



Factor-Conditioned Speaking-Style Captioning

Atsushi Ando, Takafumi Moriya, Shota Horiguchi, Ryo Masumura

NTT Corporation, Japan

atsushi.ando@ntt.com

Abstract

This paper presents a novel speaking-style captioning method that generates diverse descriptions while accurately predicting speaking-style information. Conventional learning criteria directly use original captions that contain not only speaking-style factor terms but also syntax words, which disturbs learning speaking-style information. To solve this problem, we introduce factor-conditioned captioning (FCC), which first outputs a phrase representing speaking-style factors (e.g., gender, pitch, etc.), and then generates a caption to ensure the model explicitly learns speaking-style factors. We also propose greedy-then-sampling (GtS) decoding, which first predicts speaking-style factors deterministically to guarantee semantic accuracy, and then generates a caption based on factor-conditioned sampling to ensure diversity. Experiments show that FCC outperforms the original caption-based training, and with GtS, it generates more diverse captions while keeping style prediction performance.

Index Terms: speaking-style captioning, large language models, pre-trained speech encoder, factor-conditioned captioning

1. Introduction

Understanding the speaking styles of speakers, including the speaker's attributes and prosodic information, is important in speech applications. These applications include flexible control in speech synthesis [1–3] or voice conversion [4], understanding and controlling speaking behavior for human-like spoken dialogue systems [5, 6], and supporting other speech processing technologies, such as using auxiliary labels for recognizing non-linguistic or para-linguistic information [7, 8].

Traditional studies formulate speaking-style recognition as audio classification into pre-defined speaking-style classes such as style categories, e.g., spontaneous or infant-directed [9, 10]. However, their applicability is limited due to the difficulty of predefining diverse speaking styles in real environments. In response, recent studies have tackled speaking-style captioning, which predicts speaking styles in a free descriptive format [11–13]. This eliminates the need to predefine speaking-style classes, since detailed speaking-style information, such as intensity and detailed expressions, can be used instead. Conventional style captioning models are composed of a speech encoder, bridge network, and text decoder. The speech encoder first extracts speech features from the input audio; the bridge network then converts these features into representations suitable for input to the text decoder, which generates a speaking-style caption from the representations. Pre-trained autoregressive large language models (LLMs) are used as the text decoder to generate natural and diverse captions.

One of the challenges of speaking-style captioning is to

generate captions with accurate speaking-style factors, such as information of gender, pitch, power, and speaking speed. However, using the conventional training method based on cross-entropy with ground-truth captions can make it difficult to ensure that the captioning model acquires speaking-style knowledge. The ground-truth captions contain not only speaking-style factor terms but also syntax words, which disturbs to focus solely on learning the correct speaking styles. Furthermore, even in the inference step, the conventional models suffer from errors in the prediction of speaking style factors if sampling-based decoding [14] is used. While such decoding is widely used to ensure caption diversity, it can lead to the occurrence of less probable words, potentially generating words that represent incorrect style factors.

This paper presents a simple but effective model training method named *factor-conditioned captioning (FCC)* that can predict speaking-style factors accurately. The key idea of FCC is to make the captioning model explicitly predict only style factors as intermediate reasoning steps before generating style captions, like the Chain-of-Thought Prompting approach [15]. The ground truth of FCC contains a factor phrase, the phrase representing speaking style factors, whose format is fixed, e.g., “*male, low pitch, high volume, normal speed*”, in the first part, and the original caption in the second part. The factor phrase can be generated by predicting the speaking-style factors from the original caption using GPT [16], or by the ground-truth factors if available. Our FCC-based training proposal forces the text decoder to consider style factors in generating captions since the decoder predicts the next word based on the preceding text. One of the advantages of FCC is that it is applicable simply by changing the output sentence in the training step, regardless of the model structure or loss function.

For decoding, we introduce *greedy-then-sampling (GtS)* to generate diverse captions while considering the style factors. GtS predicts the factor phrase by using a maximum likelihood criterion, followed by random sampling for the caption part. This enables the generation of diverse captions in terms of expressions and syntax while the conditioning factors are determined by the hypothesis with the highest probability, which prevents the generalization of incorrect style factors.

