

# Human-CLAP: Human-perception-based contrastive language–audio pretraining

Taisei Takano\*, Yuki Okamoto\*, Yusuke Kanamori\*, Yuki Saito\*, Ryotaro Nagase<sup>†</sup>, Hiroshi Saruwatari\*

\* The University of Tokyo, Japan

E-mail: y-okamoto@ieee.org

<sup>†</sup> Ritsumeikan University, Japan

**Abstract**—Contrastive language–audio pretraining (CLAP) is widely used for audio generation and recognition tasks. For example, CLAPScore, which utilizes the similarity of CLAP embeddings, has been a major metric for the evaluation of the relevance between audio and text in text-to-audio. However, the relationship between CLAPScore and human subjective evaluation scores is still unclarified. We show that CLAPScore has a low correlation with human subjective evaluation scores. Additionally, we propose a human-perception-based CLAP called Human-CLAP by training a contrastive language-audio model using the subjective evaluation score. In our experiments, the results indicate that our Human-CLAP improved the Spearman’s rank correlation coefficient (SRCC) between the CLAPScore and the subjective evaluation scores by more than 0.17 compared with the conventional CLAP.

## I. INTRODUCTION

Various tasks focusing on audio–text relationships, such as text-to-audio (TTA) [1]–[3], language-queried audio source separation (LASS) [4], [5], and automated audio captioning (AAC) [6]–[8], have been developed. Constructing text representations of audio is a common effective approach across such audio–text tasks, making audio more interpretable for humans [9]. Representing audio as text descriptions improves the searchability of audio data and allows more people, including those with hearing difficulties, to enjoy audio contents.

CLAP [10], [11] has been proposed as a foundation model for connecting audio and text by contrastive learning. By learning the correspondence of audio–text pairs, CLAP creates a shared embedding space between audio and text, enabling wide applications to audio recognition and generation tasks [12]. Moreover, CLAP is used not only for these tasks, but also for the evaluation of relevance between audio and text in TTA [13]–[17]. Concretely, the distance between the audio and text embedding vectors obtained by CLAP can be interpreted as a score that quantifies the degree of relevance between the audio and the text, which is called CLAPScore. Although CLAPScore is used for the objective evaluation in TTA tasks [13], [14], the relationship between CLAPScore and human subjective evaluation scores has not been investigated. Considering that TTA aims to generate sounds humans hear, it is necessary to investigate how well the CLAPScore aligns with human subjective evaluation scores. In addition, although CLAP is trained with the assumption that the audio–text pairs in the dataset match exactly, the audio–text pairs also include noisy data, such as text that does not contain all the content

in the audio. Therefore, there is a possibility that the distance between audio and text will be close in the embedding space of CLAP, even for inappropriate audio–text pairs. It is also possible to collect cleaner audio–text pairs for model training, but data collection is extremely costly.

In this paper, we first investigate the relationship between CLAPScore and subjective evaluation scores, focusing on the relevance between audio and text (REL). We also propose a human-perception-based CLAP called Human-CLAP<sup>1</sup> by fine-tuning a pretrained CLAP model using a small number of subjective scores, consisting of approximately 1/320th the training data of conventional CLAP, as an approximation of the ideal relevance of an audio-text pair. By using scores that evaluate the degree of relevance of text and audio as determined by humans, we aim to construct a CLAP model that matches human perception.

In summary, our contributions are as follows:

- We perform a subjective evaluation of the relevance between texts and audio samples. By comparing CLAPScore with collected subjective evaluation scores, we show that CLAPScore has a low correlation with subjective scores.
- We propose a human-perception-based CLAP called “Human-CLAP”, which is trained using subjective scores. We experimentally show that the CLAPScore calculated by Human-CLAP has a higher correlation with subjective scores than the conventional CLAP.

## II. RELATIONSHIP BETWEEN CLAPSCORE AND HUMAN PERCEPTION

Although the distance between CLAP embedding vectors is interpreted as the audio–text relevance, its alignment with human perception is still unclarified. We examine how well CLAP matches human perception by investigating the correlation between CLAPScore and subjective evaluation scores.

