



# Transcript-Prompted Whisper with Dictionary-Enhanced Decoding for Japanese Speech Annotation

Rui Hu<sup>1</sup>, Xiaolong Lin<sup>1</sup>, Jiawang Liu<sup>1</sup>, Shixi Huang<sup>1</sup>, Zhenpeng Zhan<sup>1</sup>

<sup>1</sup>Baidu Inc., Shenzhen, China

hurui05@baidu.com

## Abstract

In this paper, we propose a method for annotating phonemic and prosodic labels on a given audio-transcript pair, aimed at constructing Japanese text-to-speech (TTS) datasets. Our approach involves fine-tuning a large-scale pre-trained automatic speech recognition (ASR) model, conditioned on ground truth transcripts, to simultaneously output phrase-level graphemes and annotation labels. To further correct errors in phonemic labeling, we employ a decoding strategy that utilizes dictionary prior knowledge. The objective evaluation results demonstrate that our proposed method outperforms previous approaches relying solely on text or audio. The subjective evaluation results indicate that the naturalness of speech synthesized by the TTS model, trained with labels annotated using our method, is comparable to that of a model trained with manual annotations.

**Index Terms:** speech annotation, fine-tuning, error correction, text-to-speech, Japanese

## 1. Introduction

The field of text-to-speech (TTS) has seen significant advancements in the past few years. Recent studies [1, 2, 3] have demonstrated that large models trained on massive datasets can synthesize natural speech directly from grapheme sequences, eliminating the need for more fine-grained inputs such as phonemes. This is particularly notable in alphabetic languages like English. However, challenges persist for Japanese TTS due to the presence of numerous polyphones and the critical role that pitch accent plays in affecting the naturalness of speech [4, 5].

Most TTS datasets consist solely of audio-transcript pairs, as obtaining manually annotated phonemic and prosodic labels is prohibitively expensive. As a result, many TTS systems rely on text-based frontends [6, 7, 8] during data pre-processing. However, these modules can be inaccurate due to the one-to-many mapping problem between text and audio, where certain characters or words can have multiple pronunciations depending on the context. This issue is especially pronounced in the case of Japanese because: 1) there are a lot of polyphones; 2) each word has its own accent nucleus (the mora at which the pitch-downstep occurs), and adjacent words often combine to form an accent phrase and create a new accent nucleus, based on the speaker’s interpretation of the context [8].

Several recent studies have explored the use of audio information during data pre-processing. Researchers in [9] proposed a sequence-to-sequence model capable of simultaneously transcribing and annotating audio with linguistic information more accurately than methods based solely on text. Building upon this, work in [10] further demonstrated the promising performance of fine-tuning a large-scale pre-trained automatic speech recognition (ASR) model to predict phonemic and prosodic la-

els from unlabeled speech data. However, through preliminary experiments, we found that annotation models only conditioned on audio tend to produce errors related to phonetically similar sounds. Some studies have attempted to leverage both text and audio information to obtain more precise annotations. [11] achieved automatic prosody labeling using text-audio data via a speech-text model while [12] proposed a Grapheme-to-Phoneme (G2P) module utilizing both text and aligned audio to overcome pronunciation ambiguity. However, these methods either only focus on prosody annotation or rely on complicated pipelines.

In the field of ASR, Whisper [13] is a Transformer-based [14] robust foundation model pre-trained on multiple tasks and large-scale data. Several recent works [15, 16, 17, 18] have demonstrated that fine-tuning Whisper with prompts can achieve impressive performance in recognizing domain-sensitive speech, spoken name entity, target speaker speech and disordered speech, respectively.