Experiments on the PromptTTS dataset [1] demonstrate that FCC significantly improves the performance of both the speaking-style captioning and individual speaking-style factor classification derived from captions compared to the conventional method. Furthermore, FCC with GtS decoding yields more diverse captions while preserving accuracy of speaking style factor recognition. Note that we conduct comprehensive comparisons of the components of the captioning model to make a strong baseline for speaking style captioning, and thus validate the superior effectiveness of the proposed methods.

2. Related work

Speaking-style captioning is often included as a task within speech understanding, enabling a single model to address multiple speech-related tasks. Several prior studies proposed speech understanding models [12, 13, 17–22]. Though the previous studies employed various types of modules in the individual components, few have reported comparisons of the components. To rectify this omission, this study conducts comprehensive evaluations of multiple methods for each of the three components in speaking-style captioning to investigate the optimal model structure and how each component affects the final result. Note that the proposed methods, including FCC and GtS, show significant improvements in the best model.

3. Proposed speaking-style captioning

Let S, I be an input speech and an input instruction sentence to specify the task, with O is being corresponding output sentence. Speaking-style captioning is formulated as the problem of predicting output sentence O from S and I ,

$$\hat{O} = \arg \max_{O} P(O | S, I; \Theta), \quad (1)$$

where Θ is the set of captioning model parameters. The instruction sentence is fixed to learn a single task, e.g., speaking style captioning.

3.1. Captioning model

Similar to the conventional studies [11, 12, 18, 19, 22], the captioning model consists of three blocks: speech encoder, bridge network, and text decoder, as shown in Figure 1.

The speech encoder extracts multiple sequences of speech representation vectors $Z = [z_1, \dots, z_K]$ from the input speech,

$$Z = \text{SpeechEnc}(S; \Theta_s), \quad (2)$$

where $z_k \in \mathbb{R}^{V \times T}$ denotes the T -length sequence of V -dimensional vectors output from the k -th layer of the speech encoder, which consists of K layers. Θ_s denotes the set of parameters of the speech encoder that are pre-trained and frozen during the training of the captioning model. The use of multiple intermediate layers enables the model to extract different types of speech information [23, 24].

The bridge network converts the speech representations into speech embeddings suitable for input to the text encoder:

$$e = \text{Bridge}(Z; \Theta_b), \quad (3)$$

where $e \in \mathbb{R}^{D \times N}$ is the D -length N -dimensional vector. Θ_b represents the set of bridge network parameters. As the bridge network, we use a time and layer-wise Transformer (TLTR) [25] so as to combine multi-level information from different layers of the speech encoders into a V -dimensional embedding of length 1, i.e., $N=1$. Finally, the embedding is projected into the same dimensions as the input of the text decoder by a fully connected layer.

The text decoder predicts the posterior probabilities of the output sentence from the speech embeddings e and embedded input instruction sentence v as

$$v = \text{TextEmb}(\text{Tokenizer}(I); \Theta_t), \quad (4)$$

$$P(O) = \text{TextDec}(v, e; \Theta_d), \quad (5)$$

where $v \in \mathbb{R}^{D \times M}$ and M is the embedding sequence length. Θ_t, Θ_d represent the set of parameters for the text embedding

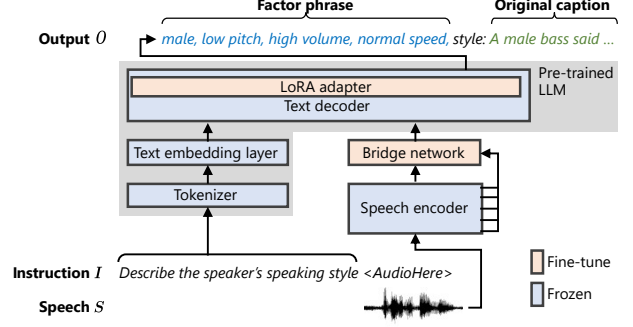


Figure 1: The outline of the proposal, FCC.




