

# COCO-NUT: CORPUS OF JAPANESE UTTERANCE AND VOICE CHARACTERISTICS DESCRIPTION FOR PROMPT-BASED CONTROL

Aya Watanabe, Shinnosuke Takamichi, Yuki Saito, Wataru Nakata, Detai Xin, Hiroshi Saruwatari

The University of Tokyo, Japan.

## ABSTRACT

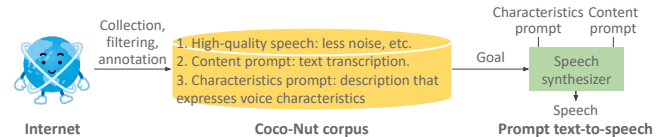
In text-to-speech, controlling voice characteristics is important in achieving various-purpose speech synthesis. Considering the success of text-conditioned generation, such as text-to-image, free-form text instruction should be useful for intuitive and complicated control of voice characteristics. A sufficiently large corpus of high-quality and diverse voice samples with corresponding free-form descriptions can advance such control research. However, neither an open corpus nor a scalable method is currently available. To this end, we develop Coco-Nut, a new corpus including diverse Japanese utterances, along with text transcriptions and free-form voice characteristics descriptions. Our methodology to construct this corpus consists of 1) automatic collection of voice-related audio data from the Internet, 2) quality assurance, and 3) manual annotation using crowdsourcing. Additionally, we benchmark our corpus on the prompt embedding model trained by contrastive speech-text learning.

**Index Terms**— Speech synthesis, speech dataset, voice characteristics, text prompt, crowdsourcing

## 1. INTRODUCTION

In human speech production, the speaker’s voice carries not only linguistic content but also unique vocal characteristics. Text-to-speech (TTS) tasks that imitate the speech production involve two significant challenges: synthesizing highly intelligible speech from the provided text (referred to as “content prompt” in this paper) and controlling the voice characteristics. This is because the characteristics greatly influence the listener’s perception, affecting their understanding of the speaker’s personality, emotion, and overall impression. Several methods of voice characteristics control have been proposed, such as a speaker index [1], speaker attributes [2, 3], personality [4], and so on [5, 6, 7, 8]. However, these methods only enable control over a narrow and simplistic range of voice characteristics, limiting their applicability in various contexts.

There has been significant advancement in techniques for synthesizing media using free-form text descriptions (text prompts). This progress is evident in various fields, such as text-to-image [9], text-to-audio [10], text-to-music [11], and text-to-video [12]. The potential of prompt-based media generation is to manipulate complicated media components, with benefits accruing from the ongoing advancements in large language models (LLMs) [13, 10]. Following these trends, we believe that voice characteristics control by a free-form description opens new doors for TTS tasks. Hence, our goal is to develop TTS capable of controlling vocal characteristics through free-form descriptions, leading to the construction of a dedicated



**Fig. 1.** Our Coco-Nut corpus towards prompt TTS. Characteristics prompt and content prompt are, for example, “middle-aged man’s voice speaking in a clear and polite tone” and “Welcome to our office!” respectively. Speech synthesizer synthesizes speech of the prompted content with the prompted voice characteristics.

corpus. We refer to this free-form description and TTS synthesizer as the “characteristics prompt” and “Prompt TTS,” respectively. As depicted in Figure 1, the aim of Prompt TTS is to synthesize speech that aligns with the prompted linguistic content and voice characteristics. The corpus designed for this purpose should encompass a wide array of vocal characteristics, unlike the existing TTS corpora [14, 15] which tend to cover only a limited range of voice attributes. However, neither an open corpus nor a scalable methodology to construct the corpus is currently available.

In this paper, we propose a methodology for constructing a corpus toward Prompt TTS. Our methodology consists of 1) machine-learning-based automatic collection of voice-related audio data from the Internet, 2) quality assurance to enhance the quality of content prompts and speech in the corpus, and 3) manual annotation of characteristics prompt using crowdsourcing. With this methodology, we construct an open corpus, *Coco-Nut*<sup>1</sup>, which is available at our project page<sup>2</sup>. This paper also benchmarks our Coco-Nut corpus. The Coco-Nut corpus is used for training a contrastive speech-text training model that embeds characteristics prompt and speech into a same space. Experimental evaluation gives results of the corpus construction and performance of the benchmark system both on objective and subjective aspects.

## 2. RELATED WORK

### 2.1. Dataset for text-to-image

Model training for text-to-image requires pairs of an image and text prompt that describes the image content. DALL-E [9], known as a pioneer in text-to-image, is trained using the MS-COCO dataset [16] (image captioning dataset) and web data [17]. MS-COCO is a dataset used for image captioning, which involves manual annotation of texts that describe the image content. In addition to MS-COCO, the use of diverse data from the Internet in training significantly contributes to the synthesis of diverse images [9]. Although HTML images and their accompanying alt-tag texts provide

This work is supported by JSPS KAKENHI 21H04900, 22H03639, 23H03418 (practical experiment), JST FOREST JPMJFR226V, and Moonshot R&D Grant Number JPMJPS2011 (algorithm development). We also appreciate Takaaki Saeki and Yuta Matsunaga of the University of Tokyo for their help.

<sup>1</sup>Corpus of connecting Nihongo utterance and text. “Nihongo” means the Japanese language in Japanese.

