

Enhancing Age-Related Robustness in Children Speaker Verification

Vishwas M. Shetty^{1,*}, Jiusi Zheng^{1,*}, Steven M. Lulich², Abeer Alwan¹

¹*Department of Electrical and Computer Engineering, University of California, Los Angeles, USA*

²*Speech Language and Hearing Sciences, Indiana University, Bloomington, Indiana, USA*
{shettyvishwas, zheng94}@ucla.edu, slulich@iu.edu, alwan@ee.ucla.edu

Abstract—One of the main challenges in children’s speaker verification (C-SV) is the significant change in children’s voices as they grow. In this paper, we propose two approaches to improve age-related robustness in C-SV. We first introduce a Feature Transform Adapter (FTA) module that integrates local patterns into higher-level global representations, reducing overfitting to specific local features and improving the inter-year SV performance of the system. We then employ Synthetic Audio Augmentation (SAA) to increase data diversity and size, thereby improving robustness against age-related changes. Since the lack of longitudinal speech datasets makes it difficult to measure age-related robustness of C-SV systems, we introduce a longitudinal dataset to assess inter-year verification robustness of C-SV systems. By integrating both of our proposed methods, the average equal error rate was reduced by 19.4%, 13.0%, and 6.1% in the one-year, two-year, and three-year gap inter-year evaluation sets, respectively, compared to the baseline.

Index Terms—Children Speaker Verification, Inter-Year Child Speaker Verification, Data augmentation, Domain adaptation

I. INTRODUCTION

With the advent of online learning, the number of children using online learning applications has increased exponentially. Speaker verification can enable personalized experiences for children in educational applications, virtual assistants, and interactive learning tools by recognizing individual voices. This can enhance engagement and tailor learning materials to each child. The current state of the art speaker verification systems [1]–[4] have achieved significant success on various benchmark datasets [5]–[8]. While speaker verification for adults has seen significant improvements in accuracy due to large datasets and optimized model architectures, speaker verification for children faces greater difficulties.

The challenges in children speaker verification are more pronounced compared to adults in the following aspects. First, the limited availability of children’s speech datasets makes it difficult for models to reach their full potential, particularly because insufficient data can lead to overfitting, which in turn reduces the model’s robustness. Many studies have addressed this issue in children’s speaker verification by employing various data augmentation techniques, including SpecAug [9], Speed Perturbation [10], Noise and RIR [11], [12], Voice Conversion [13]–[15], and a combination of several signal processing approaches [16]. Second, the anatomical changes in children’s vocal tracts during growth lead to more

pronounced variations in their voices [17]–[19], compared to the relatively stable vocal characteristics of adults [20]. Due to the scarcity of longitudinal speech datasets, research on the impact of age-related changes on the performance of speaker verification systems is limited [21], with a few studies addressing inter-year C-SV. However, understanding and improving age robustness in children’s speaker verification systems is crucial in practical applications because children’s voices can change significantly over short periods, leading to potential degradation in verification accuracy over time. Addressing this issue is essential to ensure consistent and reliable performance of these systems in real-world scenarios where a child may use the verification system over several years.

In this paper, we propose techniques to improve the inter-age C-SV through two approaches. Drawing inspiration from recent advancements in feature adapters [22], [23], we first propose the Feature Transform Adapter (FTA) which aggregates the original local features at different time scales and resolutions to eliminate disturbances caused by age-related changes in the speaker’s audio characteristics, thereby forming more robust features. We then present a data augmentation method using HiFi-GAN [24] to generate in-domain synthetic audio from speech spectrograms. Although this approach has been previously applied in other fields [25], [26], we have adapted this data augmentation technique for the children SV task. As a final step, we introduce a longitudinal children’s speech dataset designed to support the evaluation of age-related speech changes and to advance the development and testing of robust speaker processing systems for children. A comprehensive set of ablation studies is conducted to validate the efficacy of the proposed methods.

