

LAMB: LLM-BASED AUDIO CAPTIONING WITH MODALITY GAP BRIDGING VIA CAUCHY-SCHWARZ DIVERGENCE

Hyeongkeun Lee*, Jongmin Choi*, KiHyun Nam, Joon Son Chung

Korea Advanced Institute of Science and Technology, South Korea

ABSTRACT

Automated Audio Captioning (AAC) aims to describe the semantic content of input audio. Recent works have employed large language models (LLMs) as a text decoder to leverage their reasoning capabilities. However, prior approaches that project audio features into the LLM embedding space without considering cross-modal alignment fail to fully utilize these capabilities. To address this, we propose *LAMB*, an LLM-based audio captioning framework that bridges the modality gap between audio embeddings and the LLM text embedding space. LAMB incorporates a Cross-Modal Aligner that minimizes Cauchy–Schwarz (CS) divergence while maximizing mutual information, yielding tighter alignment between audio and text at both global and token levels. We further design a Two-Stream Adapter that extracts semantically enriched audio embeddings, thereby delivering richer information to the Cross-Modal Aligner. Finally, leveraging the aligned audio embeddings, a proposed Token Guide directly computes scores within the LLM text embedding space to steer the output logits of generated captions. Experimental results confirm that our framework strengthens the reasoning capabilities of the LLM decoder, achieving state-of-the-art performance on AudioCaps. Code is available at <https://github.com/Hyeongkeun/LAMB>.

Index Terms— automated audio captioning, large language model, modality gap, Cauchy-Schwarz divergence

1. INTRODUCTION

Automated Audio Captioning (AAC) is a multimodal task that aims to generate corresponding descriptions for given audio content [1]. Capturing acoustic semantic information and aligning it with textual features are key challenges in AAC, as these processes enable the model to generate semantically consistent and contextually appropriate descriptions.

Prior works [2–11] typically follow a two-stage framework, using an audio encoder [12–14] and a text decoder [15–17]. More recently, with the emergence of large language models (LLMs), studies [3–7, 18] have employed LLMs [19–21] as a text decoder to leverage their contextual reasoning.

However, because LLMs are trained on large-scale text corpora, their reasoning is most effectively exploited when the conditioning inputs (e.g. audio modality) are represented in the LLM text embedding space [22]. Existing methods [3, 6, 7] map audio features into this space, typically via linear projection or Q-Former [23], without an explicit objective that bridges the audio–text modality gap. As a result, a modality misalignment can remain between audio representations and the LLM embedding space, limiting the ability of the decoder that interprets audio semantics.

To overcome this limitation, we introduce *LAMB*, an LLM-based Audio Captioning framework for Modality Bridging. Our approach lies in a *Cross-Modal Aligner* that leverages Cauchy-Schwarz (CS) divergence [24]. This symmetric metric, combined with mutual information, provides a robust distance estimation that effectively captures both global structures and pairwise semantic correspondences between different modalities [25]. To the best of our knowledge, this is the first use of CS divergence to mitigate the modality gap for LLM-driven audio captioning. Ensuring effective cross-modal alignment, the audio features are processed with *Two-Stream Adapter*, which extracts both semantically and temporally rich context. With these high-quality audio features, the *Cross-Modal Aligner* can more accurately align them to the text embedding space. Finally, our novel *Token Guide* directly steers the text generation process. It handles scores between the aligned audio and token embeddings using only the LLM’s own token dictionary, avoiding reliance on auxiliary modules that constrained previous works [7].

Experiments validate the effectiveness of each component of the proposed method. Furthermore, our framework achieves superior performance over state-of-the-art methods on AudioCaps, and shows comparable results with some improvements on Clotho.

2. METHOD

Fig. 1 shows the overall architecture of LAMB, and the subsequent sections provide detailed explanations of each component and the training objectives.