<sup>2</sup>[https://sites.google.com/site/shinnosuketakamichi/research-topics/cocanut\\_corpus](https://sites.google.com/site/shinnosuketakamichi/research-topics/cocanut_corpus)

a massive amount of text-image pairs, data filtering is necessary due to noisiness of the Internet data. The pre-trained CLIP (contrastive language-image pretraining) model [13] is often used for data filtering purposes. The importance of data diversity and contrastive learning should be considered in other generation tasks, e.g., voice characteristics in this paper.

## 2.2. Dataset for text-to-audio and text-to-music

As the same to text-to-image, datasets for captioning are also available for text-to-audio. The typical examples are AudioCaps [18] and Clotho [19]. Additionally, the text-audio version of CLIP, CLAP (contrastive language-audio pretraining) model [10], is also used for data filtering [20] before the training. MuLan [11] in text-to-music proposes a method of retrieving music videos on the web and builds a machine learning model to identify whether the text attached to the video describes the music. This methodology has the potential to be applied to other than music.

Unlike the text-to-audio and text-to-music cases, datasets for Prompt TTS are very limited<sup>3</sup>. Existing studies have added characteristics prompts to small in-house and private datasets [21, 22]. However, typical TTS corpora [14, 15] contain only limited voice characteristics. Considering the contribution of Internet data described in Section 2.1, it is necessary to establish a methodology of corpus construction from the Internet data. Also, there is no open corpus that everyone can access.

## 2.3. Sequence generation from text

In sequence generation tasks such as text-to-video and text-to-audio, it is necessary to determine 1) *overall concepts* that represent characteristics of the entire sequence and 2) *sequence concepts* that represent characteristics of changes in the sequence. There are two kinds for describing these concepts using text.

The first is to describe both concepts in a single text, e.g., “wooden figurine surfing on a surfboard in space” [12] in text-to-video and “hip-hop features rap with an electronic backing” [11] in text-to-music<sup>4</sup>. The second is to describe each concept in separate texts. Examples of this include “bat hitting” (overall concept) and “ki-i-i-n” (sequence concept) in text-to-audio [24]<sup>5</sup> and “A toy fireman is lifting weights” (sequence concept) in text-to-video [26]<sup>6</sup>. This kind of methods is suitable for applications that require fine-grained control over the sequence, such as TTS, where the linguistic content and voice characteristics are often controlled separately [21, 22]. Therefore, we aim to collect content and characteristics prompts separately.

# 3. CORPUS CONSTRUCTION

## 3.1. Corpus composition

The corpus for Prompt TTS should include:

1. **High-quality speech.** Speech data for TTS. Unlike data in speech-to-text corpora [27, 28, 29], it should be high-quality, e.g., less noise. Also, it is paired with the content prompt and characteristics prompt.

<sup>3</sup>Audio captioning datasets [18, 19] include human voices as an environmental sound, but the voices do not strongly specify linguistic content.

<sup>4</sup>MusicLM [23] uses a variation of this kind by switching the description at fixed intervals (15 seconds in the paper) to allow for more fine-grained control of changes. This method is suitable for applications that generate sequences from rough descriptions.

<sup>5</sup>LAION-Audio-630K [25] uses text of overall concept for non-speech environmental sounds and that of sequence concept for speech-related environmental sounds.

<sup>6</sup>Overall concept is given by an image in the paper.

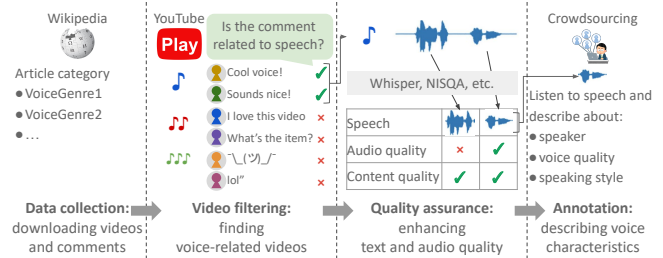


Fig. 2. Procedure of corpus construction.

2. **Content prompt.** Text transcriptions of speech. This corresponds to the sequence concept described in Section 2.3.
3. **Characteristics prompt.** Free-form descriptions that express characteristics of speech. This corresponds to the overall concept described in Section 2.3.

Existing approaches [21, 22] for constructing this kind of corpora are to add characteristics prompts to existing TTS corpora consisting of high-quality speech and content prompts, e.g., [15, 14]. However, as described in Section 2, such corpora often lack diversity of voice characteristics. Therefore, we propose a methodology that builds a corpus from very noisy Internet data.

Our methodology consists of the following four steps. Figure 2 illustrates these steps. Although the target language of this paper is Japanese, the process of these steps is language-independent; our methodology can be implemented in other languages than Japanese.

1. **Data collection.** Speech data candidates are searched out and obtained from the Internet.
2. **Video filtering.** Impressive voice data are filtered from the candidates. The “impressive voice” refers to those that have received a large number of responses on the Internet. Such data are expected to be characteristic voices and therefore suitable for the construction of the corpus.
3. **Quality assurance.** Speech and its transcriptions (content prompts) are further filtered to guarantee quality of the corpus.
4. **Manual annotation.** Characteristics prompts are manually annotated to the speech data.

The subsequent subsections describe the details of these steps.