Figure 2: Generation of the ground-truth output in FCC.

layer and the text decoder, respectively. The proposed method employs an autoregressive text decoder and the corresponding text tokenizer derived from a pre-trained LLM such as LLaMA-2 [26]. Furthermore, in the text decoder, an adapter based on low-rank adaptation (LoRA) [27] is inserted into the original decoder to adapt the output sentences of the training data. That is, the set of the text decoder parameters composed of the parameter sets of the adapter Θ_{d_a} and the original LLM decoder Θ_{d_o} , where $\Theta_d = \Theta_{d_o} \cup \Theta_{d_a}$. The parameters of Θ_t and Θ_{d_o} are frozen during training, while Θ_{d_a} are updated.

3.2. Training by factor-conditioned captioning

To help the captioning model learn speaking-style factors explicitly in addition to syntactic information, FCC uses the extended version of the original captions as ground truth. Each extended caption consists of two parts: factor phrase, i.e., comma-separated terms, each of which represents a speaking-style factor, as the first part, and the original caption as the second part. The use of this output enables the captioning model to predict the speaking-style information derived from speech first and generate the caption with explicit consideration of speaking style factors since the text decoder is autoregressive.

Readers may find it difficult to pre-define such factors, but this study avoids that difficulty by automatically generating them from the original captions alone. As shown in Figure 2, speaking-style factors are predicted by GPT [16] with a specific prompt¹, and so a template is used to create the factor phrase. The factors unspecified in the caption are regarded as *normal*. The template of the factor phrase is fixed during training, which makes the captioning model focus on predicting the factor information. Of course, if ground-truth style factor labels are available, they can be used instead of predictions by GPT.

¹“Please select the information contained in the input sentence. The options are as follows:
Gender: male or female
Pitch: low, normal, high (normal if unspecified)
Volume: low, normal, high (normal if unspecified)
Speaking Speed: slow, normal, fast (normal if unspecified)

Input: {caption}.”, where
 represents line break.

3.3. Greedy-then-sampling decoding

The common sampling strategies are not suitable for FCC because there is a risk of outputting style factors with low probability, which will yield caption errors. Thus, we also propose GtS, which is a decoding strategy for FCC in generating captions with diverse expressions. In GtS, decoding begins with a greedy search criterion until a delimiter token that connects the factors and caption (*'style:'* in Figure 1) is generated. Then, it switches to a sampling criterion [14] to generate the main speaking style caption. The greedy search ensures the correctness of the style factors by simply using those of the highest posterior probabilities, while the latter sampling allows the generation of diverse captions constrained on the predicted style factors.

4. Experiments

4.1. Setup

The PromptSpeech dataset provided by PromptTTS [1] was used for our evaluations. The audio is from the English Audiobook corpus, LibriTTS [28], and the labels include annotated speaking-style captions, as well as discrete speaking-style factors for gender (male/female), pitch (high/normal/low), volume (high/normal/low), and speaking speed (fast/normal/slow). The dataset was split into the official speaker-open subsets as in the conventional work of [11]. The training set consists of 24,953 utterances by 1,113 speakers, the development set of 857 utterances by 40 speakers, and the test set of 778 utterances by 38 speakers. The distributions of each class of the style factors were roughly equal. Note that from the total of 26,588 utterances, there were 154 unique captions.

The speech encoder was Whisper large-v3 [29] (denote as Whisper-L), consisting of 32 intermediate layers with 1,280 units each; the input was fixed to 30 seconds, i.e., 1,500 frames. As TLTR in the bridge network (TLTR-utt), initial pooling of intermediate layer outputs was performed every 20 frames, and both the time and layer-wise Transformers were single-layer Transformers with 512 hidden units. The text decoder was LLaMA-2 7B-chat [26], and the LoRA was the attention-based method [27] with $r=8$ and $\alpha=32$ (7B-chat-LoRA). The trainable/total parameters in each component were 0M/637M in the speech encoder, 90.7M/90.7M in the bridge network, and 4.2M/6.7B in the text decoder, respectively. The conventional method was StyleCap [11] consisting of WavLM-base-plus [23] speech encoder, weighted-sum + BLSTM + Q-Former [30]-based bridge network, and LLaMA-2-7B decoder without LoRA. The batch size was 16. The optimizer was AdamW [31] with learning rate of 1×10^{-4} . The number of epochs was set to 5, with early stopping based on the development set. Note that we did not pre-train the bridge network using the close-ended style classification tasks [12]². The baseline methods were the models trained by the original caption with greedy or sampling-based decoding. Two methods of creating FCC were evaluated: the case when factor phrases were created using the predicted labels using GPT (denoted as w/ predicted factors), and those created using discrete speaking-style factor labels provided in the dataset (w/ golden factors). The *gpt-4-0613* model with the same prompt as described in Section 3.2 was used for the former FCC. We evaluated the proposed FCC with greedy, sampling, and GtS decoding. In the inference step, the thresholds of the probability and the number were set to 0.9