### A. CLAPScore

CLAPScore is an evaluation metric for the relevance of texts and audio samples in TTA. CLAPScore utilizes CLAP [11] embeddings and is defined as

$$\text{CLAPScore} = \max \left( \frac{\mathbf{e}^{\text{audio}} \cdot \mathbf{e}^{\text{text}}}{\|\mathbf{e}^{\text{audio}}\| \|\mathbf{e}^{\text{text}}\|}, 0 \right), \quad (1)$$

<sup>1</sup><https://github.com/sarulab-speech/Human-CLAP>

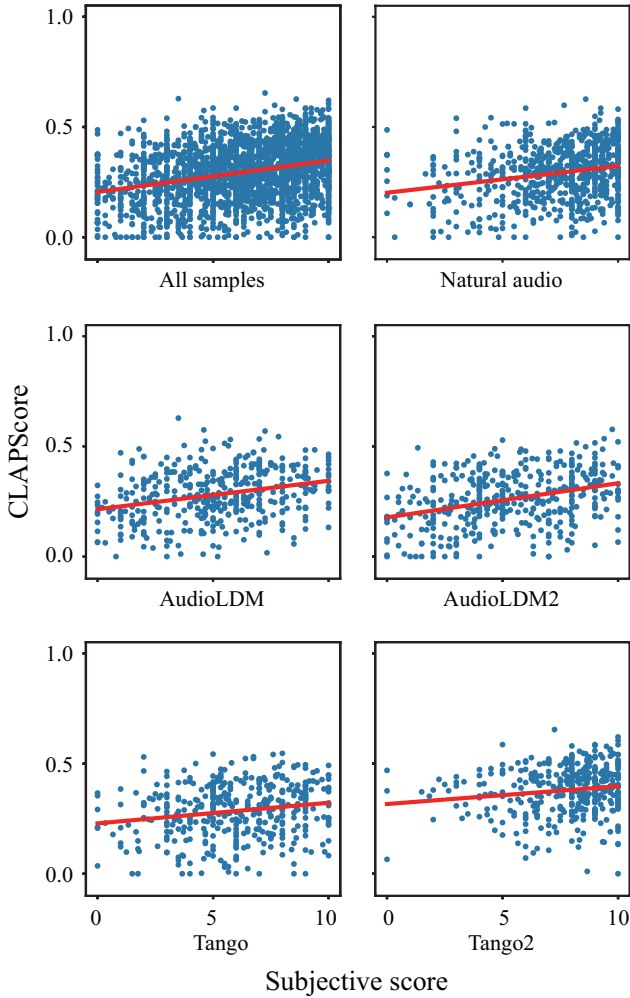


Fig. 1: Relationship between CLAPScore and subjective scores (Model: LAION CLAP)

where  $e^{\text{text}}$  and  $e^{\text{audio}}$  denote the CLAP embedding of text and audio, respectively [13]. Although CLAPScore is sometimes calculated to take values between  $-1$  and  $1$ , in this paper, we defined negative similarity scores as  $0$  to compare correlation with human evaluation scores. This definition follows the precedence set by the conventional work on CLIPScore [18].

### B. Dataset of subjective scores for audio–text pairs

We used the RELATE dataset [19], which contains subjective evaluation scores of the relevance between audio–text pairs. The RELATE dataset includes natural audio samples from AudioCaps [20] and synthesized audio samples generated using AudioLDM [1], AudioLDM2 [13], Tango [15], and Tango2 [14]. In the RELATE dataset, audio–text pairs were presented to listeners, who were then asked, “How well does the text description match the audio?” Listeners rated each audio–text pair on an 11-point scale, where  $0$  indicated no match and  $10$  indicated a perfect match. To suppress individual differences when rating, the corpus developer provided the following guidance when asking the question:

- 0 - Does not match at all




Fig. 2: Overview of Human-CLAP

- 2 - Has significant discrepancies
- 5 - Has several minor discrepancies
- 8 - Has a few minor discrepancies
- 10 - Matches exactly

To select high-quality listeners, we conducted a screening of listeners using anchor samples in the listening test. The RELATE dataset provides three anchor samples, which were mismatched audio–text pairs, in each listening test. We removed listeners who scored higher than  $2$  on average for anchor samples. This screening method is based on the subjective evaluation conducted by Choi et al. in DCASE 2023 Challenge Task 7 [21].

Each audio–text pair was evaluated by four listeners on average. After the screening,  $2,383$  audio–text pairs from the RELATE training set and  $2,405$  audio–text pairs from the RELATE test set were obtained.