## 3.2. Data collection

To obtain speech data candidates, we make search phrases and input it into the search engine of video-sharing websites, e.g., YouTube. We select article categories related to speech from Wikipedia<sup>7</sup> in the target language and use the titles of Wikipedia articles belonging to those categories as search phrases. In addition, we add related phrases that are thought to be related to the search phrase (e.g., “[article title] short clip”). After searching, we obtain a video ID, audio data, video title, and viewers’ comments of the found videos.

## 3.3. Video filtering

By filtering through the above obtained video data, videos containing “impressive voice” are acquired. In this paper, we extract videos in which many people commented about voices in the videos. Two-stage filtering are conducted, and the voices of the filtered videos are forwarded to the next “quality assurance” step.

1. **Keyword matching-based pre-filtering.** The obtained data contains many videos without audio or with nondescript voices. First, a rule-based video filter is applied. We use a set of keywords related to voice characteristics (e.g., “listen”) to distinguish whether a viewer’s comment on a video contain those

<sup>7</sup>For example, [https://en.wikipedia.org/wiki/List\\_of\\_YouTubers](https://en.wikipedia.org/wiki/List_of_YouTubers) in English.

keywords. If the number of comments containing the keywords in that video is greater than the threshold, the video is adopted.

2. **Machine learning-based filtering.** Machine learning is used to determine whether viewers' comments mention the voice in the video, obtaining in videos with "impressive voices". We create training data for this machine learning. We randomly extract viewers' comments from videos and perform crowdsourcing-based annotations on the comment mentions. A title and comment of the video are presented to the crowdworkers<sup>8</sup>. The crowdworkers answer whether the comment is 1) related to speaking voice, 2) related to singing voice, or 3) others. Before the annotation, we instruct crowdworkers that "1)" includes comments mentioning about the voice characteristics but does not include comments about the linguistic contents.

A comment content classifier is trained using the above annotated data. The classifier model is BERT [30] followed by a linear layer. The input is a video title and comment, joined by "[SEP]" token that represents "sentence separation" in BERT. The output target is binary: 1) speech-related comment and 2) singing-related comment and others. To improve classification performance, we decided to use the aforementioned keywords auxiliary. A subset of the keyword set was used, and only comments that matched one of the subsets were used to train and evaluate the classifier.

### 3.4. Quality assurance

Due to the collection of Internet data, there are some text and speech data samples that are of low quality and difficult to use. In order to ensure the quality of the audio data included in the corpus, the following processes are used to filter the data.

#### 3.4.1. Audio quality

To ensure the quality of the sound, the following operations are performed.

1. **Voice activity detection (VAD).** VAD is performed to extract only the segments containing voices from the entire video. We use `inaSpeechSegmenter` [31] to detect individual speech segments in the video.
2. **Denoising.** To enhance audio quality, we use `Demucs`<sup>9</sup>, which is a powerful source separation model based on deep neural networks, to extract the voices from noise-contaminated voices.
3. **Audio quality assessment.** There is a variety of audio quality of speech, e.g., recording device quality and effective frequency band. Also, denoising process well eliminates background noise but sometimes drops speech component. To quantify quality degradation caused by these factors, we use `NISQA` [32], a multidimensional speech quality predictor. The `NISQA` score is calculated for each speech segment, and we filtered out the segments with the score lower than the pre-determined threshold<sup>10</sup>.
4. **Threshold for duration and audio volume.** We set the acceptable duration ranges to eliminate too long and too short voices. We also set the volume threshold and filtered out inaudible (low-volume) speech.
5. **Detection of multi-speaker voice and singing voice.** Data not intended for TTS, specifically singing voices and multi-speaker voices (e.g., cheering), are manually excluded.

<sup>8</sup>For example, "Video title: My daily voice training method. Comment: Cool Voice!" The answer will be "1) related to speaking voice." Presenting title makes it easier for the crowdworkers to judge the comment content by having the crowdworkers imagine the content of the video.

<sup>9</sup><https://github.com/facebookresearch/demucs>

<sup>10</sup>We found that speech component drop can be quantified by the `NISQA` score.

6. **Voice characteristics variation.** It is desirable for the corpus to include a variety of voice characteristics. To achieve this, we perform hierarchical clustering based on Ward's method [33] using distances of  $x$ -vectors [34], which reflects not only voice quality but also speech style as suggested by [35]. The  $x$ -vector is extracted for each speech segment by a pretrained  $x$ -vector extractor. Since speech segments with similar voice characteristics are expected to be grouped, we randomly sample one speech segment as the representative of each cluster.

#### 3.4.2. Content quality

To select appropriate speech contents, the following processing steps are performed.

1. **Speech-to-text and language identification.** To obtain content prompts of speech, we use pre-trained `Whisper` speech-to-text models [36]. Jointly with speech-to-text, we identify language of speech by `Whisper` and filtered out speech of non-target language. Furthermore, manual identification is conducted to enhance the corpus quality<sup>11</sup>.
2. **NSFW (not safe for work) word detection.** We filtered out content prompts that include NSFW words. We adopt keyword matching-based NSFW word detection; the text is filtered out if the lemmatized word is found in the NSFW word dictionary. Additional manual detection is conducted to enhance the corpus quality.
3. **Non-verbal voice detection.** Since TTS does not handle non-verbal voices, e.g., scream, we filter out non-verbal voices using a large language model and content prompt texts. Masked language model (MLM) scores [37] based on BERT [30] are calculated for each segment's transcription. Since the masked tokens of content prompt text is highly predictable from the adjacent tokens<sup>12</sup>, the MLM score of the non-verbal voice becomes higher. We manually set a threshold against to the MLM score and filtered out speech with the MLM score higher than the threshold.

