

# STORYTTS: A HIGHLY EXPRESSIVE TEXT-TO-SPEECH DATASET WITH RICH TEXTUAL EXPRESSIVENESS ANNOTATIONS

Sen Liu, Yiwei Guo, Xie Chen, <sup>†</sup>Kai Yu

MoE Key Lab of Artificial Intelligence, AI Institute  
X-LANCE Lab, Department of Computer Science and Engineering  
Shanghai Jiao Tong University, Shanghai, China  
{sen.liu, yiwei.guo, chenxie95, kai.yu}@sjtu.edu.cn

## ABSTRACT

While acoustic expressiveness has long been studied in expressive text-to-speech (ETTS), the inherent expressiveness in text lacks sufficient attention, especially for ETTS of artistic works. In this paper, we introduce StoryTTS, a highly ETTS dataset that contains rich expressiveness both in acoustic and textual perspective, from the recording of a Mandarin storytelling show. A systematic and comprehensive labeling framework is proposed for textual expressiveness. We analyze and define speech-related textual expressiveness in StoryTTS to include five distinct dimensions through linguistics, rhetoric, etc. Then we employ large language models and prompt them with a few manual annotation examples for batch annotation. The resulting corpus contains 61 hours of consecutive and highly prosodic speech equipped with accurate text transcriptions and rich textual expressiveness annotations. Therefore, StoryTTS can aid future ETTS research to fully mine the abundant intrinsic textual and acoustic features. Experiments are conducted to validate that TTS models can generate speech with improved expressiveness when integrating with the annotated textual labels in StoryTTS.

**Index Terms**— expressive text-to-speech, TTS dataset, textual expressiveness, large language models

## 1. INTRODUCTION

The advancement of deep learning has significantly enhanced the quality of text-to-speech (TTS), enabling TTS models [1, 2, 3, 4, 5, 6] to produce speech that closely resembles human speech. However, these models tend to excel in synthesizing speech with relatively simple emotional characteristics. When it comes to expressive performance genres such as novels, poems, talk shows, and others, these methods still fall short of delivering the desired level of expressiveness. Such performances are usually rich in expressive text. This form of text influences the cadence and rhythm of the speaker’s delivery and refers to a form of written language that can bring to life the meaning and feeling of the subject matter to be conveyed. A common approach is mine semantic and syntactic features from the expressive text. [7] introduced a TTS model that explicitly incorporates text-context semantic information extracted from a pre-trained BERT [8] model. This approach effectively enhances the expressiveness of the synthesized speech. In [9], a syntax graph is constructed for each input sentence based on its dependency tree. Following this, they employed a graph encoder to extract syntactic features, resulting in improved duration and pitch prediction.

However, these studies often rely on the coarse-grained semantic representations of pre-trained language models or basic syntactic structures, and they have not conducted a comprehensive and thorough exploration of textual expressiveness, especially speech-related expressiveness. To achieve natural and expressive TTS synthesis, it’s essential to convey the emotional stance of the text, which requires identifying speech-related textual expressive features in the text. This is also supported by literature [10], where a thorough investigation was conducted into the relationship between prosody and linguistics. This study underscored that elements like the emotional content of the text, sentence patterns, and syntax have a direct impact on reading expressiveness. Hence, understanding how to generalize, summarize, and characterize speech-related expressive features from expressive text might be crucial for expressive TTS.

In this paper, we introduce StoryTTS, a highly expressive TTS dataset with rich expressiveness from both acoustic and textual perspectives. We initially construct the dataset from the recording of a Mandarin storytelling show with careful revision of transcripts and punctuations. Then we establish a systematic and comprehensive labeling framework for textual expressiveness. Specifically, we analyze and define speech-related textual expressiveness in StoryTTS to include five distinct dimensions through linguistics, rhetoric, and literary studies. These dimensions include rhetorical devices, sentence patterns, scenes, imitated characters, and emotional colors. Then we employ large language models (LLMs) and prompt them with a few manual annotation examples for batch annotation. The resulting corpus contains 61 hours of consecutive and highly prosodic speech equipped with accurate text transcriptions and rich textual expressiveness annotations. We further conduct experiments to validate that TTS models can produce speech with enhanced expressiveness when integrating the annotations from StoryTTS. Our contributions can be summarized as follows:

- We construct StoryTTS, the first TTS dataset that contains rich expressiveness in both speech and texts and is also equipped with comprehensive annotations for speech-related textual expressiveness. This dataset is also of high acoustic quality, organized by consecutive chapters, and of sufficient size. We release the StoryTTS dataset online<sup>1</sup>.
- We establish a framework powered by LLMs to annotate textual expressiveness in five distinct dimensions.
- We conduct experiments to validate that TTS models can produce speech with enhanced expressiveness when integrating the annotated textual expressiveness labels.

