

NIST SRE CTS Superset: A large-scale dataset for telephony speaker recognition

Omid Sadjadi

August 2021

1 Introduction

This document provides a brief description of the National Institute of Standards and Technology (NIST) speaker recognition evaluation (SRE) [1] conversational telephone speech (CTS) Superset. The CTS Superset has been created in an attempt to provide the research community with a large-scale dataset along with uniform metadata that can be used to effectively train and develop telephony (narrowband) speaker recognition systems. It contains a large number of telephony speech segments from more than 6800 speakers with speech durations distributed uniformly in the [10s, 60s] range. The segments have been extracted from the source corpora used to compile prior SRE datasets (SRE1996-2012), including the Greybeard corpus as well as the Switchboard and Mixer series collected by the Linguistic Data Consortium (LDC¹). In addition to the brief description, we also report speaker recognition results on the NIST 2020 CTS Speaker Recognition Challenge, obtained using a system trained with the CTS Superset. The results will serve as a reference baseline for the challenge.

2 CTS Superset (LDC2021E08)

The NIST SRE CTS Superset is the largest most comprehensive dataset available to date for telephony speaker recognition. It has been extracted from the source corpora (see Table 1) used to compile prior SRE datasets (SRE1996-2012). Table 2 summarizes the data statistics for the CTS Superset. There are a total of 605,760 segments originating from 6867 speakers (2885 male and 3992 female). Each segment contains approximately 10 to 60 seconds of speech ², and each speaker has at least three sessions/calls (hence at least 3 segments). Note that some speakers appear in more than one source corpus, therefore the sum of the speakers in the table is greater than 6867. Although the vast majority of the segments in the CTS Superset are spoken in English (including both native and accented English), there are more than 50 languages represented in this dataset.

¹<https://www ldc upenn edu/>

²As determined using a speech activity detector (SAD).

Table 1: Original source corpora used to create the NIST SRE CTS Superset

Source corpus	LDC Catalog ID	corpusid
Switchboard1 release2	LDC97S62 [2]	swb1r2
Switchboard2 Phase I	LDC98S75 [3]	swb2p1
Switchboard2 Phase II	LDC99S79 [4]	swb2p2
Switchboard2 Phase III	LDC2002S06 [5]	swb2p3
Switchboard Cellular Part 1	LDC2001S13 [6]	swbcellp1
Switchboard Cellular Part2	LDC2004S07 [7]	swbcellp2
Mixer3	LDC2021R03 [8]	mx3
Mixer4–5	LDC2020S03 [9]	mx45
Mixer6	LDC 2013S03 [10]	mx6
Greybeard	LDC2013S05 [11]	gb1

Table 2: Data statistics for the NIST SRE CTS Superset

corpusid	#segments	#speakers	#sessions
swb1r2	26,282	442	4757
swb2p1	33,746	566	7134
swb2p2	41,982	649	8895
swb2p3	22,865	548	5187
swbcellp1	13,496	216	2560
swbcellp2	20985	378	3966
mx3	317,950	3033	37,759
mx45	40,313	486	4997
mx6	70,174	526	8727
gb1	17,967	167	2188

The procedure for extracting segments from the original sessions/calls is as follows; given a session of arbitrary duration (typically 5–12 minutes) and speech time marks generated using a speech activity detector (SAD), we start extracting non-overlapping segments by repeatedly sampling a speech duration from the uniform distribution [10, 60], until we exhaust the duration of that session.

The LDC2021E08 package contains the following items:

- Audio segments from 6867 subjects located in the `data/{subjectids}/` directories
- Associated metadata located in the `docs/` directory

The metadata file contains information about the audio segments and includes the following fields:

- `filename` (segment filename including the relative path)
- `segmentid` (segment identifier)
- `subjectid` (LDC speaker id)
- `speakerid` (zero indexed numerical speaker id)

- `speech_duration` (segment speech duration)
- `sessionid` (segment session/call identifier)
- `corpusid` (corpus identifier as defined in Table 1)
- `phoneid` (anonymized phone number)
- `gender` (male or female)
- `language` (language spoken in the segment)

For future releases of the CTS Superset we plan to extend the source corpora to include Mixer 1, 2, and 7 as well.

