

LAVCap: LLM-based Audio-Visual Captioning using Optimal Transport

Kyeongha Rho^{1*}, Hyeongkeun Lee^{1*}, Valentio Iverson², Joon Son Chung¹

¹KAIST, South Korea, ²University of Waterloo, Canada

{khrho325, lhk528, joonson}@kaist.ac.kr, viverson@uwaterloo.ca

Abstract—Automated audio captioning is a task that generates textual descriptions for audio content, and recent studies have explored using visual information to enhance captioning quality. However, current methods often fail to effectively fuse audio and visual data, missing important semantic cues from each modality. To address this, we introduce LAVCap, a large language model (LLM)-based audio-visual captioning framework that effectively integrates visual information with audio to improve audio captioning performance. LAVCap employs an optimal transport-based alignment loss to bridge the modality gap between audio and visual features, enabling more effective semantic extraction. Additionally, we propose an optimal transport attention module that enhances audio-visual fusion using an optimal transport assignment map. Combined with the optimal training strategy, experimental results demonstrate that each component of our framework is effective. LAVCap outperforms existing state-of-the-art methods on the AudioCaps dataset, without relying on large datasets or post-processing. Code is available at <https://github.com/NAVER-INTEL-Co-Lab/gaudi-lavcap>.

Index Terms—multimodal captioning, LLM, optimal transport

I. INTRODUCTION

Automated audio captioning (AAC) is a task that aims to generate textual descriptions of a given audio input. In contrast to audio tagging and sound event classification tasks that focus on understanding low-level sound characteristics, AAC also requires a comprehension of high-level semantics contained in the audio. It has recently garnered significant attention from the deep learning community due to its wide range of applications, such as providing audio descriptions for broadcasts and movies, text query-based audio retrieval, and developing more human-like conversational AI systems [1].

Most recent works [2]–[9] for AAC adopt the encoder-decoder framework and have explored different architectures. They utilize various audio encoders [10]–[14] to extract semantically rich acoustic features. Pre-trained language models are used as text decoders because of their sequence modeling and text generation capabilities. In particular, some recent works [4], [8] leverage LLMs for their strong ability to understand contexts and generate text sequences.

Recently, several studies [15], [16] have incorporated visual modality to improve audio captioning. Visual information helps distinguish sounds in complex scenes. For example, when a man speaks while drilling, visual cues clarify the presence of both elements, allowing for generating more accurate captions.

To this end, we introduce **LAVCap**, an LLM-based Audio-Visual Captioning framework. This work focuses on exploring various audio-visual fusion methods and identifying the optimal training strategy for the LLM-based audio-visual captioning framework. First, we find out that naively combining visual features with audio features does not enable the LLM to take full advantage of the visual information. We attribute this shortfall to the significant modality gap between audio and visual feature spaces.

To bridge the modality gap, we consider the alignment of audio and visual tokens as an optimal transport (OT) problem. While previous work [17] applies OT algorithms to align video and text, this study pioneers their application to audio-visual alignment. Specifically, we introduce an optimal transport-based alignment loss (OT loss) tailored for this purpose.

OT loss encourages the encoders to extract features that are rich in semantics while ensuring that they are well-aligned. In addition, our experiment reveals that the existing cross-attention mechanism [15] struggles to integrate cross-modal features effectively. To tackle this, we propose an optimal transport attention module (*OT-Att*) that leverages the OT assignment map as attention weights for audio-visual fusion. This approach provides more effective fusion compared to other audio-visual fusion methods.

Experimental results demonstrate that LAVCap outperforms previous state-of-the-art methods on the AudioCaps benchmark, without pre-training the model on large datasets or applying post-processing to generated captions. It is noteworthy that our approach achieves such high performance although the ground-truth captions in AudioCaps are audio-centric. Furthermore, our user study shows that LAVCap obtains mean opinion scores (MOS) even higher than the ground-truth captions. This underscores the importance of utilizing the visual modality to distinguish the various sounds and understand the scene.

