all setups, 1sess-1utt and 1sess-1sess for 438-spkr set are considered as the main evaluation setups for text-independent case. Table 7 shows number of trials for these setups.

For text-prompted SV with random text, the same setup as text-independent case together with corresponding utterance transcriptions can be used.

3.4. Part3 - Speech Recognition

As explained before, Part3 of the DeepMine database can be used for Persian read speech recognition. There are only a few databases for speech recognition in Persian [7, 8]. Hence, this part can at least partly address this problem and enable robust speech recognition applications in Persian. Additionally, it can be used for speaker recognition applications, such as training deep neural networks (DNNs) for extracting bottleneck features [9], or for collecting sufficient statistics using DNNs for i-vector training.

We have randomly selected 50 speakers (25 for each gender) from the all speakers in the database which have net speech (without silence parts) between 25 minutes to 50 minutes as test speakers. For each speaker, the utterances in the first 5 sessions are included to (small) test-set and the other utterances of test speakers are considered as a large-test-set. The remaining utterances of the other speakers are included in the training set. The test-set, large-test-set and train-set contain 5.9, 28.5 and 450 hours of speech respectively.

There are about 8300 utterances in Part3 which contain only Persian full names (i.e. first and family name pairs). Each phrase consists of several full names and their phoneme transcriptions were ex-

Table 5. Numbers of trials in each different evaluation sets for all possible trial types in text-dependent SV for 1-sess enrollment.

Set name	Male				Female			
	TC	TW	IC	IW	TC	TW	IC	IW
100-spkr	98.5K	394K	7.61M	7.61M	155K	621K	14.2M	14.2M
200-spkr	111K	446K	15.2M	15.2M	187K	749K	30.5M	30.5M
300-spkr	119K	474K	22.5M	22.5M	199K	798K	43.9M	43.9M

Table 6. Different evaluation sets of text-independent part of the DeepMine database.

Set name	Male			Female		
	back	dev	eval	back	dev	eval
438-spkr	689	91	181	376	85	257
798-spkr	540	59	365	220	65	433

tracted automatically using a trained Grapheme-to-Phoneme (G2P). These utterances can be used to evaluate the performance of a systems for name recognition, which is usually more difficult than the normal speech recognition because of the lack of a reliable language model.

4. EXPERIMENTS AND RESULTS

Due to the space limitation, we present results only for the Persian text-dependent speaker verification and speech recognition.

4.1. Speaker Verification Experiments

We conducted an experiment on text-dependent speaker verification part of the database, using the i-vector based method proposed in [10, 11] and applied it to the Persian portion of Part1. In this experiment, 20-dimensional MFCC features along with first and second derivatives are extracted from 16 kHz signals using HTK [12] with 25 ms Hamming windowed frames with 15 ms overlap.

The reported results are obtained with a 400-dimensional gender independent i-vector based system. The i-vectors are first length-normalized and are further normalized using phrase- and gender-dependent Regularized Within-Class Covariance Normalization (RWCCN) [11]. Cosine distance is used to obtain speaker verification scores and phrase- and gender-dependent s-norm is used for normalizing the scores. For aligning speech frames to Gaussian components, monophone HMMs with 3 states and 8 Gaussian components in each state are used [11]. We only model the phonemes which appear in the 5 Persian text-dependent phrases.

For speaker verification experiments, the results were reported in terms of Equal Error Rate (EER) and Normalized Detection Cost Function as defined for NIST SRE08 ($NDCF_{0.01}^{\min}$) and NIST SRE10 ($NDCF_{0.001}^{\min}$). As shown in Table 8, in text-dependent SV there

Table 7. Numbers of trials in each different evaluation sets for all possible trial types in text-dependent SV for 3-sess enrollment. “cross” means enrollment using Persian utterances and test using 5 English utterances. K stands for thousand and M stands for million trials.

Set name	Male		Female	
	Target	Imposter	Target	Imposter
1sess-1utt	537K	61.7M	1013K	193.1M
1sess-1utt-cross	459K	51.3M	890K	164.4M
1sess-1sess	135K	15.6M	254K	48.7M
1sess-1sess-cross	229K	25.6M	445K	82.2M

Table 8. Trial types in text-dependent speaker verification [3].