2.1. Two-Stream Adapter

The input audio x_a is first encoded by the audio encoder E_a into embeddings $h_a = E_a(x_a) \in \mathbb{R}^{N_a \times D}$, where N_a is the number of audio tokens and D the embedding dimension.

*These authors contributed equally to this work.

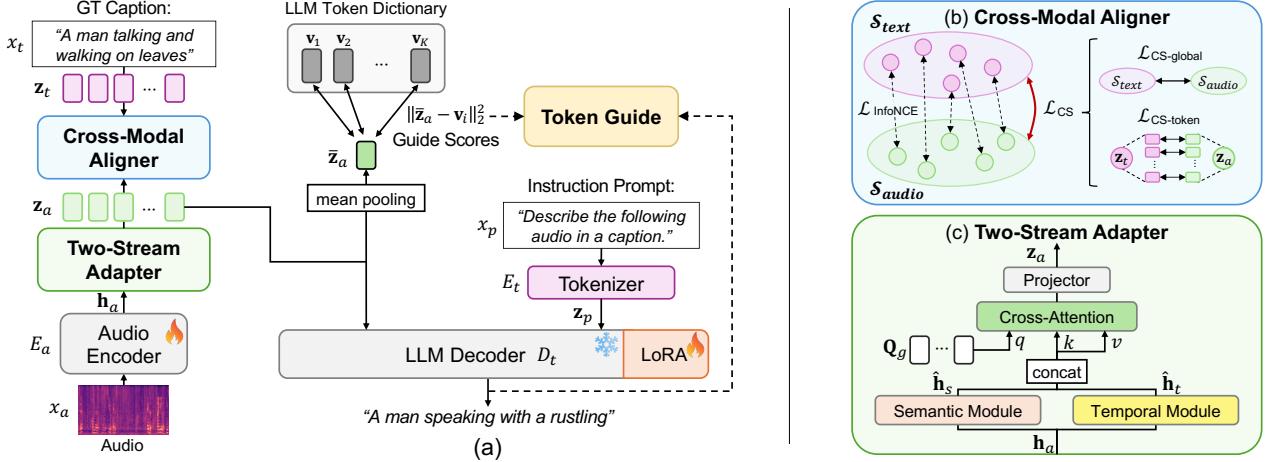


Fig. 1. (a) Overview of the proposed LAMB framework with (b) Cross-Modal Aligner and (c) Two-Stream Adapter.

We then introduce the *Two-Stream Adapter* to capture semantic and temporal context of \mathbf{h}_a through two parallel modules.

Semantic Module. Semantic features are extracted by N_s learnable queries $\mathbf{Q}_s \in \mathbb{R}^{N_s \times D}$ attending to \mathbf{h}_a through multi-head attention (MHA), followed by residual connection and LayerNorm (LN):

$$\text{MHA}(\mathbf{Q}_s, \mathbf{h}_a) = \sigma \left(\frac{\mathbf{Q}_s W_Q \cdot (\mathbf{h}_a W_K)^\top}{\sqrt{d_H}} \right) \mathbf{h}_a W_V, \quad (1)$$

$$\hat{\mathbf{h}}_s = \text{LN}(\text{MHA}(\mathbf{Q}_s, \mathbf{h}_a) + \mathbf{Q}_s), \quad (2)$$

where σ is softmax, $W_Q, W_K, W_V \in \mathbb{R}^{D \times H d_H}$ are projection matrices, H the number of heads, and d_H the head size.

Temporal Module. Temporal features are extracted using a multi-scale 1D convolutional front-end, a two-layer bidirectional GRU [26], and MHA with residual connection and LayerNorm. The convolutional front-end captures local patterns from \mathbf{h}_a , which are projected to dimension D with LayerNorm and GELU activation. The GRU encodes contextual dependencies, and N_t learnable queries $\mathbf{Q}_t \in \mathbb{R}^{N_t \times D}$ attend to its outputs to form temporal representations:

$$\hat{\mathbf{h}}_t = \text{LN}(\text{MHA}(\mathbf{Q}_t, \text{GRU}(\text{Conv1D}(\mathbf{h}_a))) + \mathbf{Q}_t), \quad (3)$$

where $\text{Conv1D}(\cdot)$ and $\text{GRU}(\cdot)$ denote the convolutional front-end and GRU encoder.