²Our preliminary experiments found there were no statistically significant improvements with bridge-network pre-training.

and 40, respectively, for both sampling and GtS decoding.

The evaluation tasks included not only speaking-style captioning but also style factor classification. The style factor classification involved using GPT to estimate factor classes for gender, pitch, volume, and speaking rate from captions, and then evaluating the accuracy against annotated discrete labels. As the mapping from caption to four speaking style factors, the predicted caption was given to the *gpt-4-0613* model with the same prompt as described in Section 3.2, then the output sentence was parsed and mapped to several non-target terms like *'unspecified'* to *'normal'* in post-processing. Note that we created a codebook of GPT input/output pairs to ensure that the same factor prediction results for the same captions were returned in all the experiments to avoid the effect of sampling-based decoding on GPT. The evaluation metrics for speaking-style captioning were BLEU-4 (B@4), ROUGE-L (ROU), METEOR (MET), BERTScore with *distilbert-base-uncased* [32] (BS), and distinct-1 and 2 (dis-1/2), as a subset of the baseline study [11]. The first four metrics were implemented by Hugging Face evaluate [33].

4.2. Results

4.2.1. Main results

The results on speaking-style captioning and style factor classification are listed in Table 1. First, as shown in the bottom line, it was found that the accuracies of factor class predictions from ground-truth captions using GPT were more than 99 % in all the factors. This suggests that GPT can extract speaking style factors from captions almost perfectly. Compared to the captioning model, the first and the second rows, our model significantly outperformed the conventional model in all tasks³. As discussed in the following sections, all the components in our model were improved. With the same model structures shown in the second, fourth, and sixth rows, FCC showed improvements in BLEU-4, ROUGE-L, METEOR, and BERTScore. These differences were statistically significant according to the bootstrap-based test ($p < .05$, $n=1000$) except for FCC with predicted factors in ROUGE-L. On the other hand, the performance of style factor classification was also improved but the differences were not significant. These indicate that both the conventional captioning system and FCC can learn information of speaking-style factors from input speech, but FCC yields better captions. One possible reason is that the factor prediction in the first step helps the model to focus on solving syntactic tasks in the caption generation step. Comparing greedy and sampling-based search on caption and FCC-based models, the sampling search yielded improved output diversity as shown by distinct-1/2, but degraded both captioning and factor classification accuracies. On the other hand, FCC with GtS showed no statistical degradation in either captioning or factor classification metrics from greedy search, while statistically improving caption diversity as shown by distinct-1/2 ($p < .05$, $n=1000$). These results show that the proposed FCC with GtS generates accurate and diverse captions. Note that FCC with sampling-based search exhibited significant degradations on both tasks, which confirms that the sampling-based prediction of style factors is inappropriate. It is finally worth mentioning that in all metrics, there is no statistical degradation between the proposed FCC with predicted factors and that with golden factors.

³The StyleCap results were with sentence rephrasing augmentation, while ours were not. Our preliminary experiments found that rephrasing yielded no significant improvement in most setups.

Table 1: Performance of speaking-style captioning and accuracy (%) of style factor classification from captions using GPT. The baselines are our model, trained by original captions with greedy or sampling-based decoding. The best results are highlighted in **bold**. Those metrics calculated using the ground-truth captions are also shown at the bottom for reference.