The main contributions of this work are summarized as follows:

- We present **LAVCap**, an LLM-based audio-visual captioning framework that leverages visual information to complement audio modality in audio captioning tasks.
- We employ optimal transport for effectively bridging the modality gap and fusing audio-visual features.
- Our approach outperforms existing state-of-the-art methods on AudioCaps without the need for pre-training the model on large datasets or applying post-processing.

*These authors contributed equally to this work.

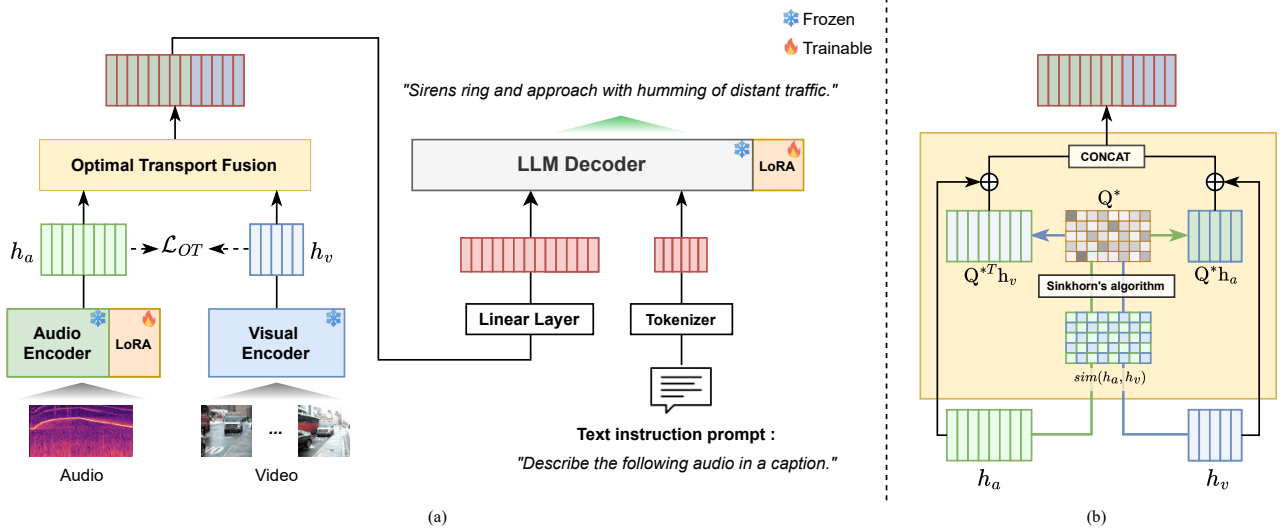


Fig. 1. (a) Overview of the proposed LAVCap Framework. (b) Detail of the Optimal Transport Fusion module.

II. METHOD

The overall architecture of LAVCap is illustrated in Fig. 1. Each component of our framework and the training objectives are explained in detail in the following sections.

A. Audio-Visual Encoding

Given an audio-visual input pair (x_a, x_v) , the audio encoder E_a and the visual encoder E_v encode them into modality-specific features with C dimensions as follows:

$$h_a = E_a(x_a), h_v = E_v(x_v) \quad (1)$$

where $h_a \in \mathbb{R}^{N_a \times C}$, $h_v \in \mathbb{R}^{N_v \times C}$ represent the audio and visual features, respectively, with N_a and N_v denoting the number of tokens for audio-visual features.

B. Audio-Visual Alignment Based on Optimal Transport

To effectively bridge the modality gap, we apply OT loss to align the audio-visual features. First of all, we compute the similarity matrix $\mathbf{S} \in \mathbb{R}^{N_a \times N_v}$ of audio and visual features. Similar to the prior work [17], we optimize the audio-visual assignment map $\mathbf{Q} \in \mathbb{R}^{N_a \times N_v}$ to maximize the global similarity between the cross-modal features:

$$\begin{aligned} \max_{\mathbf{Q} \in \mathcal{Q}} \quad & \text{tr}(\mathbf{Q}^T \mathbf{S}) + \epsilon H(\mathbf{Q}) \\ \text{s.t. } \mathcal{Q} = \{ \mathbf{Q} \mid & \mathbf{Q} \mathbf{1}_{N_v} = \frac{1}{N_a} \mathbf{1}_{N_a}, \mathbf{Q}^T \mathbf{1}_{N_a} = \frac{1}{N_v} \mathbf{1}_{N_v} \}, \end{aligned} \quad (2)$$