3 Speaker Recognition System

In this section, we describe the baseline speaker recognition system setup including speech and non-speech data used for training the system components as well as the hyper-parameter configurations used. Figure 1 shows a block diagram of the x-vector baseline system. The embedding extractor is trained using Pytorch³, while the NIST SLRE toolkit is used for front-end processing and back-end scoring.




Figure 1: Block diagram of the baseline system.

3.1 Data

The baseline system is developed using the CTS Superset described in the previous section. In order to increase the diversity of the acoustic conditions in the training set, two different data augmentation strategies are adopted. The first strategy uses noise-degraded (using babble, general noise, and music) versions of the original recordings, while the second strategy uses spectro-temporal masking applied directly on spectrograms (aka spectrogram augmentation [12]). The noise samples for the first augmentation approach are extracted from the MUSAN corpus [13]. For spectrogram augmentation, the mild and strong policies described in [12] are used.

³<https://github.com/pytorch/pytorch>

3.2 Configuration

For speech parameterization, we extract 64-dimensional log-mel spectrograms from 25 ms frames every 10 ms using a 64-channel mel-scale filterbank spanning the frequency range 80 Hz–3800 Hz. After dropping the non-speech frames using SAD, a short-time cepstral mean subtraction is applied over a 3-second sliding window.

For embedding extraction, an extended TDNN [14] with 11 hidden layers and parametric rectified linear unit (PReLU) non-linearities is trained to discriminate among the nearly 6800 speakers in the CTS Superset set. A cosine loss with additive margin [15] is used in the output layer (with $m = 0.2$ and $s = 40$). The first 9 hidden layers operate at frame-level, while the last 2 operate at segment-level. There is a 3000-dimensional statistics pooling layer between the frame-level and segment-level layers that accumulates all frame-level outputs from the 9th layer and computes the mean and standard deviation over all frames for an input segment. The model is trained using Pytorch and the stochastic gradient descent (SGD) optimizer with momentum (0.9), an initial learning rate of 10^{-1} , and a batch size of 512. The learning rate remains constant for the first 5 epochs, after which it is halved every other epoch.

To train the network, a speaker-balanced sampling strategy is implemented where in each batch 512 unique speakers are selected, without replacement, from the pool of training speakers. Then, for each speaker, a random speech segment is selected from which a 400-frame (corresponding to 4 seconds) chunk is extracted for training. This process is repeated until the training samples are exhausted.

After training, embeddings are extracted from the 512-dimensional affine component of the 10th layer (i.e., the first segment-level layer). Prior to dimensionality reduction through linear discriminant analysis (LDA) to 250, 512-dimensional embeddings are centered, whitened, and unit-length normalized. The centering and whitening statistics are computed using the CTS Superset data. For backend scoring, a Gaussian probabilistic LDA (PLDA) model with a full-rank Eigenvoice subspace is trained using the embeddings extracted from only the original (as opposed to degraded) speech segments in the CTS Superset. No parameter/domain adaptation or score normalization/calibration is applied.

4 Results

In this section, we present the experimental results on the NIST 2020 CTS Challenge progress and test sets obtained using the baseline system. Results are reported in terms of the equal error rate (EER) and the minimum cost (denoted as \min_C) and defined in the CTS Challenge evaluation plan [16].

Table 3 summarizes the baseline results on the 2020 CTS Challenge progress and test sets. Note that no calibration is applied to the baseline system output. It is also worth emphasizing here that one could potentially exploit publicly available and/or proprietary data such as the VoxCeleb to further improve the performance; nevertheless, this is beyond the scope of the baseline system, and therefore not considered in this report.

Table 3: NIST baseline system performance using on the 2020 CTS Challenge progress and test sets.

System	Approach	Training Data	Backend	Set	EER [%]	min_C
NIST baseline	x-vector	CTS Superset	Cosine	Progress	4.37	0.190
				Test	4.74	0.206
			PLDA	Progress	4.62	0.221
				Test	4.67	0.224

5 Acknowledgement

Experiments and analyses were performed, in part, on the NIST Enki HPC cluster.