### 3.5. Manual annotation

Finally, we use crowdsourcing to add characteristics prompts to the collected voices. The employed crowdworkers listen to the presented voice and describe the voice characteristics. They are instructed to include speaker attributes, voice quality, and speaking style in their descriptions<sup>13</sup>. Only descriptions with more than the threshold number of characters are accepted.

After collecting characteristics prompts, we manually filtered out characteristics prompts that include proper nouns and persons' name, e.g., "The voice is similar to [celebrity's real name]." This is done to prevent models trained on this corpus from generating the voices of actual individuals' name. We also perform text normalization to cleanse the descriptions.

## 4. EXPERIMENTS

### 4.1. Data collection

The target language was Japanese. The data collection period was from July 2022 to March 2023. The number of comments per video

<sup>11</sup>We found that language identification by `Whisper` alone would result in the inclusion of many voices of non-target languages.

<sup>12</sup>For example, let consider "aa[MASK]aaaa," a partially masked content prompt of scream. The masked token "[MASK]" will be "aa."

<sup>13</sup>The actual English-translated instruction is "Describe what kinds of speaker (age, gender, etc.), voice quality (brisk, low voice, etc.), and speaking style (angry, fast, etc.) in a free-form description of at least 20 characters. Do not include the linguistic content of the speech, and do not use expressions that indicate personal likes and dislikes (e.g., my favorite voice and disliked way of speaking)."

Retrieved item	Value
#article-categories	180 categories
#search-phrases	0.10M phrases
#videos found in the search	1.14M videos
Audio duration	0.30M hours
#comments	24.2M sentences

was limited to the top 100 comments with the highest number of “Likes.” After extracting comments in the target language by rule-based language identification, comments with less than 3 characters or more than 50 characters were excluded. Table 1 lists the results of the data collection.

## 4.2. Filtering

For keyword matching-based pre-filtering, we used eight words: “声”, “ボイス”, “ヴォイス” (voice), “響” (resonance), “音” (sound), “聴” (listen), “聞” (hear), and “歌” (song). The threshold for the number of keyword-matching comments per video was 10.

For machine learning-based filtering, we used pre-trained BERT [30] model<sup>14</sup>. We collected 32,453 labels for comments, out of which 11,647 were “speech-related.” 80% and 20% labels were used for training and evaluation, respectively. We attempted to improve the performance by using a keyword subset. We examined using all combinations of the subsets. Finally, seven different subsets and classifiers with high precision were selected. The choice of precision for the selection is to ensure the accurate extraction of “speech-related” comments. The average precision of the seven classifiers was 54.3%. In comparison, when using only the BERT-based classifier without the keyword subsets, the precision was 38.6%. This confirms the effectiveness of using the keyword subsets in combination, as it improves the precision. After training the classifiers, we classified unlabeled comments. Videos were selected if they had 10 or more comments identified as “speech-related” by any of the seven classifiers. Hereinafter, a subset of selected videos, including 1,523, was used for further processing.

## 4.3. Quality assurance

Through VAD, we obtained 54,610 speech segments. Figure 3 shows the distribution of NISQA-predicted mean opinion scores (MOSs) on audio quality. We set the threshold to 2, which is the most frequent score. Also, segments with a duration between 2 seconds and 10 seconds were retained. The audio volume was checked using Pydub<sup>15</sup>, and segments with a volume of  $-55$ dB or lower were excluded.

For transcription using Whisper [36], both the tiny and large model were employed, because the former tends to excel in fidelity to the speech while the latter excels in grammatical correctness<sup>16</sup>.

The NSFW detection was performed by MeCab<sup>17</sup> and the Japanese NSFW dictionary<sup>18</sup>.

The MLM score threshold for non-verbal voice detection was determined to be  $-0.01$ . Figure 4 is the histogram of MLM scores calculated by the transcriptions by Whisper large model. The analysis reveals that the MLM scores of collected segments are distributed around the peak of  $-3$ , with a range of approximately  $\pm 2$  interval.

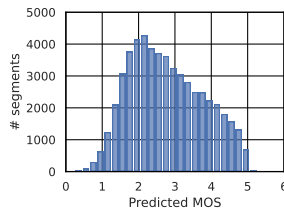
<sup>14</sup><https://huggingface.co/cl-tohoku/bert-base-japanese>

<sup>15</sup><https://github.com/jiaaro/pydub>


<sup>16</sup>The average Word Error Rate (WER) of transcriptions from Whisper large model was 22.1%. Upon final publication, we will provide manually corrected transcriptions to ensure a WER of 0.

<sup>17</sup><https://taku910.github.io/mecab/>


<sup>18</sup><https://github.com/MosasoM/inappropriate-words-ja>




**Fig. 3.** Histogram of NISQA-predicted MOS on speech quality.



**Fig. 4.** Histograms of MLM scores.



**Fig. 5.** Categories of speech.



**Fig. 6.** Gender of speech.

We observe that the percentage of segments whose MLM scores exceeded the threshold  $-0.01$  was approximately 0.05%, which was extremely low frequency.