II. METHOD

A. Feature Transform Adapter

Different from inter-year speaker verification for adult speech [21] (referred to as cross-age in [21]), age significantly impacts the performance of C-SV systems. In this context, global features, such as formant distribution [27], refer to the stable and consistent aspects of a speaker’s voice that remain relatively unchanged over time, making them crucial for identifying the speaker’s identity. In contrast, local features, such as segmental durations [17], exhibit transient, short-term

*These authors contributed equally to this work.

variations that are sensitive to slight changes in the input signal, such as age-related shifts. We hypothesize that as a child ages, local features undergo changes in both the time and frequency domains, whereas global features remain relatively stable. To address this, we introduce an innovative structure referred to as the Feature Transform Adapter (FTA), as depicted in Fig 1. This adapter aggregates features in a manner that obscures age-related local information while preserving essential global information related to the speaker’s identity. The primary goal of this adapter is to enhance the robustness of the children’s speaker verification system, making it more effective in handling age-related variability and improving overall verification accuracy.

The details of the FTA are as follows. Initially, the 80-dimensional filterbank features are layer-normalized and passed through a fully connected layer. The features are then processed through two 1D convolutional layers to aggregate local features along the frequency dimension. The aim is to capture the average variations along the frequency dimension while maintaining the global speaker identity information. The ReLU activation function is employed between the two convolutional layers to introduce non-linearity into the adapter. Finally, the processed features are integrated with the original features via a residual connection. These features can be then used to train the SV system.

B. Synthetic Audio Augmentation

Synthetic Audio Augmentation (SAA) involves using artificially generated audio samples to diversify training datasets. We assume that such synthesis methods preserve the speaker’s identity but may not accurately retain age-related information due to challenges in capturing age-specific features during synthesis. In this paper, we introduce an SAA method based on HiFi-GAN [24]. HiFi-GAN includes one generator and two discriminators: the generator converts mel spectrograms into raw waveforms through the convolutional neural network. Among the discriminators, the Multi-Period Discriminator (MPD) captures diverse periodic patterns by analyzing various segments of the audio, while the Multi-Scale Discriminator (MSD) focuses on both short-term and long-term dependencies across multiple resolutions. During training, the generator learns to produce realistic audio waveforms that can deceive both discriminators, while the discriminators simultaneously refine their ability to distinguish between real and generated audio. In this work we directly use the pre-trained open source HiFi-GAN generator [28] to generate the synthetic audio. For each child’s speech signal from a specific year, we input the corresponding spectrogram into the HiFi-GAN generator, generating synthetic audio to augment the training dataset.

C. Child Longitudinal Test Set

To address the dataset gap in inter-year children speaker verification, we present a longitudinal evaluation set. Derived from an existing longitudinal child speech dataset [19], [27], [29] and referred to as the IU dataset, it includes data from eight children (four boys, four girls) recorded annually from

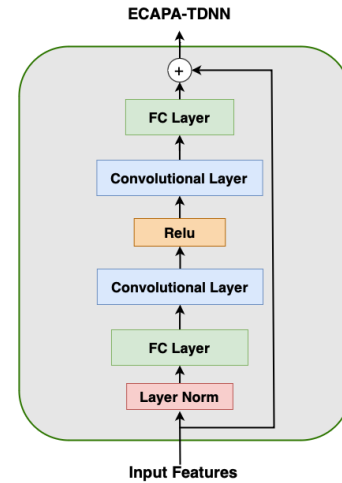


Fig. 1: Overview of the proposed Feature Transform Adapter

Grade 1 to Grade 4. The speech utterances were elicited through tasks commonly used by speech-language pathologists (Goldman and Fristoe Test of Articulation, 3rd Edition, aka GFTA-3) and were recorded at Indiana University, Bloomington. We applied Voice Activity Detection (VAD) [30] to extract child speech segments of at least 2 seconds, creating the IU evaluation database.

IU inter-year verification pairs were created by pairing enrollment speech from one grade with test segments from subsequent grades. For example, Grade 1 (G1) enrollment speech was paired with Grade 2 (G2), Grade 3 (G3), and Grade 4 (G4) segments, forming G1-G2, G1-G3, and G1-G4 target pairs. Negative pairs were sourced from the same year to increase task complexity. Other sets like G2-G3, G2-G4, and G3-G4 were constructed similarly. Additionally, IU intra-year sets, i.e., G1-G1, G2-G2, G3-G3, and G4-G4 were created. Each IU evaluation set contains 2,000 verification pairs—1,000 positive and 1,000 negative.

III. EXPERIMENTAL SETTINGS

A. Databases

The CSLU [31] dataset is utilized in all the training and fine-tuning experiments proposed in this paper. The CSLU speech corpus, also referred to as the OGI kids’ speech corpus, includes both spontaneous and scripted speech data from approximately 1,100 children, ranging from kindergarten (*K00*) through grade 10 (*K10*). Each age group consists of approximately 100 speakers recorded uttering single words, sentences, and digit sequences. For our experiments, we utilized only the scripted speech data.