Model	Output	Decoding	Speaking-style captioning						Style factor classification				
			B@4↑	ROU↑	MET↑	BS↑	dis-1↑	dis-2↑	Gender	Pitch	Speed	Volume	Avg.
StyleCap [11]	Caption	Greedy	0.273	0.497	0.469	0.855	0.023	0.073	91.0	61.0	85.2	69.9	76.8
Ours (Whisper-L, TLTR-utt, 7B-chat-LoRA)	Caption	Greedy	0.613	0.710	0.698	0.925	0.017	0.051	97.9	91.4	93.3	97.9	95.1
		Sampling	0.594	0.691	0.679	0.920	0.020 [†]	0.069 [‡]	97.8	88.7	91.7	98.6	94.2
	FCC (w/ predicted factors)	Greedy	0.638*	0.726	0.714*	0.930*	0.016	0.047	98.3	92.2	94.2	98.5	95.8
		Sampling	0.607	0.694	0.687	0.922	0.020 [‡]	0.070 [‡]	98.7	89.9	94.5	98.6	95.4
		GtS	0.640 ⁺	0.727 ⁺	0.716 ⁺	0.929 ⁺	0.020 [‡]	0.070 [‡]	99.0	93.2 ⁺	94.5 ⁺	99.2	96.5
	FCC (w/ golden factors)	Greedy	0.638*	0.727 [*]	0.716 [*]	0.931 [*]	0.016	0.048	98.3	92.3	94.2	98.5	95.8
		Sampling	0.607	0.693	0.687	0.922	0.020 [‡]	0.070 [‡]	98.7	89.9	94.5	98.6	95.4
		GtS	0.640 ⁺	0.727 ⁺	0.716 ⁺	0.929 ⁺	0.020 [‡]	0.067 [‡]	99.0	93.2 ⁺	94.5 ⁺	99.2	96.5
Ground-truth captions			1.000	1.000	1.000	1.000	0.020	0.071	99.1	99.1	99.6	99.4	99.3

^{*/+} Statistically significant improvement from the model trained using the original captions with greedy search (*) / sampling (+).

^{†/‡} Statistically significant improvement from greedy search with the model trained using the original captions (†) / FCC (‡).

Table 2: Comparison of the speech encoders. The component used in StyleCap [11] is underlined.

	Style captioning		Style-factor classification			
	B@4↑	ROU↑	Gender	Pitch	Speed	Volume
WavLM-B-plus	0.419	0.513	97.8	<u>77.6</u>	<u>92.4</u>	<u>76.2</u>
HuBERT-L	0.517	0.626	98.2	80.1	87.5	99.0
WavLM-L	0.597	0.693	99.0	88.4	92.5	97.8
Whisper-L	0.613	0.710	97.9	91.4	93.3	97.9

Table 3: Comparison of the bridge networks.

	Style captioning		Style-factor classification			
	B@4↑	ROU↑	Gender	Pitch	Speed	Volume
W.sum + AvP	0.505	0.615	98.1	85.7	81.4	95.5
W.sum + CNN	0.383	0.501	98.1	74.3	79.4	88.6
W.sum + QF	<u>0.561</u>	<u>0.668</u>	98.6	88.3	91.7	97.0
TLTR-seg	0.578	0.666	97.8	88.4	89.9	98.2
TLTR-utt	0.613	0.710	97.9	91.4	93.3	97.9

4.2.2. Ablation study: Encoder, Bridge network, and Decoder

To investigate the best components of the speaking-style captioning model, we compared several widely used modules in each component. The base structure uses Whisper-L as the speech encoder, TLTR-utt as the bridge network, and 7B-chat-LoRA as the text decoder. Each model was trained to output the original speaking-style captions and greedy search was used for decoding, which corresponds to the second row in Table 1. The training hyperparameters were the same as those in the main experiment. Due to space limitations, we report here only BLEU-4, ROUGE-L, and four classification accuracies.

Comparisons of the speech encoders with WavLM-base-plus [23] (WavLM-B-plus), WavLM-large (WavLM-L), and HuBERT-large [24] (HuBERT-L) are detailed in Table 2. WavLM-B-plus showed lower accuracy than the other encoders in pitch and volume prediction, same as the previous work [11]. With regard to the remaining three encoders, Whisper-L achieved the best performance in both captioning and factor classification tasks. This indicates that style captioning performance is highly dependent on the capability of the speech encoder, and that well trained large-scale models will

Table 4: Comparison of the text decoders.