### C. Correlation between CLAPScore and Subjective Evaluation Scores

We investigated the relationship between CLAPScore and subjective evaluation score using  $2,405$  audio–text pairs from test data collected in Section II-B. Two types of CLAP model were used, models by Wu et al. (LAION CLAP) [10] and Elizalde et al. (MS CLAP) [11]. To evaluate the correlation between CLAPScore and subjective evaluation scores, we used Spearman’s rank correlation coefficient (SRCC), linear correlation coefficient (LCC), and Kendall’s rank correlation coefficient (KTAU).

The CLAPScore calculated with LAION CLAP had an SRCC of  $0.280$  and an LCC of  $0.294$ . Similarly, those calculated with MS CLAP had an SRCC of  $0.278$  and an LCC of  $0.296$ . From these results, CLAPScore has a low correlation with subjective evaluation scores. Such a low correlation is also visualized in the scatter plot of the relationship between CLAPScore and subjective score shown in Fig. 1. These results indicate that CLAPScore is not necessarily suitable for an objective evaluation metric in TTA.

## III. PROPOSED METHOD

### A. Framework and Motivation

We propose Human-CLAP, which utilizes human subjective scores for training the CLAP model. Figure 2 shows the overview of our proposed method. The conventional CLAP was trained to maximize the cosine similarity of embedding vectors for paired text and audio, while minimizing it for

unpaired ones. However, this approach heavily relies on the quality of the text labels in the dataset. Even though some low-quality text labels might be included, the conventional method equally maximizes the similarity of each audio–text pair. On the other hand, we expect the CLAP model to train the relationship between audio and text on the basis of human perception using subjective evaluation scores as target scores. When using it as a target score, we rescaled the subjective score to the range of 0 to 1.

### B. Loss Function

Our Human-CLAP model is trained to minimize a loss function that combines contrastive learning and regression approaches. Regression losses such as the mean square error (MSE) and mean absolute error (MAE) are often used to train models that predict subjective evaluation scores for speech–text pairs or audio–text pairs [22]–[24]. In this study, we used the regression loss between the target score  $a_i \in [0, 1]$  and the predicted cosine similarity  $y_i$ . We utilize MSE and MAE to calculate such regression losses. MSE is calculated as

$$L_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (a_i - y_i)^2, \quad (2)$$

$$y_i = \text{ReLU} \left( \frac{\mathbf{e}_i^{\text{text}} \cdot \mathbf{e}_i^{\text{audio}}}{\|\mathbf{e}_i^{\text{text}}\| \|\mathbf{e}_i^{\text{audio}}\|} \right), \quad (3)$$

where  $a_i \in [0, 1]$  is the rescaled subjective score of the  $i$ th audio–text pair in a batch,  $\mathbf{e}_i^{\text{text}}$  and  $\mathbf{e}_i^{\text{audio}}$  denote the embeddings of  $i$ th data in batch,  $\text{ReLU}(\cdot)$  denotes the rectified linear unit, and  $N$  denotes the batch size. Similarly to MSE, we calculated MAE as

$$L_{\text{MAE}} = \frac{1}{N} \sum_{i=1}^N |a_i - y_i|. \quad (4)$$

We also trained the model with the contrastive learning approach performed in the conventional CLAP model [10], [11]. The conventional CLAP is trained by using symmetric cross entropy (SCE) loss [10], [25] to construct its audio–text embedding space. We propose the weight SCE (wSCE) loss by adding weights based on the subjective score to the SCE loss:

$$L_{\text{wSCE}} = -\frac{1}{2N} \sum_{i=1}^N a_i \left( \log \frac{\exp(\mathbf{e}_i^{\text{text}} \cdot \mathbf{e}_i^{\text{audio}} / \tau)}{\sum_{j=1}^N \exp(\mathbf{e}_i^{\text{text}} \cdot \mathbf{e}_j^{\text{audio}} / \tau)} \right) + \log \frac{\exp(\mathbf{e}_i^{\text{audio}} \cdot \mathbf{e}_i^{\text{text}} / \tau)}{\sum_{j=1}^N \exp(\mathbf{e}_i^{\text{audio}} \cdot \mathbf{e}_j^{\text{text}} / \tau)}, \quad (5)$$

where  $a_i$  denotes the rescaled subjective score of  $i$ th audio–text pair and  $\tau$  is a learnable temperature parameter for scaling the loss. The SCE loss minimizes the distance between relevant audio–text embeddings and maximizes the distance between unpaired embeddings. By weighting the SCE loss with subjective scores, we believe that the wSCE loss function will construct a better embedding space, reflecting human perception.