<sup>†</sup>Kai Yu is the corresponding author.

<sup>1</sup><https://goarsenal.github.io/StoryTTS>

## 2. STORYTTS

The construction of StoryTTS is introduced in detail in this section, including data selection and retrieval, audio quality analysis, speech segmentation and automatic recognition, manual correction of recognition errors, and the enhancement of punctuation. The detailed statistics of StoryTTS are shown in Table 1.

**Table 1.** Detailed statistics of StoryTTS.

Domain	Feature	Statistics
Speech	Number of speakers	1
	Total duration (hours)	60.9
	Mean duration per utterance (seconds)	6.8
	Sampling rate (kHz)	16
	Estimated SNR (dB)	32
Text	Chapters	160
	Sentences	33108

**Table 2.** Comparison of often used publicly available TTS datasets. Here, ‘**Lang**’ refers to the respective language, ‘**Num of spks**’ means the number of speakers, ‘**Hours / spk**’ indicates the average hours per speaker, ‘**Pitch Std. / spk**’ represents the average standard deviation of pitch per speaker, and ‘**Expr. Annots.**’ denotes whether expressiveness annotations are available.

Dataset	Lang	Num of spks	Hours / spk	Pitch Std. / spk	Expr. Annots.
LJSpeech [11]	EN	1	24	49.93	N
Blizzard-2013 [12]	EN	1	319	49.32	N
HiFi-TTS [13]	EN	10	29.2	55.09	N
LibriTTS [14]	EN	2456	4.2	48.51	N
Aishell3 [15]	ZH	218	0.39	39.63	N
Biaobei [16]	ZH	1	12	58.57	N
StoryTTS	ZH	1	60.9	<b>98.22</b>	<b>Y</b>

### 2.1. Data selecting and retrieval

Storytelling, also known as “Pingshu”, is a traditional Chinese oral art form where performers narrate stories, imitate various voices, and portray characters to enthrall audiences. This form of oral art is usually based on historical novels, which makes storytelling not only rich in speech prosody but also diverse in textual expressiveness, such as linguistic structures, figures of speech, role-playing, etc. Hence, storytelling shows satisfy our goal to a great extent. We selected a storytelling show titled “Zhi Sheng Dongfang Shuo”, recounting the legend of Dongfang Shuo, a key figure in the development of the Han Dynasty in ancient China. This performance is skillfully delivered by a female artist, Lian Liru. To construct the dataset, we retrieved the recorded speech data from a public website, which is organized into 160 consecutive chapters. These chapters have an approximate duration of 24 minutes each, hence amounting to 64 hours in total, including interval breaks.

### 2.2. Audio quality analysis

We also estimated the signal-to-noise ratio (SNR) of the speech data, where the noise power was computed using the silence segments predicted by a voice activity detection (VAD) tool. As can be seen in Table 1, the SNR is estimated to be 32dB, indicating the high audio quality of the waveforms. Subsequently, we conducted statistical analysis on multiple common Mandarin (ZH) and English (EN) datasets, as shown in Table 2. The results reveal that StoryTTS exhibits a significantly higher pitch standard deviation than other datasets, providing compelling evidence of its substantial acoustic expressiveness. Furthermore, StoryTTS includes expressiveness annotations, which will be further explained in the next section.

### 2.3. Speech segmentation and automatic recognition

To process the originally coarsely segmented speech data, we implemented a three-step approach. We first employed a VAD tool to segment the chapter-level speech into utterances based on the duration of silence segments. Long silences were also removed in this step, resulting in 60.9 hours of speech. Subsequently, given the absence of matching text transcripts, we obtained text transcripts using Whisper [17], a popular speech recognition model. We observed that there still exist speech segments that remained excessively long after the VAD process. To address this issue, we identified these speech segments and their corresponding texts. We manually divided the prolonged text into smaller sentences and then utilized the Aeneas<sup>2</sup> tool for synchronizing text fragments with speech. This alignment allowed us to accurately cut the speech, producing a final dataset of 33108 pairs of speech and text.