We employ two distinct train-eval splits from the CSLU database, as detailed in Table 1. To achieve optimal results through fine-tuning and our proposed methods, we use CSLU-S, a split with a large training set, providing ample data for fine-tuning. As a result, the evaluation set in CSLU-S contains fewer speakers, with 866 speakers in the training set and 255

TABLE I: *Train-eval splits used to train/fine-tune the models discussed in this paper. #Spks refers to number of speakers, #Hrs represents duration in hours, and #Trials stands for number of SV trials in the evaluation set.*

	Train		Eval	
	# Spks	# Hrs	# Spks	# Trials
CSLU-S [32]	866	24.00	255	44000
CSLU-L [16]	120	2.65	993	190972

speakers in the evaluation set, covering 44,000 intra-year SV trials across all grades. To validate the effectiveness of our methods with a larger number of speakers in intra-year SV and to compare our approaches with existing state of the art C-SV models, we utilize the CSLU-L dataset in some of the experiments. In this split, the training set contains 120 speakers, and the evaluation set includes 993 speakers with 190,972 SV trials. There is no overlap between the speakers in the train and evaluation subsets of both splits. Along with the evaluation sets from CSLU database, we also evaluate the models on IU dataset discussed in section II-C.

B. Baseline System

An ECAPA-TDNN [1] network was trained for a speaker identification (SID) task using the CSLU dataset, with the training process conducted via the SpeechBrain toolkit [33]. The input features consisted of 80-dimensional filter bank features, extracted every 10 ms using a 25 ms window. The ECAPA-TDNN architecture includes frame-level convolutional layers with 1024 channels and 128 dimensional attention channels. The output of the model is the 192 dimensional speaker vectors. We then compare these vectors to examine speaker similarity. In terms of the training details, all models are trained for 15 epochs with a batch size of 16. The Adam optimizer is utilized with a learning rate of $1e-3$. Additionally, four online data augmentation methods were employed during training, with detailed parameter configurations provided in III-C. This model is used as the baseline in our experiments and is referred to as *Baseline* in this paper.

C. Augmentation Setup

To construct a strong baseline system, we incorporated four distinct data augmentation techniques, each with the following parameter configurations. For Noise Augmentation, the signal-to-noise ratio (SNR) of the added noise was varied between 0 and 15 dB. Additionally, Room Impulse Response (RIR) augmentation was employed to enhance the system’s robustness in reverberant environments. The frequency drop augmentation allowed for the random removal of 1 to 3 frequency bands at a time, across the full frequency range, with a frequency band width of 0.05. In time drop augmentation, audio chunks ranging from 1000 to 2000 samples were randomly removed, with between 1 and 5 chunks being dropped per audio sample. The baseline model described in III-B was trained using the above discussed augmentation techniques on the CSLU train data. In the Synthetic Audio Augmentation process discussed

in Section II-B, we utilized the open-sourced pretrained HiFi-GAN vocoder and denoiser [28], where the denoising strength was set to 0.005.

IV. RESULTS AND DISCUSSION

A. Performance of Feature Transform Adapter

To evaluate the robustness of the proposed C-SV systems, Table II shows the Equal Error Rates (EER) for different SV systems trained using CSLU-S and evaluated under intra-year and inter-year verification (i.e., longitudinal evaluation) on the CSLU and IU evaluation sets. Fine-tuning significantly reduces the EER for intra-year verification compared to the baseline but does not consistently improve inter-year performance, likely due to insufficient representation of inter-year scenarios during training, which leads to poor generalization and variable performance. The FTA approach mitigates this limitation by leveraging convolutional layers to capture generalizable patterns, resulting in improved results across most CSLU evaluation sets and all IU evaluation sets compared to the conventional fine-tuning. In contrast, replacing the convolutional layers with fully connected layers (+RA) shows no improvement, hence supporting our hypothesis on the ability of the convolutional layers to aggregate local features as a key to enhancing robustness.