	Target Speaker	Imposter Speaker
Correct Pass-Phrase	Target-Correct	Imposter-Correct
Wrong Pass-Phrase	Target-Wrong	Imposter-Wrong

are 4 types of trials: *Target-Correct* and *Imposter-Correct* refer to trials when the pass-phrase is uttered correctly by target and imposter speakers respectively, and in same manner, *Target-Wrong* and *Imposter-Wrong* refer to trials when speakers uttered a wrong pass-phrase. In this paper, only the correct trials (i.e. Target-Correct as target trials vs Imposter-Correct as non-target trials) are considered for evaluating systems as it has been proved that these are the most challenging trials in text-dependent SV [9, 13].

Table 9 shows the results of text-dependent experiments using Persian 100-spkr and 3-sess setup. For filtering trials, the respondents’ mobile brand and model were used in this experiment. In the table, the first two letters in the filter notation relate to the target trials and the second two letters (i.e. right side of the colon) relate for non-target trials. For target trials, the first Y means the enrolment and test utterances were recorded using a device with the same brand by the target speaker. The second Y letter means both recordings were done using exactly the same device model. Similarly, the first Y for non-target trials means that the devices of target and imposter speakers are from the same brand (i.e. manufacturer). The second Y means that, in addition to the same brand, both devices have the same model. So, the most difficult target trials are “NN”, where the speaker has used different a device at the test time. In the same manner, the most difficult non-target trials which should be rejected by the system are “YY” where the imposter speaker has used the same device model as the target speaker (note that it does not mean physically the same device because each speaker participated in the project using a personal mobile device). Hence, the similarity in the recording channel makes rejection more difficult.

The first row in Table 9 shows the results for all trials. By comparing the results with the best published results on RSR2015 and RedDots [11, 9, 13], it is clear that the DeepMine database is more challenging than both RSR2015 and RedDots databases. For RSR2015, the same i-vector/HMM-based method with both RWCCN and s-norm has achieved EER less than 0.3 % for both genders (Table VI in [11]). The conventional Relevance MAP adaptation with HMM alignment without applying any channel-compensation techniques (i.e. without applying RWCCN and s-norm due to the lack of suitable training data) on RedDots Part1 for the male has achieved EER around 1.5 % (Table XI in [11]). It is worth noting that EERs for DeepMine database without any channel-compensation techniques are 2.1 and 3.7 % for males and females respectively.

One interesting advantage of the DeepMine database compared to both RSR2015 and RedDots is having several target speakers with more than one mobile device. This allows us to analyse the effects of channel compensation methods. The second row in Table 9 corresponds to the most difficult trials where the target trials come from mobile devices with different models while imposter trials come from the same device models. It is clear that severe degradation was caused by this kind of channel effects (i.e. decreasing within-speaker similarities while increasing between-speaker similarities), especially for females.

The results in the third row show the condition when target

Table 9. i-vector/HMM results on text-dependent part (i.e. Part1) on 100-spk, 3-sess setup of Persian language. The results are only for correct trials (i.e. Target-Correct vs Imposter-Corrects). The first column shows filtering of the trials in target:imposter format. The first Y letter on each side of the colon shows the condition where the same mobile brand and the second Y shows exactly the same device models were used for recording the enrollment and test utterances.

Filter	Male					Female				
	# Target	# Imposter	EER [%]	$NDCF_{0.01}^{\min}$	$NDCF_{0.001}^{\min}$	# Target	# Imposter	EER [%]	$NDCF_{0.01}^{\min}$	$NDCF_{0.001}^{\min}$
-- : --	408951	30742420	1.54	0.189	0.633	666946	59953740	3.01	0.163	0.318
NN : YY	55950	233560	4.61	0.297	0.432	71459	434464	15.54	0.562	0.701
YY : --	346799	30742420	1.09	0.159	0.568	560827	59953740	1.28	0.107	0.247
YY : YY	346799	233560	1.48	0.098	0.153	560827	434464	1.79	0.138	0.217
YY : YN	346799	8216620	1.13	0.084	0.142	560827	20902608	1.33	0.097	0.158
Y - : --	353001	30742420	1.11	0.163	0.578	595487	59953740	1.49	0.123	0.271
- N : Y -	62152	8450180	3.49	0.258	0.408	106119	21337072	10.45	0.439	0.590

speakers at the test time use exactly the same device that was used for enrollment. Comparing this row with the results in the first row proves how much improvement can be achieved when exactly the same device is used by the target speaker.

The results in the fourth row show the condition when imposter speakers also use the same device model at test time to fool the system. So, in this case, there is no device mismatch in all trials. By comparing the results with the third row, we can see how much degradation is caused if we only consider the non-target trials with the same device.

The fifth row shows similar results when the imposter speakers use device of the same brand as the target speaker but with a different model. Surprisingly, in this case, the degradation is negligible and it means that mobiles from a specific brand (manufacturer) have different recording channel properties.