### 2.4. Manual correction of recognition errors

Given the extremely variable pitch and speaking rate in the storytelling performances, the speech recognition results exhibited a higher error rate compared to standard speech. In response to this challenge, we meticulously reviewed every speech segment line by line and rectified the recognition errors. Furthermore, we have made diligent efforts to replace onomatopoeic elements in the speech with appropriate words from the corresponding text.

### 2.5. Punctuation enhancement

Punctuation plays a vital role in text presentation, conveying emotions like surprise or shock through exclamation points and indicating character dialogue or inner thoughts through double quotes. Although Whisper can identify some punctuation marks, it still falls far short of expectations. Thus, during our text review process, we made careful punctuation corrections and additions to ensure precise punctuation usage whenever possible. This attention to punctuation accuracy also significantly benefited our subsequent work on textual sentiment analysis.

## 3. LLM-DRIVEN EXPRESSIVENESS ANNOTATION

The expressive texts within StoryTTS exhibit a high degree of colloquialism and are rich in role-playing, psychological, action, and environmental descriptions. Inspired and guided by [10], we introduced a systematic and comprehensive labeling framework that harnesses the power of LLMs.

### 3.1. Exploring textual expressiveness

In our investigation of speech-related textual expressiveness, we classify it into five dimensions drawn from the fields of literary studies, linguistics, and rhetoric. These dimensions include rhetorical devices, sentence patterns, scenes, imitated characters, and emotional colors.

**Rhetorical devices**, like hyperbole, and **sentence patterns**, such as declarative sentences, are commonly employed textual expressive devices. For instance, using an exclamatory sentence or incorporating rhetorical devices like hyperbole can evoke emotions of excitement or surprise. We also employed **scenes** such as role-playing, taking into account the characteristics of StoryTTS. For example, a role-playing scene often carries strong emotional content, while an aside typically lacks emotional elements. The specific categorization regarding sentence patterns, scenes, and rhetorical devices is shown in Table 3.

<sup>2</sup><https://github.com/readbeyond/aeneas>

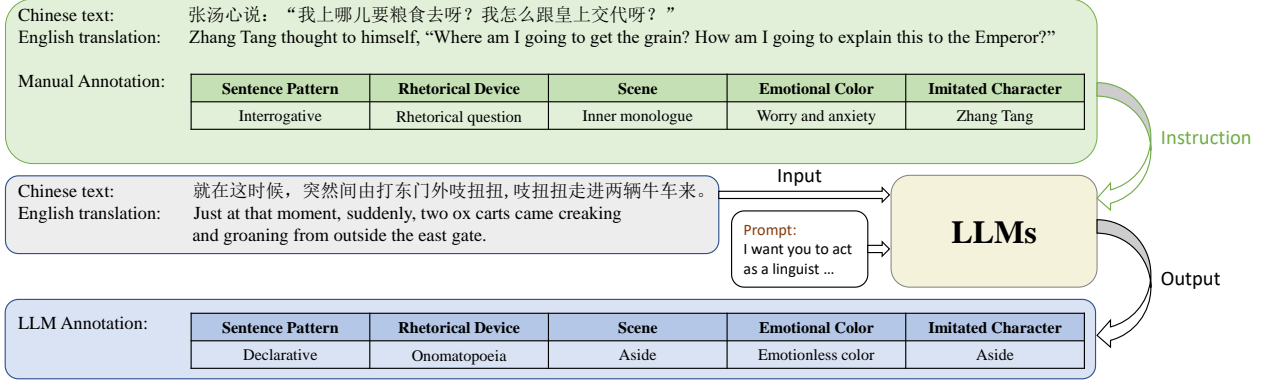


Fig. 1. An Illustration of Expressiveness Annotation Using LLMs.

**Table 3.** Specific categories of Sentence Pattern, Scene, and Rhetoric Device in StoryTTS. Here, ‘Num’ denotes the number of categories.

Dimension	Num	Category
Sentence Pattern	4	Declarative, Interrogative, Imperative, Exclamatory
Scene	3	Role-playing, Aside, Inner monologue
Rhetoric Device	11	Hyperbole, Antithesis, Onomatopoeia, Simile, Personification, Quotation, Irony, Rhetorical question, Hypophora, Rhetorical repetition, Transition

**Emotional colors** often have a direct impact on the performer’s expression. Instead of categorizing them into various polarities or predefined categories, we have chosen a more precise approach: summarizing the emotional color of a sentence using several words. This method can provide a more accurate description of the text’s emotion compared to traditional categorization.