B. Effect of Synthetic Audio Augmentation

Overall, using SAA to fine-tune along with FTA, referred to as ++SAA in Table II, further reduces the EER on IU inter-year evaluation sets compared to +FTA, although SAA carries the risk of lowering intra-year verification accuracy. In inter-year verification, as shown in Table II, SAA achieved the lowest EER in 5 out of 6 IU inter-year evaluation sets, demonstrating a clear improvement. Specifically, the average EER were reduced by 19.4%, 13.0%, and 6.1% in the one-year, two-year, and three-year gap evaluation sets, respectively, compared to the baseline. In contrast, in intra-year verification, SAA resulted in a noticeable increase in verification EER in 7 out of 11 age groups on the CSLU dataset, and 2 out of 4 intra-year evaluation sets on the IU dataset also showed EER degradation compared to using only the FTA method. We noticed that this trade-off between inter-year and intra-year verification accuracy aligns with the previous study on inter-year adult speech verification [21] (inter-year is referred to as cross-age in [21]).

C. The Impact of Number of Speakers

To evaluate the generalizability of our proposed methods on a larger intra-year evaluation set, we retrained the models using the CSLU-L train-eval split, with results summarized in the first four rows of Table III. Our methods demonstrate continued effectiveness despite the substantial increase in the number of speakers in the CSLU intra-year evaluation set. Specifically, the model fine-tuned with both SAA and FTA (i.e., ++SAA[†]) reduced the EER by 28.3% on the CSLU intra-year set, 29.3% on the IU intra-year set, and 10.8%

TABLE II: *Equal Error Rate (EER) results for various techniques across different grades (K00–K10 and G1–G4, representing different age groups). ‘Intra-Year’ refers to enrollment and test segments from the same grade, while ‘Inter-Year’ indicates segments from different grades (e.g., G1-G2). Techniques evaluated include the baseline model, fine-tuning, FTA (Feature Transform Adapter), RA (Residual Adapter), and SAA (Synthetic Audio Augmentation), all using the ‘CSLU-S’ dataset. The lower EER on the CSLU test set compared to the IU test set for the same age group is likely due to CSLU’s longer average audio duration, which provides more detailed speaker information.*

CSLU Intra-Year											
	K00	K01	K02	K03	K04	K05	K06	K07	K08	K09	K10
<i>Baseline</i>	13.2	12.85	10.45	9	9.55	7.05	4.7	4.9	4.5	3.45	4.4
<i>Finetune</i>	9.55	10.65	9.05	6.95	7.25	4.1	4.05	2.95	2.75	1.85	3.3
+ RA	9.7	11.5	8.75	7.5	7.1	4.35	3.85	3.2	2.6	1.95	3.45
+ FTA	10.5	10.35	9.05	6.85	5.8	4.2	3.45	2.85	2.55	1.7	3.25
++ SAA	10.6	12.1	9.65	6.8	5.75	4.4	3	2.85	2.6	2.65	3.6

IU Intra-Year				IU Inter-Year						
	G1-G1	G2-G2	G3-G3	G4-G4	G1-G2	G1-G3	G1-G4	G2-G3	G2-G4	G3-G4
<i>Baseline</i>	22.5	15.5	17.5	16.2	29.5	32.8	34.7	39.5	40.3	23.4
<i>Finetune</i>	19	13.9	15.8	11	28.7	36.1	36.7	37.1	39	24.1
+ RA	19.4	16.7	13	13.1	27.9	38	37.1	39.5	38.3	22.3
+ FTA	16.4	12.9	12.2	10.6	26.6	34.2	35.8	34.2	35.2	21
++ SAA	19.9	12.2	8.1	13.1	27.1	31.15	32.6	30.6	32.3	16.8

on the IU inter-year set compared to the *Baseline*[†] model. In the CSLU-L case, fine-tuning with only the FTA adapter (+*FTA*[†]) showed no improvements. One possible explanation is that the CSLU-L train set has a significantly lower number of speakers and hence the limited size of the training set becomes the dominant factor constraining model performance. This hypothesis is supported by the considerable performance gains from incorporating the SAA data augmentation method during fine-tuning (i.e., ++*SAA*[†]), which mitigated the impact of data scarcity.

