



Text-only domain adaptation for end-to-end ASR using integrated text-to-mel-spectrogram generator

Vladimir Bataev^{1,2}, Roman Korostik^{1,3}, Evgeny Shabalin^{1,4}, Vitaly Lavrukhin¹, Boris Ginsburg¹

¹NVIDIA

²University of London, UK; ³ITMO University, Russia; ⁴Higher School of Economics, Russia
{vbataev, rkorostik, eshabalin, vlavrukhin, bginsburg}@nvidia.com

Abstract

We propose an end-to-end Automatic Speech Recognition (ASR) system that can be trained on transcribed speech data, text-only data, or a mixture of both. The proposed model uses an integrated auxiliary block for text-based training. This block combines a non-autoregressive multi-speaker text-to-mel-spectrogram generator with a GAN-based enhancer to improve the spectrogram quality. The proposed system can generate a mel-spectrogram dynamically during training. It can be used to adapt the ASR model to a new domain by using text-only data from this domain. We demonstrate that the proposed training method significantly improves ASR accuracy compared to the system trained on transcribed speech only. It also surpasses cascade TTS systems with the vocoder in the adaptation quality and training speed.

1. Introduction

Modern end-to-end automatic speech recognition (ASR) systems are superior to traditional HMM-DNN systems in both academic benchmarks and commercial applications [1]. But customizing such models, especially adapting to a new domain, is still challenging. The standard approach of using audio-text pairs from a new domain has a high cost of collecting and transcribing speech. In many scenarios, text-only adaptation is preferable due to the wide availability of text data. In traditional HMM-DNN systems, a WFST recognition graph is built using a statistical language model (LM), which can be easily constructed from the new text-only data and can significantly improve the performance on a new domain. End-to-end ASR systems can also benefit from an external language model [2], but this approach requires significantly more computational resources than greedy decoding.

To use text data for training or finetuning, audio can be synthesized from the text to use in a conventional training or finetuning pipeline [3, 4, 5, 6, 7] or the unpaired text input itself can be integrated into an end-to-end ASR neural system [8, 9, 10, 11, 12]. Furthermore, some works propose modifying end-to-end model architectures to make the decoder behave like an actual language model for encoder-decoder with attention [13] and RNN-Transducer [14] models, and by such behavior enable text-only adaptation for the decoder part. These approaches require changing the training pipeline significantly, in most cases making it incompatible with existing models. Using a pretrained text-to-speech (TTS) system to synthesize audio for audio-text pairs is simpler and does not require modification of existing models.

The main disadvantages of TTS-based ASR customization are a requirement for large storage for generated data, the high computational cost of generating speech, and a mismatch be-

tween natural and synthetic audio. By generating speech on the fly using a multi-speaker TTS system, it is possible to produce a practically infinite amount of data without limitations of required space. Each text can be generated with different speakers, and the number of speakers can also be unlimited by sampling from speaker embedding space [15]. Conversely, dynamic audio generation using the complete TTS system will significantly slow down the training process and impact used memory.

Modern TTS pipelines are usually composed of a mel spectrogram generator followed by a vocoder to transform the spectrogram into a speech signal, and the vocoder is usually the most computationally intensive part of the TTS pipeline. Since ASR models can also use mel spectrogram features as input, the vocoder part can be omitted. On the other side, training on spectrograms computed from vocoder output leads to higher ASR performance [4, 16].

Some solutions are proposed to solve the mismatch problem, e.g., training ASR and TTS models with shared components [17] or with consistency loss [18, 19, 20, 21], using discrete representations instead of mel spectrogram [22] or training an additional input block for the combination of ASR and TTS system to improve performance with a frozen ASR backbone [23]. Most of these methods make ASR and TTS models dependent on each other, and we are trying to avoid this. The rejection sampling algorithm can be used to improve distribution mismatch [24], but this exacerbates the problem of computing resources. Also, recently an approach of using an additional block for text-to-mel-spectrogram generator was presented in [16] with 5-10% relative WER improvements.