Fusion. The outputs of the two parallel modules $\hat{\mathbf{h}}_s$ and $\hat{\mathbf{h}}_t$ are concatenated along the token axis, and then processed with N_g global queries \mathbf{Q}_g via a cross-attention layer with residual and LayerNorm. The resulting representations are projected into the D_{llm} -dimensional LLM text embedding space via $W_p \in \mathbb{R}^{D \times D_{llm}}$, yielding $\mathbf{z}_a \in \mathbb{R}^{N_g \times D_{llm}}$:

$$\mathbf{z}_a = \left(\text{LN}(\text{MHA}(\mathbf{Q}_g, [\hat{\mathbf{h}}_s; \hat{\mathbf{h}}_t]) + \mathbf{Q}_g) \right) W_p. \quad (4)$$

2.2. Cross-Modal Aligner

A *Cross-Modal Aligner* employs Cauchy–Schwarz (CS) divergence with mutual information to align audio embeddings \mathbf{z}_a with caption embeddings \mathbf{z}_t in a batch B , where \mathbf{z}_t is the encoding of caption x_t with the LLM tokenizer E_t .

The representative vectors $\bar{\mathbf{z}}_a$ and $\bar{\mathbf{z}}_t$ are computed via mean pooling to quantify the modality gap between audio and the LLM text distributions. The CS divergence, D_{CS} , provides a symmetric and robust measure of the distance between the distributions of audio p_a and text p_t as follows:

$$D_{CS}(p_a; p_t) = -\log \frac{\int p_a(\omega) p_t(\omega) d\omega}{\sqrt{\int p_a^2(\omega) d\omega \int p_t^2(\omega) d\omega}}, \quad (5)$$

with $0 \leq D_{CS} < \infty$ and $D_{CS} = 0$ if and only if $p_a = p_t$.

Global-level CS Divergence. Since Eq. 5 is intractable, we approximate D_{CS} via non-parametric KDE[28] using i.i.d. samples $\{\bar{\mathbf{z}}_a^{(i)}\}_{i=1}^B \sim p_a(\bar{\mathbf{z}}_a)$ and $\{\bar{\mathbf{z}}_t^{(i)}\}_{i=1}^B \sim p_t(\bar{\mathbf{z}}_t)$, where each batch B spans the entire dataset, to measure the *global-level* distributional distance, as follows:

$$\hat{D}_{CS}(p_a(\bar{\mathbf{z}}_a); p_t(\bar{\mathbf{z}}_t)) = -\log \frac{\sum_{i,j=1}^B \kappa_{at}^{ij}}{\sqrt{\left(\sum_{i,j=1}^B \kappa_{aa}^{ij} \right) \left(\sum_{i,j=1}^B \kappa_{tt}^{ij} \right)}}, \quad (6)$$

where $\kappa_{mn}^{ij} = \kappa(\bar{\mathbf{z}}_m^{(i)}, \bar{\mathbf{z}}_n^{(j)})$, κ is Gaussian kernel $\kappa_\sigma(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|_2^2/2\sigma^2)$ with kernel width σ . The global distribution distance is reduced by the loss $\mathcal{L}_{CS-global} = \hat{D}_{CS}$.