Inspired by previous works, this paper proposes fine-tuning Whisper conditioned on ground truth transcripts to simultaneously produce phrase-level graphemes along with their phonemic and prosodic labels (hereinafter referred to as *TTS labels*) from speech data. To further improve the performance of phonemic labeling, we employ a decoding strategy that leverages dictionary prior knowledge. We conduct two experiments to evaluate the effectiveness of our method. The first one is an objective evaluation. We begin by verifying that the graphemes of each phrase align faithfully with the given transcript. Then we calculate the character error rate (CER) for phonemic labeling, as well as the accuracy of accent phrase boundary prediction and the  $F_1$  scores for mora-level pitch sequence in prosodic labeling. The second one is a subjective evaluation, where human raters assess the naturalness of speech synthesized by TTS models trained on labels generated by different annotation methods. Experimental results show that the proposed method produces more accurate TTS labels and helps enhance the performance of TTS models in terms of naturalness. Audio samples are available on our demo page<sup>1</sup>.

## 2. Method

### 2.1. Notation method for TTS labels

Japanese is a pitch-accent language with a writing system that combines logograms (Kanji) with phonograms (Hiragana and Katakana). As depicted in Figure 2, annotation for Japanese utterances typically consists of transcript and pitch status aligned at the mora level. We adopt the notation method proposed in [5] to construct TTS labels from the annotations following the rules

<sup>1</sup><https://hurui0318.github.io/japanese-annotator-demo/>

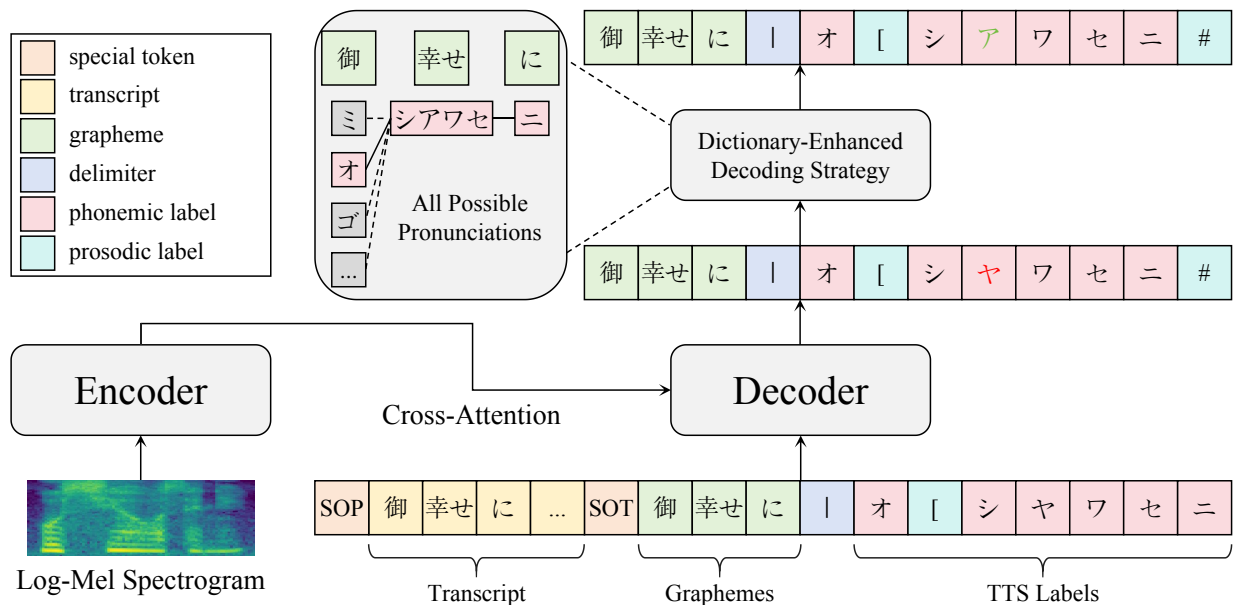


Figure 1: Overview of proposed annotation model with a decoding strategy that leverages dictionary prior knowledge.