In our work, we are building a lightweight modular text-only adaptation system based on a text-to-mel-spectrogram generator¹. We augment an end-to-end ASR model with an additional module for the on-the-fly generation of spectrograms from the text during training. We use FastPitch [25] modified to produce synthetic mel spectrograms with the same STFT parameters that are used in the ASR front-end. This approach does not significantly affect the training speed and is fully compatible with existing pretrained models. We address the problem of the mismatch between generated and real spectrograms by applying a small GAN-based enhancer directly to the generated spectrograms. Our main contributions are:

- We extend the text-to-spectrogram (TTS) module with an additional block based on StyleGAN 2 to mitigate the mismatch between real and synthetic spectrograms.
- We show that fine-tuning the ASR model on combined

¹We limit the investigated task to adaptation to a new text domain. We are building a modular system, therefore training ASR and TTS models together is undesirable. We also assume that the model will be applied in the target domain and do not mitigate performance drop on the source domain.

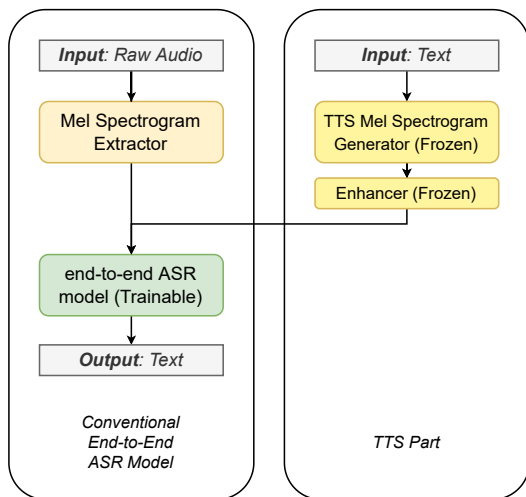


Figure 1: ASR system with text-to-mel-spectrogram frontend

speech-text and text-only datasets using the proposed system leads to a significant accuracy improvement. We also demonstrate that one can improve the acoustic end-to-end ASR model by utilizing textual data, usually used for training a language model in hybrid HMM-DNN models.

Models and code are released in the NeMo [26] framework.

2. ASR Model with integrated text-to-spectrogram front-end

To enable text-only adaptation, we add a text-to-spectrogram front-end to the ASR model as shown in Fig. 1. The model can take text or audio as input during training. If the input is audio, the mel spectrogram is extracted and passed directly to the ASR network. If the input is text, the mel spectrogram is produced on the fly using the pretrained frozen mel spectrogram generator and then fed into the ASR model, as in the previous case. At inference time, only the standard speech front-end is used.

2.1. Text-to-spectrogram frontend

We use multi-speaker FastPitch [25] as our base text-to-spectrogram model. Neural TTS models produce blurry spectrograms, so we add an Enhancer block to add fine details to synthesized spectrograms. Our Enhancer is based on the StyleGAN 2 [27] architecture².

We modify the original StyleGAN 2 to operate on 80-band mel spectrograms of arbitrary length L , which are treated as grayscale images. The network starts from a $5 * L / 16$ fixed random image and outputs an $80 * L$ detailed spectrogram. We add an appropriately downsampled and broadcasted spectrogram to the input and output of each generator block. This way, the generative process for the residual becomes spatially conditioned on the input spectrogram, and the network learns only to generate the details. Within the discriminator, we average across the time axis before projecting to logits.

The Enhancer is trained adversarially. We generate blurry spectrograms by passing TTS training data through the corresponding FastPitch model using ground-truth F_0 and speaker IDs. Blurry spectrograms passed through the generator are

²As a starting point, we used StyleGAN 2 implementation from <https://github.com/lucidrains/stylegan2-pytorch>

