
Few-Shot Synthetic Accented Speech for ASR Fine-Tuning: What Helps and When?

Yurii Halychanskyi^{1 2} Nimet Beyza Bozdogan¹ Mark Hasegawa-Johnson¹ Dilek Hakkani-Tür¹
Volodymyr Kindratenko^{1 2}

Abstract

Synthetic accented speech is a promising way to improve automatic speech recognition (ASR) when real accented recordings are scarce. We ask what makes such data useful for ASR fine-tuning: target-accent phoneme edits that expose the recognizer to accent-specific pronunciations, or random phoneme perturbations that act as augmentation in phoneme space. In a few-shot TTS pipeline, we compare LLM-generated accent edits with matched-rate random substitutions and oracle controls using ground-truth accented phonemes and prosody. Random substitutions recover much of the ASR gain: LLM target-accent edits improve over random by only a small margin, ground-truth phonemes stay close to the random baseline and nearly converge with it as the synthetic ASR fine-tuning set grows larger, and adding ground-truth prosody yields only a modest further gain. Mixing synthetic with real accented speech also stabilizes low-resource fine-tuning, but a fixed synthetic budget can later dilute the information in real data, showing that the real–synthetic ratio matters.

1. Introduction

ASR systems often underperform on accented speech, with accent-related performance gaps documented across benchmarks and bias analyses (Koenecke et al., 2020; Feng et al., 2021; Shi et al., 2021). Existing work addresses this problem from two directions. Model-centric approaches adapt the recognizer itself, for example through accent embeddings, adversarial objectives, or meta-learning (Jain et al., 2018; Sun et al., 2018; Das et al., 2021; Winata et al., 2020). Data-centric approaches instead augment training with accented speech generated by voice conversion or TTS (Zhang et al., 2022; Klumpp et al., 2023; Nespoli et al., 2024).

¹University of Illinois Urbana-Champaign, USA ²National Center for Supercomputing Applications, USA. Correspondence to: Yurii Halychanskyi <yuriih2@illinois.edu>.

While effective, both directions often require minutes to hours of accented speech (Zhou et al., 2023; Zhang et al., 2023; Nespoli et al., 2024), motivating few-shot techniques for settings where only a handful of accented utterances are available. Few-shot accented synthesis is difficult because accent variation spans several factors, including segmental pronunciations, prosody, and speaker-correlated acoustics (Zhong et al., 2025; Lee et al., 2023; Du et al., 2023). Prior work models accents through local/global representations, pronunciation rules, or multilingual resources (Zhou et al., 2025; Goronzy et al., 2004; Schaden, 2003; Inoue et al., 2025; Karakasidis et al., 2023); these approaches improve synthesis, but can still depend on substantial accented data, handcrafted rules, or multilingual corpora.

In this work, we study few-shot synthetic accented speech as a resource for ASR fine-tuning¹. We test two hypotheses about why synthetic accented data helps: **(i)** target-accent phoneme edits improve ASR by exposing the recognizer to accent-specific pronunciation patterns; **(ii)** random phoneme perturbations improve ASR by acting as noise augmentation in phoneme space and increasing robustness to pronunciation variation. We use an LLM as the practical structured-edit generator: unlike oracle accented phonemes or prosody, it can propose target-accent pronunciation changes for new source utterances from general phonetic knowledge and a few examples, without requiring ground-truth accented phonetic annotations for every sentence. We compare these LLM edits with matched-rate random substitutions, and use ground-truth accented phonemes and prosody as diagnostic oracle bounds. We also study the budget question that arises when real and synthetic accented speech are combined: how much synthetic data helps at a given real-data budget, when it stabilizes fine-tuning, and when it begins to dilute the information carried by real examples. Finally, we examine how little target-accent speech is needed to build the synthetic generator itself.

