

CN-CELEB: A CHALLENGING CHINESE SPEAKER RECOGNITION DATASET

Y. Fan¹, J.W. Kang², L.T. Li², K.C. Li², H.L. Chen², S.T. Cheng², P.Y. Zhang², Z.Y. Zhou², Y.Q. Cai⁴, D. Wang³

¹Key Laboratory of Transient Physics, Nanjing University of Science and Technology, China

²Center for Speech and Language Technologies, Tsinghua University, China

³Beijing National Research Center for Information Science and Technology, China

⁴Department of Computer Science and Technology, Tsinghua University, China

wangdong99@mails.tsinghua.edu.cn

ABSTRACT

Recently, researchers set an ambitious goal of conducting speaker recognition in unconstrained conditions where the variations on ambient, channel and emotion could be arbitrary. However, most publicly available datasets are collected under constrained environments, i.e., with little noise and limited channel variation. These datasets tend to deliver over-optimistic performance and do not meet the request of research on speaker recognition in unconstrained conditions.

In this paper, we present *CN-Celeb*, a large-scale speaker recognition dataset collected ‘in the wild’. This dataset contains more than 130,000 utterances from 1,000 Chinese celebrities, and covers 11 different genres in real world. Experiments conducted with two state-of-the-art speaker recognition approaches (i-vector and x-vector) show that the performance on *CN-Celeb* is far inferior to the one obtained on *VoxCeleb*, a widely used speaker recognition dataset. This result demonstrates that in real-life conditions, the performance of existing techniques might be much worse than it was thought. Our database is free for researchers and can be downloaded from <http://project.csl.org>.

Index Terms— speaker recognition, Chinese, dataset

1. INTRODUCTION

Speaker recognition including identification and verification, aims to recognize claimed identities of speakers. After decades of research, performance of speaker recognition systems has been vastly improved, and the technique has been deployed to a wide range of practical applications. Nevertheless, the present speaker recognition approaches are still far from reliable in unconstrained conditions where uncertainties within the speech recordings could be arbitrary. These uncertainties might be caused by multiple factors, including free text, multiple channels, environmental noises, speaking styles, and physiological status. These uncertainties make the speaker recognition task highly challenging [1, 2].

This work was supported by the National Natural Science Foundation of China No.61633013 and the Postdoctoral Science Foundation of China No.2018M640133. Dong Wang is the corresponding author.

Researchers have devoted much effort to address the difficulties in unconstrained conditions. Early methods are based on probabilistic models that treat these uncertainties as an additive Gaussian noise. JFA [3, 4] and PLDA [5] are the most famous among such models. These models, however, are shallow and linear, and therefore cannot deal with the complexity of real-life applications. Recent advance in deep learning methods offers a new opportunity [6, 7, 8, 9]. Resorting to the power of deep neural networks (DNNs) in representation learning, these methods can remove unwanted uncertainties by propagating speech signals through the DNN layer by layer and retain speaker-relevant features only [10]. Significant improvement in robustness has been achieved by the DNN-based approach [11], which makes it more suitable for applications in unconstrained conditions.

The success of DNN-based methods, however, largely relies on a large amount of data, in particular data that involve the true complexity in unconstrained conditions. Unfortunately, most existing datasets for speaker recognition are collected in constrained conditions, where the acoustic environment, channel and speaking style do not change significantly for each speaker [12, 13, 14]. These datasets tend to deliver over optimistic performance and do not meet the request of research on speaker recognition in unconstrained conditions.

To address this shortage in datasets, researchers have started to collect data ‘in the wild’. The most successful ‘wild’ dataset may be *VoxCeleb* [15, 16], which contains millions of utterances from over thousands of speakers. The utterances were collected from open-source media using a fully automated pipeline based on computer vision techniques, in particular face detection, tracking and recognition, plus video-audio synchronization. The automated pipeline is almost costless, and thus greatly improves the efficiency of data collection.

In this paper, we re-implement the automated pipeline of *VoxCeleb* and collect a new large-scale speaker dataset, named *CN-Celeb*. Compared with *VoxCeleb*, *CN-Celeb* has three distinct features:

- *CN-Celeb* specially focuses on Chinese celebrities, and contains more than 130,000 utterances from 1,000 persons.