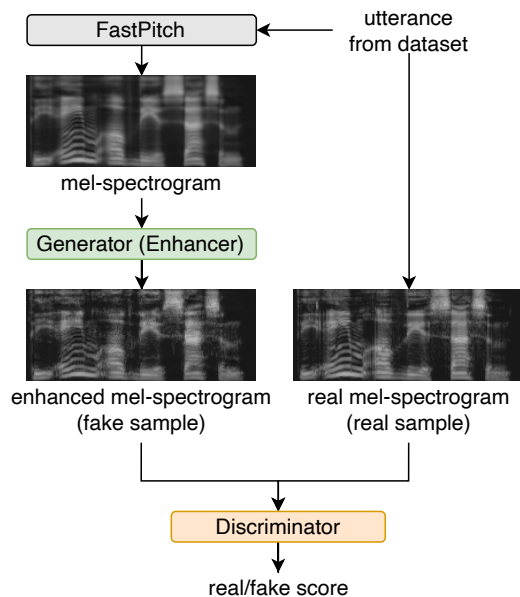


Figure 2: Spectrogram Enhancer adversarial training setup

considered “fake” and real spectrograms are considered “real” (Fig. 2). The training procedure is the usual alternation between gradient steps for the discriminator and the generator. Both are trained using hinge loss [28]. The gradient penalty loss [29, 27] is used every 4 steps for the discriminator.

Ablation runs showed that the enhancer is prone to generating artifacts. To fix this, we add a consistency loss during generator training: L1 distance between real and fake spectrograms, both downsampled 4x along the frequency axis.

2.2. ASR backbone

In the proposed pipeline, any end-to-end model can be used as an ASR backbone, which takes in a spectrogram and produces the corresponding transcription. For experiments we use Conformer-L [30] (120M), and Conformer-M (32M parameters) Transducers³ implemented with NeMo [26]. The encoder is a convolution-augmented transformer network; the decoder is a single-layer LSTM network with 640 hidden units. The encoder contains either BatchNorm (BN) [31] layers following the original architecture, or LayerNorm (LN) [32]. Since BN layers in inference mode use statistics accumulated at training time, due to the mismatch between synthetic and real audio additional issues are introduced when using synthetic audio for finetuning. To avoid this, we find it useful either to use models with LayerNorm, which normalizes the input samples independently and doesn’t have a mismatch between inference and training, or to fuse the BatchNorm layer (Fused BN) into a trainable projection. This is equivalent to removing BN (it can be further fused with a convolutional layer), but it allows us to use pretrained BN-based models for adaptation. Separating statistics in BN for real and synthetic data was proposed in [24], but this approach can be applied only when the real data is present. We believe that fusing BN is a more universal alternative for finetuning since we did not observe any meaningful quality degradation between pure and fused BN even for real audio.

³https://github.com/NVIDIA/NeMo/blob/main/examples/asr/conf/conformer/conformer_transducer_bpe.yaml

Table 1: **Training overhead** for Conformer-M on text vs real audio. Relative average time per training batch, measured on synthetic LibriSpeech data, NVIDIA V100 GPU.

Model	Input	Relative Training Time
Conformer-M	Audio	1x
+ FastPitch	Text	1.03x
+ FastPitch + Enhancer	Text	1.10x
+ FastPitch + Vocoder	Text	1.35x

3. Experiments

3.1. Experimental setup

Conformer was trained and finetuned with the AdamW [33] optimizer with a weight decay of 10^{-3} . In training from scratch, we use the Noam annealing learning rate (LR) scheduler [34] and global batch size of 2048 and train the model for 600 epochs (1920 GPU hours). For finetuning, we use the Cosine Annealing [35] LR scheduler with a warmup of 20% of training steps, and maximum LR of 10^{-4} . We finetune Conformer-Medium for 10K (290 GPU hours without enhancer) and Large for 15K steps (420 GPU hours) on SLURP [36] data with global batch size of 512. For finetuning with WSJ [37] data we use batch size of 256, and finetune for 20K steps on usual training data (160 GPU hours), and for 40K when mixing texts and audio-text pairs. In all scenarios SpecAugment [38] is applied for both real and synthetic data. All experiments are done on the cluster of NVIDIA DGX Stations. Conformer models are trained on 4 nodes, each containing 16 V100 GPUs.