Our contributions are:

¹Audio samples, the full LLM prompt, and code are available at https://claussss.github.io/few_shot_accent_synthesis_demo

- Comparing matched-rate random substitutions, LLM target-accent edits, and oracle inputs reveals that phoneme perturbation alone explains much of the ASR fine-tuning gain. LLM target-accent edits and ground-truth accented phonemes both stay close to the random baseline and nearly converge with it at larger budgets, while adding ground-truth prosody yields only a modest further improvement.
- Mixing synthetic and real accented speech sharply reduces sensitivity to which real utterances are sampled at low real-data budgets, leading to substantially lower run-to-run variance. However, using much more synthetic than real speech can later dilute the benefit of additional real data; the crossover point differs by accent, highlighting the importance of choosing the real-synthetic ratio.
- Supporting analyses show that ASR trained on one synthetic speaker transfers to other speakers of the same accent, and that the synthetic set can be created from as little as $K=3$ target-accent references while LLM editing remains nearly unchanged even with no in-context examples.

2. Method

Our pipeline starts from a Standard American source utterance, its transcript, and a small set of target-accent reference utterances. We build on the phoneme-conditioned TTS acoustic model of [Zaïdi et al. \(2022\)](#). We represent phonemes using the ARPAbet symbol set. For each phoneme i , we associate prosodic controls (d_i, p_i, e_i) , where d_i is duration in frames using a 256-sample hop at 22.05 kHz, approximately 11.6 ms per frame, p_i is mean $\log-F_0$ in log-Hz over voiced frames, and e_i is mean log-energy. These controls are extracted from source speech using transcript-based phoneme alignments rather than predicted internally. We reuse source phoneme-level prosody to preserve natural prosodic diversity from real speech and broaden the synthetic data distribution. To encourage the decoder to follow these externally supplied controls, we add auxiliary pitch and energy reconstruction losses against the extracted frame-level contours during pretraining.

The pipeline introduces target-accent information in two stages. First, we perform target-speaker decoder adaptation: the TTS decoder is fine-tuned on the target-accent speaker’s reference utterances, using speaker and style embeddings extracted from the same target speaker. This shifts the acoustic renderer toward the target speaker and accent before any symbolic editing: from a few reference utterances, the decoder can learn speaker-level and accent-correlated realization patterns, so the same phoneme-prosody inputs are rendered closer to how the target speaker

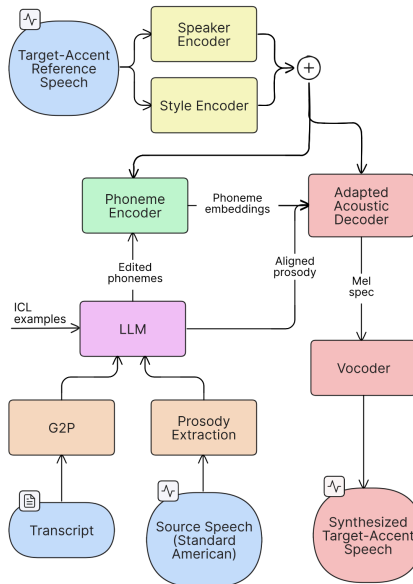


Figure 1. Few-shot synthetic accented speech pipeline used for ASR fine-tuning.

produces them. After adaptation, the system generates only the adapted target speaker’s voice. Second, an LLM edits the source phoneme sequence to simulate target-accent pronunciations while preserving phoneme-prosody alignment. Insertions, deletions, splits, and merges are allowed; prosody is modified only when needed to keep the sequence aligned. For example, the Standard American realization of *will*, W IH1 L | d:10,7,7 | p:5.3,5.3,5.2 | e:0.8,3.6,3.1, may be edited to the Indian-English pronunciation V IH1 L; because the sequence length is unchanged, the phoneme-level prosody is copied directly from the source. In-context examples contain paired source and target-accent phoneme-prosody sequences.

3. Experimental Setup

The TTS backbone is pretrained on LJSpeech and the English subset of ESD, both Standard American English only ([Ito & Johnson, 2017](#); [Zhou et al., 2022](#)). Accented speech comes from L2-ARCTIC; CMU Arctic speaker CLB supplies matched Standard American source utterances ([Zhao et al., 2018](#); [Kominek & Black, 2004](#)). We evaluate Indian English (TNI, RRBI, SVBI) and Korean English (HKK, YDCK, YKWK). For each accent, one adaptation speaker (TNI or HKK) supplies reference speech for embedding extraction, decoder adaptation, LLM prompting, and synthetic-data generation. Unless otherwise stated, we use $K=10$ reference utterances (about 36 seconds), disjoint from evaluation data.