TABLE I: Evaluation results of CLAPScore and subjective evaluation scores

Model	SRCC $\uparrow$	LCC $\uparrow$	KTAU $\uparrow$	MSE $\downarrow$
<b>Human-CLAP (ours)</b>				
wSCE + MSE	0.447	0.471	0.313	0.059
wSCE + MAE	<b>0.457</b>	<b>0.481</b>	<b>0.320</b>	0.057
wSCE	0.383	0.410	0.265	0.063
MSE	0.439	0.462	0.307	0.056
MAE	0.453	0.472	0.317	<b>0.051</b>
<b>Baseline</b>				
LAION CLAP	0.280	0.294	0.192	0.068
MS CLAP	0.278	0.296	0.192	0.078

The overall loss function  $L$  is formulated as a linear combination of the wSCE and MSE or MAE losses:

$$L = \lambda_1 L_{\text{wSCE}} + \lambda_2 L_{\text{reg}}, \quad (6)$$

where  $L_{\text{reg}}$  is the regression loss, which is  $L_{\text{MSE}}$  or  $L_{\text{MAE}}$ ,  $\lambda_1$ , and  $\lambda_2$  are hyperparameters for controlling the effects of the loss functions.

## IV. EXPERIMENTS

We compared the relationship between CLAPScores and subjective evaluation scores under different conditions to confirm whether our proposed method is effective in improving the correlation relationship.

### A. Experimental Setup

We fine-tuned LAION CLAP [10] using the subjective evaluation score collected in Section II. The learning rate was set to  $10^{-5}$  and AdamW [26] was used for optimization, and the batch size was 8. We trained the model for 50 epochs and selected the best model that had the lowest validation loss. The hyperparameters of  $L$  (Eq. (6)) were empirically determined to be  $[\lambda_1, \lambda_2] = [0.1, 1]$  after multiple rounds of hyperparameter tuning. We used RoBERTa [27] and hierarchical token-semantic audio Transformer (HTS-AT) [28] for the text and audio encoders, respectively. We split the train data collected in Section II into 1,925 and 458 audio–text pairs for model training and validation.

We evaluated the correlation between CLAPScore and subjective evaluation scores using 2,405 audio–text pairs collected from the AudioCaps [20] test set in Section II. We used SRCC and LCC to calculate the correlation between CLAPScore and subjective score, and used MSE to evaluate the error between CLAPScore and subjective score. LAION CLAP [10] and MS CLAP [11] are used as the comparison methods.

### B. Results

**Overall results.** Table I shows the results. The proposed method improved the SRCC by 0.15 or more compared with baseline models. In particular, the model fine-tuned with wSCE + MAE achieved the best scores across almost all metrics. From these results, we confirmed that our proposed method is effective in improving the correlation between CLAPScore and subjective evaluation scores.