In the case of **imitated characters** in StoryTTS, the performer often mimics the speech patterns of characters when delivering their lines. For example, she deliberately lowers the pitch and slows down when portraying an old man while raising the pitch and speeding up when mimicking a villain. We categorized the characters into **19** role types based on characteristics like age, gender, and status, which also served as the basis for the performer’s mimicry. The six most frequent role types are shown in Figure 2.




Fig. 2. Top 6 role types in StoryTTS.

### 3.2. Batch annotation via LLMs

LLMs have found applications in a wide range of studies. [18] demonstrated that GPT3 [19] performs well in data annotation tasks at a relatively low cost, making it a viable choice for individuals and organizations with limited budgets. To expedite the labeling process and reduce costs, we utilized GPT4 [20] and Claude2 [21], both of which are more powerful than GPT3, for batch annotation. During our annotation process, Claude2 was used for annotating sentence patterns, rhetorical devices, scenes, and imitated characters. However, when it came to summarizing emotional colors in the text, we found Claude2’s performance to be suboptimal. Consequently, we turned to GPT4, which proved to be more proficient in this aspect. In our prompt, we initiated the persona of a linguist for the LLMs. We then proceeded to guide the model, informing it that the input texts were consecutive and required labeling with contextual information. For instance, two consecutive sentences might belong to the same role-play scenario. Subsequently, we provided details about the sources and features of the text, emphasizing their richness in elements such as onomatopoeia, inner monologue, and role-plays. Finally, we instructed the model to annotate each sentence in a prescribed format, adhering to the prompts and requirements outlined. These annotations align with the guidelines detailed in Section 3, where sentence patterns, scenes, rhetorical devices and imitated characters must be assigned specific categories, and each emotional color should be summarized in several words.

Initially, we attempted labeling in the zero-shot setting, but the results exhibited low accuracy. Consequently, we transitioned to the few-shot setting, where we provided the model with comprehensive and diverse labeled text and explained the rationale behind each labeling decision. In this setting, the model demonstrates improved accuracy and meets our labeling requirements. An illustration of annotation using LLMs is shown in Figure 1.

## 4. EXPERIMENTS AND RESULTS

In this section, we build TTS models to analyze the impact of annotated textual expressiveness labels on synthetic speech.

### 4.1. Model Architecture

#### 4.1.1. Baseline Model

Our baseline model is implemented based on VQTTS [22], which utilizes self-supervised vector-quantized (VQ) [23, 24, 25, 26], acoustic features rather than traditional mel-spectrogram. Specifically, it consists of an acoustic model, t2v, and a vocoder, v2w. T2v accepts the phoneme sequence and then outputs the VQ Acoustic

feature and Auxiliary feature, which consist of pitch, energy, and probability of voice [27], and v2w receives them and thus synthesizes the waveform.

#### 4.1.2. Expressiveness Encoder

To fully leverage our expressiveness annotations, we developed an expressiveness encoder. We employed four separate learnable embedding tables to supply information to the model for the four labels: sentence pattern, scene, rhetorical method, and imitated character. For each sentence, we assigned four category numbers based on these four expressive labels. We then inputted these numbers into the corresponding embedding tables, with vector dimensions of 32, 32, 64, and 256, respectively.

Regarding the modeling of emotional color, we employed distinct model structures. Given that emotion descriptions typically condense into several words, representing the overall sentiment of a sentence, while emotions may change within a single sentence. For instance, in an exclamatory sentence, emotion often intensifies towards the end. We initially extracted word-level embeddings for the entire sentence using a pre-trained BERT. Subsequently, we extracted the embedding of emotional color using a Sentence BERT. Through cross-attention [28] between these embeddings, we aimed to capture the distribution of emotions at different locations within the text, enhancing their expressive accuracy. Following this, we up-sampled the results to the phoneme level based on the word-to-phoneme correspondence and added them to the encoder output, along with the previous four embeddings.

#### 4.2. Experimental setup

We conducted experiments to evaluate the impact of each of the five textual expressiveness labels on the expressiveness of the synthesized speech. Additionally, we assessed the cumulative effect of utilizing all these labels together. For these experiments, we trained an acoustic model separately for 300 epochs using a batch size equal to 8. The vocoder was shared, and we trained 100 epochs on StoryTTS using a batch size of 8. The remaining model configurations and parameters remained consistent with those in [29]. Each experiment was performed on a single 2080 Ti GPU. To preprocess the text data, we utilized our internal Grapheme-to-Phoneme (G2P) tool for text-to-phoneme conversion. We also set aside 5% of the text for test and validation sets, where the test set consists of 3 consecutive chapters. To obtain ground truth phoneme duration, we employed the Montreal Forced Aligner [30], which conducts forced alignment using Kaldi [31].