Figure 2: Example of Japanese speech annotation for “輪廻は廻って、” (The cycle of reincarnation turns, and...). The yellow squares denote the transcript segmented into accent phrases, with each accent phrase further divided into individual words. The white circles in gray area indicate the pitch status (high or low) of each mora. The red and blue squares denote phonemic and prosodic labels, respectively.

of Tokyo dialect. The TTS labels use the following symbols to represent phonemic and prosodic information: (1) Katakana characters indicating pronunciation; (2) “[” for pitch rising; (3) “]” for pitch falling; (4) “#” for accent phrase boundary; (5) “-” for pause.

## 2.2. Annotation model

Whisper is a robust ASR model based on the encoder-decoder Transformer architecture, pre-trained on a large amount of data. The encoder processes log-Mel spectrograms to extract high-level hidden representations, while the decoder incorporates encoder output through cross-attention mechanism and generates transcripts in an auto-regressive manner. In the decoder input, we can use a <startofprev> token to represent the beginning of previous context and a <startoftranscript> token to indicate the beginning of transcription.

Researchers in [10] first proposed fine-tuning pre-trained Whisper model to produce TTS labels conditioned on audio. However, in our preliminary experiments, we observed that the model tends to lose some semantic knowledge acquired during pre-training. Consequently, it is prone to making mistakes such

Table 1: Examples of errors produced by the annotation model. “Content” denotes the ground truth transcript of speech content. “Reference” and “Prediction” refer to the annotations from the human annotator and the annotation model, respectively.

Source	Example
Content	成功してもしなくても
Reference	セ[-コーシテ]モ#シ[ナ]クテモ#
Prediction	セ[-コーシテ]モ#シ[ナ]ク <del>タ</del> モ#
Content	漁業と商業
Reference	ギョ]ギョ-ート#シヨ]-ギョ-ーテ#
Prediction	ギョ]-ギョ-ート#シヨ]-ギョ-ーテ#
Content	新しい家
Reference	ア[タラシ]-#イ[エ]オ#
Prediction	ア[タラシ]-# <del>エ</del> オ#

as failing to distinguish between phonetically similar sounds, omitting or adding certain pronunciations. Some examples are shown in Table 1. We speculate that this is likely because the model relies solely on speech as input, with training targets limited to TTS labels. To mitigate this, inspired by [9], we also attempted to have the model simultaneously output phrase-level graphemes and TTS labels, expecting that the model would understand the speech content before annotation. We found that the model has shown some improvements but errors still persist. This is reasonable since the base model, Whisper, is not a perfect transcriber. Therefore, during inference, the preceding graphemes may be inaccurate and could provide insufficient, even misleading guidance for the subsequent decoding of TTS labels.

Our proposed annotation model, however, takes both audio and its transcript as input and is required to output phrase-level graphemes along with their TTS labels. As illustrated in Figure 1, the transcript is tokenized and placed be-

tween <startofprev> and <startoftranscript> as a prompt to the decoder. Symbol “|” is the delimiter between graphemes and TTS labels of each phrase. In this approach, the decoder can attend to both the acoustic information from input speech and the semantic information of the transcript when generating labels.

### 2.3. Decoding strategy

Upon close examination, we still found the aforementioned errors present in the phonemic labeling of Kanji characters. Since the model is requested to sequentially output the graphemes and TTS labels of each accent phrase during decoding, we can utilize the preceding graphemes and a dictionary to correct the following phonemic labels.

As shown in Figure 1, when the annotation mode decodes an accent phrase boundary symbol “#”, we first get the graphemes and TTS labels for the phrase based on the delimiter “|”. The predicted TTS labels are converted into pronunciation in Katakana and accent type (the position of mora where accent nucleus appears) following the rules of Tokyo dialect. Then we use MeCab [19] with UniDic [20] to segment the graphemes into words and look up each word’s candidate pronunciations to get all possible pronunciations for the phrase. Next, we select the one with the shortest edit distance to the pronunciation predicted by the annotation model as the corrected pronunciation. After this, we need to restore it to the form of TTS labels with some modifications to the accent type. As shown in Equation 1, the accent type remains unchanged for phrases whose mora size stays the same after pronunciation correction or when the original accent type is 0 (flat), 1 (head high) or  $n$  (tail high), where  $n$  is the mora size. For the rest (middle high), the modified accent type  $A_{\text{mod}}$  is calculated by adding the difference between the mora size of the corrected phrase  $M_{\text{mod}}$  and that of the original phrase  $M_{\text{orig}}$  to the original accent type  $A_{\text{orig}}$ . Finally, we replace the original TTS labels with the corrected ones and continue with the decoding of subsequent phrases.

