



Speaker discrimination in humans and machines: Effects of speaking style variability

Amber Afshan¹, Jody Kreiman², and Abeer Alwan¹

¹Dept. of Electrical and Computer Engineering, University of California, Los Angeles, USA

²Depts. of Head and Neck Surgery and Linguistics, University of California, Los Angeles, USA

amberafshan@g.ucla.edu, jkreiman@g.ucla.edu, alwan@g.ucla.edu

Abstract

Does speaking style variation affect humans' ability to distinguish individuals from their voices? How do humans compare with automatic systems designed to discriminate between voices? In this paper, we attempt to answer these questions by comparing human and machine speaker discrimination performance for read speech versus casual conversations. Thirty listeners were asked to perform a same versus different speaker task. Their performance was compared to a state-of-the-art x-vector/PLDA-based automatic speaker verification system. Results showed that both humans and machines performed better with style-matched stimuli, and human performance was better when listeners were native speakers of American English. Native listeners performed better than machines in the style-matched conditions (EERs of 6.96% versus 14.35% for read speech, and 15.12% versus 19.87%, for conversations), but for style-mismatched conditions, there was no significant difference between native listeners and machines. In all conditions, fusing human responses with machine results showed improvements compared to each alone, suggesting that humans and machines have different approaches to speaker discrimination tasks. Differences in the approaches were further confirmed by examining results for individual speakers which showed that the perception of distinct and confused speakers differed between human listeners and machines.

Index Terms: speaker perception, speaking style, automatic speaker verification, human assisted speaker discrimination

1. Introduction

Speaking style variations are prevalent in everyday life, changing as we move from talking to a friend to reading aloud, from public speaking to talking to an infant. Regardless of these variations, humans are often able to recognize a familiar voice after hearing it for a few seconds [1]. Previous research suggests that for humans, recognizing familiar talkers entails matching a sample to stored voice templates, whereas recognizing unfamiliar talkers is a much more involved process requiring acoustic feature comparisons [2]. In this work, we are interested in understanding the effects of speaking style variations on the abilities of humans to distinguish between unfamiliar voices from short duration (~3 s), text-independent utterances. We are also interested in comparing human performance against state-of-the-art automatic systems.

A speaker rarely says an utterance twice in the exact same manner. This variability maybe intentional—for example, to hide one's identity or to communicate irony—or it may be introduced without conscious intention due to changes in emotion, social context, or physiological state [3]. These voice variations have

been extensively studied in the forensic literature [4, 5, 6, 7]. For example, style variability confuses ear witnesses hearing a criminal shouting vs reading aloud during a voice lineup [8]. Human and machine speaker discrimination performances have been compared when style changed from read to pet-directed speech, which is characterized by exaggerated prosody [9]. In both examples, differences in style were extreme, and little is known about how moderate variations in style, for example between read and conversational speech, affect the relative performance of humans vs machines in speaker discrimination performance. Evidence from voice sorting tasks indicates that humans vary their perceptual strategies when “telling people together” versus “telling people apart” [10, 11], while machines apply the same classification approach in target and non-target trials. Given that humans outperform machines in some tasks [e.g., [12]], this suggests that machines can adopt strategies from humans, and humans might do better with machine assistance in certain situations. Additionally, it has previously been observed that there is a correlation between the linguistic proficiency of the listener and their speaker identification accuracy [13, 14]. There is uncertainty, however, whether the effects of linguistic proficiency are consistent across different speaking styles.

In this study, we employed an unfamiliar *speaker discrimination* task in which the listener decides if two samples are from the same speaker or not. Using read and conversational speech stimuli, we compared the discrimination performance and reliability of humans and machines when confronted with moderate differences in speaking style. We hypothesized that such variations will have different effects on automatic speaker verification (ASV) systems than they do on human performance. We also examined the relationship between listeners' native languages and their accuracy in performing speaker discrimination tasks for different speaking styles.

The paper is organized as follows. Section 2 describes the databases used. Methods for machine and human experiments are detailed in Section 3. Results and discussion are presented in Section 4, and we conclude the paper with Section 5.