where $\mathbf{1}_{N_a}$, $\mathbf{1}_{N_v}$ denotes the vector of ones in dimension N_a and N_v . $H(\mathbf{Q})$ is an entropy regularization term with ϵ controlling the smoothness. Then the optimal solution \mathbf{Q}^* is obtained by iterative Sinkhorn-Knopp algorithm [18]:

$$\mathbf{Q}^* = \text{Diag}(\boldsymbol{\mu}) \exp(\mathbf{S}/\epsilon) \text{Diag}(\boldsymbol{\nu}) \quad (\boldsymbol{\mu} \in \mathbb{R}^{N_a}, \boldsymbol{\nu} \in \mathbb{R}^{N_v}) \quad (3)$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are the non-negative scaling vectors.

Based on the optimal solution \mathbf{Q}^* and the similarity matrix \mathbf{S} , OT loss is computed as follows:

$$\mathcal{L}_{\text{OT}} = \sum_{m \in \{a, v\}} \left[- \sum_{i=1}^{N_m} \log \frac{\exp(\langle \mathbf{q}_i^m, \mathbf{s}_i^m \rangle / \tau)}{\sum_{j=1}^{N_m} \exp(\langle \mathbf{q}_i^m, \mathbf{s}_j^m \rangle / \tau)} \right], \quad (4)$$

where $\{\mathbf{q}_i^m\}_{i=1}^{N_m}$, $\{\mathbf{s}_i^m\}_{i=1}^{N_m}$ denote the row vectors of \mathbf{Q} and \mathbf{S} when $m = a$, and the column vectors when $m = v$. OT loss encourages maximizing the similarity between audio-visual token sequences within a sample, rather than between averaged features in a mini-batch. This provides more fine-grained supervision and thus a more effective way to align the cross-modal features.

C. Audio-Visual Fusion and Projection

In addition to using the optimal transport assignment map for loss computation, we also employ it for the fusion of audio-visual features. We refer to this fusion module as **OT-Att** module in Eq. 5. OT-Att module operates similarly to the conventional cross-attention but utilizes the optimal transport assignment map as the attention weight. The visually-attended audio feature and audio-attended visual feature are computed as follows:

$$\begin{aligned} \hat{h}_a &= \text{OT-Att}(h_a, \mathbf{Q}^*, h_v) = h_a + \mathbf{Q}^* h_v \\ \hat{h}_v &= \text{OT-Att}(h_v, \mathbf{Q}^{*T}, h_a) = h_v + \mathbf{Q}^{*T} h_a \end{aligned} \quad (5)$$

Then these two features are concatenated token-wise and then projected to the LLM latent space through the linear projector:

$$h_{av} = W_{av \rightarrow t} \text{Concat}(\hat{h}_a, \hat{h}_v) \quad (6)$$

D. Text Decoding

Since we leverage an LLM as the text decoder, the text instruction prompt x_t needs to be transformed into text embeddings through the LLM tokenizer E_t . Then the fused audio-visual features and text embeddings are concatenated and fed into the LLM decoder D_t to produce the output z_t :

$$z_t = D_t(\text{Concat}(h_{av}, E_t(x_t)), M_{att}) \quad (7)$$

where M_{att} denotes attention masks for reflecting the autoregressive property of LLM.

TABLE I

AAC RESULTS ON AUDIOCAPS TEST SPLIT FOR VARIOUS MODELS. BOLD INDICATES THE BEST AMONG AUDIO-VISUAL METHODS, WHILE UNDERLINED REPRESENTS THE OVERALL TOP PERFORMERS. NOTE THAT GROUND-TRUTH CAPTIONS ARE BASED SOLELY ON AUDIO. [†]RESULTS REPRODUCED ON OUR ENVIRONMENT. AS: AUDIOSET, AC: AUDIOCAPS, WC: WAVCAPS, CL: CLOTHO, MA: MULTI-ANNOTATOR CAPTIONED SOUNDSCAPES.