$$A_{\text{mod}} = \begin{cases} A_{\text{orig}}, & A_{\text{orig}} \in \{0, 1, n\}, \\ A_{\text{orig}} + (M_{\text{mod}} - M_{\text{orig}}), & A_{\text{orig}} \in \{2, \dots, n - 1\}. \end{cases} \quad (1)$$

Note that the restoration of TTS labels is not perfect but we show that the impact is minimal in the subsequent experiments.

## 3. Experimental setup

### 3.1. Evaluation metrics

In the experiments, we aim to evaluate our proposed annotation method in three aspects:

- The alignment of predicted phrase-level graphemes with the given transcript.
- The performance of phonemic and prosodic labeling.
- The effectiveness of its application in TTS systems.

We first segment the predicted utterance-level TTS labels into accent phrases based on the accent phrase boundary symbol “#”. For each accent phrase, we divide it into graphemes and TTS labels using the delimiter “|”. Following the rules of Tokyo dialect, the phrase-level TTS labels are converted into phonemic labels in Katakana and mora-level high-low pitch sequences as prosodic labels. The phrase-level graphemes, phonemic labels and prosodic labels are then concatenated back to the utterance

level, respectively. CERs are calculated for both the grapheme sequences and phonemic labels.

Utterances where all models correctly predict the phonemic labels are used for evaluating prosodic labeling. This is because comparing ground truth prosodic labels with predicted ones becomes infeasible when the predicted phonemic labels are incorrect [10]. To evaluate the accuracy of accent phrase boundary prediction, we calculate a metric defined as the number of correctly predicted accent phrases divided by the total number of accent phrases in the manual annotations. For pitch labeling, we compute  $F_1$  scores on the mora-level high-low pitch sequences.

Finally, we directly use the utterance-level TTS labels produced by the annotation model to train TTS models and conduct subjective listening test based on the mean opinion score (MOS). Ten raters proficient in Japanese are asked to assess 50 randomly selected speech samples for pronunciation accuracy and prosodic naturalness.

### 3.2. Data

We use three datasets for model training and evaluation:

**JSUT**: a public Japanese speech corpus with a single female voice [21]. The *basic5000* subset and their manually annotated TTS labels<sup>2</sup> are used. The TTS labels are at the accent phrase level but there is no alignment information with the text. Therefore, we feed them into GPT-4 [22] to obtain phrase-level aligned annotations. The data are used for training annotation models.

**PRIV**: a private Japanese speech corpus recorded by seven male and four female professional native Japanese speakers. It comprises 104,183 text samples and 121.82 hours of speech, with annotations provided at the mora level. Data from six male and four female speakers are used to train the annotation models, while the remaining data from one male speaker, consisting of 8,888 samples, are reserved for testing. The held-out data are further randomly split into 7,110, 889 and 889 samples for training, validation and testing in TTS experiments, respectively.

**JVS**: a public Japanese speech corpus containing audio-transcript pairs of 100 professional speakers [23]. We only use the *parallel100* and *nonpara30* subsets, with a total of 13,000 samples. These samples are randomly divided into 10,400, 1300 and 1300 samples for training, validation and testing in TTS experiments, respectively.