We used the  $x$ -vectors extracted by `xvector_jtubespeech`<sup>19</sup> for voice characteristics variation. We performed hierarchical clustering and made 11,000 clusters based on voice characteristics similarity. From each cluster, a single audio segment was randomly selected. After the selection, we further conducted manual annotation for whether the segments include NSFW words, non-target language, and multi-speakers. Finally, 7,667 segments, with a total length of 30,661 seconds, were selected.

## 4.4. Annotation

We hired workers through the crowdsourcing platform, Lancers<sup>20</sup>. Each worker annotated 10 segments. There were a total of 1,318 workers, and each worker was paid 200 yen as reward.

Before the annotation, in preparation for the machine learning experiments described at Section 4.6.1, we designed the training, validation, and test sets. To avoid data leakage caused by similar voice characteristics within the same video or YouTube channel, we ensured that the sets have no overlap in YouTube channel and included a diverse range of segments. As a result, we created training, validation, and test sets with 6,463, 593 and 611 segments respectively.

We designed our corpus to include variations introduced by workers. Specifically, for the training set, we included one characteristics prompt per segment, while five prompts per segment for the other sets, following the existing studies [19, 38].

## 4.5. Corpus analysis

We analyze the constructed corpus. Particularly, we investigate data diversity, which is our aim of the corpus.

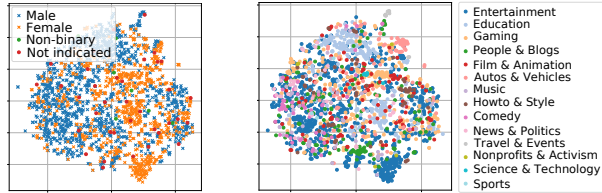
### 4.5.1. Video categories

We investigated which video category speech segments in the corpus belonged to. The source video of each segment was classified according to the YouTube video category. Figure 5 shows the results. The corpus contains 14 categories, indicating that it covers a

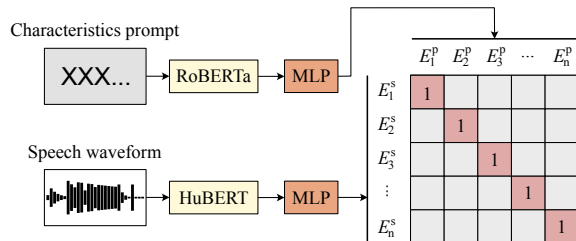
<sup>19</sup>[https://github.com/sarulab-speech/xvector\\_jtubespeech](https://github.com/sarulab-speech/xvector_jtubespeech)

<sup>20</sup><https://www.lancers.jp>





**Fig. 7.**  $x$ -vector distributions colored by gender. **Fig. 8.**  $x$ -vector distributions colored by category.



**Fig. 9.** Baseline model architecture. MLP means multi-layer perceptron.  $E^s$  and  $E^p$  mean the  $n$ -dimensional embeddings of speech and characteristics prompt, respectively.

wide variety of categories. The top three categories (entertainment, education, gaming) account for approximately 70%, and minor categories such as Science & Technology are also included.

#### 4.5.2. Gender distributions

We manually annotated gender to the characteristics prompts to analyze gender diversity. Figure 6 presents the distribution of gender. While the majority of characteristics prompts are labeled as male or female, non-binary and some prompts that don’t mention gender (not-indicated) still exist. Similar to a typical TTS corpus, clusters can be observed for male and female voices. Non-binary and not indicated categories do not form distinct clusters but are scattered throughout. To provide a detailed analysis, we present the t-SNE visualization of  $x$ -vectors colored by gender in Figure 7. There are clusters for male and female voices. However, non-binary and not indicated categories do not form clusters but appear scattered.

#### 4.5.3. Voice characteristics of video categories

To examine the relation between  $x$ -vector and video categories, we present the t-SNE visualization of  $x$ -vectors colored by the video category in Figure 8. In the Entertainment and Education categories, specific clusters can be observed, particularly in the bottom-right and top-central regions. This suggests that typical voice characteristics are gathered within each category. On the other hand, for the majority of the scatter plot, no prominent clusters are observed. This indicates that the speech in this corpus encompass both typical voices within categories and voices that are shared across categories.

### 4.6. Machine learning baseline

Using the constructed corpus, we conduct machine learning experiments to align speech and characteristics prompts. These gives future directions of Prompt TTS.

#### 4.6.1. Model construction

We constructed a baseline model that aligns speech and characteristics prompts. The model was inspired by CLAP [39], the model that embeds both audio and text into the same space by contrastive learning. While HTS-AT [40] is used as the audio encoder for CLAP, we changed it into HuBERT [41] to grasp speech features well.

**Table 2.** Zero-shot gender classification

		Classification result	
		Male	Female
Actual gender	Male	3442	1456
	Female	1048	4051

We used `japanese-roberta-base`<sup>21</sup> and `japanese-hubert-base`<sup>22</sup> as pre-trained models of RoBERTa [42] and HuBERT, respectively. Figure 9 shows the overview of baseline model architecture. Most of hyperparameters followed official implementation of CLAP<sup>23</sup>. The batch size was set to 48, and the learning rate was 0.0001. We used 8 GPUs, NVIDIA A100 for NVLink 40GiB HBM2. The training process took approximately 1 hour.

#### 4.6.2. Evaluation tasks

Following the CLAP paper [10], we evaluate the trained model and obtained embeddings.