Conditions. **American TTS** uses neither adaptation nor phoneme editing. **Adapt-only** applies target-

speaker decoder adaptation without symbolic edits. **Adapt+LLM** adds LLM-generated target-accent phoneme edits. **Adapt+Random** is the perturbation control, replacing uniformly sampled phoneme positions with uniformly sampled ARPAbet phonemes at the same average edit rate as the LLM for each accent: 19% for Indian English and 35% for Korean English. **Adapt+GT-phoneme** replaces source phonemes with accented sequences derived from L2-ARCTIC perceived phoneme labels (PPLs), while retaining aligned source prosody, and **Adapt+GT-phoneme+prosody** additionally uses prosody extracted from the corresponding real accented audio and aligned to the accented phoneme sequence. **Real** uses real accented fine-tuning utterances; **Real+Synth** combines real utterances with LLM-generated synthetic utterances.

Evaluation. Acoustic evaluation uses 50 held-out synthesized utterances per speaker (150 total per accent) and reports Whisper-small WER as an intelligibility proxy, UT-MOS for naturalness, and SpeechBrain accent-embedding cosine similarity (AccSim) (Radford et al., 2022; Saeki et al., 2022; Ravanelli et al., 2021). ASR experiments fine-tune wav2vec 2.0 Base with CTC (Baevski et al., 2020). ASR scaling varies the number of fine-tuning utterances $N \in [1, 500]$. Synthetic-only uses N synthetic utterances, while Real+Synth uses N real utterances plus a fixed set of 500 LLM-generated synthetic utterances. All ASR results are evaluated on real speech from all three speakers per accent (500 utterances each; 1,500 total) and averaged over 7 runs. Oracle comparisons use a separate annotated Indian-English set containing 96 utterances. Additional implementation details for acoustic features, TTS pretraining, adaptation, alignment, LLM prompting and model choice, and ASR fine-tuning are provided in Appendix A.1.

Supporting analyses. For cross-speaker transfer, ASR models fine-tuned on 500 synthetic utterances from the adaptation speaker are evaluated on the other speakers of the same accent. For synthetic-data sample efficiency, we vary $K \in \{1, 3, 5, 10, 15\}$ on Indian English (TNI). In each sweep, one component uses only K reference utterances while the remaining components stay fixed at 15: decoder adaptation utterances, speaker-embedding utterances, style-embedding utterances, or LLM in-context examples. The LLM sweep also includes $K=0$.

4. Results

4.1. What drives ASR fine-tuning gains from synthetic accented speech?

Table 1 shows that Adapt-only substantially increases AccSim relative to American TTS for both accents (0.27→0.69 for Indian; 0.32→0.61 for Korean), indicating that target-speaker adaptation shifts the rendered speech toward the

Table 1. Accent synthesis quality under different input conditions. Results are pooled across the three speakers of each accent.

Accent	Condition	WER (%)↓	UTMOS↑	AccSim↑
Indian English	American TTS	6.4	3.78	0.27
	Adapt-only	11.7	2.70	0.69
	Adapt + LLM	14.8	2.63	0.72
	Adapt + Random	47.2	2.31	0.68
	Adapt + GT-phon.+pros.	20.5	2.58	0.77
	Real accent	8.6	3.89	0.86
Korean English	American TTS	7.3	3.72	0.32
	Adapt-only	11.9	2.63	0.61
	Adapt + LLM	33.8	2.51	0.61
	Adapt + Random	93.4	2.12	0.58
	Adapt + GT-phon.+pros.	21.6	2.65	0.62
	Real accent	14.1	3.81	0.72




Figure 2. ASR WER versus fine-tuning budget N for Indian (TNI) and Korean (HKK) English. Error bars denote standard deviation over runs. The shaded band shows the gap between Adapt+Random and Adapt+LLM; the band’s lower edge corresponds to Adapt+LLM and the upper edge to Adapt+Random. Real+Synth mixes N real utterances with 500 LLM-generated synthetic utterances.

target-accent acoustics. Adding LLM phoneme edits improves Indian AccSim further (0.69→0.72) with moderate WER degradation, but does not improve Korean AccSim and raises acoustic WER from 11.9% to 33.8%. Korean edits are denser than Indian edits (35% versus 19%), suggesting that even after adaptation the decoder struggles to realize the larger symbolic shift intelligibly. Random substitutions are much less acoustically plausible still, reaching 47.2% WER for Indian English and 93.4% for Korean English without improving AccSim.