### 3.3. Model details

The proposed annotation model is fine-tuned from *whisper-large-v3-turbo*<sup>3</sup> for its robustness and inference speed, using the transformers library [24]. The model is fine-tuned for 100k steps with a batch size of 32 on a single A100 GPU. The learning rate is initialized at 0.00001 with 500 warm-up steps. During training, the labels for the prompt tokens are set to  $-100$  to ignore loss calculation. Notably, in our preliminary experiments, we found that directly using symbols, such as “#” could lead to severe hallucinations during inference. We conjecture that this is because these symbols have already been assigned different meanings during pre-training, which are difficult to alter through fine-tuning. Therefore, we use Unicode characters U+2191, U+2193, U+2460, U+2462 and U+2223, which are not present in the Whisper tokenizer’s vocabulary, to represent

<sup>2</sup><https://github.com/sarulab-speech/jsut-label>

<sup>3</sup><https://huggingface.co/openai/whisper-large-v3-turbo>

“[”, “[”, “#”, “.” and “[”, respectively. The following summarizes the annotation methods used in our experiments:

**NLP:** Open JTalk<sup>4</sup>, a text-based annotation method. We use its Python wrapper<sup>5</sup>.

**Annt-v1:** Whisper model that is conditioned solely on audio and is fine-tuned to output TTS labels. We train two models on JSUT and PRIV, referred to as **Annt-v1-pub** and **Annt-v1-priv**.

**Annt-v2:** Whisper model that is conditioned solely on audio and is fine-tuned to output phrase-level graphemes and TTS labels. We train two models on JSUT and PRIV, referred to as **Annt-v2-pub** and **Annt-v2-priv**.

**Annt-v3:** Proposed method where Whisper model is conditioned on both audio and its corresponding transcript, fine-tuned to output phrase-level graphemes and TTS labels. We train two models on JSUT and PRIV, referred to as **Annt-v3-pub** and **Annt-v3-priv**.

**Annt-v4:** **Annt-v3** incorporating proposed decoding strategy. We apply it to both **Annt-v3-pub** and **Annt-v3-priv**, referred to as **Annt-v4-pub** and **Annt-v4-priv**.

The TTS model is built upon VITS [25], an end-to-end architecture that can synthesize high-quality speech. We use the official implementation<sup>6</sup> and train all models for 500k steps on a single A100 GPU, adhering to the original configuration, using TTS labels generated by the text-based method and annotation models trained on PRIV. Additionally, we incorporate two more sources to ensure a fairer comparison:

**Reference:** Recorded speech samples.

**Oracle:** TTS model trained on manually annotated labels. Since JVS lacks manual annotations, we only train a model on PRIV.

For each TTS model, we input the labels produced by its corresponding annotation method to generate speech samples for evaluation, ensuring consistency between training and inference.

## 4. Experimental results

We calculate the CER of grapheme sequences predicted by **Annt-v3-priv** on the testing data of PRIV and the result is 0.16%. Interestingly, most of the errors occur when the model changes graphemes to another form identical in pronunciation, such as “プイ” and “ピ”, or when it normalizes numbers, such as converting “1週” to “一週”. One possible reason is the inconsistency of annotation style across human annotators. Therefore, we can assume that the predicted graphemes of each accent phrase align well with the given transcript, ensuring the feasibility of our proposed decoding strategy.

Table 2 shows the objective evaluation results on the PRIV testing data. All annotation models utilizing audio information significantly outperform the text-based method (i.e., **NLP**). Among the models conditioned solely on audio, those that produce both phrase-level graphemes and TTS labels (i.e., **Annt-v2-pub** and **Annt-v2-priv**) perform better than those that only generate TTS labels (i.e., **Annt-v1-pub** and **Annt-v1-priv**). Our proposed models (i.e., **Annt-v3-pub** and **Annt-v3-priv**), relying on audio and the corresponding transcript, are more accurate in terms of both phonemic and prosodic labeling. The

<sup>4</sup><https://open-jtalk.sp.nitech.ac.jp/>

<sup>5</sup><https://github.com/r9y9/pyopenjtalk>

<sup>6</sup><https://github.com/jaywalnut310/vits>

Table 2: Objective evaluation results on each metric. CER is for phonemic labeling while accuracy and  $F_1$  are for prosodic labeling