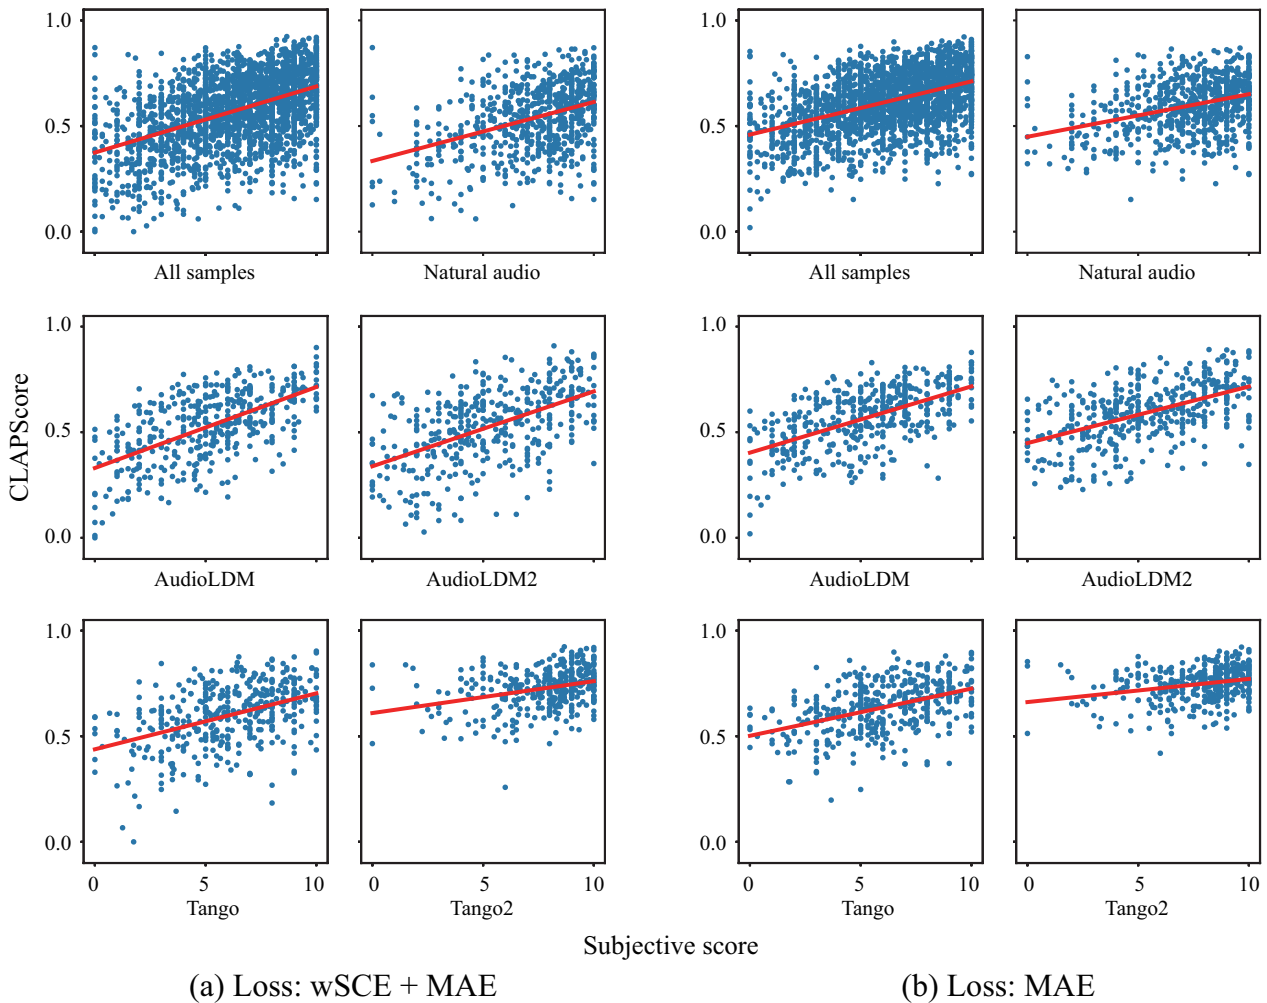


Fig. 3: Relationship between CLAPScore and subjective evaluation score of Human-CLAP

Figure 3(a) shows the relationship between CLAPScore and subjective evaluation scores of the best-performing model trained with wSCE + MAE. Compared with the CLAPScore before fine-tuning in Fig. 1, it can be confirmed that the scores are predicted consistently regardless of whether the target subjective evaluation scores are high or low. These results demonstrate that fine-tuning the CLAP model with the proposed loss function can predict scores that are closely aligned with human subjective evaluations.

Figure 3(b) shows the relationship between CLAPScore fine-tuned with only MAE and human subjective score. As shown in the figure, it can be confirmed that the model trained using only MAE assigns a relatively higher CLAPScore than the model fine-tuned with wSCE + MAE. This indicates that the model fine-tuned with only MAE struggles to assign low CLAPScores to low subjective scores, particularly those close to 0. These results suggest that the Human-CLAP fine-tuned with wSCE has more potential to score low-relevant audio-text pairs and to assign low scores to low subjective scores.

**Results on natural and synthesized audio.** We divided the test data into natural and synthesized audio and calculated

TABLE II: SRCC of CLAPScore and subjective evaluation scores on natural and synthesized audio

Model	Natural	Synthesized				
		AudioLDM	AudioLDM 2	Tango	Tango 2	All
<b>Human-CLAP (ours)</b>						
wSCE + MSE	0.331	0.549	0.487	0.385	0.275	0.548
wSCE + MAE	<b>0.358</b>	<b>0.589</b>	<b>0.538</b>	0.425	<b>0.310</b>	<b>0.583</b>
wSCE	0.303	0.485	0.465	0.326	0.230	0.488
MSE	0.338	0.513	0.485	0.354	0.278	0.527
MAE	0.333	0.567	0.513	<b>0.448</b>	0.282	0.574
<b>Baseline</b>						
LAION CLAP	0.201	0.288	0.332	0.178	0.144	0.338

the SRCC between CLAPScore and subjective score. For the synthesized audio samples, we further calculated the SRCC for each TTA model. Table II shows the results. We confirmed that our model has successfully improved the SRCC for all synthesized audio samples. This result suggests that the CLAPScore obtained using Human-CLAP evaluates synthesized audio samples in TTA tasks more closely to human evaluation, showing its effectiveness for the evaluation in TTA tasks.