6 Disclaimer

The NIST baseline system was developed to support speaker recognition research. Comparison of systems and results against this system and its results are not to be construed or represented as endorsements of any participant's system or commercial product, or as official findings on the part of NIST or the U.S. Government. The reader of this report acknowledges that changes in the data domain and system configurations, or changes in the amount of data used to build a system, can greatly influence system performance.

Because of the above reasons, this system should not be used for commercial product testing and the results should not be used to make conclusions regarding which commercial products are best for a particular application.

References

- [1] C. S. Greenberg, L. P. Mason, S. O. Sadjadi, and D. A. Reynolds, "Two decades of speaker recognition evaluation at the national institute of standards and technology," *Computer Speech & Language*, vol. 60, p. 101032, 2020.
- [2] J. Godfrey and E. Holliman, "Switchboard-1 Release 2," <https://catalog.ldc.upenn.edu/LDC97S62>, 1993, [Online; accessed 07-August-2021].
- [3] D. Graff, A. Canavan, and G. Zipperlen, "Switchboard-2 Phase I," <https://catalog.ldc.upenn.edu/LDC98S75>, 1998, [Online; accessed 07-August-2021].
- [4] D. Graff, K. Walker, and A. Canavan, "Switchboard-2 Phase II," <https://catalog.ldc.upenn.edu/LDC99S79>, 1999, [Online; accessed 07-August-2021].
- [5] D. Graff, D. Miller, and K. Walker, "Switchboard-2 Phase III," <https://catalog.ldc.upenn.edu/LDC2002S06>, 2002, [Online; accessed 07-August-2021].
- [6] D. Graff, K. Walker, and D. Miller, "Switchboard Cellular Part 1 Audio," <https://catalog.ldc.upenn.edu/LDC2001S13>, 2001, [Online; accessed 07-August-2021].
- [7] D. Graff, K. Walker, and D. Miller, "Switchboard Cellular Part 2 Audio," <https://catalog.ldc.upenn.edu/LDC2004S07>, 2004, [Online; accessed 07-August-2021].
- [8] C. Cieri, L. Corson, D. Graff, and K. Walker, "Resources for new research directions in speaker recognition: The Mixer 3, 4 and 5 corpora," in *Proc. INTERPSEECH*, Antwerp, Belgium, August 2007.
- [9] L. Brandschain, K. Walker, D. Graff, C. Cieri, A. Neely, N. Mirghafori, B. Peskin, J. Godfrey, S. Strassel, F. Goodman, G. R. Doddington, and M. King, "Mixer 4 and 5 speech," <https://catalog.ldc.upenn.edu/LDC2020S03>, 2020, [Online; accessed 07-August-2021].
- [10] L. Brandschain, D. Graff, K. Walker, and C. Cieri, "Mixer 6 speech," <https://catalog.ldc.upenn.edu/LDC2013S03>, 2013, [Online; accessed 07-August-2021].
- [11] L. Brandschain and D. Graff, "Greybeard," <https://catalog.ldc.upenn.edu/LDC2013S05>, 2013, [Online; accessed 07-August-2021].
- [12] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," 2019, pp. 2613–2617.
- [13] D. Snyder, G. Chen, and D. Povey, "MUSAN: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [14] D. Snyder, D. Garcia-Romero, G. Sell, A. McCree, D. Povey, and S. Khudanpur, "Speaker recognition for multi-speaker conversations using x-vectors," in *Proc. IEEE ICASSP*, May 2019, pp. 5796–5800.

- [15] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "Cosface: Large margin cosine loss for deep face recognition," in *Proc. IEEE CVPR*, 2018, pp. 5265–5274.
- [16] S. O. Sadjadi, C. Greenberg, E. Singer, Douglas, and L. Mason, "NIST 2020 CTS Speaker Recognition Challenge Evaluation Plan," <https://www.nist.gov/document/nist-cts-challenge-evaluation-plan>, July 2020, [Online; accessed 22-July-2021].