The downstream ASR results tell a different story. Figure 2 shows that Adapt+Random approaches Adapt+LLM for Indian English and is only modestly worse for Korean English. This indicates that a large part of the ASR gain comes from pronunciation-space augmentation: the recognizer benefits from exposure to symbolic variation even when the generated audio is not a faithful target-accent realization. Target-accent edits still help beyond random perturbation, possibly because they remain more linguistically coherent and easier for the adapted decoder to render than arbitrary substitutions, but the margin is small.

Figure 3 sharpens that interpretation. GT-phoneme remains close to the random baseline across the annotated Indian subset, with a gap that is typically under 1 abso-




Figure 3. Oracle analysis on an Indian-English subset with ground-truth annotations. GT-phoneme uses accented phonemes with aligned source prosody; GT-phoneme+prosody uses both accented phonemes and forced-aligned target-accent prosody.




Figure 4. Sample efficiency on Indian English (TNI). (a) Accent similarity and (b) UTMOS versus the number of real target-accent references K .

lute WER point and nearly vanishes at larger N . Even GT-phoneme+prosody reaches only a modest additional advantage over random, about 2 points at the largest budget despite using oracle symbolic and prosodic inputs. Thus, target-accent structure matters, especially at small budgets, but much of the synthetic-only ASR benefit in this setup is explained by phoneme perturbation acting as augmentation.

4.2. When does synthetic data help alongside real speech?

Synthetic data provides an additional benefit when combined with real accented speech, but only if the real-synthetic balance is appropriate. In Fig. 2, Real+Synth improves over real-only fine-tuning at the smallest budgets and sharply stabilizes training: at $N=3$, cross-run standard deviation falls from 3.11 to 0.49 for Indian English and from 2.71 to 0.09 for Korean English. The useful range differs by accent. For Indian English, adding 500 synthetic utterances helps mainly at $N=3-7$ (19.48%→16.81% at $N=3$), after which real-only adaptation quickly overtakes the fixed mixture near $N\approx 8$. For Korean English, the improvement is larger and persists longer: 20.40%→15.83% at $N=5$, with crossover only near $N\approx 25$.

These crossovers suggest a mixing effect rather than a universal synthetic-data recipe. When real data are extremely scarce, synthetic utterances reduce the dependence on which particular real examples are sampled and lower variance

across runs. Once the real set becomes informative, however, a fixed 500-utterance synthetic pool can dilute the stronger signal from newly added real speech. The accent-specific crossover may reflect differences in both base-model mismatch and synthetic-data quality, but the main conclusion is that the synthetic-to-real ratio changes the outcome. We do not optimize mixture weighting here, but the results show that the synthetic-to-real ratio is itself an important design variable.

4.3. Supporting analyses

Cross-speaker transfer. Although the adapted TTS decoder generates only one speaker, ASR models trained on 500 synthetic utterances from that speaker generalize to unseen speakers of the same accent (Table 2). This indicates that the synthetic set contains accent-relevant variation useful beyond the adaptation speaker.

Table 2. Per-speaker ASR generalization results. Values are WER (%) \pm standard deviation.