TABLE III: SRCC depending on the subjective evaluation score. Nat. and Syn. denote natural and synthesized audio, respectively.

Model	Sub. score $\leq 5$			Sub. score $> 5$		
	All $\uparrow$	Nat. $\uparrow$	Syn. $\uparrow$	All $\uparrow$	Nat. $\uparrow$	Syn. $\uparrow$
<b>Human-CLAP (ours)</b>						
wSCE + MSE	0.327	0.246	0.356	0.275	0.205	0.382
wSCE + MAE	<b>0.357</b>	<b>0.287</b>	<b>0.392</b>	<b>0.280</b>	<b>0.229</b>	<b>0.408</b>
wSCE	0.281	0.156	0.322	0.213	0.184	0.306
MSE	0.327	0.269	0.348	0.272	0.207	0.368
MAE	0.336	0.262	0.367	0.277	0.208	0.406
<b>Baseline</b>						
LAION CLAP	0.205	0.072	0.244	0.162	0.120	0.213

**Results depending on the values of subjective scores.** We also divided the test data, depending on the values of subjective evaluation scores, to assess the effectiveness of Human-CLAP in evaluating highly related or less related audio-text pairs. Table III shows the results. An increase in SRCC was observed on both high-score and low-score data. It is noteworthy that there was a significant improvement in SRCC for natural audio, which was notably low in conventional CLAP. This result indicates that the Human-CLAP can be more useful for the evaluation of natural audio compared with the conventional CLAP.

## V. CONCLUSION

In this paper, we investigated the correlation between subjective evaluation scores and CLAPScore, which uses the embedding space of CLAP. Results showed that the CLAPScore had a low correlation with subjective evaluation scores. We also proposed Human-CLAP, a human-perception-based CLAP, fine-tuned by using a small amount of subjective evaluation scores. CLAPScore obtained using Human-CLAP indicated a higher SRCC than the conventional CLAP models by 0.17 or higher. In particular, the model fine-tuned with wSCE + MAE obtained the best scores across almost all metrics. In addition, for only synthesized audio, an improvement of approximately 0.25 in SRCC was confirmed compared with the conventional CLAP model. We confirmed that Human-CLAP can achieve higher SRCC on both natural and synthesized audio samples than conventional CLAP. In the future, we will collect subjective evaluation scores for other audio datasets and investigate the performance of Human-CLAP for out-of-domain data.

## ACKNOWLEDGMENT

The work was supported by JSPS KAKENHI Grant Number 24K23880, 25K21221, ROIS NII Open Collaborative Research 2025-(251S4-22735), JST Moonshot Grant Number JPMJMS2237.

## REFERENCES

- [1] H. Liu, Z. Chen, Y. Yuan, *et al.*, “AudioLDM: Text-to-audio generation with latent diffusion models,” *Proceedings of the International Conference on Machine Learning*, pp. 21 450–21 474, 2023.
- [2] D. Yang, J. Yu, H. Wang, *et al.*, “Diffsound: Discrete diffusion model for text-to-sound generation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 1720–1733, 2023.
- [3] F. Kreuk, G. Synnaeve, A. Polyak, *et al.*, “AudioGen: Textually guided audio generation,” in *Proc. the Eleventh International Conference on Learning Representations*, 2023.
- [4] X. Liu, H. Liu, Q. Kong, *et al.*, “Separate what you describe: Language-queried audio source separation,” in *Proc. INTERSPEECH 2022*, 2022.
- [5] H. Ma, Z. Peng, X. Li, M. Shao, X. Wu, and J. Liu, “Clapsep: Leveraging contrastive pre-trained model for multi-modal query-conditioned target sound extraction,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 4945–4960, 2024.
- [6] K. Drossos, S. Adavanne, and T. Virtanen, “Automated audio captioning with recurrent neural networks,” in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017.
- [7] G. Karanth, N. Rao, S. Subramanian, and A. Shah, “DCASE 2023 task 6 automated audio captioning and language-based retrieval,” DCASE2023 Challenge, Tech. Rep., 2023.
- [8] K. Greeshma, N. Rao, S. Subramanian, and A. Shah, “DCASE 2023 task 6 automated audio captioning and language-based retrieval,” DCASE2023 Challenge, Tech. Rep., 2023.
- [9] C.-C. Wang, J. Du, and J.-S. R. Jang, “DCASE 2023 task 6b: Text-to-audio retrieval using pretrained models,” DCASE2023 Challenge, Tech. Rep., 2023.
- [10] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, “Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [11] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, “Clap learning audio concepts from natural language supervision,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [12] R. Huang, J. Huang, D. Yang, *et al.*, “Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models,” in *Proc. the 40th International Conference on Machine Learning*, vol. 202, 2023, pp. 13 916–13 932.
- [13] H. Liu, Y. Yuan, X. Liu, *et al.*, “Audioldm 2: Learning holistic audio generation with self-supervised pretraining,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 2871–2883, 2024.
- [14] N. Majumder, C.-Y. Hung, D. Ghosal, W.-N. Hsu, R. Mihalcea, and S. Poria, “Tango 2: Aligning diffusion-based text-to-audio generations through direct prefer-