2. Databases

2.1. The UCLA speaker variability database

The UCLA Speaker Variability Database [15, 16] captures commonly occurring variations in speech from 103 female and 105 male speakers. These variations were due to phonetic content, speaking style, and affect conditions. Speech was recorded in a sound-attenuated booth at a sampling rate of 22kHz. The experiments reported in this study used a subset of recordings from 40 female talkers who were self-reported native speakers of American English (confirmed *post hoc* by two linguists). In order to avoid gender dependent cues, only female talkers were

This study was supported in part by NSF

chosen. Additionally, female talkers had discernible prosody changes between speaking styles in comparison to male talkers. Two types of speech samples were selected for each talker. The first included five phonetically-rich Harvard sentences [17], read twice in random order (clear read speech); the second comprised the talker’s side of a 2-minute telephone conversation with a family member or friend (casual conversational speech).

2.2. NIST SRE and Switchboard databases

We used the NIST SRE 04, 05, 06, 08 and 10 databases [18, 19, 20] and the Switchboard II corpus, phase 2 [21] to train the ASV system. These databases provide more than 3,000 hours of speech samples from 3,408 female and 1,832 male talkers, sampled at 8 kHz. Recordings from the UCLA database were downsampled during the ASV experiments to match this rate.

3. Method

3.1. Perceptual speaker discrimination

Perceptual tasks included trials with two different read sentences, trials with two different sentences excerpted from a conversation, and trials with one read sentence and one conversational sentence. In each case, equal numbers of “same speaker” (target) and “different speaker” (non-target) trials were included, for a total of six kinds of trials. The stimuli were all ~3 sec long, and long silences before and after the sentences were removed.

Distinct stimuli were selected for each session so that a listener never heard the same stimulus twice. In the case of read speech, there were only five sentences, so a second recording of one of the sentences was repeated at random in each session. Among the conversations, six different snippets were chosen at random. Selections were carefully made to ensure that semantic cues would not bias responses. All non-speech vocalizations (laughing, giggling, sighing) were deleted.

To minimize fatigue, each listener heard a subset of 24 talkers selected at random from the pool of 40, for a total of 144 trials/listener (6 trial types x 24 talkers). Fifteen normal-hearing subjects heard each subset. Here we report results from two groups of listeners who heard 2 different subsets of the voices. Twenty-four/30 listeners were native English speakers; 22 were female and 8 were male. They ranged in age from 17-21. The non-native English speakers consisted of 3 native speakers of Spanish, 2 of Mandarin, and 1 of Hindi.

Stimuli were randomized prior to each presentation. During the listening experiments, each pair of stimuli could be played only once in each presentation order (AB/BA). Listeners were asked to decide if the stimuli represented the same talker or two different talkers. They also reported their confidence in their response on a scale of 0 to 5 (0=wild guess and 5=very confident). Listeners were not aware of the number of talkers included in each experiment, and were encouraged to complete the experiments at their own pace, and take breaks as necessary. Testing time averaged about 45 minutes.

3.2. Automatic Speaker Discrimination

An x-vector [22]/PLDA [23] (probabilistic linear discriminant analysis) based ASV system was used, and the PLDA was adapted [24] using in-domain data with both read and conversational styles to achieve the best possible machine performance. In order to ensure a fair comparison between humans and the ASV system, trials matched the stimulus pairs presented to lis-

teners. We followed the Kaldi [25] SRE16 recipe to train the x-vector embedding extractor. Mel-frequency cepstral coefficients of 23 dimensions calculated using a frame-length of 25ms were used as features for the x-vectors. The system was trained using the SRE and Switchboard databases.

3.3. Evaluation Metric

3.3.1. Calculation of scores

Human responses were unfolded to obtain a similarity score between each stimuli pair. Confidence ratings (0 to 5) were multiplied by the decision (different=-1 and same= 1) to provide continuous scores between -5 (highly confident that the voices are different) and 5 (highly confident that the voices are same). This ensured that the similarity score reflected the decision as well as the confidence of human responses. For ASV systems, the PLDA score acted as the similarity measure. The PLDA score represents the ratio of the likelihood that a given pair of stimuli is from the same speaker to the likelihood that the pair is from two different speakers.