Token-level CS Divergence. To achieve fine-grained alignment between audio and caption embeddings, we present a token-level loss based on token-wise CS divergence. For each pair of token-level embeddings $\{\mathbf{z}_a^{(i)}, \mathbf{z}_t^{(i)}\}_{i=1}^B$, the CS divergence is estimated as $\hat{D}_{CS}(p_a(\mathbf{z}_a^{(i)}); p_t(\mathbf{z}_t^{(i)}))$ and the token-level loss is defined as the average of per-sample divergences:

$$\mathcal{L}_{CS-token} = \frac{1}{B} \sum_{i=1}^B \hat{D}_{CS}(p_a(\mathbf{z}_a^{(i)}); p_t(\mathbf{z}_t^{(i)})). \quad (7)$$

Lastly, we adopt the InfoNCE loss [29] to maximize lower bounds of mutual information between modalities, defined as:

$$\mathcal{L}_{\text{InfoNCE}} = \frac{1}{2} (\mathcal{L}_{a \rightarrow t} + \mathcal{L}_{t \rightarrow a}), \quad (8)$$

$$\mathcal{L}_{m \rightarrow n} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\text{sim}(\mathbf{z}_m^i, \mathbf{z}_n^i)/\tau)}{\sum_{j=1}^B \exp(\text{sim}(\mathbf{z}_m^i, \mathbf{z}_n^j)/\tau)}. \quad (9)$$

Here, τ is the temperature and $\text{sim}(\cdot, \cdot)$ the cosine similarity.

Table 1. AAC results on AudioCaps and Clotho. Bold numbers denote the best scores and underlined numbers the second best. Gray-shaded rows indicate results with additional pretraining datasets. *Reported from previous work [7]. Pre-training datasets include AudioCaps (AC), Clotho (CL), augmented Clotho (CL_P [7]), WavCaps (WC), MACS (MA), and AutoACD (AA). Evaluation metrics are METEOR (MT), CIDEr (CD), SPICE (SC), SPIDEr (SD), SPIDEr-FL (SF), and FENSE (FS).

Model	PT Dataset	AudioCaps						Clotho					
		MT	CD	SC	SD	SF	FS	MT	CD	SC	SD	SF	FS
HTSAT-BART [27]	AC+CL+WC	25.0	78.7	18.2	48.5	48.3*	64.2*	18.5	48.8	13.3	31.0	29.6*	50.1*
EnCLAP-large [8]	AC+CL	25.5	80.3	18.8	49.5	49.9*	65.5*	18.6	46.4	13.3	29.9	28.9*	50.7*
AutoCap [9]	AC+CL+WC	25.6	80.4	19.0	49.7	-	-	-	-	-	-	-	-
LOAE [6]	AC+CL+WC	26.7	81.6	19.3	50.5	50.4	66.4	19.7	51.3	14.7	33.0	33.0	53.8
CLAP-ART [10]	- / AC	25.6	80.7	18.8	49.8	-	65.5	18.7	47.5	13.3	30.4	-	51.1
MQ-Cap [11]	AC+CL+WC+AA	-	84.5	19.4	51.9	-	-	-	49.6	14.3	31.9	-	-
SLAM-AAC [7]	AC+CL _P +WC+MA	26.8	84.1	19.4	51.8	51.5	66.8	19.7	51.5	14.8	33.2	33.0	54.0
LAMB (Ours)	AC+CL+WC	27.1	91.1	19.7	55.4	55.3	67.7	19.8	52.3	14.7	33.4	33.1	53.4

Table 2. Ablation of alignment methods on AudioCaps.

Align Method	$L_2 \downarrow$	Cost \uparrow	MT \uparrow	CD \uparrow	SC \uparrow	SD \uparrow
Q-former	111.8	0.04	24.4	66.9	18.1	42.5
Linear Layer	62.6	-.01	25.3	77.2	18.1	47.7
CMA (ours)	10.9	0.58	27.1	91.1	19.7	55.4

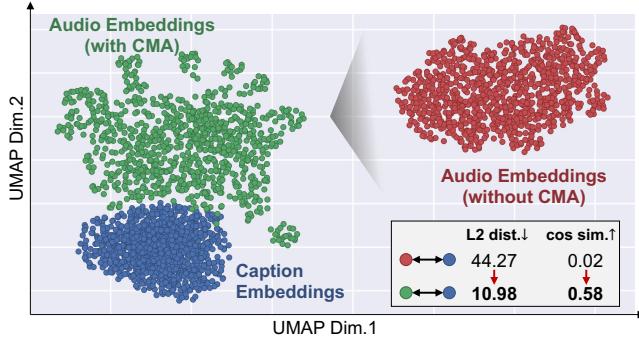


Fig. 2. UMAP Visualization of audio–text embedding alignment with/without Cross-Modal Aligner (CMA).

