

A Unit-based System and Dataset for Expressive Direct Speech-to-Speech Translation

Anonymous submission to Interspeech 2024

Abstract

Current research in speech-to-speech translation (S2ST) primarily concentrates on translation accuracy and speech naturalness, often overlooking key elements like paralinguistic information, which is vital for conveying emotions and attitudes in human communication. This oversight is particularly significant in applications requiring expressive speech, such as video dubbing. To remedy these gaps, our research introduces a novel, meticulously assembled multilingual dataset from various movie audio tracks. Each dataset pair is carefully matched for paralinguistic content and duration. We enhance this approach by integrating multiple prosody transfer techniques, aiming for translations that are not only accurate and natural-sounding but also rich in paralinguistic detail. Our experimental results confirm that our model successfully retains more paralinguistic information from the source speech while upholding high standards of translation accuracy and naturalness.

Index Terms: Expressive speech-to-speech translation, controllable text-to-speech, prosody transfer

1. Introduction

Speech-to-speech translation (S2ST) enables the translation of spoken language into another spoken language, significantly enhancing communication between different language speakers. Traditional S2ST systems [1, 2, 3] rely on a pipeline of automatic speech recognition (ASR), machine translation (MT), or speech-to-text translation (S2T), followed by text-to-speech synthesis (TTS). While speech translation traditionally involves converting speech to text or vice versa in different languages, recent developments have shifted towards an end-to-end S2ST system [4, 5, 6, 7, 8]. These systems minimize error propagation between ASR and MT, resulting in a streamlined process with reduced computational costs, particularly advantageous for languages without a written form.

In the realm of expressive speech-to-speech translation (S2ST) [9, 5, 10, 3, 11], some research focuses on intonation transfer, utilizing statistical word alignment to transfer source intonation characteristics to the target language. Many methods [12, 13, 14] evolved to include word emphasis transfer, ultimately leading to sequence-to-sequence models for simultaneous emphasis and content translation. Despite the progress, these approaches only focused on individual expression elements and did not fully capture the emotional aspects of speech.

However, a notable challenge in direct S2ST, especially with style transfer, is the scarcity of paired data where the source and target speech have the same speaker. There are some works [15, 16] developing methods to utilize non-parallel data. However, they lack ground truth for human evaluation, limiting further study. To address this, we introduce a novel,

carefully curated multilingual dataset from diverse movie audio tracks. This dataset, primarily consisting of paired Spanish-English data from clear, emotionally rich dialogues in movies and TV shows, offers a unique opportunity to capture nuanced emotional variations often missed in standard speech synthesis data.

To extract fine-grained emotional information from these paired datasets, we propose a model structure that learns both global style features and local pitch features, recognizing that discrete speech representations may lose some information. In our acoustic unit modelling, we introduce a direct S2ST model that translates between languages without intermediate text, leveraging self-supervised learning representations. Our approach differs by simultaneously integrating multiple expressive aspects, representing a significant advancement in S2ST. This methodology facilitates real-time improvements and ensures duration alignment without relying on text intermediaries.

Our contributions are as follows:

- We introduce the first dataset designed for training paired speech emotion translation from movies and TV shows with multiple audio tracks.
- We propose a novel approach to direct S2ST with style transfer, integrating pitch and global style transferring. This method preserves intricate emotional characteristics without the need for text as an intermediary while maintaining translation accuracy.
- Our experiments demonstrate that our method produces high-quality translations while maintaining stylistic fidelity to the source speech.




Figure 1: Direct speech-to-speech translation system compared with cascaded speech-to-speech translation system

2. Movie Dataset

In this section, we detail the construction and processing of the movie dataset, which is crucial for advancing research in speech-to-speech translation between English and Spanish and sets the stage for future research in producing emotionally paired multilingual speech datasets.

2.1. Dataset Source

Our dataset, a substantial collection of approximately 300 hours of paired English-Spanish television series and movie audio, is carefully curated to facilitate advanced translation model development. Our dataset comprises content from the following sources: “Money Heist” seasons 1-5, “Elite” seasons 1-4, 59 Disney movies, “Dragon Ball Z”, 24 “James Bond” Collection movies, 10 superhero series movies, “Shrek” movies 1-4, “Harry Potter” movies 1-8, and “Poltergeist” movies 1-3. These videos contain both English and Spanish audio tracks, and the diligent dubbing by actors ensures alignment of emotion and timing. Our dataset maintains gender consistency between English actors and their Spanish dubbing counterparts. This uniformity ensures standardization, which is critical for voice recognition and translation models.