Base	LoRA	Style captioning		Style-factor classification			
		B@4↑	ROU↑	Gender	Pitch	Speed	Volume
<u>7B</u>	✓	0.578	0.684	<u>97.7</u>	89.0	93.6	98.6
		0.602	0.695	98.1	91.5	92.8	94.6
7B-chat	✓	0.575	0.675	98.8	88.3	92.9	96.5
		0.613	0.710	97.9	91.4	93.3	97.9

provide better results.

Various bridge networks including the weighted sum of the intermediate layer outputs and average pooling (W.sum + AvP), CNN (W.sum + CNN), or BLSTM+Q-former (W.sum + QF), and TLTR with segment-level outputs (TLTR-seg, removing the temporal pooling as LTU-AS [12]) are also compared in Table 3. The results indicate that capturing temporal context in the bridge is effective (1, 2 vs. 3, 4, 5 rows), and that outputting a single-length embedding is more suitable than sequential embeddings in speaking style captioning and style factor classification (4 vs. 5 rows).

Comparisons of the text decoders, LLaMA-2 7B or 7B-chat with and without LoRA, are shown in Table 4. LoRA improved speaking-style captioning and pitch classification, but the performance difference between text decoders is smaller than that of the speech encoders and bridge networks. This is because the variation in style captions was limited in this dataset, and so did not require a strong text generation module.

5. Conclusion

This paper presented FCC, a novel speaking-style captioning method. FCC uses factor and caption terms that support the generation of captions conditioned by the predicted speech factors. GtS, greedy-then-sampling, decoding initially evaluates the speech factors by maximum likelihood criteria and then generates captions, which yields diverse outputs while preventing factor prediction errors. We also investigated the optimal Speech-Text model components. Experiments showed that our method significantly outperformed the conventional style captioning model, and that FCC with GtS can generate accurate and diverse style captions while identifying style factors.