The overall cross-modal alignment loss is as follows:

$$\mathcal{L}_{\text{cma}} = \alpha_1 \mathcal{L}_{\text{CS-global}} + \alpha_2 \mathcal{L}_{\text{CS-token}} + \alpha_3 \mathcal{L}_{\text{InfoNCE}}. \quad (10)$$

2.3. Text Decoding with Token Guide

We propose *Token Guide* to steer the output logits using solely the LLM token dictionary without auxiliary module or external corpora. At each decoding step l , the logits \mathbf{z}_l of LLM decoder D_t are obtained from the concatenation of the audio embedding \mathbf{z}_a and the instruction prompt embedding $\mathbf{z}_p = E_t(x_p)$, where E_t is the LLM tokenizer:

$$\mathbf{z}_l = D_t([\mathbf{z}_a; E_t(x_p)], M_{\text{att}}), \quad (11)$$

where M_{att} is the attention mask enforcing the autoregressive property of the LLM. At each step l , the guided logit $\hat{z}_{l,i}$ is obtained by correcting the original logit $z_{l,i}$ with *Guide Scores*, squared L_2 distance between global mean \bar{z}_a and token embeddings $\{\mathbf{v}_i\}_{i=1}^K$, scaled by a learnable coefficient β :

$$\hat{z}_{l,i} = z_{l,i} - \beta \cdot \|\bar{z}_a - \mathbf{v}_i\|_2^2, \quad i = 1, \dots, K. \quad (12)$$

Given the audio embeddings \mathbf{z}_a and prompt embeddings \mathbf{z}_p , the autoregressive cross-entropy loss is defined as

$$\mathcal{L}_{\text{ce}} = -\frac{1}{L} \sum_{l=1}^L \log [p(y_l | y_{<l}, \mathbf{z}_a, \mathbf{z}_p)], \quad (13)$$

where $p(\cdot)$ denotes the predicted probability distribution at step l . Both the original logits ($p_{\text{orig}} = \text{Softmax}(\mathbf{z}_l)$) and the guided logits ($p_{\text{tg}} = \text{Softmax}(\hat{\mathbf{z}}_l)$) are utilized to compute the loss functions \mathcal{L}_{dec} and \mathcal{L}_{tg} , respectively.

2.4. Loss Functions

The final objective combines all components with weights:

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{cma}} + \lambda_2 \mathcal{L}_{\text{tg}} + \lambda_3 \mathcal{L}_{\text{dec}}. \quad (14)$$

3. EXPERIMENTS

3.1. Experimental Settings

3.1.1. Datasets

We pre-train the model on the combined dataset of AudioCaps [30], Clotho [31], and WavCaps [27], and fine-tune it separately on AudioCaps and Clotho. AudioCaps provides 48,595 training and 944 test clips (10 s each) with audio-based annotations, retrieved from available links. Clotho v2.1 contains 3,839 development, 1,045 validation, and 1,045 evaluation clips (15–30 s), each with five captions. WavCaps comprises 403,050 clips (a few seconds to over one minute) collected from AudioSet [32], BBC Sound Effects,¹ FreeSound,² and SoundBible.³

3.1.2. Implementation Details and Metrics

We adopt the consistent ensemble distillation [33] model as the audio encoder and LLaMA 2 (7B) [20] as the text decoder, fine-tuned with LoRA [34]. Query numbers are set to $N_s=8$, $N_t=8$, and $N_g=32$. All experiments used AdamW with

¹<https://sound-effects.bbcrewind.co.uk>

²<https://freesound.org>

³<https://soundbible.com>

Table 3. Ablation of Guide Scores on AudioCaps.