2.2. Dataset Construction

We initiate our dataset construction by converting subtitle SRT files into more structured CSV files. This conversion process was enhanced with sophisticated filtering rules aimed at eliminating irrelevant or inconsistent data. Additionally, we merge continuous sentences from the same speaker using [17] to confirm speaker consistency, consolidating them into a single data point. This enhances dataset coherence and contextual relevance, crucial for training models on realistic dialogue patterns and maintaining narrative continuity.

Next, we subject all audio files to a denoising process using a noise suppression library¹ based on a recurrent neural network, significantly improving audio clarity. This clarity is vital for the subsequent step of automatic speech recognition (ASR) using Azure², as cleaner audio leads to more accurate transcription. We select segments where the ASR output is closely aligned with the subtitles, choosing the top 80% of segments with a word error rate more significant than 0.6 to ensure dataset accuracy. Additionally, we carefully filter segments based on appropriate sentence lengths, optimizing the dataset for practical training while maintaining contextual richness. Specifically, we exclude segments shorter than 3 seconds or longer than 15 seconds. All audio is processed at the sampling rate of 16000.

The dataset was further enriched by extracting speaker (Spk) embeddings from the sentences and calculating the cosine similarities for these embeddings. These embeddings and their corresponding cosine similarities were systematically organized and saved in specific files. This organization not only enhances the accessibility of the data but also provides a rich source for analyzing and comparing vocal features across languages.

2.3. Dataset analysis

Additionally, we analyze the English audio segments. This involved matching English segments with corresponding Spanish segments within the same TV show, ensuring consistency in contextual and emotional content. We prioritize segments where the cosine similarity between the matched English and Spanish audio was below 0.5, ensuring diversity in the dataset that challenges and thus strengthens the robustness of the translation models. Furthermore, we impose a criterion that only lists at least five remaining segments that would be saved, ensuring a meaningful sample size for model training.

¹<https://github.com/xiph/rnnoise>

²<https://azure.microsoft.com/en-us/products/ai-services/ai-speech>




Figure 2: This is the length Distribution of the Utterances, and the yellow ones denote the utterances that have a word error rate under 40%. There are 12610 utterances in total. The maximum duration is 244.250s, while the minimum duration is 0.833s. The average duration of utterances is 5.096s.

By ensuring gender consistency, enhancing audio quality, meticulously matching segments, and enriching the dataset with detailed speaker embeddings and cosine similarity analyses, we provide a comprehensive and robust foundation for developing advanced translation models.

3. Method

To extract fine-grained emotional information from these paired datasets, we develop a model to transfer the emotion of the reference speech. Our S2ST system comprises three components. Initially, speech from one language is converted into discrete units for direct speech-to-speech translation. Subsequently, speaker identification is extracted from the speech. Illustrated in Figure 3, we introduce a unit-hifigan-based emotion transfer model, where speech in the target language, enriched with the appropriate emotions, is synthesized.

3.1. Obtain discrete Units

In the first stage, we extract discrete units using a process inspired by the HuBERT [18] framework as in [5], which employs self-supervised learning techniques for speech representation. HuBERT leverages K-means clustering on its intermediate representations or Mel-frequency cepstral coefficient (MFCC) features [19] in the initial iteration to categorize masked audio segments into discrete labels. By pre-training a HuBERT model on an unlabelled speech corpus in the target language, we can encode target speech into continuous representations for every 20ms frame. Subsequently, a K-means algorithm is applied to these representations to generate K cluster centroids. These centroids are instrumental in encoding target utterances into sequences of cluster indices at the same 20ms interval.

In our implementation, the target speech is encoded into a vocabulary of 1000 discrete units. The models for HuBERT and K-means are derived from a combination of unlabeled English, Spanish, and French speech data sourced from the Vox-Populi [20] corpus. Our focus is solely on encoding English and Spanish target speech.