Calibrated log-likelihood ratios (LLR; L) were obtained from similarity scores using a calibration system based on standard logistic regression [26]. The resulting LLRs represent scalar responses by humans and machines.

3.3.2. Analysis of performance errors

System-level speaker discrimination performance was evaluated in terms of equal error rates (EER) and the log-likelihood-ratio cost function (C_{lr}) [27]. While the EER is a widely used measure, it does not measure calibration, the ability to set good decision thresholds. Hence, C_{lr} , an application-independent measure for evaluating soft decisions, was also used. It can be interpreted as a measure of loss of information; thus, the lower the C_{lr} , the more the average information per trial (in bits) increases by applying the system (humans or machines).

To perform calibration and to calculate the evaluation measures, we used the Bosaris toolkit [26]. As a result of the limited amount of data, the calibration parameters were trained and applied on the same set of scores.

3.3.3. System Fusion

In order to better understand the similarities and differences between responses from humans and machines on the speaker discrimination task, logistic regression-based [28] system fusion was performed on the log-likelihood-ratios from humans and machines using the Bosaris toolkit [26]. If humans and machine use different strategies in performing this task, system fusion should outperform either system alone.

3.3.4. Speaker-Level Analysis

This section describes speaker-level equivalents of the measures described in the previous sections. The log-likelihood-ratio L_t , as outlined in Section 3.3.1, was obtained for each trial t . L_t represents the scalar response by the system for the given trial. To compare the scores between target and non-target trials for each speaker, the L^{tar} for target and L^{non} for non-target trials were calculated separately.

L^{tar} indicates within-speaker variability: a large L^{tar} means small within-speaker variability (i.e., these target trials are easy). L^{tar} for a speaker was obtained by averaging the L_t values over the target trials that included that particular speaker. On the other hand, L^{non} represents between-speaker variability,

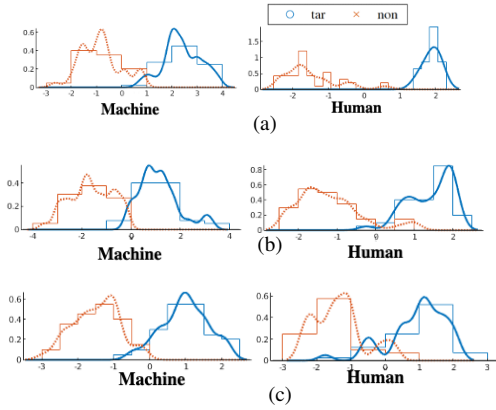


Figure 1: Histograms of L^{tar} and L^{non} per speaker for human and machine responses for trial types: (a) read-read (b) conversation-conversation. (c) read-conversation.




Figure 2: Confusion matrix on the difficulty level of speakers for machines versus humans based on the three subsets: easy, average, and hard. Numbers in parentheses represent the count of speakers in the subsets (easy, average, hard) for fused scores.

so a large L^{non} value indicates that the speaker did not differ very much from others, so that it is difficult for the system to distinguish her from others.

A speaker-level aggregation of the log-likelihood-ratio cost function (C_{llr} ; section 3.3.2) was also computed, and represents the confidence the system has when identifying a speaker. Speaker-level C_{llr} values for target trials ($C_{\text{llr}}^{\text{tar}}$) and non-target trials ($C_{\text{llr}}^{\text{non}}$) were also computed.

4. Results and Discussion

4.1. Human and Machine Performance

Table 1 presents results for the three speaking-style conditions (read speech – read speech, conversation – conversation and read speech – conversation). Results for native and non-native English listeners are shown separately in Table 1 as “native” and “non-native”. Statistical significance was analyzed using a two-sample Kolmogorov-Smirnoff test [29] and a paired sample McNemar’s test [30] as appropriate.