For setup with Enhancer, we train FastPitch using NeMo toolkit [26] with configuration `fastpitch_align.v1.05` (50.7 M parameters). The Enhancer has a latent dimension of 192, depth of the style network is 4, "network capacity" is 16, and the upper bound on the number of feature maps in convolutional layers is 192. The generator has 3.5M parameters, and the discriminator has 4.5M parameters. We trained model on 8 V100 GPUs for 20 epochs using Adam [39] optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.9$ and LR of $2 \cdot 10^{-4}$. Batch size is 128, and the consistency loss has a weight of 0.1. We use greedy decoding for WER evaluation. All experiments done on LibriTTS train-clean-100 (53 hours) and train-960 (585 hours) subsets. After filtering out utterances longer than 15s, we are left with 45 hours and 468 hours respectively. Spectrogram parameters follow ASR: 25ms window, and 10ms hop, with mel bands ending at 8 kHz.

For the text-to-waveform TTS setup we use FastPitch with UnivNet [40] vocoder trained on LibriTTS clean-100 subset. Intermediate acoustic features are 80-band mel-spectrograms, calculated using 1024-sample windows with 256-sample hop calculated from a 22.05 kHz signal. Mel bands end at 8 kHz. Effectively this means ~ 46 ms window size with ~ 12 ms hop.

The Enhancer-based TTS frontend is much faster than the Conformer models. During training, we observe an overhead of $\leq 10\%$ for Conformer-M (Table 1). The benchmarking was done on V100 GPU with batch size of 16.

3.2. Training ASR model using synthetic only data

We start our experiments with a synthetic but representative setup, to demonstrate the problem of mismatch between real and synthetic data for training ASR models. In such a setup, a model trained only on real audio typically fails to generalize, requiring a mixture of real and synthetic audio to achieve

Table 2: **Text-only training on LibriSpeech texts.** Model: Conformer-M with LN. Greedy WER[%].

Training setup	dev		test	
	clean	other	clean	other
TTS 45h	27.0	59.2	27.6	61.8
+ vocoder	16.3	44.5	17.3	46.2
+ enhancer	11.6	38.9	12.6	40.8
TTS 468h	16.5	43.2	17.1	45.5
+ enhancer	8.9	22.3	9.3	22.1
Oracle (real audio-text)	2.7	6.6	2.9	6.6

Table 3: **Finetuning model with LayerNorm on SLURP text-only data.** Base model: Conformer-M with LN, trained on LibriSpeech. Greedy WER[%].

Model	dev	test
base: Conformer-M with LN	48.9	49.3
+ SLURP texts \rightarrow TTS 45h	37.5	38.1
+ vocoder	36.9	37.9
+ enhancer	36.1	36.8
+ SLURP texts \rightarrow TTS 468h	33.9	34.7
+ enhancer	31.3	32.4

good performance [3]. We pretrain the FastPitch using subsets of LibriTTS [41], which consists of the same speakers and mostly the same data as LibriSpeech. We then only use text from the LibriSpeech [42] audio-text paired dataset to generate spectrograms with random speakers, to train a Conformer-M model with LayerNorm. We can see that using more data for TTS model leads to a significantly lower WER on both "clean" and "other" parts of the dataset (see Table 2), but the difference compared to real audio (6.3% vs 45.5% WER on test-other) is still considerable. Using the enhancer model to generate more realistic spectrogram allows to achieve significantly better results (22.1% WER on test-other), and also surpasses cascade TTS system with vocoder.

3.3. Text-only finetuning on SLURP dataset

We use the SLURP [36] dataset to study text-only adaptation using the proposed approach. SLURP is a spoken language understanding dataset, which contains diverse verbose commands for smart home control. Since its text significantly differs from audio books and conventional dialog data, it is suitable for investigating text-only adaptation.

Base models are trained on LibriSpeech data, and we use text-only data from SLURP training set (~ 11 K utterances) to adapt the model. We report WER for the original dev and test subsets containing each ~ 10 h of text-audio pairs. Similar to the previous setup, ablation studies in Table 3 with finetuning Conformer-M with LN show improvement from using more data for mel spectrogram generator (train-clean-100 vs full LibriTTS), and improvement from the enhancer.