The second phase involves processing the discrete units obtained from the first phase. Due to the high length of the original unit, we follow [5] to adopt a strategy to condense continuous repetitions of the same unit into a single unit. This approach not

only streamlines the dataset but also aids in reducing computational complexity. In the third and final phase, these condensed units are expanded back to their original form during the Unit-to-Waveform conversion process.

3.2. Get speaker ID

In our methodology, a critical step is the enhancement of our dataset through the extraction of speaker embeddings from each utterance. These embeddings capture the unique vocal characteristics of the speakers, which are crucial for our analysis. To further enrich this data, we compute the cosine similarities (cos sim) between these embeddings. This computation serves as a measure of similarity between different vocal features, facilitating a deeper understanding of language-specific vocal nuances.

We meticulously organize these speaker embeddings and their corresponding cosine similarities into dedicated files. This systematic organization not only improves the accessibility of the data but also creates a robust framework for comparative analysis across different languages. By arranging the data in this manner, we ensure that the vocal features can be easily retrieved and analyzed for further research purposes.

Utilizing this approach, we are able to identify and extract speaker IDs from the dataset effectively. The speaker ID provides a unique identifier for each speaker, enabling us to track and analyze individual vocal characteristics across different linguistic contexts. This is particularly valuable in speech-to-speech translation research, where understanding and preserving individual speaker characteristics is essential for generating accurate and natural translations.

The process of extracting speaker embeddings and computing their cosine similarities, therefore, plays a pivotal role in our methodology. It not only aids in the accurate identification of speakers but also contributes significantly to our understanding of how vocal features vary across languages. This understanding is crucial for the development of advanced speech-to-speech translation systems that are capable of handling the complexities inherent in human speech.

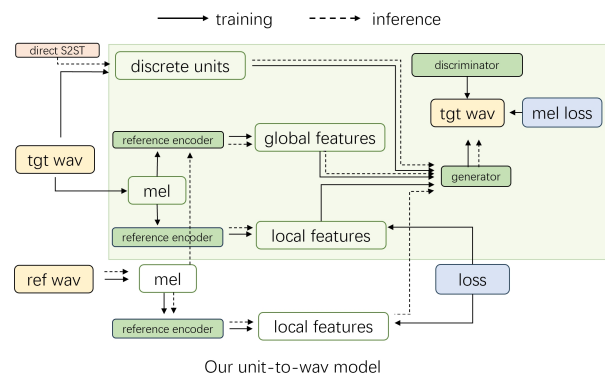


Figure 3: *unit-hifigan based, voice style transfer model at the training stage and inference stage.*

3.3. Unit2Wav synthesis

In the third part of our method, we focus on synthesizing speech in a different language from the voice and translating discrete representations. This presents two primary challenges: firstly, maintaining high audio quality after dataset denoising, and secondly, addressing the tonal differences between matched audio

in different languages, which create a gap that complicates the direct computation of Mel-spectral loss.

To tackle these issues, as shown in Figure 3, we employ a unit-based variant of HiFi-GAN [21], termed “unit-HiFiGAN”. The structure is inspired mainly by HiFi-GAN, which excels in handling signals of varying periodicity in speech by employing multiple smaller sub-discriminators. These sub-discriminators individually process different periodic patterns, resulting in superior performance. Additionally, this model architecture allows for parallel processing of these patterns, enhancing computational efficiency. Controllable Text-to-Speech (TTS) has developed in two main directions: global and fine-grained style transfer [22, 23, 24, 25]. Global style transfer, encapsulating overall speech attributes into a single embedding, contrasts with fine-grained style transfer, which captures local prosodic features but faces alignment challenges. Global style transfer is more adaptable to non-parallel scenarios, so we leverage it in our S2ST framework following [23].

Our innovation primarily lies in two areas. First, addressing the challenge of preserving high audio quality post-denoising, our model can be trained on high-quality monolingual datasets. This initial phase establishes a foundation for quality. Subsequently, we train the model and discriminator predictors on a mixed dataset, further refining the system.