- *CN-Celeb* covers more genres of speech. We intentionally collected data from 11 genres, including entertainment, interview, singing, play, movie, vlog, live broadcast, speech, drama, recitation and advertisement. The speech of a particular speaker may be in more than 5 genres. As a comparison, most of the utterances in *VoxCeleb* were extracted from interview videos. The diversity in genres makes our database more representative for the true scenarios in unconstrained conditions, but also more challenging.
- *CN-Celeb* is not fully automated, but involves human check. We found that more complex the genre is, more errors the automated pipeline tends to produce. Ironically, the error-pron segments could be highly valuable as they tend to be boundary samples. We therefore choose a two-stage strategy that employs the automated pipeline to perform pre-selection, and then perform human check.

The rest of the paper is organized as follows. Section 2 presents a detailed description for *CN-Celeb*, and Section 3 presents more quantitative comparisons between *CN-Celeb* and *VoxCeleb* on the speaker recognition task. Section 4 concludes the entire paper.

2. THE CN-CELEB DATASET

2.1. Data description

The original purpose of the *CN-Celeb* dataset is to investigate the *true* difficulties of speaker recognition techniques in unconstrained conditions, and provide a resource for researchers to build prototype systems and evaluate the performance. Ideally, it can be used as a standalone data source, and can be also used with other datasets together, in particular *VoxCeleb* which is free and large. For this reason, *CN-Celeb* tries to be distinguished from but also complementary to *VoxCeleb* from the beginning of the design. This leads to three features that we have discussed in the previous section: Chinese focused, complex genres, and quality guarantee by human check.

In summary, *CN-Celeb* contains over 130,000 utterances from 1,000 Chinese celebrities. It covers 11 genres and the total amount of speech waveforms is 274 hours. Table 1 gives the data distribution over the genres, and Table 2 presents the data distribution over the length of utterances.

2.2. Challenges with CN-Celeb

Table 3 summarizes the main difference between *CN-Celeb* and *VoxCeleb*. Compared to *VoxCeleb*, *CN-Celeb* is a more complex dataset and more challenging for speaker recognition research. More details of these challenges are as follows.

- Most of the utterances involve real-world noise, including ambient noise, background babbling, music, cheers and laugh.

Table 1. The distribution over genres.

Genre	# of Spks	# of Utters	# of Hours
Entertainment	483	22,064	33.67
Interview	780	59,317	135.77
Singing	318	12,551	28.83
Play	69	4,245	4.95
Movie	62	2,749	2.20
Vlog	41	1,894	4.15
Live Broadcast	129	8,747	16.35
Speech	122	8,401	36.22
Drama	160	7,274	6.43
Recitation	41	2,747	4.98
Advertisement	17	120	0.18
Overall	1,000	130,109	273.73

Table 2. The distribution over utterance length.

Length (s)	# of Utterances	Proportion
<2	41,658	32.0%
2-5	38,629	30.0%
5-10	23,497	18.0%
10-15	10,687	8.0%
15-20	5,334	4.0%
20-25	3,218	2.5%
25-30	1,991	1.5%
>30	5,095	4.0%

Table 3. Comparison between *CN-Celeb* and *VoxCeleb*.

	<i>CN-Celeb</i>	<i>VoxCeleb</i>
Source media	bilibili.com	youtube.com
Language	Chinese	Mostly English
Genre	11	Mostly interview
# of Spks	1,000	7,363
# of Utters	130,109	1281,762
# of Hours	274	2,794
Human Check	Yes	No

- A certain amount of utterances involve strong and overlapped background speakers, especially in the drama and movie genres.
- Most of speakers have different genres of utterances, which results in significant variation in speaking styles.
- The utterances of the same speaker may be recorded at different time and with different devices, leading to serious cross-time and cross-channel problems.
- Most of the utterances are short, which meets the scenarios of most real applications but leads to unreliable decision.

2.3. Collection pipeline

CN-Celeb was collected following a two-stage strategy: firstly we used an automated pipeline to extract potential segments of the Person of Interest (POI), and then applied a human check to remove incorrect segments. This process is much faster than purely human-based segmentation, and reduces errors caused by a purely automated process.