We also compared the proposed approaches with other state-of-the-art methods in C-SV. The *Proposed 3/11* model from [16] was trained using 11 different data augmentation techniques on the out-of-domain dataset VoxCeleb2. We observe that incorporating in-domain child speech CSLU data significantly outperforms using only out-of-domain data augmentation for both intra-year and inter-year verification. Additionally, to assess our proposed approaches against an in-domain model, we compared our models with the *Baseline 3/5* from [16], which was trained using 315 hours of MyST child speech data and five data augmentation techniques. Notably, our model fine-tuned with only 24 hours of CSLU-S training data with ++*SAA*^{*} method in row 6 of Table III achieved an EER of 28.72% on the IU inter-year evaluation set, outperforming the *Baseline 3/5* model, which achieved an EER of 30.83%. This highlights that conventional training and augmentation methods alone are insufficient for addressing robustness issues in inter-year C-SV, whereas larger datasets of child speech show more pronounced benefits in intra-year verification. Specifically, we observed that the *Baseline 3/5* model performs notably better on the IU intra-year evaluation set. This is likely because the MyST data (Grades 3 to 5) used for *Baseline 3/5* aligns better with the IU dataset (Grades 1 to 4), while the OGI evaluation sets span a broader age range (Grades 0 to 10).

TABLE III: *Equal Error Rate (EER) comparisons across three test sets are shown. CSLU is an intra-year test set obtained by combining K00 to K10 test sets from CSLU-L evaluation split. IU Intra-Year and IU Inter-Year are combined ‘‘Intra-Year’’ and ‘‘Inter-Year’’ IU test sets from Table II. † indicates train data from CSLU-L, and * indicates train data from CSLU-S.*

	CSLU	IU Intra-Year	IU Inter-Year
<i>Baseline</i> [†]	12.15	27.02	38.7
<i>Finetune</i> [†]	11.1	19.53	35.15
+ <i>FTA</i> [†]	11.58	20.30	36.97
++ <i>SAA</i> [†]	8.71	19.1	34.5
+ <i>FTA</i> [*]	-	12.97	31.2
++ <i>SAA</i> [*]	-	13.8	28.72
<i>Proposed 3/11</i> [16]	16.77	24.03	41.65
<i>Baseline 3/5</i> [16]	8.15	10.24	30.83

V. CONCLUSIONS

In this paper, we propose two approaches to enhance the robustness of children speaker verification (C-SV) systems, focusing on reducing age-related impact. The first approach introduces a Feature Transform Adapter (FTA) to make input acoustic features age-invariant. The second approach enhances data diversity through Synthetic Audio Augmentation (SAA) using HiFi-GAN. We present a new C-SV evaluation dataset with both intra-year and inter-year child speech data. We also evaluate our proposed models against the existing C-SV systems and show that in the inter-year evaluation scenario, our best proposed system performs better than the existing C-SV systems.

VI. ACKNOWLEDGEMENTS

This work was funded in part by National Science Foundation (NSF) collaborative research grants 1551131 and 1551113.