- ence optimization,” in *Proc. the 32nd ACM International Conference on Multimedia*, 2024, pp. 564–572.
- [15] D. Ghosal, N. Majumder, A. Mehrish, and S. Poria, “Text-to-audio generation using instruction guided latent diffusion model,” in *Proc. the 31st ACM International Conference on Multimedia*, 2023, pp. 3590–3598.
- [16] J. Xue, Y. Deng, Y. Gao, and Y. Li, “Auffusion: Leveraging the power of diffusion and large language models for text-to-audio generation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 4700–4712, 2024.
- [17] H. Liao, H. Han, K. Yang, *et al.*, “Baton: Aligning text-to-audio model using human preference feedback,” in *Proc. the Thirty-Third International Joint Conference on Artificial Intelligence*, 2024, pp. 4542–4550.
- [18] J. Hessel, A. Holtzman, M. Forbes, R. Le Bras, and Y. Choi, “CLIPScore: A reference-free evaluation metric for image captioning,” in *Proc. 2021 Conference on Empirical Methods in Natural Language Processing*, M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih, Eds., 2021, pp. 7514–7528.
- [19] Y. Kanamori, Y. Okamoto, T. Takano, S. Takamichi, Y. Saito, and H. Saruwatari, “Relate: Subjective evaluation dataset for automatic evaluation of relevance between text and audio,” in *Proc. INTERSPEECH*, 2025.
- [20] C. D. Kim, B. Kim, H. Lee, and G. Kim, “Audiocaps: Generating captions for audios in the wild,” in *Proc. the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 119–132.
- [21] K. Choi, J. Im, L. Heller, *et al.*, “Foley sound synthesis at the DCASE 2023 challenge,” in *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2023.
- [22] A. Tjandra, Y.-C. Wu, B. Guo, *et al.*, “Meta audiobox aesthetics: Unified automatic quality assessment for speech, music, and sound,” *arXiv preprint arXiv:2502.05139*, 2025.
- [23] K. Baba, W. Nakata, Y. Saito, and H. Saruwatari, “The t05 system for the VoiceMOS Challenge 2024: Transfer learning from deep image classifier to naturalness MOS prediction of high-quality synthetic speech,” in *Proc. IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 818–824.
- [24] C.-C. Lo, S.-W. Fu, W.-C. Huang, *et al.*, “Mosnet: Deep learning-based objective assessment for voice conversion,” in *Proc. INTERSPEECH*, 2019, pp. 1541–1545.
- [25] A. Radford, J. W. Kim, C. Hallacy, *et al.*, “Learning transferable visual models from natural language supervision,” in *Proc. the 38th International Conference on Machine Learning (ICML)*, vol. 139, 2021, pp. 8748–8763.
- [26] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *Proc. International Conference on Learning Representations*, 2019.
- [27] Y. Liu, M. Ott, N. Goyal, *et al.*, “Roberta: A robustly optimized BERT pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [28] K. Chen, X. Du, B. Zhu, Z. Ma, T. Berg-Kirkpatrick, and S. Dubnov, “Hts-at: A hierarchical token-semantic audio transformer for sound classification and detection,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2022, pp. 646–650.