Briefly, the automated pipeline we used is similar to the one used to collect *VoxCeleb1* [15] and *VoxCeleb2* [16], though we made some modification to increase efficiency and precision. Especially, we introduced a new face-speaker double check step that fused the information from both the image and speech signals to increase the recall rate while maintaining the precision.

The detailed steps of the collection process are summarized as follows.

- **STEP 1. POI list design.** We manually selected 1,000 Chinese celebrities as our target speakers. These speakers were mostly from the entertainment sector, such as singers, drama actors/actresses, news reporters, interviewers. Region diversity was also taken into account so that variation in accent was covered.
- **STEP 2. Pictures and videos download.** Pictures and videos of the 1,000 POIs were downloaded from the data source (<https://www.bilibili.com/>) by searching for the names of the persons. In order to specify that we were searching for POI names, the word ‘human’ was added in the search queries. The downloaded videos were manually examined and were categorized into the 11 genres.
- **STEP 3. Face detection and tracking.** For each POI, we first obtained the portrait of the person. This was achieved by detecting and clipping the face images from all pictures of that person. The RetinaFace algorithm was used to perform the detection and clipping [17]. Afterwards, video segments that contain the target person were extracted. This was achieved by three steps: (1) For each frame, detect all the faces appearing in the frame using RetinaFace; (2) Determine if the target person appears by comparing the POI portrait and the faces detected in the frame. We used the ArcFace face recognition system [18] to perform the comparison; (3) Apply the MOSSE face tracking system [19] to produce face streams.
- **STEP 4. Active speaker verification.** As in [15], an active speaker verification system was employed to verify if the speech was really spoken by the target person. This is necessary as it is possible that the target person appears in the video but the speech is from other persons. We used the SyncNet model [20] as in [15] to perform the task. This model was trained to detect if a stream of mouth movement and a stream of speech are synchronized. In our implementation, the stream of mouth movement was derived from the face stream produced by the MOSSE system.
- **STEP 5. Double check by speaker recognition.** Although SyncNet worked well for videos in simple genres, it failed for videos of complex genres such as movie and vlog. A possible reason is that the video content of these genres may change dramatically in time, which leads to unreliable estimation for the stream of the mouth movement, hence unreliable synchronization detection. In order to improve the robustness of the active speaker verification in complex genres, we introduced a double check procedure based on speaker recognition. The idea is simple: whenever the speaker recognition system states a *very low* confidence for the target speaker, the segment will be discarded even if the confidence from SyncNet is high; vice versa, if the speaker recognition system states a *very high* confidence, the segment will be retained. We used an off-the-shelf speaker recognition system [21] to perform this double check. In our study, this double check improved the recall rate by 30% absolutely.
- **STEP 6. Human check.** The segments produced by the above automated pipeline were finally checked by human. According to our experience, this human check is rather efficient: one could check 1 hour of speech in 1 hour. As a comparison, if we do not apply the automated pre-selection, checking 1 hour of speech requires 4 hours.

3. EXPERIMENTS ON SPEAKER RECOGNITION

In this section, we present a series of experiments on speaker recognition using *VoxCeleb* and *CN-Celeb*, to compare the complexity of the two datasets.

3.1. Data

VoxCeleb: The entire dataset involves two parts: *VoxCeleb1* and *VoxCeleb2*. We used *SITW* [22], a subset of *VoxCeleb1* as the evaluation set. The rest of *VoxCeleb1* was merged with *VoxCeleb2* to form the training set (simply denoted by *VoxCeleb*). The training set involves 1,236,567 utterances from 7,185 speakers, and the evaluation set involves 6,445 utterances from 299 speakers (precisely, this is the *Eval. Core* set within *SITW*).

CN-Celeb: The entire dataset was split into two parts: the first part *CN-Celeb(T)* involves 111,260 utterances from 800 speakers and was used as the training set; the second part *CN-Celeb(E)* involves 18,849 utterances from 200 speakers and was used as the evaluation set.

3.2. Settings

Two state-of-the-art baseline systems were built following the Kaldi *sitw* and *cncelb* recipe [23]: an i-vector system [4] and

an x-vector system [11].