Model	Pre-training Dataset	BLEU ₁ ↑	BLEU ₄ ↑	ROUGE _L ↑	METEOR ↑	CIDEr ↑	SPICE ↑	SPIDEr ↑
Audio-Based								
ACT [2]	AS	64.7	25.2	46.8	22.2	67.9	16.0	42.0
BART-tags [3]	AS	69.9	26.6	49.3	24.1	75.3	17.6	46.5
CNN-GPT2 [4]	-	71.3	30.9	50.3	24.0	73.3	17.7	45.5
EnCLAP-large [5]	-	-	-	-	25.5	80.3	18.8	49.5
HTSAT-BART [6]	AC+CL+WC	70.7	28.3	50.7	25.0	78.7	18.2	48.5
CNext-trans [7]	AC+CL+MA+WC	-	-	-	25.2	80.6	18.4	49.5
LOAE [†] [8]	-	69.8	25.8	49.2	24.8	75.6	18.2	46.9
LOAE [8]	AC+CL+WC	-	-	-	<u>26.7</u>	81.6	<u>19.3</u>	50.5
AutoCap (audio+text) [9]	AC	72.1	28.6	51.5	25.6	80.0	18.8	49.4
AutoCap (audio+text) [9]	AC+CL+WC	<u>72.3</u>	<u>29.7</u>	<u>51.8</u>	25.3	83.2	18.2	50.7
Audio-Visual								
V-ACT [15]	-	69.8	28.1	49.4	23.7	71.1	17.2	44.2
AVCap (freeze) [16]	-	70.8	29.5	49.8	22.8	74.4	16.2	45.5
AVCap (finetuning) [16]	-	68.1	28.7	49.1	24.3	75.8	17.8	46.8
LAVCaps (ours)	-	72.3	29.7	51.0	26.2	84.9	18.5	51.7

TABLE II
ABLATION ON THE USE OF VISUAL MODALITY AND \mathcal{L}_{OT} .

Audio	Visual	\mathcal{L}_{OT}	METEOR ↑	CIDEr ↑	SPICE ↑	SPIDEr ↑
✓	✗	✗	24.6	77.9	17.4	47.6
✓	✓	✗	25.4	78.1	18.8	48.5
✓	✓	✓	26.3	83.1	18.9	51.0

TABLE III
ABLATION ON AUDIO-VISUAL FUSION METHOD. \mathcal{L}_{OT} IS ALL APPLIED.

Method	METEOR ↑	CIDEr ↑	SPICE ↑	SPIDEr ↑
Q-Former	24.4	77.0	18.0	47.5
Joint encoder	24.9	76.7	18.4	47.5
Cross attention	24.8	80.1	17.9	49.0
OT-Att (ours)	26.2	84.9	18.5	51.7

E. Training Objectives

As well as OT loss, the conventional autoregressive cross-entropy loss is also used for training:

$$\mathcal{L}_{CE} = -\frac{1}{T} \sum_{i=1}^T \log p(y_i | y_{1:i-1}, h_{av}, h_t) \quad (8)$$

where y_i is a i -th text token. The final training objective is the weighted sum of two losses:

$$\mathcal{L} = \lambda_{CE} \mathcal{L}_{CE} + \lambda_{OT} \mathcal{L}_{OT} \quad (9)$$

III. EXPERIMENTS

A. Experimental settings

1) *Datasets*: We utilize the AudioCaps dataset [19] for training and evaluation, where each 10-second clip is annotated based on its audio component. Due to the limited accessibility of some YouTube links, we acquired 48,595 clips for the training and 944 clips for the test set. For audio pre-processing, we applied Short-Time Fourier Transform using a 25-ms Hanning window with a 10-ms hop size to each 10-second waveform sampled at 16 kHz, resulting in a 1024×64 spectrogram. For visual input, 20 frames are uniformly selected from a 10-second clip at 2 FPS, then center-cropped to 224×224 pixels and normalized.