REFERENCES

- [1] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in tdn based speaker verification," in *Interspeech*, 2020, pp. 3830–3834.
- [2] J. Villalba, N. Chen, D. Snyder, D. Garcia-Romero, A. McCree, G. Sell, J. Borgstrom, L. P. García-Perera, F. Richardson, R. Dehak, P. A. Torres-Carrasquillo, and N. Dehak, "State-of-the-art speaker recognition with neural network embeddings in NIST SRE18 and speakers in the wild evaluations," *Computer Speech Language*, vol. 60, p. 101026, 2020.
- [3] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu, H. yi Lee, and H. Meng, "MFA-Conformer: Multi-scale feature aggregation conformer for automatic speaker verification," in *Interspeech*, 2022, pp. 306–310.
- [4] X. Xiang, S. Wang, H. Huang, Y. Qian, and K. Yu, "Margin matters: Towards more discriminative deep neural network embeddings for speaker recognition," in *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2019, pp. 1652–1656.
- [5] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [6] K. Hechmi, T. N. Trong, V. Hautamäki, and T. Kinnunen, "Voxceleb enrichment for age and gender recognition," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 687–693.
- [7] R. A. Cole, M. Noel, and V. Noel, "The CSLU speaker recognition corpus," in *Fifth International Conference on Spoken Language Processing*, 1998.
- [8] S. Zheng, L. Cheng, Y. Chen, H. Wang, and Q. Chen, "3d-speaker: A large-scale multi-device, multi-distance, and multi-dialect corpus for speech representation disentanglement," *arXiv preprint arXiv:2306.15354*, 2023.
- [9] 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," in *Interspeech*, 2019, pp. 2613–2617.
- [10] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Interspeech*, 2015, pp. 3586–3589.
- [11] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *ICASSP*, 2017, pp. 5220–5224.
- [12] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [13] S. Shahnawazuddin, W. Ahmad, N. Adiga, and A. Kumar, "In-domain and out-of-domain data augmentation to improve children's speaker verification system in limited data scenario," in *ICASSP*, 2020, pp. 7554–7558.
- [14] S. Aziz and S. Shahnawazuddin, "Enhancing children's short utterance-based asv using inverse gamma-tone filtered cepstral coefficients," *Circuits, Systems, and Signal Processing*, vol. 43, no. 5, pp. 3020–3041, 2024.
- [15] S. Aziz, Ankita, and S. Shahnawazuddin, "Short-utterance-based children's speaker verification in low-resource conditions," *Circuits, Systems, and Signal Processing*, vol. 43, no. 3, pp. 1715–1740, 2024.
- [16] V. P. Singh, M. Sahidullah, and T. Kinnunen, "ChildAugment: Data augmentation methods for zero-resource children's speaker verification," *The Journal of the Acoustical Society of America*, vol. 155, no. 3, pp. 2221–2232, 2024.
- [17] S. Lee, A. Potamianos, and S. Narayanan, "Acoustics of children's speech: Developmental changes of temporal and spectral parameters," *The Journal of the Acoustical Society of America*, vol. 105, no. 3, pp. 1455–1468, 1999.
- [18] G. Yeung and A. Alwan, "On the difficulties of automatic speech recognition for kindergarten-aged children," in *Interspeech*, 2018, pp. 1661–1665.
- [19] V. M. Shetty, S. M. Lulich, and A. Alwan, "Developmental articulatory and acoustic features for six to ten year old children," in *Interspeech*, 2023, pp. 4598–4602.
- [20] A. O. Santos, J. Godoy, K. Silverio, and A. Brasolotto, "Vocal changes of men and women from different age decades: an analysis from 30 years of age," *Journal of Voice*, vol. 37, no. 6, pp. 840–850, 2023.
- [21] X. Qin, N. Li, W. Chao, D. Su, and M. Li, "Cross-age speaker verification: Learning age-invariant speaker embeddings," in *Interspeech*, 2022, pp. 1436–1440.
- [22] J. Tian, X. Hu, and X. Xu, "Learning emotion-invariant speaker representations for speaker verification," in *ICASSP*, 2024, pp. 10611–10615.
- [23] W. Huang, B. Han, S. Wang, Z. Chen, and Y. Qian, "Robust cross-domain speaker verification with multi-level domain adapters," in *ICASSP*, 2024, pp. 11781–11785.
- [24] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in neural information processing systems*, vol. 33, pp. 17022–17033, 2020.
- [25] N. Rossenbach, A. Zeyer, R. Schlüter, and H. Ney, "Generating synthetic audio data for attention-based speech recognition systems," in *ICASSP*, 2020, pp. 7069–7073.
- [26] A. Madhu and S. Kumaraswamy, "Data augmentation using generative adversarial network for environmental sound classification," in *2019 27th European signal processing conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [27] S. M. Lulich and S. D. Charles, "Development of formant frequency distributions in American English-speaking elementary school-aged children: A longitudinal study," in *Proceedings of Meetings on Acoustics*, vol. 42, no. 1. AIP Publishing, 2020.
- [28] <https://github.com/NVIDIA/DeepLearningExamples/tree/master/PyTorch/SpeechSynthesis/HiFiGAN>.
- [29] S. Charles and S. M. Lulich, "An ultrasound study of American English laterals produced by children," in *Proceedings of Meetings on Acoustics*, vol. 36, no. 1. AIP Publishing, 2019.
- [30] S. Team, "Silero VAD: pre-trained enterprise-grade voice activity detector (vad), number detector and language classifier," <https://github.com/snakers4/silero-vad>, 2024.
- [31] K. Shobaki, J.-P. Hosom, and R. Cole, "The OGI kids' speech corpus and recognizers," in *Proc. of ICSLP*. Citeseer, 2000, pp. 564–567.
- [32] R. Fan, N. Balaji Shankar, and A. Alwan, "Benchmarking children's asr with supervised and self-supervised speech foundation models," in *Interspeech*, 2024, pp. 5173–5177.
- [33] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawlatatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, arXiv:2106.04624.