The style-matched, read speech – read speech condition had the best results for both machines and humans in terms of EERs (see Table 1). Although conversation – conversation trials were also style matched, both humans and machines performed less well than for read speech – read speech trials, possibly because casual conversations can vary in style depending on the context, topic, and speaker. As hypothesized, the worst performance for both humans and machines was obtained for style mismatched read speech – conversation trials. Results were

consistent with the hypothesis of our previous study [9] of read and pet-directed speech from the same set of speakers. That study showed that humans consistently performed better than machines in both read speech – read speech (EER = 19.02% versus 30.31%) and read speech – pet-directed speech (EER = 39.23% versus 44.17%) trials. Comparing those results with the present study, we can see that the performance gap between humans and machines decreased with a decreasing amount of style-mismatch (read speech – conversation). However, caution must be applied, as the EERs could have been affected by score calibration and an improved ASV system used in the present study.

Native listeners performed better than non-natives in all conditions (EERs = 6.96% versus 12.39% for read speech – read speech, 15.12% versus 23.22% for conversation – conversation, and 20.68% versus 31.46% for read speech – conversation, for native versus non-native listeners respectively). All differences were statistically significant ($p < 0.05$). It can, therefore, be assumed that, across both style-matched and style-mismatched trials, a listener’s fluency in the language being spoken affects perception. Note that in the case of native speakers, humans outperformed machines for style-matched cases (EERs of 6.96% versus 14.35%; $p < 0.05$ and 15.12% versus 19.87%; $p < 0.05$). In the style-mismatched condition, however, there was no significant difference in the performances between native listeners and machines. Consistent with [9, 12], native listeners performed better than machines in style-matched trials. There was no significant difference between the performances of non-native listeners and machines on read speech – read speech and read speech – conversation trials. But in the case of conversation – conversation, machines performed better than non-natives (EER of 19.47% versus 23.22%, $p < 0.05$).

Finally, the fusion of human and machine scores improved performance significantly in the majority of conditions ($p < 0.05$). One exception was read speech – read speech fusion of native listeners and machine scores ($p = 0.05$ for humans versus fusion). The other exception was the conversation – conversation condition for the fusion of non-native listeners and machine scores ($p = 0.053$ for machine versus fusion). The overall improvement due to fusion was consistent with previous reports [31, 12], and with our hypothesis that humans and machines use different approaches to speaker discrimination. The small size of the dataset meant that it was not possible to split it into development and evaluation sets while performing fusion, potentially resulting in some over-fitting.

Differences between the same speaker versus different speaker tasks can be studied by comparing $C_{\text{llr}}^{\text{tar}}$ and $C_{\text{llr}}^{\text{non}}$ values in Table 1. The comparison is highlighted in each condition by showing the better value between $C_{\text{llr}}^{\text{tar}}$ and $C_{\text{llr}}^{\text{non}}$ in purple and the worse value in orange. We see that target (same speaker) trials were easier for native listeners when compared to the non-target (different speaker) trials ($p < 0.05$). In contrast, for machines the non-target trials were easier ($p < 0.05$). On the other hand, non-native listeners were not consistent; they found non-target trials easier in the read speech – conversation condition and target trials easier in the other two conditions. Recall that C_{llr} is a measure of soft decisions, and represents the reliability of log-likelihood scores. Table 1 shows that native listeners were more reliable than non-natives ($p < 0.05$). In subsequent analyses, we only use scores from native listeners.

Female versus male listeners comparison showed that females did better in read speech – read speech and read speech – conversation conditions whereas males did better in conver-

Table 1: Speaker discrimination performance in terms of equal error rates (EER, %) and log-likelihood-ratio cost function: combined (C_{lr}), target trials C_{lr}^{tar} and non-target trials C_{lr}^{non} . A comparison of C_{lr}^{tar} and C_{lr}^{non} is also shown by using purple for the better of the two values in each condition and orange for the worse value.