Guide Scores	MT	CD	SC	SD	SF	FS
Cosine sim.	26.6	81.6	19.0	50.3	50.2	66.8
L_1 dist.	25.9	88.6	17.5	53.1	52.4	65.5
L_2 dist. (ours)	27.1	91.1	19.7	55.4	55.3	67.7

Table 4. Ablation of LAMB components of on AudioCaps.

Element Ablation	MT	CD	SC	SD	SF	FS
LAMB (Ours)	27.1	91.1	19.7	55.4	55.3	67.7
w/o TSA	26.7	87.6	18.9	53.3	53.2	65.9
w/o \mathcal{L}_{CS} ($\mathcal{L}_{\text{CS-global+CS-token}}$)	24.4	84.8	16.6	50.7	49.4	63.0
w/o CMA (\mathcal{L}_{cma})	25.8	78.6	18.9	48.7	48.6	65.8
w/o TG	26.8	82.2	19.2	50.7	50.6	66.9

weight decay 1e-6, trained for 30 epochs (2 warmup epochs) under a cosine annealing schedule. Pre-training employed a learning rate of 5e-5 with batch size 32, and fine-tuning used 3e-6 with batch size 8. For evaluation, we adopt common AAC metrics, including METEOR [35], CIDEr [36], SPICE [37], SPIDER [38], SPIDER-FL [39], and FENSE [39].

3.2. Main Results

Tab. 1 summarizes the performance of LAMB on the AAC benchmarks. LAMB surpasses prior work on *all* evaluation metrics, establishing clear state-of-the-art results on AudioCaps. On AudioCaps, it achieves notable gains, particularly on CIDEr, SPIDER, and SPIDER-FL, while also improving METEOR and SPICE, reflecting stronger semantic and structural fidelity. On Clotho, which is more challenging due to diverse captions, LAMB shows competitive performance with improvements on several metrics, demonstrating robustness to annotation variability. Notably, while several prior works rely on additional pre-training data, our model attains comparable or better results across the datasets.

These results show that bridging the modality gap enhances the LLM’s reasoning and produces captions that are more faithful, coherent, and aligned with human preference, reflected in improvements in FENSE and SPIDER-FL.

3.3. Ablation Studies

Modality Gap Alignment Strategies. We first assess the effectiveness of the Cross-Modal Aligner (CMA). As shown in Tab. 2, CMA yields the lowest L_2 distance and highest cosine similarity, surpassing linear and Q-Former baselines. Moreover, Fig. 2 visualizes with UMAP that CMA clearly reduces the modality gap compared to the case without CMA.

Token Guide Metrics. Tab. 3 presents the ablation results on identifying effective methods for computing Guide Scores in Token Guide (TG). Squared L_2 distance proves most effective, stably steering decoder logits and achieving the highest scores, while L_1 distance slightly degrades performance and cosine similarity yields results inferior to the baseline without guidance. Even without auxiliary modules or external knowl-

Table 5. Qualitative results without and with Token Guide. Underlined text indicates audio cues captured more explicitly.

w/o Token Guide	w Token Guide
“ <i>Basketballs are dribbled and shoes squeak as a man speaks</i> ”	“ <i>Sneakers squeaking and basketballs bouncing on a hard surface as a group of people talk in the background</i> ”
“ <i>A loud crash followed by a man speaking and a woman screaming</i> ”	“ <i>Glass shattering followed by a man speaking then a woman speaking before a rooster crows and birds chirping</i> ”

edge, Guide Scores derived from the LLM text embedding space effectively guide the decoder.