**Speech retrieval from characteristics prompt.** We calculate the cosine similarity between the embeddings of the input prompt and the set of embeddings of target speech segments. A higher cosine similarity value indicates a higher-ranked retrieval result. We evaluate whether the proper segment can be retrieved by the prompt.

**Zero-shot speech classification.** We automatically generate characteristics prompts using categorical labels, such as “a voice of [label].” Then a prompt closest to the audio segment in the embedding space is selected. The label associated with that prompt is considered as the classification label for that speech. We evaluate whether the correct label can be obtained without additional training.

#### 4.6.3. Objective evaluation

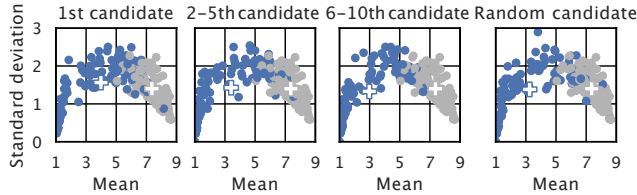
We evaluated our model using mean average precision at top 10 retrieval (mAP@10) following [10]. mAP@10 is an evaluation metric that measures how accurately the speech corresponding to each characteristics prompt is retrieved within the top 10 retrievals. At the best epoch, the text to speech mAP@10 on the test set reached 8.63%, while it was 0.54% before the training. In comparison to what was trained specifically for environmental sounds [10], the obtained value may appear lower. However, it is important to note that mAP@10 will be 10% when the 10th candidate in every retrieval is the correct pair. Therefore, an 8% value can be considered a reasonable indication of learning to a certain extent.

To test whether the model recognizes the simple characteristics of the speech with the unseen data, we conducted the gender classification on JVS [15] parallel100 set, which consists of 49 male speakers and 51 female speakers, with each speaker having 100 speech samples. Using labels of two gender, we made two characteristics prompts: 男性の声 and 女性の声, which meant “a male voice” or “a female voice” and retrieved one prompt closest to the JVS speech in the embedding space. The gender of the retrieved prompt is considered as the gender of the speech. For example, if a male speaker’s speech retrieved the prompt “a male voice,” then the classification would be correct. Table 2 shows the confusion matrix of the result. It is observed that both genders’ data are correctly identified at around 70%, indicating that the model had effectively learned to associate the speech of one gender with the text indicating same gender.

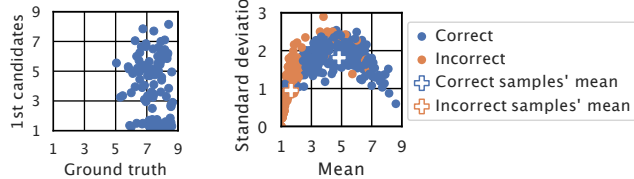
<sup>21</sup><https://huggingface.co/rinna/japanese-roberta-base>

<sup>22</sup><https://huggingface.co/rinna/japanese-hubert-base>

<sup>23</sup><https://github.com/LAION-AI/CLAP>



**Fig. 10.** Mean and standard deviation of subjective evaluation on each prompt-speech pair. Blue circles indicate retrieved pairs of the figure title, and gray ones indicates ground-truth. “+” marks indicate average mean and standard deviation of same color plots.



**Fig. 11.** Mean relationship of ground truth results and 1st retrieved results, both of the same characteristics prompt.

**Fig. 12.** Mean and standard deviation of subjective evaluation on each of prompt and retrieved speech, colored by whether the retrieved speech possesses gender described in the prompt.

#### 4.6.4. Subjective evaluation

To perceptually evaluate the retrieved speech by the prompt, we conducted subjective evaluations.

We randomly selected 100 characteristics prompts from the test set and retrieved speech from the whole test set. From retrieval results, we created four kinds of speech paired with the characteristics prompt: 1st candidate, 2–5th candidates, 6–10th candidates, and random candidates. The first means the 1st candidate retrieved by the prompt, and the last is randomly selected from the test set.

We presented the prompt and speech to crowdworkers and let them evaluate how much the prompt represents the speech characteristics on a nine-point scale, 9 is the best matching and 1 is the worst. For comparison, we added ground-truth speech (paired with the prompt in the test set) to the listening test set. Each worker was presented with a total of 20 pairs. We employed 500 workers and obtained 20 evaluations for each pair.

Figure 10 shows the results. Each figure illustrates mean and standard deviation of the scores of each prompt-speech pair. Those of ground-truth pairs are illustrated for comparison.

**Q. Is free-form text truly appropriate in describing the voice characteristics?** We validate the adequacy of the free-form characteristics representation we present in this paper. As shown in Figure 10, the ground-truth pairs obtained a sufficiently high average score of 7.37 despite the difference between the writers of the free-form expressions and the evaluators. It indicates that the free-form expressions can appropriately represent voice characteristics regardless of the writers or evaluators. Note that, compared to conventional categorization (e.g., gender), the scores tend to be more variable.

**Q. Does the baseline model retrieve perceptually good speech from the given prompt?** We compared between retrieval rank (1st, 2–5th, 6–10th candidates and random). The average scores for each method were 3.98 (1st), 3.42 (2–5th), 2.98 (6–10th) and 3.25 (random). Statistically significant differences were observed between

**Table 3.** Retrieved pair examples. “Score” means the average of subjective evaluated appropriateness. “Rank” indicates the rank of retrieved candidate. Text of retrieved candidate means one characteristics prompt of retrieved candidate. Each text has been translated to English.