For the i-vector system, the acoustic feature involved 24-dimensional MFCCs plus the log energy, augmented by the first- and second-order derivatives. We also applied the cepstral mean normalization (CMN) and the energy-based voice active detection (VAD). The universal background model (UBM) consisted of 2,048 Gaussian components, and the dimensionality of the i-vector space was 400. LDA was applied to reduce the dimensionality of the i-vectors to 150. The PLDA model was used for scoring [5].

For the x-vector system, the feature-learning component was a 5-layer time-delay neural network (TDNN). The slicing parameters for the five time-delay layers were: $\{t-2, t-1, t, t+1, t+2\}$, $\{t-2, t, t+2\}$, $\{t-3, t, t+3\}$, $\{t\}$, $\{t\}$. The statistic pooling layer computed the mean and standard deviation of the frame-level features from a speech segment. The size of the output layer was consistent with the number of speakers in the training set. Once trained, the activations of the penultimate hidden layer were read out as x-vectors. In our experiments, the dimension of the x-vectors trained on *VoxCeleb* was set to 512, while for *CN-Celeb*, it was set to 256, considering the less number of speakers in the training set. Afterwards, the x-vectors were projected to 150-dimensional vectors by LDA, and finally the PLDA model was employed to score the trials. Refer to [11] for more details.

3.3. Basic results

We first present the basic results evaluated on *SITW* and *CN-Celeb(E)*. Both the front-end (i-vector or x-vector models) and back-end (LDA-PLDA) models were trained with the *VoxCeleb* training set. Note that for *SITW*, the averaged length of the utterances is more than 80 seconds, while this number is about 8 seconds for *CN-Celeb(E)*. For a better comparison, we resegmented the data of *SITW* and created a new dataset denoted by *SITW(S)*, where the averaged lengths of the enrollment and test utterances are 28 and 8 seconds, respectively. These numbers are similar to the statistics of *CN-Celeb(E)*.

The results in terms of the equal error rate (EER) are reported in Table 4. It can be observed that for both the i-vector system and the x-vector system, the performance on *CN-Celeb(E)* is much worse than the performance on *SITW* and *SITW(S)*. This indicates that there is big difference between these two datasets. From another perspective, it demonstrates that the model trained with *VoxCeleb* does not generalize well, although it has achieved reasonable performance on data from a similar source (*SITW*).

3.4. Further comparison

To further compare *CN-Celeb* and *VoxCeleb* in a quantitative way, we built systems based on *CN-Celeb* and *VoxCeleb*, respectively. For a fair comparison, we randomly sampled 800 speakers from *VoxCeleb* and built a new dataset *VoxCeleb(L)*

Table 4. EER(%) results of the i-vector and x-vector systems trained on *VoxCeleb* and evaluated on three evaluation sets.

System	Training Set		Evaluation Set		
	Front-end	Back-end	SITW	SITW(S)	CN-Celeb(E)
i-vector	VoxCeleb	VoxCeleb	5.30	7.30	19.05
x-vector	VoxCeleb	VoxCeleb	3.75	4.78	15.52

whose size is comparable to *CN-Celeb(T)*. This data set was used for back-end (LDA-PLDA) training.

The experimental results are shown in Table 5. Note that the performance of all the comparative experiments show the same trend with the i-vector system and the x-vector system, we therefore only analyze the i-vector results.

Firstly, it can be seen that the system trained purely on *VoxCeleb* obtained good performance on *SITW(S)* (1st row). This is understandable as *VoxCeleb* and *SITW(S)* were collected from the same source. For the pure *CN-Celeb* system (2nd row), although *CN-Celeb(T)* and *CN-Celeb(E)* are from the same source, the performance is still poor (14.24%). More importantly, with re-training the back-end model with *VoxCeleb(L)* (4th row), the performance on *SITW* becomes better than the same-source result on *CN-Celeb(E)* (11.34% vs 14.24%). All these results reconfirmed the significant difference between the two datasets, and indicates that *CN-Celeb* is more challenging than *VoxCeleb*.

Table 5. EER(%) results with different data settings.