Speaker	Zero-shot	Adapt-only	Adapt + LLM	Real
Indian English				
TNI	25.4	19.6 \pm 0.5	18.9 \pm 0.4	12.2 \pm 0.2
RRBI	25.3	15.5 \pm 0.3	14.6 \pm 0.4	10.5 \pm 0.1
SVBI	26.1	17.4 \pm 0.4	17.1 \pm 0.7	11.4 \pm 0.3
Korean English				
HKK	28.1	22.2 \pm 0.5	20.3 \pm 0.4	14.1 \pm 0.3
YDCK	20.5	15.3 \pm 0.4	13.2 \pm 0.2	11.6 \pm 0.2
YKWK	21.5	15.6 \pm 0.6	13.9 \pm 0.2	12.3 \pm 0.3

Synthetic-data sample efficiency. Figure 4 studies how much target-accent reference data is needed to create the synthetic set. Under joint adaptation, quality improves sharply from $K=1$ to $K=3$ (AccSim 0.67→0.74; UTMOS 1.85→2.46) and then stays near plateau, so three reference utterances are sufficient in this setup. Decoder fine-tuning is the most sample-sensitive sweep, whereas speaker and style embeddings change only mildly across K . The LLM curve is nearly flat from $K=0$ to 15, suggesting that pronunciation editing can rely largely on pretrained accent knowledge and reasoning even without in-context examples.

5. Conclusion

Few-shot synthetic accented speech improves ASR fine-tuning largely through robustness to pronunciation-space perturbation. In our experiments, much of the gain comes from phoneme-space augmentation, while structured target-accent edits and oracle prosody provide only modest additional benefit beyond random perturbation. In mixed training, synthetic data is valuable because it reduces variance when real data are extremely scarce, yet a synthetic-heavy mixture can dilute increasingly informative real data; choosing the real-synthetic ratio is therefore central. Supporting

analyses show that single-speaker synthetic training transfers within accent and that the generator can operate from as little as three reference utterances, with LLM editing remaining robust even without in-context examples. Future work should extend these controls to more accents and learn better strategies for mixing real and synthetic speech.

Acknowledgements

This research used the DeltaAI advanced computing and data resource, which is supported by the National Science Foundation (award OAC 2320345) and the State of Illinois. DeltaAI is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications.

Impact Statement

This work studies why, and under what conditions, synthetic accented speech can improve ASR fine-tuning for accented speakers. The intended impact is to support more inclusive speech recognition systems by reducing performance gaps for speakers whose accents are underrepresented in existing datasets. By analyzing when synthetic data helps, when it mainly acts as pronunciation-space augmentation, and when it should be balanced with real accented speech, the paper aims to make accent-focused ASR adaptation more reliable and data-efficient. At the same time, synthetic accented speech should be used carefully: accent groups are heterogeneous, and simplified synthetic representations may fail to capture the full diversity of real speakers. Responsible deployment therefore requires evaluation across speakers and accents, transparent reporting of data sources and limitations, and attention to possible bias amplification.

References

- Baevski, A., Zhou, H., Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations, 2020. URL <https://arxiv.org/abs/2006.11477>.
- Das, N., Bodapati, S., Sunkara, M., Srinivasan, S., and Chau, D. H. Best of both worlds: Robust accented speech recognition with adversarial transfer learning, 2021. URL <https://arxiv.org/abs/2103.05834>.
- Du, C., Guo, Y., Chen, X., and Yu, K. Speaker adaptive text-to-speech with timbre-normalized vector-quantized feature. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:3446–3456, 2023. doi: 10.1109/TASLP.2023.3308374.
- Feng, S., Kudina, O., Halpern, B. M., and Scharenborg, O. Quantifying bias in automatic speech recognition, 2021. URL <https://arxiv.org/abs/2103.15122>.
- Goronzy, S., Rapp, S., and Kompe, R. Generating non-native pronunciation variants for lexicon adaptation. *Speech Communication*, 42(1):109–123, 2004. ISSN 0167-6393. doi: <https://doi.org/10.1016/j.specom.2003.09.003>. URL <https://www.sciencedirect.com/science/article/pii/S0167639303001158>. Adaptation Methods for Speech Recognition.
- Inoue, S., Wang, S., Wang, W., Zhu, P., Bi, M., and Li, H. Macst: Multi-accent speech synthesis via text transliteration for accent conversion, 2025. URL <https://arxiv.org/abs/2409.09352>.
- Ito, K. and Johnson, L. The lj speech dataset, 2017. URL <https://keithito.com/LJ-Speech-Dataset/>.
- Jain, A., Upreti, M., and Jyothi, P. Improved accented speech recognition using accent embeddings and multi-task learning. In *Interspeech*, 2018. URL <https://api.semanticscholar.org/CorpusID:52187265>.
- Karakasidis, G., Robinson, N., Getman, Y., Ogayo, A., Al-Ghezi, R., Ayasi, A., Watanabe, S., Mortensen, D., and Kurimo, M. *Multilingual TTS Accent Impressions for Accented ASR*, pp. 317–327. 08 2023. ISBN 978-3-031-40497-9. doi: 10.1007/978-3-031-40498-6_28.
- Klumpp, P., Chitkara, P., Sari, L., Serai, P., Wu, J., Veliche, I.-E., Huang, R., and He, Q. Synthetic cross-accent data augmentation for automatic speech recognition, 2023. URL <https://arxiv.org/abs/2303.00802>.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Toups, C., Rickford, J. R., Jurafsky, D., and Goel, S. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689, 2020. doi: 10.1073/pnas.1915768117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1915768117>.
- Kominek, J. and Black, A. W. The cmu arctic speech databases. In *Proc. 5th ISCA Workshop on Speech Synthesis (SSW5)*, pp. 223–224, 2004. URL https://www.isca-archive.org/ssw_2004/kominek04b_ssw.html.
- Kong, J., Kim, J., and Bae, J. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis, 2020. URL <https://arxiv.org/abs/2010.05646>.