Model	CER (↓)	Acc. (↑)	$F_1$ (↑)
NLP	3.56%	69.33%	83.69%
Annt-v1-pub	2.41%	81.49%	92.03%
Annt-v2-pub	1.92%	81.79%	92.15%
Annt-v3-pub (ours)	1.59%	<b>81.95%</b>	<b>92.27%</b>
Annt-v4-pub (ours)	<b>1.31%</b>	<b>81.95%</b>	92.25%
Annt-v1-priv	1.13%	86.04%	94.35%
Annt-v2-priv	0.80%	86.40%	94.55%
Annt-v3-priv (ours)	0.63%	<b>87.32%</b>	<b>94.59%</b>
Annt-v4-priv (ours)	<b>0.57%</b>	<b>87.32%</b>	94.58%

Table 3: MOS test results for TTS models trained using TTS labels from different annotation models with 95% confidence intervals.

Model	JVS	PRIV
Reference	4.43 ± 0.08	4.68 ± 0.07
Oracle	-	4.27 ± 0.08
NLP	3.63 ± 0.09	3.85 ± 0.10
Annt-v1-priv	3.87 ± 0.08	4.06 ± 0.09
Annt-v2-priv	3.92 ± 0.09	4.11 ± 0.09
Annt-v3-priv (ours)	4.12 ± 0.09	4.23 ± 0.09
Annt-v4-priv (ours)	<b>4.16 ± 0.08</b>	<b>4.29 ± 0.08</b>

models applying proposed decoding strategy (i.e., **Annt-v4-pub** and **Annt-v4-priv**) achieve the best performance of phonemic labeling with a slight degradation in pitch labeling, which is expected as mentioned in Section 2.3. The accuracy of accent phrase boundary prediction remains unchanged since the decoding strategy is performed within each accent phrase.

Table 3 presents the results of subjective evaluation. Consistent with the objective evaluation results, all methods leveraging audio information significantly outperform the text-based method. Furthermore, our proposed method not only outperforms the baseline methods across all datasets but also achieves scores slightly higher than manual annotation (i.e., **Oracle**) on PRIV. This is probably due to the inconsistent annotation style across human annotators, which is also reported in [10, 11]. The results demonstrate that the proposed method is capable of generating high-quality annotations from audio-transcript pairs, thereby enhancing the performance of Japanese TTS models.

## 5. Conclusion

In this paper, we propose an annotation model fine-tuned from a pre-trained ASR model to produce graphemes along with their phonemic and prosodic labels for given audio-transcript pairs. To address the phonemic labeling issue associated with Kanji characters, we adopt a decoding strategy that incorporates dictionary prior knowledge. Experimental results demonstrate that the proposed annotation method is effective in building high-quality Japanese TTS models. Since the model can generate phrase-level graphemes and TTS labels, it holds potential for constructing polyphone disambiguation datasets and training text-based accent sandhi estimation models in future work.