Retrieval text	Score	Rank	Retrieved candidate	Score
A young woman in her twenties is speaking slowly with a sweet, clinging voice.	8.40	1st	A young woman is speaking in a cloyingly sweet tone.	8.15
A young woman in her twenties is speaking with a cheerful voice.	8.10	1st	A young man is speaking in a gentle voice, whispering softly.	1.05
A middle-aged cheerful man is speaking in a clear voice, addressing in a questioning manner.	7.75	2–5th	A young man is speaking in a high-pitched voice, as if he is excited.	2.35

the 1st candidate and random ( $p < 0.05$ ), indicating that the model can retrieve appropriate speech. However, there is still room for improvement in the trained model to reach the ground-truth score, and the samples in the 6–10th positions fall below the random method, indicating the need for improvement in the retrieval method.

**Q. Is the low score for 1st candidate due to the low ground-truth score?** As mentioned above, the scores of the ground-truth samples vary among the samples. We investigated whether this variability in ground-truth scores affects the retrieval performance of the model. Figure 11 illustrates the ground-truth scores along with the 1st candidate’s corresponding to each prompt. From this figure, there is no clear correlation between the two, indicating that the variability in ground-truth scores has little impact on the retrieval performance. Therefore, the low scores of the 1st candidate primarily will reflect the performance of the retrieval model itself.

**Q. What happens when scores are extremely low?** Retrieved candidates includes extremely low scores as observed in the bottom left of Figure 10. To investigate the reason behind this, we examined the correspondence between the gender of the prompt and the gender mentioned in the ground-truth prompt associated with the retrieved speech. As shown in Figure 12, in the samples with low scores, it frequently occurs that the gender is misaligned. For example, there are cases where the input prompt includes the term “female,” but our model retrieves a male voice. To address this issue, we need a training method that embeds the same gender samples close.

**Q. What are the actual examples?** Finally, we provide examples of the input prompt and the corresponding ground-truth prompt for the retrieved speech in Table 3. In the case of high scores (8.15), we can observe that not only the gender and age (“young women”) but also the style (“sweet”) are aligned. As mentioned earlier, when the gender is different, the scores significantly drop (1.05). On the other hand, even when the gender is aligned, if there are differences in the age group (“middle-aged” vs. “young”) or style (“questioning manner” vs. “excited”), we can see the low score (2.35).

## 5. CONCLUSION

In this paper, we developed a paired corpus of speech and characteristics prompts and conducted evaluations of both the corpus itself and a baseline model. This corpus will promote research on Prompt TTS, where the speaker is controlled by characteristics prompts. The consideration of Prompt TTS architecture and the expansion of the corpus itself are tasks for future work.