The degraded female results in the sixth row as compared to the third row show the effect of using a different device model from the same brand for target trials. For males, the filters brings almost the same subsets of trials, which explains the very similar results in this case.

Looking at the first two and the last row of Table 9, one can notice the significantly worse performance obtained for the female trials as compared to males. Note that these three rows include target trials where the devices used for enrollment do not necessarily match the devices used for recording test utterances. On the other hand, in rows 3 to 6, which exclude such mismatched trials, the performance for males and females is comparable. This suggest that the degraded results for females are caused by some problematic trials with device mismatch. The exact reason for this degradation is so far unclear and needs a further investigation.

In the last row of the table, the condition of the second row is relaxed: the target device should have different model possibly from the same brand and imposter device only needs to be from the same brand. In this case, as was expected, the performance degradation is smaller than in the second row.

4.2. Speech Recognition Experiments

In addition to speaker verification, we present several speech recognition experiments on Part3. The experiments were performed with the Kaldi toolkit [14]. For training HMM-based MonoPhone model, only 20 thousands of shortest utterances are used and for other models the whole training data is used. The DNN based acoustic model

Table 10. Factorized TDNN architecture for acoustic modeling using Kaldi toolkit [14]. Note that Batch-Norm applied after each ReLU is omitted for brevity.

Layer	Layer type	Context factor1	Context factor2	Skip conn. from layer	Size	Inner size
1	TDNN-ReLU	t			1536	
2	F-TDNN-ReLU	t-1, t	t, t+1	0 (input)	1536	256
3	F-TDNN-ReLU	t-1, t	t, t+1	1	1536	256
4	F-TDNN-ReLU	t-1, t	t, t+1	2	1536	256
5	F-TDNN-ReLU	t	t	3	1536	256
6	F-TDNN-ReLU	t-3, t	t, t+3	4	1536	256
7	F-TDNN-ReLU	t-3, t	t, t+3	5	1536	256
8	F-TDNN-ReLU	t-3, t	t, t+3	6	1536	256
9	F-TDNN-ReLU	t-3, t	t, t+3	7	1536	256
10	F-TDNN-ReLU	t-3, t	t, t+3	8	1536	256
11	F-TDNN-ReLU	t-3, t	t, t+3	9	1536	256
12	F-TDNN-ReLU	t-3, t	t, t+3	10	1536	256
13	F-TDNN-ReLU	t-3, t	t, t+3	11	1536	256
14	F-TDNN-ReLU	t-3, t	t, t+3	12	1536	256
15	F-TDNN-ReLU	t-3, t	t, t+3	13	1536	256
16	F-TDNN-ReLU	t-3, t	t, t+3	14	1536	256
17	F-TDNN-ReLU	t-3, t	t, t+3	15	1536	256
18	Linear				256	
19	Dense-ReLU-Linear				256	1536
20	Dense				# Tar.	

is a time-delay DNN with low-rank factorized layers and skip connections without i-vector adaptation (a modified network from one of the best performing LibriSpeech recipes). The network is shown in Table 10: there are 16 F-TDNN layers, with dimension 1536 and linear bottleneck layers of dimension 256. The acoustic model is trained for 10 epochs using lattice-free maximum mutual information (LF-MMI) with cross-entropy regularization [15]. Re-scoring is done using a pruned trigram language model and the size of the dictionary is around 90,000 words.

Table 11 shows the results in terms of word error rate (WER) for different evaluated methods. As can be seen, the created database can be used to train well performing and practically usable Persian ASR models.

Table 11. WER [%] obtained with different ASR systems trained on DeepMine database. Test-set has 5.9 hours of speech and Large-Test-set contains 28.5 hours of speech.

	Test-set	Large-Test-set
MonoPhone	41.31	40.70
TriPhone + Deltas + Delta-Deltas	17.30	16.50
TriPhone + LDA + MLLT	14.08	13.32
TriPhone + LDA + MLLT + SAT	11.52	10.95
Kaldi F-TDNN	4.44	4.09

5. CONCLUSIONS

In this paper, we have described the final version of a large speech corpus, the DeepMine database. It has been collected using crowd-sourcing and, according to the best of our knowledge, it is the largest public text-dependent and text-prompted speaker verification database in two languages: Persian and English. In addition, it is the largest text-independent speaker verification evaluation database, making it suitable to robustly evaluate state-of-the-art methods on different conditions. Alongside these appealing properties, it comes with phone-level transcription, making it suitable to train deep neural network models for Persian speech recognition.