- Lee, J. Y., Bae, J., Mun, S., Lee, J., Lee, J.-H., Cho, H.-Y., and Kim, C. Hierarchical timbre-cadence speaker encoder for zero-shot speech synthesis. In *Interspeech*, 2023. URL <https://api.semanticscholar.org/CorpusID:260917414>.
- McAuliffe, M., Socolof, M., Mihuc, S., Wagner, M., and Sonderegger, M. Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi. In *Proc. Interspeech 2017*, pp. 498–502, 2017. doi: 10.21437/Interspeech.2017-1386.
- Nespoli, F., Barreda, D., and Naylor, P. A. Zero shot text to speech augmentation for automatic speech recognition on low-resource accented speech corpora, 2024. URL <https://arxiv.org/abs/2409.11107>.
- OpenAI. GPT-5.1 Model Documentation, 2025. Online, Available from: <https://platform.openai.com/docs/models>.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. Robust speech recognition via large-scale weak supervision, 2022. URL <https://arxiv.org/abs/2212.04356>.
- Ravaneli, M., Parcollet, T., Rouhe, A., Plantinga, P., Rastorgueva, E., Lugosch, L., Dawalatabad, N., Heba, A., Grondin, F., Arisoy, E., and et al. Speechbrain: A general-purpose speech toolkit. In *Proc. Interspeech*, pp. 96–100, 2021.
- Saeki, T., Xin, D., Nakata, W., Koriyama, T., Takamichi, S., and Saruwatari, H. Utmos: Utokyo-sarulab system for voicemos challenge 2022, 2022. URL <https://arxiv.org/abs/2204.02152>.
- Schaden, S. Rule-based lexical modelling of foreign-accented pronunciation variants. *EACL '03*, pp. 159–162, USA, 2003. Association for Computational Linguistics. ISBN 1111567890. doi: 10.3115/1067737.1067773. URL <https://doi.org/10.3115/1067737.1067773>.
- Shi, X., Yu, F., Lu, Y., Liang, Y., Feng, Q., Wang, D., Qian, Y., and Xie, L. The accented english speech recognition challenge 2020: Open datasets, tracks, baselines, results and methods, 2021. URL <https://arxiv.org/abs/2102.10233>.
- Sun, S., Yeh, C.-F., Hwang, M.-Y., Ostendorf, M., and Xie, L. Domain adversarial training for accented speech recognition, 2018. URL <https://arxiv.org/abs/1806.02786>.
- Talkin, D. Reaper: Robust epoch and pitch estimator. GitHub repository, 2015. URL <https://github.com/google/REAPER>.
- Winata, G. I., Cahyawijaya, S., Liu, Z., Lin, Z., Madotto, A., Xu, P., and Fung, P. Learning fast adaptation on cross-accented speech recognition. In *Interspeech*, 2020. URL <https://api.semanticscholar.org/CorpusID:211990418>.
- Zaïdi, J., Seuté, H., van Niekerk, B., and Carbonneau, M.-A. Daft-exprt: Cross-speaker prosody transfer on any text for expressive speech synthesis. In *Interspeech 2022*, pp. 4591–4595, 2022. doi: 10.21437/Interspeech.2022-10761.
- Zhang, M., Zhou, Y., Wu, Z., and Li, H. Zero-shot multi-speaker accent tts with limited accent data. In *2023 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 1931–1936, 2023. doi: 10.1109/APSIPAASC58517.2023.10317526.
- Zhang, Y., Zhang, Y., Halpern, B., Patel, T., and Scharenborg, O. Mitigating bias against non-native accents. pp. 3168–3172, 09 2022. doi: 10.21437/Interspeech.2022-836.
- Zhao, G., Sonsaat, S., Silpachai, A., Lucic, I., Chukharev-Hudilainen, E., Levis, J., and Gutierrez-Osuna, R. L2-arctic: A non-native english speech corpus. In *Proc. Interspeech*, pp. 2783–2787, 2018. doi: 10.21437/Interspeech.2018-1110.
- Zhong, J., Richmond, K., Su, Z., and Sun, S. Accentbox: Towards high-fidelity zero-shot accent generation. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, April 2025. doi: 10.1109/icassp49660.2025.10888332. URL <http://dx.doi.org/10.1109/ICASSP49660.2025.10888332>.
- Zhou, K., Sisman, B., Liu, R., and Li, H. Emotional voice conversion: Theory, databases and esd. *Speech Communication*, 137:1–18, 2022. ISSN 0167-6393.
- Zhou, X., Zhang, M., Zhou, Y., Wu, Z., and Li, H. Accented text-to-speech synthesis with limited data, 2023. URL <https://arxiv.org/abs/2305.04816>.
- Zhou, X., Zhang, M., Zhou, Y., Wu, Z., and Li, H. Multi-scale accent modeling and disentangling for multi-speaker multi-accent text-to-speech synthesis, 2025. URL <https://arxiv.org/abs/2406.10844>.