Effect of Components. Tab. 4 highlights the role of each module by ablating components of LAMB. Two-Stream Adapter (TSA) provides semantic and temporal features of audio to CMA, and its removal reduces both semantic coverage and structural fidelity. Within CMA, two loss terms are used: CS divergence and InfoNCE. Removing \mathcal{L}_{CS} causes a larger performance drop (rows 1 & 3) than removing $\mathcal{L}_{\text{InfoNCE}}$ (rows 3 & 4), and excluding both further lowers all metrics. TG is also important, as its absence lowers performance across all metrics. Overall, these modules offer complementary benefits, with the full model achieving the best results.

3.4. Qualitative Results

In Tab. 5, we present captions generated by LAMB trained with and without TG. Since the Guide Scores are computed as the squared L_2 distance between aligned audio embeddings and LLM token embeddings, semantically related tokens receive higher weights, enabling more accurate and precise guidance. In the first example, TG allows the model to capture not only the sound of a basketball bouncing but also people talking, generating details such as “*hard surface*” and “*a group of people talk in the background*”. In the second example, TG links “*loud crash*” to “*glass shattering*” and additionally identifies background sounds such as “*a rooster crows*” and “*birds chirping*”. These results indicate that TG can directly guide in the LLM text embedding space, enabling precise and accurate caption generation.

4. CONCLUSION

In this work, we introduced **LAMB**, an LLM-based audio captioning framework that, for the first time, applies Cauchy-Schwarz divergence to bridge the audio-text modality gap in audio captioning. LAMB integrates a Cross-Modal Aligner for fine-grained alignment and a Two-Stream Adapter to enrich audio representations, while the Token Guide steers decoder logits using Guide Scores in the LLM text embedding space. Experimental results show that LAMB outperforms prior work on AudioCaps and achieves competitive performance on Clotho, indicating that bridging the modality gap enables effective use of LLM reasoning in audio captioning.

5. ACKNOWLEDGEMENT

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2025-02215122, Development and Demonstration of Lightweight AI Model for Smart Homes).

References

[1] Xuenan Xu, Zeyu Xie, Mengyue Wu, and Kai Yu, “Beyond the status quo: A contemporary survey of advances and challenges in audio captioning,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 32, pp. 95–112, 2023.

[2] Minkyu Kim, Kim Sung-Bin, and Tae-Hyun Oh, “Prefix tuning for automated audio captioning,” in *Proc. ICASSP*, 2023.

[3] Changli Tang, Wenyi Yu, et al., “Salmonn: Towards generic hearing abilities for large language models,” *arXiv preprint arXiv:2310.13289*, 2023.

[4] Jongsuk Kim, Jiwon Shin, and Junmo Kim, “AVCap: Leveraging audio-visual features as text tokens for captioning,” in *Proc. Interspeech*, 2024.

[5] Kyeongha Rho, Hyeongkeun Lee, Valentio Iverson, and Joon Son Chung, “Lavcap: Llm-based audio-visual captioning using optimal transport,” in *Proc. ICASSP*, 2025.

[6] Jizhong Liu, Gang Li, et al., “Enhancing automated audio captioning via large language models with optimized audio encoding,” in *Proc. Interspeech*, 2024.

[7] Wenxi Chen, Ziyang Ma, et al., “SLAM-AAC: Enhancing audio captioning with paraphrasing augmentation and clap-refine through llms,” in *Proc. ICASSP*, 2025.

[8] Jaeyeon Kim, Jaeyoon Jung, et al., “EnCLAP: Combining neural audio codec and audio-text joint embedding for automated audio captioning,” in *Proc. ICASSP*, 2024.

[9] Moayed Haji-Ali, Willi Menapace, Aliaksandr Siarohin, Guha Balakrishnan, and Vicente Ordonez, “Taming data and transformers for audio generation,” *arXiv*, 2024.

[10] Daiki Takeuchi, Binh Thien Nguyen, Masahiro Yasuda, et al., “Clap-art: Automated audio captioning with semantic-rich audio representation tokenizer,” in *Proc. Interspeech*, 2025.