6. References

- [1] Z. Guo, Y. Leng, Y. Wu, S. Zhao, and X. Tan, "PromptTTS: Controllable text-to-speech with text descriptions," in *Proc. ICASSP*, 2023.
- [2] R. Shimizu, R. Yamamoto, M. Kawamura, Y. Shirahata, H. Doi, T. Komatsu, and K. Tachibana, "PromptTTS++: Controlling speaker identity in prompt-based text-to-speech using natural language descriptions," in *Proc. ICASSP (to appear)*, 2024.
- [3] S. Ji, J. Zuo, M. Fang, Z. Jiang, F. Chen, X. Duan, B. Huai, and Z. Zhao, "TextrolSpeech: A text style control speech corpus with codec language text-to-speech models," in *Proc. ICASSP (to appear)*, 2024.
- [4] J. Yao, Y. Yang, Y. Lei, Z. Ning, Y. Hu, Y. Pan, J. Yin, H. Zhou, H. Lu, and L. Xie, "PromptVC: Flexible stylistic voice conversion in latent space driven by natural language prompts," in *Proc. ICASSP (to appear)*, 2024.
- [5] J. Miehle, I. Feustel, J. Hornauer, W. Minker, and S. Ultes, "Estimating user communication styles for spoken dialogue systems," in *Proc. LREC*, 2020, pp. 540–548.
- [6] G.-T. Lin, C.-H. Chiang, and H. yi Lee, "Advancing large language models to capture varied speaking styles and respond properly in spoken conversations," 2024, arXiv:2402.12786.
- [7] H. Zhang, M. Mimura, T. Kawahara, and K. Ishizuka, "Selective multi-task learning for speech emotion recognition using corpora of different styles," in *Proc. ICASSP*, 2022, pp. 7707–7711.
- [8] M. Sharma, "Multi-lingual multi-task speech emotion recognition using wav2vec 2.0," in *Proc. ICASSP*, 2022, pp. 6907–6911.
- [9] T. Shinozaki, M. Ostendorf, and L. Atlas, "Characteristics of speaking style and implications for speech recognition," *Journal of Acoust. Soc. Am.*, vol. 126, no. 3, pp. 1500–1510, 2009.
- [10] A. Veiga, S. Candeias, D. Celorico, J. Proença, and F. Perdigão, "Towards automatic classification of speech styles," in *Proc. Computational Processing of the Portuguese Language*, 2012, pp. 421–426.
- [11] K. Yamauchi, Y. Ijima, and Y. Saito, "StyleCap: Automatic speaking-style captioning from speech based on speech and language self-supervised learning models," in *Proc. ICASSP (to appear)*, 2024.
- [12] Y. Gong, A. H. Liu, H. Luo, L. Karlinsky, and J. Glass, "Joint audio and speech understanding," in *Proc. ASRU*, 2023.
- [13] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, "Qwen-Audio: Advancing universal audio understanding via unified large-scale audio-language models," 2023, arXiv:2311.07919.
- [14] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, "The curious case of neural text degeneration," in *Proc. ICLR*, 2020.
- [15] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Advances in NeurIPS*, vol. 35, 2022, pp. 24 824–24 837.
- [16] OpenAI, "GPT-4 Technical Report," 2023, arXiv:2303.08774.
- [17] K.-W. Chang, W.-C. Tseng, S.-W. Li, and H. yi Lee, "An exploration of prompt tuning on generative spoken language model for speech processing tasks," in *Proc. Interspeech*, 2022, pp. 5005–5009.
- [18] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, "Pengi: An audio language model for audio tasks," in *Advances in NeurIPS*, vol. 36, 2023, pp. 18 090–18 108.
- [19] M. Wang, W. Han, I. Shafran, Z. Wu, C.-C. Chiu, Y. Cao, N. Chen, Y. Zhang, H. Soltau, P. K. Rubenstein, L. Zilka, D. Yu, G. Pundak, N. Siddhartha, J. Schalkwyk, and Y. Wu, "SLM: Bridge the thin gap between speech and text foundation models," in *Proc. ASRU*, 2023.
- [20] T. Wang, L. Zhou, Z. Zhang, Y. Wu, S. Liu, Y. Gaur, Z. Chen, J. Li, and F. Wei, "VioLA: Unified codec language models for speech recognition, synthesis, and translation," 2023, arXiv:2305.16107.
- [21] P. K. Rubenstein, C. Asawaroengchai, D. D. Nguyen, A. Bapna, Z. Borsos, F. de Chaumont Quitry, P. Chen, D. E. Badawy *et al.*, "AudioPaLM: A large language model that can speak and listen," 2023, arXiv:2306.12925.
- [22] Y. Xu, H. Chen, J. Yu, Q. Huang, Z. Wu, S. Zhang, G. Li, Y. Luo, and R. Gu, "SECap: Speech emotion captioning with large language model," in *Proc. AACL (to appear)*, 2024.
- [23] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [24] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [25] Y. Gong, S. Khurana, L. Karlinsky, and J. Glass, "Whisper-AT: Noise-robust automatic speech recognizers are also strong general audio event taggers," in *Proc. Interspeech*, 2023, pp. 2798–2802.
- [26] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esioiu *et al.*, "Llama 2: Open foundation and fine-tuned chat models," 2023, arXiv:2307.09288.
- [27] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. ICLR*, 2022.
- [28] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," in *Proc. Interspeech*, 2019, pp. 1526–1530.
- [29] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *Proc. ICML*, 2023, pp. 28 492–28 518.
- [30] J. Li, D. Li, S. Savarese, and S. Hoi, "BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models," in *Proc. ICML*, 2023.
- [31] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. ICLR*, 2017.
- [32] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," in *Proc. NeurIPS EMCC Workshop*, 2019.
- [33] L. Von Werra, L. Tunstall, A. Thakur, S. Luccioni, T. Thrush, A. Piktus, F. Marty, N. Rajani, V. Mustar, and H. Ngo, "Evaluate & evaluation on the hub: Better best practices for data and model measurements," in *Proc. EMNLP*, 2022, pp. 128–136.