#### 4.3. Speech Synthesis Evaluation

##### 4.3.1. Metrics

We performed a mean opinion score (MOS) listening test involving 20 native listeners who were asked to rate each sample. MOS ratings were based on a 1-5 scale with 0.5-point increments and 95% confidence intervals. During our tests, we instructed listeners to specifically assess the level of expressiveness in the synthesized speech, all while evaluating speech quality. For objective evaluations, we computed Mel-cepstral distortion (MCD) using dynamic time warping (DTW). Additionally, we analyzed the log F0 root mean square error (log-F0 RMSE), also computed with DTW. MCD measures general speech quality, while log-F0 RMSE assesses performance in terms of speech prosody. Lower values for both of these metrics indicate better sound quality and rhythm in speech performance.

##### 4.3.2. Results

Table 4 presents the evaluation results. It can be seen that when incorporating the expressiveness labels into the model, both the sub-

**Table 4.** Results of evaluation for different setups. Here, ‘GT(Voc.)’ denotes the vocoded ground truth speech, and ‘+ALL’ means we utilized all the expressiveness labels together.

Model	MCD ↓	log-F0 RMSE ↓	MOS ↑
GT(Voc.)	3.765	0.322	4.29 ± 0.06
Baseline	6.904	0.437	3.88 ± 0.07
+Sentence Pattern	6.746	0.432	3.90 ± 0.08
+Scene	6.692	0.431	3.90 ± 0.07
+Rhetoric Device	6.633	0.421	3.93 ± 0.07
+Emotional Color	6.271	0.412	3.92 ± 0.08
+Imitated Character	6.508	0.411	3.96 ± 0.07
+ALL	<b>6.181</b>	<b>0.402</b>	<b>4.09 ± 0.07</b>

jective and objective scores outperform the baseline model. Specifically, sentence patterns and scene boosts were relatively minimal. This might be attributed to the prevalence of declarative sentences in the dataset, resulting in limited information acquired by the model. Additionally, while scene types are distributed fairly evenly, their diversity is insufficient to furnish the model with adequate information, owing to the different character imitations found in role-playing and inner monologue scenes. Rhetorical devices and emotional colors bring more obvious enhancements. Among the individual expressiveness labels, the imitated characters stand out as the most effective, as they directly provide information about the characters currently being mimicked, enabling the model to efficiently learn how the mimicked characters speak and thus synthesize speech close to that of the original data.

Finally, the fusion of all expressive labels provides the most significant enhancement. It outperforms other setups significantly in both objective and subjective metrics, supplying the model with increasingly accurate information about imitated characters and scenes. This fusion also benefits from the complementary nature of sentence patterns, rhetorical devices, and emotional colors.

## 5. CONCLUSIONS

This paper presented StoryTTS, the first TTS dataset that encompasses rich expressiveness from both acoustic and textual perspectives. Derived from a high-quality recording of a Mandarin storytelling show, this dataset serves as a valuable resource for researchers aiming to investigate acoustic expressiveness on the one hand. Additionally, we conducted a comprehensive analysis of expressive text and categorized speech-related textual expressiveness into five distinct dimensions. Then we employed LLMs and provided them with a few manual annotation examples for batch annotation. The effective labeling of LLMs also offers insights for similar data labeling endeavors. The dataset is thus equipped with abundant textual expressiveness annotations. Experimental results demonstrated that TTS models can generate speech with significantly improved expressiveness when incorporated with the annotated textual expressiveness labels. Future work may focus on integrating these expressiveness annotations with acoustic expressiveness to further enhance expressive speech synthesis.

## 6. ACKNOWLEDGEMENTS

This work was supported by China NSFC Project (No.92370206), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102) and the Key Research and Development Program of Jiangsu Province, China (No.BE2022059).