[11] Changin Choi, Sungjun Lim, and Wonjong Rhee, “Enhancing retrieval-augmented audio captioning with generation-assisted multimodal querying and progressive learning,” in *Proc. Interspeech*, 2025.

[12] Qiuqiang Kong, Yin Cao, et al., “PANNs: Large-scale pre-trained audio neural networks for audio pattern recognition,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.

[13] Sanyuan Chen, Yu Wu, et al., “BEATs: audio pre-training with acoustic tokenizers,” in *Proc. ICML*, 2023.

[14] Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang, “CLAP: Learning audio concepts from natural language supervision,” in *Proc. ICASSP*, 2023.

[15] Jacob Devlin et al., “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL*, 2019.

[16] Alec Radford, Jeffrey Wu, et al., “Language models are unsupervised multitask learners,” *OpenAI blog*, 2019.

[17] Mike Lewis, Yinhan Liu, et al., “BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” in *Proc. ACL*, 2020.

[18] KiHyun Nam, Jongmin Choi, Hyeongkeun Lee, Jungwoo Heo, and Joon Son Chung, “Diffusion-link: Diffusion probabilistic model for bridging the audio-text modality gap,” *arXiv*, 2025.

[19] Wei-Lin Chiang, Zhuohan Li, et al., “Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality,” *vicuna.lmsys.org*, vol. 2, no. 3, pp. 6, 2023.

[20] Hugo Touvron, Louis Martin, et al., “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.

[21] Hyung Won Chung, Le Hou, et al., “Scaling instruction-finetuned language models,” *J. Mach. Learn. Res.*, vol. 25, no. 70, pp. 1–53, 2024.

[22] Ahmed Masry, Juan A Rodriguez, et al., “Alignvilm: Bridging vision and language latent spaces for multimodal understanding,” *arXiv*, 2025.

[23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi, “Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” in *Proc. ICML*, 2023.

[24] Jose C Principe et al., “Learning from examples with information theoretic criteria,” *J. VLSI Signal Process. Syst. Signal Image Video Technol.*, vol. 26, no. 1, pp. 61–77, 2000.

[25] Wenzhe Yin, Zehao Xiao, et al., “Distributional vision-language alignment by cauchy-schwarz divergence,” *arXiv preprint arXiv:2502.17028*, 2025.

[26] Junyoung Chung et al., “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv*, 2014.

[27] Xinhao Mei, Chutong Meng, et al., “WavCaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2024.

[28] Emanuel Parzen, “On estimation of a probability density function and mode,” *Ann. Math. Stat.*, vol. 33, no. 3, pp. 1065–1076, 1962.

[29] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” *arXiv*, 2018.

[30] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim, “Audiocaps: Generating captions for audios in the wild,” in *Proc. NAACL*, 2019.

[31] Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen, “Clotho: An audio captioning dataset,” in *Proc. ICASSP*, 2020.

[32] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, et al., “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. ICASSP*, 2017.

[33] Heinrich Dinkel et al., “Ced: Consistent ensemble distillation for audio tagging,” in *Proc. ICASSP*, 2024.

[34] Edward J Hu, Yelong Shen, et al., “LoRA: Low-rank adaptation of large language models,” in *Proc. ICLR*, 2022.

[35] Satanjeev Banerjee and Alon Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. ACL Workshop*, 2005.

[36] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh, “Cider: Consensus-based image description evaluation,” in *Proc. CVPR*, 2015.

[37] Peter Anderson et al., “Spice: Semantic propositional image caption evaluation,” in *Proc. ECCV*, 2016.

[38] Siqi Liu, Zhenhai Zhu, et al., “Improved image captioning via policy gradient optimization of spider,” in *Proc. ICCV*, 2017.

[39] Zelin Zhou et al., “Can audio captions be evaluated with image caption metrics?,” in *Proc. ICASSP*, 2022.