Finetuning models with BN is more challenging. As described in Section 2.2, we replace BN with a trainable projection initialized from the original layer parameters. Experiments in Table 4 show that finetuning Conformer-M with fused BN on text-only data leads to larger WER improvements. The relative impact of the enhancer is larger on models with pure BN, but

Table 4: *Finetuning models with BatchNorm on SLURP text-only data. Base model: Conformer-M with BN, trained on LibriSpeech. Greedy WER[%].*

Model	dev	test
base: Conformer-M with BN	49.6	49.8
+ SLURP texts	36.6	37.5
+ fused BN	34.2	35.1
+ enhancer	33.1	33.9
+ enhancer + fused BN	31.9	32.5

Table 5: *Finetuning medium and large models on SLURP text-only data. Base models: Conformer-M and Conformer-L with fused BN, trained on LibriSpeech. Greedy WER[%].*

Model	dev	test
base: Conformer-M	49.6	49.8
+ SLURP texts + enhancer	31.9	32.5
base: Conformer-L	46.7	47.1
+ SLURP texts + enhancer	27.0	27.7

the system with fused BN has better performance and is comparable with LN-based models.

Table 5 shows that larger ASR model size leads to better relative improvement, and the system trained on LibriSpeech with the baseline WER of 47.1% is able to achieve 27.7% WER on SLURP test (41.2% relative improvement).

Text-only adaptation helps even when the ASR model has been trained on a large amount of speech data. For this scenario we finetuned Conformer-L pretrained on a large corpus of 24K hours of read and conversational speech data (not including SLURP) starting from a publicly available checkpoint⁴. With text-only finetuning the system is able to achieve 15.0% WER (36.2% relative improvement), see Table 6.

3.4. Finetuning on WSJ dataset

To compare the contributions of using original audio-text pairs and synthetic data whilst using a large text corpus for improving the model, we use the WSJ [37] corpus. The audio conditions are similar to LibriSpeech, however the domain of text is different. To compare the text-only adaptation with text-audio pairs we use text-only training data. The results are in Table 7. Our system achieves 47% relative improvement on eval-92 using only texts, and surpasses text-only adaptation of MHAT [14], which achieves 24% WERR (from 9.5% to 7.2%) in the same scenario. The quality of text-only adaptation is still below than the quality of training on real data, but the difference is small. Also, we observe that the proposed method can not only re-

Table 6: *Finetuning model pretrained on large ASR corpus on SLURP text-only data. Base: Conformer-L with BN, trained on a large (24K hours) audio-text dataset. Greedy WER[%]*

Model	dev	test
base: Conformer-L	23.2	23.5
+ SLURP texts + enhancer	16.1	16.7
+ fused BN	14.7	15.0

⁴https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_transducer_large

Table 7: *Finetuning on WSJ: audio-text pairs vs text-only data. Base model: Conformer-L with fused BN trained on LibriSpeech. Greedy WER[%]*

Model	dev-93	eval-92
base: Conformer-L	9.3	7.2
texts	5.5	3.8
audio-text pairs	4.2	3.1

Table 8: *Finetuning using WSJ audio-text pairs and text from WSJ Language Model corpus. Base model: Conformer-L with fused BN trained on LibriSpeech. Greedy WER[%].*

audio:text ratio	dev-93	eval-92
base: Conformer-L	9.3	7.2
1:1	2.7	1.7
1:2	2.5	1.5

place adaptation on real data with using text-only samples, but it can gain benefits by combining them both. WSJ contains also a large LM corpus, used in traditional HMM-DNN systems to build a speech recognition graph based on a statistical LM. When using synthetic mel spectrogram from this text, we observe significant improvement compared to using speech-only data, achieving 1.5% WER on eval-92 (Table 8). Each training epoch combines speech data and sampled texts from the LM corpus with the shown ratio.