## 6. REFERENCES

- [1] Nobukatsu Hojo, Yusuke Ijima, and Hideyuki Mizuno, “DNN-based speech synthesis using speaker codes,” *IEICE TRANSACTIONS on Information and Systems*, vol. E101-D, no. 2, pp. 462–472, 2017.
- [2] Daisy Stanton, Matt Shannon, Soroosh Mariooryad, RJ Skerry-Ryan, Eric Battenberg, Tom Bagby, and David Kao, “Speaker generation,” in *Proc. ICASSP*, 2022, pp. 7897–7901.
- [3] Aya Watanabe, Shinnosuke Takamichi, Yuki Saito, Detai Xin, and Hiroshi Saruwatari, “Mid-attribute speaker generation using optimal-transport-based interpolation of gaussian mixture models,” in *Proc. ICASSP. IEEE*, 2023, pp. 1–5.
- [4] Joakim Gustafson, Jonas Beskow, and Eva Szekely, “Personality in the mix - investigating the contribution of fillers and speaking style to the perception of spontaneous speech synthesis,” in *Proc. Speech Synthesis Workshop*, 2021, pp. 48–53.
- [5] Yu Zhang, Ron J. Weiss, Heiga Zen, Yonghui Wu, Zhifeng Chen, R.J. Skerry-Ryan, Ye Jia, Andrew Rosenberg, and Bhuvana Ramabhadran, “Learning to Speak Fluently in a Foreign Language: Multilingual Speech Synthesis and Cross-Language Voice Cloning,” in *Proc. Interspeech*, Graz, Austria, Sep. 2019, pp. 2080–2084.
- [6] Rui Liu, Berrak Sisman, and Haizhou Li, “Reinforcement learning for emotional text-to-speech synthesis with improved emotion discriminability,” in *Proc. Interspeech*, Aug. 2021, pp. 4648–4652.
- [7] Kumi Ohta, Tomoki Toda, Yamato Ohtani, Hiroshi Saruwatari, and Kiyohiro Shikano, “Adaptive voice-quality control based on one-to-many eigenvoice conversion,” in *Proc. Interspeech*, 2010, pp. 2158–2161.
- [8] Tuomo Raitio, Ramya Rasipuram, and Dan Castellani, “Controllable Neural Text-to-Speech Synthesis Using Intuitive Prosodic Features,” in *Proc. Interspeech*, 2020, pp. 4432–4436.
- [9] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever, “Zero-shot text-to-image generation,” *arXiv:2102.12092*, 2021.
- [10] Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang, “CLAP: Learning audio concepts from natural language supervision,” *arXiv:2206.04769*, 2022.
- [11] Qingqing Huang, Aren Jansen, Joonseok Lee, Ravi Ganti, Judith Yue Li, and Daniel P. W. Ellis, “MuLan: A joint embedding of music audio and natural language,” *arXiv:2208.12415*, 2022.
- [12] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Salimans, “Imagen Video: High definition video generation with diffusion models,” *arXiv:2210.02303*, 2022.
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever, “Learning transferable visual models from natural language supervision,” in *Proceedings of the 38th International Conference on Machine Learning*, Marina Meila and Tong Zhang, Eds. 18–24 Jul 2021, vol. 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763, PMLR.
- [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.
- [15] Shinnosuke Takamichi, Ryosuke Sonobe, Kentaro Mitsui, Yuki Saito, Tomoki Koriyama, Naoko Tanji, and Hiroshi Saruwatari, “JSUT and JVS: Free Japanese voice corpora for accelerating speech synthesis research,” *Acoustical Science and Technology*, vol. 41, no. 5, pp. 761–768, 2020.
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár, “Microsoft COCO: Common objects in context,” *arXiv:1405.0312*, 2014.
- [17] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut, “Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Melbourne, Australia, July 2018, pp. 2556–2565, Association for Computational Linguistics.
- [18] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim, “Audiocaps: Generating captions for audios in the wild,” in *NAACL-HLT*, 2019.
- [19] Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen, “Clotho: an audio captioning dataset,” in *Proc. ICASSP*, 2020, pp. 736–740.
- [20] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov, “Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation,” *arXiv:2211.06687*, 2022.
- [21] Zhifang Guo, Yichong Leng, Yihan Wu, Sheng Zhao, and Xu Tan, “PromptTTS: Controllable text-to-speech with text descriptions,” *arXiv:2211.12171*, 2022.
- [22] Dongchao Yang, Songxiang Liu, Rongjie Huang, Guangzhi Lei, Chao Weng, Helen Meng, and Dong Yu, “InstructTTS: Modelling expressive TTS in discrete latent space with natural language style prompt,” *arXiv:2301.13662*, 2023.
- [23] Andrea Agostinelli, Timo I. Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, Matt Sharifi, Neil Zeghidour, and Christian Frank, “MusicLM: Generating music from text,” *arXiv:2301.11325*, 2023.
- [24] Hien Ohnaka, Shinnosuke Takamichi, Keisuke Imoto, Yuki Okamoto, Kazuki Fujii, and Hiroshi Saruwatari, “Visual onoma-to-wave: environmental sound synthesis from visual onomatopoeias and sound-source images,” in *Proc. ICASSP. IEEE*, 2023, pp. 1–5.
- [25] Yusong Wu, Ke Chen, Tianyu Zhang, Marianna Nezhurina, and Yuchen Hui, “LAION-Audio-630K,” 2022, <https://github.com/LAION-AI/audio-dataset>.
- [26] Eyal Molad, Eliahu Horwitz, Dani Valevski, Alex Rav Acha, Yossi Matias, Yael Pritch, Yaniv Leviathan, and Yedid Hoshen, “Dreamix: Video diffusion models are general video editors,” *arXiv:2302.01329*, 2023.
- [27] Guoguo Chen, Shuzhou Chai, Guanbo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, et al., “GigaSpeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio,” *arXiv preprint arXiv:2106.06909*, 2021.

- [28] Shinnosuke Takamichi, Ludwig Kürzinger, Takaaki Saeki, Sayaka Shiota, and Shinji Watanabe, “JTubeSpeech: corpus of Japanese speech collected from YouTube for speech recognition and speaker verification,” *arXiv:2112.09323*, 2021.
- [29] Seiji Fujimoto Yue Yin, Daijiro Mori, “ReasonSpeech: A free and massive corpus for Japanese ASR,” in *Annual meetings of the Association for Natural Language Processing*, 2023.
- [30] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [31] David Doukhan, Jean Carrive, Félicien Vallet, Anthony Larcher, and Sylvain Meignier, “An open-source speaker gender detection framework for monitoring gender equality,” in *Proc. ICASSP*. IEEE, 2018.
- [32] Gabriel Mittag, Babak Naderi, Assmaa Chehadi, and Sebastian Möller, “NISQA: A deep CNN-self-attention model for multidimensional speech quality prediction with crowdsourced datasets,” in *Proc. Interspeech*, 2021, pp. 2127–2131.
- [33] Joe H. Ward Jr., “Hierarchical grouping to optimize an objective function,” *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 236–244, 1963.
- [34] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *Proc. ICASSP*. IEEE, 2018, pp. 5329–5333.
- [35] Andrew Brown, Jaesung Huh, Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Playing a part: Speaker verification at the movies,” in *Proc. ICASSP*. IEEE, 2021, pp. 6174–6178.
- [36] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [37] Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff, “Masked language model scoring,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online, July 2020, pp. 2699–2712, Association for Computational Linguistics.
- [38] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim, “Audiocaps: Generating captions for audios in the wild,” in *NAACL-HLT*, 2019.
- [39] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov, “Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation,” in *Proc. ICASSP*. IEEE, 2023, pp. 1–5.
- [40] Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov, “Hts-at: A hierarchical token-semantic audio transformer for sound classification and detection,” in *Proc. ICASSP*. IEEE, 2022, pp. 646–650.
- [41] 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.
- [42] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.