TABLE IV
ABLATION OF THE ENCODER AND DECODER TRAINING STRATEGIES.

Encoder	Decoder	METEOR ↑	CIDEr ↑	SPICE ↑	SPIDEr ↑
Finetune	Freeze	25.3	77.7	17.3	47.5
LoRA	Freeze	24.4	73.5	17.3	45.4
Finetune	LoRA	25.5	78.9	18.1	48.5
LoRA	LoRA	26.2	84.9	18.5	51.7

2) *Implementation Details and Metrics*: During training, we utilize the AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of $1e-6$. For the first two epochs, out of a total of 100, the learning rate warms up to $5e-6$, and then it gradually decreases following a cosine annealing strategy. We adopt a pre-trained Consistent Ensemble Distillation model [20] as an audio encoder, and a pre-trained CLIP ViT-L/14 model [21] as a visual encoder. For text decoding, Llama 2 [22] with 7B parameters is employed. Both the audio encoder and text decoder are fine-tuned using low-rank adaptation (LoRA) [23], while the visual encoder is kept frozen. When evaluating our methods, we use the metrics commonly employed for AAC, including BLEU [24], ROUGE-L [25], METEOR [26], CIDEr [27], SPICE [28], and SPIDEr [29]. Experiments are conducted using Intel Gaudi 2 AI Accelerator.

B. Main Results

The audio captioning performance of LAVCap on the AudioCaps dataset is shown in Table I. LAVCap not only outperforms previous works in closely matching the lexical content of ground truth captions but also shows enhanced semantic relevance and informativeness. It is impressive that our framework achieves high performance although the ground-truth captions in AudioCaps are annotated based solely on audio. Notably, our model still achieves comparable performance to the concurrent works [8], [9] that utilize multiple and additional datasets for pre-training. The results demonstrate that using an OT mapping between audio and visual modalities enables the model to train more effectively on semantically aligned features across these modalities, compared to previous methods such as cross-attention and concatenation. A further analysis of the results can be found in Section III-C.

TABLE V
EXAMPLES OF INSTRUCTION PROMPTS FOR THE PROPOSED METHOD.

Prompts I
USER: <Audio>[Token _{a+vv}]/</Audio> Please describe the audio.\n
ASSISTANT:
Prompts II
USER: <Content>[Token _{a+vv}]/</Content> Please describe the content.\n
ASSISTANT:
Prompts III
USER: <Audio>[Token _a]/</Audio><Visual>[Token _v]/</Visual> Please describe the audio and visual elements.\n
ASSISTANT:

TABLE VI
ABLATION RESULTS BASED ON EXAMPLES OF INSTRUCTION PROMPTS.

Prompt	METEOR \uparrow	CIDEr \uparrow	SPICE \uparrow	SPIDEr \uparrow
Prompts I	26.2	84.9	18.5	51.7
Prompts II	25.8	82.5	19.0	50.8
Prompts III	25.9	81.8	19.1	50.5

C. Ablation Studies

1) *Visual Modality and OT Loss*: We conduct an ablation study to evaluate the effectiveness of leveraging visual modality and OT loss. As shown in Table. II, the performance improvement from simply adding visual modality is marginal without OT loss. This demonstrates that bridging the modality gap through OT loss is crucial for enabling LLM to comprehend multi-modal contexts and process audio-visual features.

2) *Audio-Visual Fusion Methods*: Based on the utilization of visual modality and OT loss, we explore various audio-visual fusion methods. The results in Table. III show that the cross-modal fusion methods employed by the previous works are not effective in our setting. We infer that this is due to the lack of training data to optimize the learnable parameters of these fusion methods and simultaneously align the fused audio-visual features to the LLM latent space. On the other hand, the proposed OT attention module does not need any learnable parameters and leverages an optimal transport assignment map as an attention weight, thus providing effective audio-visual fusion in a data-efficient way.

3) *Encoder-Decoder Training Strategy*: After deciding on the optimal network architecture, we seek the best training strategy for our framework. As shown in Table. IV, training the encoder with LoRA instead of fine-tuning the entire parameters is much more efficient under data-limited circumstances. In addition, freezing the weights of the LLM decoder severely degrades the performance, indicating that the LLM decoder should be adapted to the target dataset.