We provided several evaluation protocols for each part of the database. The protocols allow researchers to investigate the performance of different methods in various scenarios and study the effects of channels, duration and phrase text on the performance. We also provide two test sets for speech recognition: One normal test set with a few minutes of speech for each speaker and one large test set with more (30 minutes on average) speech that can be used for any speaker adaptation method.

As baseline results, we reported the performance of an i-vector/HMM based method on Persian text-dependent part. Moreover, we conducted speech recognition experiments using conventional HMM-based methods, as well as state-of-the-art deep neural network based method using Kaldi toolkit with promising performance. Text-dependent results have shown that the DeepMine database is more challenging than RSR2015 and RedDots databases.

6. ACKNOWLEDGMENTS

The data collection project was mainly supported by Sharif DeepMine company. The work on the paper was supported by Czech National Science Foundation (GACR) project "NEUREM3" No. 19-26934X and the National Programme of Sustainability (NPU II) project "IT4Innovations excellence in science - LQ1602".

7. REFERENCES

- [1] Arsha Nagrani, Joon Son Chung, and Andrew Senior, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv:1706.08612*, 2017.
- [2] Georg Heigold, Ignacio Moreno, Samy Bengio, and Noam Shazeer, "End-to-End text-dependent speaker verification," in *Acoustics, Speech and Signal Processing (ICASSP), IEEE International Conference on*, 2016, pp. 5115–5119.
- [3] Anthony Larcher, Kong Aik Lee, Bin Ma, and Haizhou Li, "Text-dependent speaker verification: Classifiers, databases and RSR2015," *Speech Communication*, vol. 60, pp. 56–77, 2014.
- [4] Kong Aik Lee, Anthony Larcher, Guangsen Wang, Patrick Kenny, Niko Brümmer, David van Leeuwen, Hagai Aronowitz, Marcel Kockmann, Carlos Vaquero, Bin Ma, et al., "The RedDots data collection for speaker recognition," in *InterSpeech*, 2015, pp. 2996–3000.
- [5] Hossein Zeinali, Hossein Sameti, and Themis Stafylakis, "DeepMine speech processing database: Text-dependent and independent speaker verification and speech recognition in Persian and English," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 386–392.
- [6] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *Acoustics, Speech and Signal Processing (ICASSP), IEEE International Conference on*, 2015, pp. 5206–5210.
- [7] Mahmood Bijankhan, Javad Sheikhzadegan, and Mahmood R. Roohani, "FarsDat - the speech database of Farsi spoken language," in *5th Australian international conference, Speech science and technology*, 1994, pp. 826–831.
- [8] Mahmood Bijankhan, Javad Sheikhzadegan, Mahmood R. Roohani, Rahman Zarrintare, Seyyed Z. Ghasemi, and Mohammad E. Ghasedi, "TFarsDat - the telephone Farsi speech database," in *EuroSpeech*, 2003, pp. 1525–1528.
- [9] Hossein Zeinali, Hossein Sameti, Lukáš Burget, and Jan Černocký, "Text-dependent speaker verification based on i-vectors, deep neural networks and hidden Markov models," *Computer Speech & Language*, vol. 46, pp. 53–71, 2017.
- [10] Hossein Zeinali, Lukas Burget, Hossein Sameti, Ondrej Glembek, and Oldrich Plchot, "Deep neural networks and hidden Markov models in i-vector-based text-dependent speaker verification," in *Odyssey-The Speaker and Language Recognition Workshop*, 2016, pp. 24–30.
- [11] Hossein Zeinali, Hossein Sameti, and Lukas Burget, "HMM-based phrase-independent i-vector extractor for text-dependent speaker verification," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 7, pp. 1421–1435, 2017.
- [12] Steve Young, Gunnar Evermann, Mark Gales, Thomas Hain, Dan Kershaw, Xunying Liu, Gareth Moore, Julian Odell, Dave Ollason, Dan Povey, et al., *The HTK book*, vol. 2, Entropic Cambridge Research Laboratory Cambridge, 1997.
- [13] Hossein Zeinali, Hossein Sameti, Lukas Burget, Jan Černocký, Nooshin Maghsoodi, and Pavel Matejka, "i-vector/HMM based text-dependent speaker verification system for RedDots challenge," in *InterSpeech*, 2016, pp. 440–444.
- [14] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., "The Kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society, 2011.
- [15] Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahremani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MML," in *Interspeech*, 2016, pp. 2751–2755.