## 6. References

- [1] S. E. Eskimez, X. Wang, M. Thakker, C. Li, C.-H. Tsai, Z. Xiao, H. Yang, Z. Zhu, M. Tang, X. Tan *et al.*, “E2 tts: Embarrassingly easy fully non-autoregressive zero-shot tts,” *arXiv preprint arXiv:2406.18009*, 2024.
- [2] Z. Du, Q. Chen, S. Zhang, K. Hu, H. Lu, Y. Yang, H. Hu, S. Zheng, Y. Gu, Z. Ma *et al.*, “Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens,” *arXiv preprint arXiv:2407.05407*, 2024.
- [3] S. Liao, Y. Wang, T. Li, Y. Cheng, R. Zhang, R. Zhou, and Y. Xing, “Fish-speech: Leveraging large language models for advanced multilingual text-to-speech synthesis,” *arXiv preprint arXiv:2411.01156*, 2024.
- [4] T. Fujimoto, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, “Impacts of input linguistic feature representation on japanese end-to-end speech synthesis,” in *10th ISCA Speech Synthesis Workshop. ISCA, Vienna, Austria*, 2019.
- [5] K. Kurihara, N. Seiyama, and T. Kumano, “Prosodic features control by symbols as input of sequence-to-sequence acoustic modeling for neural tts,” *IEICE Transactions on Information and Systems*, vol. 104, no. 2, pp. 302–311, 2021.
- [6] Z. Ying, C. Li, Y. Dong, Q. Kong, Q. Tian, Y. Huo, and Y. Wang, “A unified front-end framework for english text-to-speech synthesis,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 10 181–10 185.
- [7] Y. Zhang, L. Deng, and Y. Wang, “Unified mandarin tts front-end based on distilled bert model,” *arXiv preprint arXiv:2012.15404*, 2020.
- [8] B. Park, R. Yamamoto, and K. Tachibana, “A unified accent estimation method based on multi-task learning for japanese text-to-speech,” in *Interspeech 2022*, 2022, pp. 1931–1935.
- [9] M. Omachi, Y. Fujita, S. Watanabe, and M. Wiesner, “End-to-end asr to jointly predict transcriptions and linguistic annotations,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021, pp. 1861–1871.
- [10] Y. Shirahata, B. Park, R. Yamamoto, and K. Tachibana, “Audio-conditioned phonemic and prosodic annotation for building text-to-speech models from unlabeled speech data,” in *Interspeech 2024*, 2024, pp. 2795–2799.
- [11] Z. Dai, J. Yu, Y. Wang, N. Chen, Y. Bian, G. Li, D. Cai, and D. Yu, “Automatic prosody annotation with pre-trained text-speech model,” in *Interspeech 2022*, 2022, pp. 5513–5517.
- [12] X. Yang, “G2pa: G2p with aligned audio for mandarin chinese,” in *Interspeech 2024*, 2024, pp. 2800–2804.
- [13] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [14] A. Vaswani, “Attention is all you need,” *Advances in Neural Information Processing Systems*, 2017.
- [15] F.-T. Liao, Y.-C. Chan, Y.-C. Chen, C.-J. Hsu, and D.-s. Shiu, “Zero-shot domain-sensitive speech recognition with prompt-conditioning fine-tuning,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [16] X. Wei and S. McGregor, “Prompt tuning for speech recognition on unknown spoken name entities,” in *Interspeech 2024*, 2024, pp. 762–766.
- [17] H. Ma, Z. Peng, M. Shao, J. Li, and J. Liu, “Extending whisper with prompt tuning to target-speaker asr,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 12 516–12 520.
- [18] Y. Jiang, T. Wang, X. Xie, J. Liu, W. Sun, N. Yan, H. Chen, L. Wang, X. Liu, and F. Tian, “Perceiver-prompt: Flexible speaker adaptation in whisper for chinese disordered speech recognition,” in *Interspeech 2024*, 2024, pp. 2025–2029.
- [19] T. Kudo, K. Yamamoto, and Y. Matsumoto, “Applying conditional random fields to japanese morphological analysis,” in *Proceedings of the 2004 conference on empirical methods in natural language processing*, 2004, pp. 230–237.
- [20] Y. Den, J. Nakamura, T. Ogiso, and H. Ogura, “A proper approach to japanese morphological analysis: Dictionary, model, and evaluation,” in *LREC*, 2008.
- [21] R. Sonobe, S. Takamichi, and H. Saruwatari, “Jstt corpus: free large-scale japanese speech corpus for end-to-end speech synthesis,” *arXiv preprint arXiv:1711.00354*, 2017.
- [22] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [23] S. Takamichi, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, “Jvs corpus: free japanese multi-speaker voice corpus,” *arXiv preprint arXiv:1908.06248*, 2019.
- [24] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [25] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 5530–5540.