4. Conclusions

In this paper, we presented an end-to-end ASR model which can be trained both on text-audio pairs and text-only data. We enable text-only training using an integrated neural text-to-spectrogram module, composed of a modified non-autoregressive text-to-mel-spectrogram generator with a lightweight GAN-based spectrogram enhancer. Such a model does not require external storage for generated audio and also outperforms the TTS system with a vocoder in both speed and quality. The amount of generated synthetic data can be theoretically infinite when using a multi-speaker text-to-mel-spectrogram model. The Enhancer helps to mitigate the mismatch between real and synthetic mel spectrograms. This leads to significant improvement of ASR accuracy in text-only adaptation.

By using text-only LibriSpeech data to train ASR model from scratch, we show the effectiveness of the proposed approach. We also show up to the 41.2% relative improvement on SLURP data in the text-only scenario. On the WSJ dataset, we demonstrate the effectiveness of our proposed system by combining a large text corpus with text-audio pairs and achieving 1.5% WER on the test set.

We used Conformer as ASR backbone and FastPitch as text-to-mel-spectrogram generator. Since the proposed architecture is modular, the approach should generalize to other ASR and TTS models.

5. Acknowledgments

We thank our colleagues Aleksandr Laptev, Sean Narenthiran, Jocelyn Huang, and Elena Rastorgueva for the help with the code and paper review.