Regarding the second challenge of tonal differences between matched audio in various languages, our approach includes predicting the speaker (spk) attributes from the outputs processed by the reference (ref) encoder. We then calculate a loss function based on this prediction, aiming to minimize the non-timbral features in the ref encoder’s output. Acknowledging the typically deeper tones in Spanish speech, we input and output both Spanish-to-English and English-to-Spanish translations to mitigate potential model biases. Additionally, we integrate a pitch predictor and an unvoiced/voiced predictor into our system.

These strategic innovations in our method not only address the inherent challenges in cross-lingual speech synthesis but also push the boundaries of what’s achievable in terms of audio quality and linguistic versatility.

4. Experiments

4.1. Experimental setup

4.1.1. Dataset setup

In addition to the movie dataset we introduced, to enhance the audio quality, we also utilize the supplementary training material from [20, 26, 27, 28], which comprises over 400K hours of high-quality audio. As they lack paired translated audio, we employ the same audio during training as a substitute for translated audio. We use English as the target language and Spanish as the reference language.

4.1.2. Human evaluation protocol

To measure expressivity preservation, we adopt the protocols [3] proposed, focusing on four specific aspects of expressiveness. This system, based on established methodologies in the field, categorizes expressiveness into four core aspects: emphasis, intonation, rhythm, and emotion. These categories were selected based on internal qualitative research, which pinpointed them as critical for preserving expressiveness in speech translation.

Among these aspects, emphasis, intonation, and rhythm are

related to more localized or prosodic features of speech. These elements play a crucial role in conveying the subtleties of spoken language. On the other hand, emotion is treated as the most “global” aspect, encompassing the overall feeling or mood conveyed by the speech. It’s important to note that while we assess naturalness in our translations, it is conducted in a separate study and thus not included as a core expressiveness aspect in this experimental setup.

To evaluate emotion expression and objectively rate the performance of our model across the identified expressivity aspects, we recruit 10 participants to assess multiple pairs of results containing our method and baseline outcomes through questionnaires. Each pair is evaluated using four different scores, and the average rating across all 10 participants is calculated. Their evaluations are crucial in providing an unbiased assessment of our model’s capability to preserve expressivity in speech-to-speech translation.

4.1.3. Model setup

In our expressive S2ST system, we use the open-source S2T model in the Fairseq toolkit [29, 30, 31]. We use the pretrained Es-En model they provide.

As for the unit-to-wav model, our model configuration uses several key parameters during training. We train from scratch, with settings of a learning rate of $2e^{-4}$ and learning rate decay of 0.999, inverse square root learning rate scheduler with warmup, adam optimizer with specific beta values of (0.8,0.99), 0.1 dropout probability in reference encoder and prosody encoder. These parameter choices crucially influenced the model’s training behavior.

For our baseline model, we utilize a vanilla implementation of the HiFi-GAN in [32], trained on the LJ Speech dataset [28]. This choice of baseline provides a reliable foundation for evaluating our model’s performance, given the HiFi-GAN’s proven effectiveness in generating high-quality speech audio.

4.2. Main Results

4.2.1. Preserve emotion

Our experimental results provide compelling evidence of the superiority of our model over traditional vanilla unit-based TTS systems, particularly in the realms of emphasis, intonation, and rhythm. These aspects are critical in achieving natural-sounding, expressive speech synthesis, a goal that has remained elusive in many existing TTS technologies.

Table 1: *The Cascade system’s performance on various aspects of speech*

System	Emotion	Emphasis	Intonation	Rhythm
Vanilla TTS	2.034	2.684	2.462	2.297
Holistic Cascade	3.576	3.257	3.173	3.562

Emotion and emphasis: The performance of our model in replicating the emotion and the emphasis in the speech was markedly superior. This was quantitatively measured using a set of metrics designed to capture the degree of emphasis correctly copied from the source material. Our model showed an improvement of 21% over traditional vanilla TTS, indicating a more dynamic and contextually accurate speech synthesis. Regarding the baseline, due to the lack of additional linguistic information in discrete units, the speech tones are nearly uniform.

Intonation: Intonation, a vital component in conveying emotions and questions in speech, was another area where our model excelled. Using a specialized intonation accuracy index, we observe that our model’s ability to mimic the natural intonation patterns of human speech surpassed that of vanilla TTS by 29%. This improvement is indicative of the model’s sophisticated understanding of speech patterns and its ability to generate more human-like, expressive speech.