System	Training Set		Evaluation Set	
	Front-end	Back-end	SITW(S)	CN-Celeb(E)
i-vector	VoxCeleb	VoxCeleb(L)	8.34	17.43
	CN-Celeb(T)	CN-Celeb(T)	14.87	14.24
	VoxCeleb	CN-Celeb(T)	12.96	15.00
	CN-Celeb(T)	VoxCeleb(L)	11.34	15.50
x-vector	VoxCeleb	VoxCeleb(L)	5.93	13.64
	CN-Celeb(T)	CN-Celeb(T)	15.23	14.78
	VoxCeleb	CN-Celeb(T)	10.72	11.99
	CN-Celeb(T)	VoxCeleb(L)	12.68	15.62

4. CONCLUSIONS

We introduced a free dataset *CN-Celeb* for speaker recognition research. The dataset contains more than 130k utterances from 1,000 Chinese celebrities, and covers 11 different genres in real world. We compared *CN-Celeb* and *VoxCeleb*, a widely used dataset in speaker recognition, by setting up a series of experiments based on two SOTA speaker recognition models. Experimental results demonstrated that *CN-Celeb* is significantly different from *VoxCeleb*, and it is more challenging for speaker recognition research. The EER performance obtained in this paper suggests that in unconstrained conditions, the performance of the current speaker recognition techniques might be much worse than it was thought.

5. REFERENCES

- [1] Lara Lynn Stoll, *Finding difficult speakers in automatic speaker recognition*, Ph.D. thesis, UC Berkeley, 2011.
- [2] Thomas Fang Zheng and Lantian Li, *Robustness-related issues in speaker recognition*, Springer, 2017.
- [3] Patrick Kenny, Gilles Boulianne, Pierre Ouellet, and Pierre Dumouchel, “Joint factor analysis versus eigenchannels in speaker recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 4, pp. 1435–1447, 2007.
- [4] Najim Dehak, Patrick J Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet, “Front-end factor analysis for speaker verification,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [5] Sergey Ioffe, “Probabilistic linear discriminant analysis,” in *European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
- [6] Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Brian Kingsbury, et al., “Deep neural networks for acoustic modeling in speech recognition,” *IEEE Signal processing magazine*, vol. 29, 2012.
- [7] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [8] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [10] Lantian Li, Yixiang Chen, Ying Shi, Zhiyuan Tang, and Dong Wang, “Deep speaker feature learning for text-independent speaker verification,” *arXiv preprint arXiv:1705.03670*, 2017.
- [11] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust DNN embeddings for speaker recognition,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018.
- [12] William M Fisher, “The DARPA speech recognition research database: specifications and status,” in *Proc. DARPA Workshop on Speech Recognition, Feb. 1986*, 1986, pp. 93–99.
- [13] John J Godfrey, Edward C Holliman, and Jane McDaniel, “SWITCHBOARD: Telephone speech corpus for research and development,” in *1992 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1992, vol. 1, pp. 517–520.
- [14] Seyed Omid Sadjadi, Timothée Kheyrkhah, Audrey Tong, Craig S Greenberg, Douglas A Reynolds, Elliot Singer, Lisa P Mason, and Jaime Hernandez-Cordero, “The 2016 NIST speaker recognition evaluation,” in *Interspeech*, 2017, pp. 1353–1357.
- [15] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Voxceleb: a large-scale speaker identification dataset,” *arXiv preprint arXiv:1706.08612*, 2017.
- [16] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman, “Voxceleb2: Deep speaker recognition,” *arXiv preprint arXiv:1806.05622*, 2018.
- [17] Jiankang Deng, Jia Guo, Yuxiang Zhou, Jinke Yu, Irene Kotsia, and Stefanos Zafeiriou, “Retinaface: Single-stage dense face localisation in the wild,” *arXiv preprint arXiv:1905.00641*, 2019.
- [18] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [19] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui, “Visual object tracking using adaptive correlation filters,” in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 2010, pp. 2544–2550.
- [20] Joon Son Chung and Andrew Zisserman, “Out of time: Automated lip sync in the wild,” in *Asian conference on computer vision*. Springer, 2016, pp. 251–263.
- [21] Weidi Xie, Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Utterance-level aggregation for speaker recognition in the wild,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5791–5795.
- [22] Mitchell McLaren, Luciana Ferrer, Diego Castan, and Aaron Lawson, “The Speakers in the Wild (SITW) speaker recognition database,” in *Interspeech*, 2016, pp. 818–822.
- [23] 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 *Workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society, 2011, number EPFL-CONF-192584.