A. Supplementary Material

A.1. Implementation Details

Acoustic features and vocoder. Mel-spectrograms use 80 mel bins at 22,050 Hz, a 1024-point FFT, hop length 256, and mel range 0–8 kHz. Waveforms are synthesized using the universal HiFi-GAN vocoder (Kong et al., 2020).

Backbone TTS training. The backbone is trained for 72k iterations with Adam (batch size 128). The mel reconstruction loss uses ℓ_1 and ℓ_2 terms, each weighted by 1.0. Pitch and energy consistency losses use $\lambda_p=1.0$ and $\lambda_e=0.2$, and the FiLM regularization weight is 10^{-3} . Pitch consistency uses a 3-layer 1D convolutional network (kernel size 3, hidden size 256) that predicts frame-level $\log-F_0$ from mel-spectrograms. Dynamic speaker statistics are computed from randomly sampled subsets of $M \in \{1, \dots, 15\}$ utterances per speaker and refreshed every 2500 training steps.

Few-shot adaptation and alignment. We fine-tune the entire acoustic decoder for 600 steps at learning rate 0.001 using mel reconstruction loss only. Text-to-phoneme conversion and phoneme durations use Montreal Forced Aligner English resources with CMU Pronouncing Dictionary lookup and MFA’s pretrained English G2P fallback (McAuliffe et al., 2017). Frame-level pitch is extracted with REAPER; energy is computed from mel-spectrogram magnitude (Talkin, 2015). For Adapt+LLM, phoneme edits and required alignment adjustments are generated by an OpenAI GPT-5.1 model via prompt design and in-context examples (OpenAI, 2025). All ASR models are fine-tuned for 15 epochs with AdamW at 3×10^{-5} ; the wav2vec 2.0 feature extractor is frozen and SpecAugment is disabled.