Rhythm: In terms of replicating the natural rhythm of speech, our model again outperformed the vanilla unit-based TTS. The rhythm conformity score, which measures how closely the synthesized speech matches the rhythm of natural speech, was 55% higher in our model. This result underscores our model’s advanced capability in capturing and reproducing the subtle temporal characteristics of speech, which are essential for naturalness and expressiveness.

These results were further corroborated by subjective evaluations, where a panel of listeners rated our model’s outputs as significantly more natural and expressive compared to those generated by the vanilla TTS system.

4.2.2. Accuracy

Table 2: *The BLEU of different system’s output compared to the English groundtruth. “S2ST” denotes for using the translation result as the input of TTS/Unit2Wav.*

System	S2ST	GT input
Vanilla TTS	29.2	81.0
Ours	28.3	74.6
GT	-	78.2

To measure the ability of our model to keep the high audio quality while maintain both emotional expressions, we conduct Automatic Speech Recognition by Azure ¹ on the translations generated by our model and compare them to the ground truth results. The BLEU score achieved 74.6, whereas the ground truth yielded a BLEU score of 78.2. These results demonstrate the superior performance of our approach in terms of both fluency and accuracy in preserving the original content. The baseline vanilla TTS exhibits a higher BLEU score, attributed to the fact that the LJ Speech training set comprises passages read by a single speaker. In contrast, our movie dataset consists of numerous noisy clips. This disparity is reasonable and highlights a limitation of our method. Future research endeavours could work on this challenge.

5. Conclusion

Direct S2ST encounters a challenge due to data scarcity. To address this, we introduce the first training dataset for expressive speech translation and propose a model for emotion translation, advancing S2ST. Our approach integrates pitch and global style transfer, facilitating real-time improvements without relying on text intermediaries. This novel method promises high-quality translations while preserving stylistic fidelity to the source speech. The dataset comprises paired Spanish-English data from movie and TV show dialogues, capturing nuanced emotional variations often overlooked in conventional speech synthesis datasets. These contributions collectively push the boundaries of direct S2ST, paving the way for future developments in expressive speech synthesis and style transfer.