Listeners	System	read-read				conversation-conversation				read-conversation			
		EER %	C_{lr}	C_{lr}^{tar}	C_{lr}^{non}	EER %	C_{lr}	C_{lr}^{tar}	C_{lr}^{non}	EER %	C_{lr}	C_{lr}^{tar}	C_{lr}^{non}
Native	Machines	14.35	0.4421	0.4443	0.4399	19.87	0.5765	0.5888	0.5642	21.78	0.6492	0.6501	0.6484
	Humans	6.96	0.2642	0.2103	0.3182	15.12	0.5291	0.5010	0.5573	20.68	0.6911	0.6900	0.6923
	Fusion	4.92	0.1731	0.1447	0.2015	11.20	0.3790	0.3624	0.3956	16.39	0.5207	0.5213	0.5202
Non Native	Machines	13.95	0.4113	0.4147	0.4079	19.47	0.5682	0.6116	0.5248	19.64	0.6754	0.6791	0.6718
	Humans	12.39	0.4292	0.3836	0.4748	23.22	0.7026	0.6667	0.7385	31.46	0.8723	0.8730	0.8716
	Fusion	5.69	0.1953	0.1616	0.2289	13.57	0.4406	0.4349	0.4463	19.34	0.6339	0.6352	0.6327

sation – conversation condition (EERs = 9.08% versus 10.24% for read speech – read speech, 17.47% versus 15.09% for conversation – conversation, and 23.12% versus 25.16% for read speech – conversation, for female versus male listeners respectively). All the results were statistically significant ($p < 0.05$) and will be studied in detail in follow-up work.

4.2. Speaker-level Log-Likelihood-Ratio Analysis

The histograms of speaker-level log-likelihood-ratios over target trial (L^{tar}) and non-target trial (L^{non}) distributions are shown in Figure 1. For read speech – read speech trials, the L^{tar} and L^{non} distributions are skewed towards correct responses for humans. Increased separation of the target and non-target distributions makes discrimination easier, resulting in better human performance than machine performance (EER = 6.96% versus 14.35%).

In the read speech – read speech condition, the L^{tar} distribution for humans showed small variance ($variance = 0.05$) and was confined to the positive response region. This was not the case with L^{non} ($variance = 0.57$). Hence, humans were more confident when classifying same voices than different voices, similar to [12]. In contrast, the L^{tar} distributions had a larger variance and overlapped with L^{non} distributions in the other two conditions (conversation – conversation and read speech – conversation). Hence, there was a decrease in listeners’ confidence in classifying both same and different voices.

For machines, variances of L^{tar} ($variance = 0.57$) and L^{non} ($variance = 0.66$) distributions the read speech – read speech trials were large and overlapping. This indicates that the degree of uncertainty was higher in the machine scores when compared to humans. There was also a similar overlap in the score distributions for the other two conditions, which suggests that there may be no particular difference between same speaker and different speaker decisions in machines.

4.3. Speaker-level Log-Likelihood-Ratio Cost Analysis

Speaker-level C_{lr} can be interpreted as the information available to the system (humans or machines) for that speaker. The higher C_{lr} for the speaker, the more confused the system is when distinguishing that speaker from others. We divided the speaker population into three subsets for both humans and machine scores, based on the speaker-level C_{lr} . The subset “easy” (easy to distinguish speakers) comprised thirteen speakers with the lowest C_{lr} values. Likewise, the “hard” subset (difficult to distinguish speakers) included the thirteen speakers with the highest C_{lr} values. The remaining fourteen speakers (neither easy nor hard to distinguish) were classified as “average”.

The confusion matrix in Figure 2 provides a visualization

of the distribution of speakers in the three subsets. An entry $count_{m_i, h_j}$ denotes the total count of speakers from subset i of machines overlapping with subset j of humans. For example, the second entry (5) is the number of speakers who were easy for machines to distinguish but were of average difficulty for humans. A matrix in which entries fall primarily on the diagonal would mean that the difficulty of individual speakers is similar for machines and humans. This is not the case in Figure 2. Thus, the degree of speaker discrimination difficulty is different for humans and machines. The triplet in parentheses represents the distribution of speakers of the corresponding entry in the matrix into (easy, average, hard) subsets based on the fusion scores. Closer inspection shows that diagonal entries get further distributed into off-diagonal entries for the fused scores. The last entry (9) illustrates this point clearly; hard speakers for both humans and machines are predominantly average (6/9) for the fused system.