## 7. REFERENCES

- [1] Yuxuan Wang, R. J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, and et al., “Tacotron: Towards end-to-end speech synthesis,” in *Proc. Interspeech*. 2017, pp. 4006–4010, ISCA.
- [2] Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, and et al., “Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions,” in *Proc. ICASSP*. 2018, pp. 4779–4783, IEEE.
- [3] Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” in *Proc. ICLR*, 2021.
- [4] Jaehyeon Kim, Jungil Kong, and Juhee Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proc. ICML*. 2021, vol. 139, pp. 5530–5540, PMLR.
- [5] Chenpeng Du and Kai Yu, “Rich prosody diversity modelling with phone-level mixture density network,” in *Proc. Interspeech*. 2021, pp. 3136–3140, ISCA.
- [6] Yiwei Guo, Chenpeng Du, and Kai Yu, “Unsupervised word-level prosody tagging for controllable speech synthesis,” in *Proc. ICASSP*. 2022, pp. 7597–7601, IEEE.
- [7] Tomoki Hayashi, Shinji Watanabe, Tomoki Toda, Kazuya Takeda, Shubham Toshniwal, and Karen Livescu, “Pre-trained text embeddings for enhanced text-to-speech synthesis,” in *Proc. Interspeech*. 2019, pp. 4430–4434, ISCA.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. ACL*, 2019, pp. 4171–4186.
- [9] Zhenhui Ye, Zhou Zhao, Yi Ren, and Fei Wu, “Syntaspeech: Syntax-aware generative adversarial text-to-speech,” in *Proc. IJCAI*. 2022, pp. 4468–4474, ijcai.org.
- [10] Paula J Schwanenflugel and Rebekah George Benjamin, “Reading expressiveness,” *Fluency instruction: Research-based best practices*, vol. 35, 2012.
- [11] Keith Ito and Linda Johnson, “The LJ speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [12] Simon King and Vasilis Karaiskos, “The blizzard challenge 2013,” 2014.
- [13] Evelina Bakhturina, Vitaly Lavrukhin, Boris Ginsburg, and Yang Zhang, “Hi-Fi multi-speaker english TTS dataset,” in *Proc. Interspeech*. 2021, pp. 2776–2780, ISCA.
- [14] Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J. Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu, “LibriTTS: A corpus derived from librispeech for text-to-speech,” in *Proc. Interspeech*. 2019, pp. 1526–1530, ISCA.
- [15] Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li, “AISHELL-3: A multi-speaker mandarin TTS corpus,” in *Proc. Interspeech*. 2021, pp. 2756–2760, ISCA.
- [16] “Biaobei dataset,” <https://en.data-baker.com/datasets/freeDatasets>, 2017.
- [17] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. ICML*. 2023, vol. 202, pp. 28492–28518, PMLR.
- [18] Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing, “Is GPT-3 a good data annotator?,” in *Proc. ACL*, 2023, pp. 11173–11195.
- [19] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, and et al., “Language models are few-shot learners,” in *Proc. NeurIPS*, 2020.
- [20] OpenAI, “GPT-4 technical report,” *CoRR*, vol. abs/2303.08774, 2023.
- [21] Anthropic, “Model card and evaluations for claude models,” <https://www-files.anthropic.com/production/images/Model-Card-Claude-2.pdf>.
- [22] Chenpeng Du, Yiwei Guo, Xie Chen, and Kai Yu, “VQTTs: high-fidelity text-to-speech synthesis with self-supervised VQ acoustic feature,” in *Proc. Interspeech*. 2022, pp. 1596–1600, ISCA.
- [23] Alexei Baevski, Steffen Schneider, and Michael Auli, “vq-wav2vec: Self-supervised learning of discrete speech representations,” in *Proc. ICLR*. 2020, OpenReview.net.
- [24] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Proc. NeurIPS*, vol. 33, pp. 12449–12460, 2020.
- [25] Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli, “Unsupervised cross-lingual representation learning for speech recognition,” in *Proc. Interspeech*, Hynek Hermans, Honza Cernocký, Lukás Burget, Lori Lamel, Odette Scharenborg, and Petr Motlíček, Eds. 2021, pp. 2426–2430, ISCA.
- [26] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [27] Pegah Ghahremani, Bagher BabaAli, Daniel Povey, Korbinian Riedhammer, Jan Trmal, and Sanjeev Khudanpur, “A pitch extraction algorithm tuned for automatic speech recognition,” in *Proc. ICASSP*. 2014, pp. 2494–2498, IEEE.
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Proc. NeurIPS*, 2017, pp. 5998–6008.
- [29] Sen Liu, Yiwei Guo, Chenpeng Du, Xie Chen, and Kai Yu, “DSE-TTS: Dual Speaker Embedding for Cross-Lingual Text-to-Speech,” in *Proc. Interspeech*. 2023, pp. 616–620, ISCA.
- [30] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldi,” in *Proc. Interspeech*. 2017, pp. 498–502, ISCA.
- [31] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlíček, and et al., “The Kaldi speech recognition toolkit,” in *Proc. ASRU*. 2011, IEEE.