6. References

- [1] J. Li, "Recent advances in end-to-end automatic speech recognition," *APSIPA Transactions on Signal and Information Processing*, 2022.
- [2] A. Laptev, S. Majumdar, and B. Ginsburg, "CTC variations through new WFST topologies," in *Interspeech*, 2022.
- [3] J. Li, R. Gadde, B. Ginsburg, and V. Lavrukhin, "Training neural speech recognition systems with synthetic speech augmentation," *arXiv:1811.00707*, 2018.
- [4] A. Laptev, R. Korostik, A. Svischev, A. Andrusenko, I. Medennikov, and S. Rybin, "You do not need more data: Improving end-to-end speech recognition by text-to-speech data augmentation," in *CISP-BMEI*, 2020.
- [5] J. Li, R. Zhao, Z. Meng, Y. Liu, W. Wei, S. Parthasarathy, V. Mazalov, Z. Wang, L. He, S. Zhao, and Y. Gong, "Developing RNN-T Models Surpassing High-Performance Hybrid Models with Customization Capability," in *Interspeech*, 2020.
- [6] X. Zheng, Y. Liu, D. Gunceler, and D. Willett, "Using synthetic audio to improve the recognition of out-of-vocabulary words in end-to-end ASR systems," *ICASSP*, 2020.
- [7] N. Robinson, P. Ogayo, S. Gangu, D. R. Mortensen, and S. Watanabe, "When is TTS augmentation through a pivot language useful?" in *Interspeech*, 2022.
- [8] T. N. Sainath, R. Pang, R. Weiss, Y. He, C.-C. Chiu, and T. Strohmman, "An attention-based joint acoustic and text on-device end-to-end model," *ICASSP*, 2020.
- [9] S. Thomas, B. Kingsbury, G. Saon, and H.-K. J. Kuo, "Integrating text inputs for training and adapting RNN Transducer ASR models," in *ICASSP*, 2022.
- [10] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. J. Moreno, A. Bapna, and H. Zen, "Maestro: Matched speech text representations through modality matching," in *Interspeech*, 2022.
- [11] H. Sato, T. Komori, T. Mishima *et al.*, "Text-only domain adaptation based on intermediate CTC," *Interspeech*, 2022.
- [12] A. Mittal, S. Sarawagi, and P. Jyothi, "In-situ text-only adaptation of speech models with low-overhead speech imputations," in *The Eleventh International Conference on Learning Representations*, 2023.
- [13] V. T. Pham, H. Xu, Y. Khassanov, Z. Zeng, C. E. Siong, C. Ni, B. Ma, and H. Li, "Independent language modeling architecture for end-to-end ASR," *ICASSP*, 2019.
- [14] Z. Meng, T. Chen, R. Prabhavalkar, Y. Zhang, G. Wang, K. Audhkhasi, J. Emond, T. Strohmman, B. Ramabhadran, W. R. Huang, E. Variani, Y. Huang, and P. J. Moreno, "Modular hybrid autoregressive transducer," *SLT*, 2022.
- [15] D. Stanton, M. Shannon, S. Mariooryad, R. J. Skerry-Ryan, E. Battenberg, T. Bagby, and D. Kao, "Speaker generation," *ICASSP*, 2021.
- [16] S. Ueno and T. Kawahara, "Phone-informed refinement of synthesized mel spectrogram for data augmentation in speech recognition," in *ICASSP*, 2022.
- [17] S. Karita, S. Watanabe, T. Iwata, M. Delcroix, A. Ogawa, and T. Nakatani, "Semi-supervised end-to-end speech recognition using text-to-speech and autoencoders," *ICASSP*, 2019.
- [18] T. Hori, R. Astudillo, T. Hayashi, Y. Zhang, S. Watanabe, and J. Le Roux, "Cycle-consistency training for end-to-end speech recognition," in *ICASSP*, 2019.
- [19] G. Wang, A. Rosenberg, Z. Chen *et al.*, "Improving speech recognition using consistent predictions on synthesized speech," in *ICASSP*, 2020.
- [20] M. K. Baskar, L. Burget, S. Watanabe, R. F. Astudillo, and J. H. Cernocky, "Eat: Enhanced ASR-TTS for self-supervised speech recognition," *ICASSP*, 2021.
- [21] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. J. Moreno, and G. Wang, "Tts4pretrain 2.0: Advancing the use of text and speech in ASR pretraining with consistency and contrastive losses," *ICASSP*, 2022.
- [22] S. Ueno, M. Mimura, S. Sakai, and T. Kawahara, "Data augmentation for ASR using TTS via a discrete representation," *ASRU*, 2021.
- [23] G. Kurata, G. Saon, B. Kingsbury, D. Haws, and Z. Tüske, "Improving customization of neural transducers by mitigating acoustic mismatch of synthesized audio," in *Interspeech*, 2021.
- [24] T. Yao Hu, M. Armandpour, A. Shrivastava, J.-H. R. Chang, H. S. Koppula, and O. Tuzel, "Synt++: Utilizing imperfect synthetic data to improve speech recognition," *ICASSP*, 2021.
- [25] A. Łańcucki, "FastPitch: Parallel text-to-speech with pitch prediction," in *ICASSP*, 2021.
- [26] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Krivan, S. Beliaev, V. Lavrukhin *et al.*, "Nemo: a toolkit for building ai applications using neural modules," in *NeurIPS Workshop on Systems for ML*, 2019.
- [27] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of styleGAN," in *CVPR*, 2020.
- [28] J. H. Lim and J. C. Ye, "Geometric GAN," *arXiv:1705.02894*, 2017.
- [29] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for GANs do actually converge?" in *ICML*, 2018.
- [30] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Interspeech*, 2020.
- [31] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015.
- [32] J.-L. Ba, J.-R. Kiros, and G.-E. Hinton, "Layer normalization," *arXiv:1607.06450*, 2016.
- [33] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, 2017.
- [35] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," in *ICLR*, 2017.
- [36] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, "SLURP: A spoken language understanding resource package," in *EMNLP*, 2020.
- [37] D. B. Paul and J. Baker, "The design for the Wall Street Journal-based CSR corpus," in *Speech and Natural Language Workshop*, 1992.
- [38] 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," *Interspeech*, 2019.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [40] W. Jang, D. C. Y. Lim, J. Yoon, B. Kim, and J. Kim, "Univnet: A neural vocoder with multi-resolution spectrogram discriminators for high-fidelity waveform generation," in *Interspeech*, 2021.
- [41] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech," in *Interspeech*, 2019.
- [42] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: an ASR corpus based on public domain audio books," in *ICASSP*, 2015.