Dissimilarities in the distributions of speakers seem to suggest that humans and machines perform the speaker discrimination tasks differently. They might be using different discrimination strategies while performing these tasks. The distributions based on the fused scores show a redistribution that is skewed towards correct responses, emphasizing the complementary nature of human and machine scores. This highlights instances when machines could assist humans in speaker discrimination.

5. Conclusion

The present study was designed to investigate the effects of style variability on speaker discrimination performance for humans and machines. These effects were evaluated using a short-utterance, text-independent speaker discrimination task with read and conversational speech. Experimental results show that both humans and machines performed better in the style-matched condition of read speech – read speech, followed by conversation – conversation. This suggests that conversations introduce more variability in the speakers’ acoustic spaces. Both humans and machines have the lowest performance for read speech – conversation, a style-mismatched condition. Overall, this implies that speaking style variability affects the performances of both humans and machines. Native English-speaking listeners showed higher reliability in performing the tasks than non-natives, and performed better than machines in style-matched conditions. Despite speaking-style variability, listeners were consistently more confident when they are performing the same speaker tasks versus different speaker tasks, but their overall confidence was the highest for read speech – read speech trials. Speakers who populate the subsets that humans and machines found easy or difficult to distinguish were not entirely the same.

6. References

- [1] S. J. Wenndt and R. L. Mitchell, "Machine recognition vs human recognition of voices," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 4245–4248.
- [2] D. Van Lancker and J. Kreiman, "Voice discrimination and recognition are separate abilities," *Neuropsychologia*, vol. 25, no. 5, pp. 829–834, 1987, iSBN: 0028-3932.
- [3] J. Kreiman and D. Sidtis, *Foundations of voice studies: An interdisciplinary approach to voice production and perception*. John Wiley & Sons, 2011.
- [4] H. Saslove and A. D. Yarmey, "Long-term auditory memory: speaker identification," *The Journal of Applied Psychology*, vol. 65, no. 1, pp. 111–116, Feb. 1980.
- [5] H. Blatchford and P. Foulkes, "Identification of voices in shouting," *International Journal of Speech, Language and the Law*, vol. 13, no. 2, pp. 241–254, Aug. 2006, number: 2.
- [6] R. González Hautamäki, V. Hautamäki, and T. Kinnunen, "On the limits of automatic speaker verification: Explaining degraded recognizer scores through acoustic changes resulting from voice disguise," *The Journal of the Acoustical Society of America*, vol. 146, no. 1, pp. 693–704, Jul. 2019.
- [7] R. González Hautamäki, T. Kinnunen, V. Hautamäki, and A.-M. Laukkanen, "Automatic versus human speaker verification: The case of voice mimicry," *Speech Communication*, vol. 72, pp. 13–31, Sep. 2015.
- [8] M. Jessen, "Forensic Phonetics," *Language and Linguistics Compass*, vol. 2, no. 4, pp. 671–711, 2008, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1749-818X.2008.00066.x>.
- [9] S. J. Park, G. Yeung, N. Vesselinova, J. Kreiman, P. A. Keating, and A. Alwan, "Towards understanding speaker discrimination abilities in humans and machines for text-independent short utterances of different speech styles," *JASA*, vol. 144, no. 1, pp. 375–386, 2018.
- [10] N. Lavan, L. F. K. Burston, and L. Garrido, "How many voices did you hear? Natural variability disrupts identity perception from unfamiliar voices," *British Journal of Psychology*, vol. 110, no. 3, pp. 576–593, 2019.