6. References

- [1] A. Lavie, A. Waibel, L. Levin, M. Finke, D. Gates, M. Gavalda, T. Zeppenfeld, and P. Zhan, "Janus-iii: Speech-to-speech translation in multiple languages," in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1. IEEE, 1997, pp. 99–102.
- [2] S. Nakamura, K. Markov, H. Nakaiwa, G.-i. Kikui, H. Kawai, T. Jitsuhiro, J.-S. Zhang, H. Yamamoto, E. Sumita, and S. Yamamoto, "The atr multilingual speech-to-speech translation system," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 2, pp. 365–376, 2006.
- [3] W.-C. Huang, B. Peloquin, J. Kao, C. Wang, H. Gong, E. Salesky, Y. Adi, A. Lee, and P.-J. Chen, "A holistic cascade system, benchmark, and human evaluation protocol for expressive speech-to-speech translation," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [4] A. Lee, P.-J. Chen, C. Wang, J. Gu, S. Popuri, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino, and W.-N. Hsu, "Direct speech-to-speech translation with discrete units," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, S. Muresan, P. Nakov, and A. Villavicencio, Eds. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 3327–3339. [Online]. Available: <https://aclanthology.org/2022.acl-long.235>
- [5] S. Popuri, P.-J. Chen, C. Wang, J. Pino, Y. Adi, J. Gu, W.-N. Hsu, and A. Lee, "Enhanced direct speech-to-speech translation using self-supervised pre-training and data augmentation," 2022.
- [6] Y. Jia, M. T. Ramanovich, T. Remez, and R. Pomerantz, "Translatotron 2: High-quality direct speech-to-speech translation with voice preservation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 10 120–10 134.
- [7] Y. Jia, R. J. Weiss, F. Biadsy, W. Macherey, M. Johnson, Z. Chen, and Y. Wu, "Direct speech-to-speech translation with a sequence-to-sequence model," *arXiv preprint arXiv:1904.06037*, 2019.
- [8] A. Lee, H. Gong, P.-A. Duquenne, H. Schwenk, P.-J. Chen, C. Wang, S. Popuri, Y. Adi, J. Pino, J. Gu *et al.*, "Textless speech-to-speech translation on real data," *arXiv preprint arXiv:2112.08352*, 2021.
- [9] A. Lee, P.-J. Chen, C. Wang, J. Gu, S. Popuri, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino, and W.-N. Hsu, "Direct speech-to-speech translation with discrete units," 2022.
- [10] R. Huang, J. Liu, H. Liu, Y. Ren, L. Zhang, J. He, and Z. Zhao, "Transpeech: Speech-to-speech translation with bilateral perturbation," 2023.
- [11] L. Barrault, Y.-A. Chung, M. C. Meglioli, D. Dale, N. Dong, M. Duppenhaler, P.-A. Duquenne, B. Ellis, H. Elshar, J. Haasheim *et al.*, "Seamless: Multilingual expressive and streaming speech translation," *arXiv preprint arXiv:2312.05187*, 2023.
- [12] Q. T. Do, T. Toda, G. Neubig, S. Sakti, and S. Nakamura, "Preserving word-level emphasis in speech-to-speech translation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 3, pp. 544–556, 2016.
- [13] Q. T. Do, S. Sakti, and S. Nakamura, "Sequence-to-sequence models for emphasis speech translation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1873–1883, 2018.
- [14] A. Waibel, M. Behr, D. Yaman, F. I. Eyiokur, T.-N. Nguyen, C. Mullov, M. A. Demirtas, A. Kantarci, S. Constantin, and H. K. Ekenel, "Face-dubbing++: Lip-synchronous, voice preserving translation of videos," in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*. IEEE, 2023, pp. 1–5.
- [15] A. Diwan, A. Srinivasan, D. Harwath, and E. Choi, "Unit-based speech-to-speech translation without parallel data," *arXiv preprint arXiv:2305.15405*, 2023.
- [16] E. Nachmani, A. Levkovitch, Y. Ding, C. Asawaroengchai, H. Zen, and M. T. Ramanovich, "Translatotron 3: Speech to speech translation with monolingual data," *arXiv preprint arXiv:2305.17547*, 2023.
- [17] B. Desplanques, J. Thienpondt, and K. Demuyne, "Ecapadnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *arXiv preprint arXiv:2005.07143*, 2020.
- [18] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," 2021.
- [19] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE transactions on acoustics, speech, and signal processing*, vol. 28, no. 4, pp. 357–366, 1980.
- [20] C. Wang, M. Rivière, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," 2021.
- [21] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," 2020.
- [22] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," 2018.
- [23] R. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. J. Weiss, R. Clark, and R. A. Saurous, "Towards end-to-end prosody transfer for expressive speech synthesis with tacotron," 2018.
- [24] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech 2: Fast and high-quality end-to-end text to speech," 2022.
- [25] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, and Y. Wu, "Fully-hierarchical fine-grained prosody modeling for interpretable speech synthesis," 2020.
- [26] J. Kahn, M. Riviere, W. Zheng, E. Kharitonov, Q. Xu, P. Mazare, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, "Libri-light: A benchmark for asr with limited or no supervision," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, May 2020. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP40776.2020.9052942>
- [27] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "English multi-speaker corpus for cstr voice cloning toolkit," *The Centre for Speech Technology Research (CSTR), University of Edinburgh*, 2017.
- [28] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [29] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," *arXiv preprint arXiv:1904.01038*, 2019.
- [30] C. Wang, Y. Tang, X. Ma, A. Wu, S. Popuri, D. Okhonko, and J. Pino, "Fairseq s2t: Fast speech-to-text modeling with fairseq," *arXiv preprint arXiv:2010.05171*, 2020.
- [31] C. Wang, W.-N. Hsu, Y. Adi, A. Polyak, A. Lee, P.-J. Chen, J. Gu, and J. Pino, "fairseq s²: A scalable and integrable speech synthesis toolkit," *arXiv preprint arXiv:2109.06912*, 2021.
- [32] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux, "Speech resynthesis from discrete disentangled self-supervised representations," *arXiv preprint arXiv:2104.00355*, 2021.