- [11] J. Johnson, C. McGettigan, and N. Lavan, "Comparing unfamiliar voice and face identity perception using identity sorting tasks," *PsyArXiv*, 2019.
- [12] S. J. Park, A. Afshan, J. Kreiman, G. Yeung, and A. Alwan, "Target and Non-target Speaker Discrimination by Humans and Machines." Brighton, United Kingdom: IEEE, May 2019, pp. 6326–6330.
- [13] T. K. Perrachione and P. C. Wong, "Learning to recognize speakers of a non-native language: Implications for the functional organization of human auditory cortex," *Neuropsychologia*, vol. 45, no. 8, pp. 1899–1910, 2007.
- [14] S. R. Schweinberger, H. Kawahara, A. P. Simpson, V. G. Skuk, and R. Zäske, "Speaker perception," *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 5, no. 1, pp. 15–25, 2014, iSBN: 1939-5078.
- [15] P. Keating, J. Kreiman, and A. Alwan, "A New Speech Database For Within- and Between-Speaker Variability," *Proc of the 19th ICPhS*, p. 4, 2019.
- [16] J. Kreiman, S. J. Park, P. A. Keating, and A. Alwan, "The Relationship Between Acoustic and Perceived Intraspeaker Variability in Voice Quality," in *Interspeech*, Dresden, Germany, 2015.
- [17] "IEEE Recommended Practice for Speech Quality Measurements," *IEEE No 297-1969*, pp. 1–24, Jun. 1969.
- [18] M. Przybocki and A. Martin, "NIST Speaker Recognition Evaluation Chronicles," in *Proc. Odyssey*, 2004, pp. 12–22.
- [19] M. Przybocki, A. Martin, and A. Le, "NIST Speaker Recognition Evaluation Chronicles - Part 2," in *Proc. Odyssey*, 2006, pp. 1–6.
- [20] A. F. Martin and C. S. Greenberg, "NIST 2008 Speaker Recognition Evaluation: Performance across Telephone and Room Microphone Channels," in *Proc. Interspeech*, Brighton, UK, 2009, pp. 2579–2582, iSBN: 19909772.
- [21] D. Graff, K. Walker, and A. Canavan, "Switchboard-2 phase ii," *LDC 99S79*—<http://www.ldc.upenn.edu/Catalog>, 1999.
- [22] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *ICASSP*, 2018.
- [23] P. Kenny, T. Stafylakis, P. Ouellet, M. J. Alam, and P. Dumouchel, "PLDA for Speaker Verification with Utterances of Arbitrary Duration," in *Proc. ICASSP*, 2013, pp. 7649–7653, iSBN: 15206149.
- [24] D. Garcia-Romero, A. McCree, S. Shum, N. Brummer, and C. Vaquero, "UNSUPERVISED DOMAIN ADAPTATION FOR I-VECTOR SPEAKER RECOGNITION," *Odyssey*, p. 5, 2014.
- [25] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, and others, "The Kaldi speech recognition toolkit," *IEEE Signal Processing Society, Tech. Rep.*, 2011.
- [26] N. Brümmer and E. De Villiers, *The BOSARIS Toolkit User Guide: Theory, Algorithms and Code for Binary Classifier Score Processing*. arXiv preprint arXiv:1304.2865, 2011.
- [27] D. A. van Leeuwen and N. Brümmer, "An Introduction to Application-Independent Evaluation of Speaker Recognition Systems," in *Speaker Classification I: Fundamentals, Features, and Methods*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 330–353.
- [28] N. Brümmer, L. Burget, J. H. Černocký, O. Glembek, F. Grézl, M. Karafiát, D. A. Van Leeuwen, P. Matějka, P. Schwarz, and A. Strasheim, "Fusion of heterogeneous speaker recognition systems in the STBU submission for the NIST Speaker Recognition Evaluation 2006," in *Proc. ICASSP*, 2007, iSBN: 15587916.
- [29] N. Smirnov, "Table for Estimating the Goodness of Fit of Empirical Distributions," *The Annals of Mathematical Statistics*, vol. 19, no. 2, pp. 279–281, 1948.
- [30] Q. McNemar, "Note on the sampling error of the difference between correlated proportions or percentages," *Psychometrika*, vol. 12, no. 2, pp. 153–157, 1947, publisher: Springer.
- [31] R. G. Hautamäki, V. Hautamäki, P. Rajan, and T. Kinnunen, "Merging Human and Automatic System Decisions to Improve Speaker Recognition Performance," in *Proc. Interspeech*, Lyon, France, 2013, pp. 2519–2523, iSBN: 19909772.