4) *Instruction Prompts*: To ensure that the LLM text decoder understands the input tokens properly, we explore various instruction prompts detailed in Table V. While *prompts I* and *II* handle audio and visual tokens together, *prompt III* is specifically designed to understand them separately. Since AAC focuses primarily on audio-based captioning, *prompt I* performs slightly better than the others, as shown in Table VI.

D. Qualitative Results

In Fig. 2 we visualize the captions generated by models trained with audio-only, visual-only, and combined audio-



Fig. 2. Qualitative results of captions generated from models trained solely on audio, only on visual, and on both audio and visual.

TABLE VII
MOS WITH 95% CONFIDENCE INTERVALS.

Input modality	MOS \uparrow (1~5)
ground truth	3.81 \pm 0.12
audio-only	2.89 \pm 0.12
visual-only	2.71 \pm 0.12
audio-visual	4.08 \pm 0.10

visual data. While using a single modality may capture incorrect elements, utilizing both audio and visual inputs provides a more detailed view of the video. Furthermore, we conduct a user study with 20 participants, using MOS to rate each of the 35 generated captions from 1 to 5 based on how well they describe the video. The results in Table VII show that the captions generated by LAVCap concretely represent the audio content. Remarkably, MOS of LAVCap is even higher than the ground-truth captions. This highlights the significant benefit of incorporating the visual modality to better distinguish the various sounds and comprehend the scene context.

IV. CONCLUSION

In this work, we propose LAVCap, an LLM-based audio-visual captioning framework designed to incorporate visual information into the audio modality using an optimal transport-based strategy. Specifically, *OT loss* and *OT-Att* are introduced to align the modality gap and promote effective fusion of audio and visual features. The proposed model outperforms previous captioning methods on the AudioCaps dataset, without the need for pre-training on large datasets or post-processing, highlighting its promise in designing new audio-visual fusion methods.

V. ACKNOWLEDGMENT

This research was supported in part by the NAVER-Intel Co-Lab. The work was conducted by Korea Advanced Institute of Science and Technology and reviewed by both NAVER and Intel.

REFERENCES

- [1] Q. Huang, Y. Zhang, T. Ko, X. Liu, B. Wu, W. Wang, and H. Tang, "Personalized dialogue generation with persona-adaptive attention," in *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 12916–12923, 2023.
- [2] X. Mei, X. Liu, Q. Huang, M. D. Plumbley, and W. Wang, "Audio captioning transformer," in *Proceedings of the 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop*, pp. 211–215, 2021.
- [3] F. Gontier, R. Serizel, and C. Cerisara, "Automated audio captioning by fine-tuning bart with audioset tags," in *DCASE 2021 - 6th Workshop on Detection and Classification of Acoustic Scenes and Events*, 2021.
- [4] M. Kim, K. Sung-Bin, and T.-H. Oh, "Prefix tuning for automated audio captioning," in *International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2023.
- [5] J. Kim, J. Jung, J. Lee, and S. H. Woo, "EnCLAP: Combining neural audio codec and audio-text joint embedding for automated audio captioning," in *International Conference on Acoustics, Speech and Signal Processing*, pp. 6735–6739, 2024.
- [6] X. Mei, C. Meng, H. Liu, Q. Kong, T. Ko, C. Zhao, M. D. Plumbley, Y. Zou, and W. Wang, "WavCaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [7] E. Labb, T. Pellegrini, J. Pinquier, *et al.*, "CoNeTTE: An efficient audio captioning system leveraging multiple datasets with task embedding," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [8] J. Liu, G. Li, J. Zhang, H. Dinkel, Y. Wang, Z. Yan, Y. Wang, and B. Wang, "Enhancing automated audio captioning via large language models with optimized audio encoding," in *Interspeech*, pp. 1135–1139, 2024.
- [9] Y. Zhang, Q. Chen, M. Li, W. Che, and L. Qin, "AutoCAP: Towards automatic cross-lingual alignment planning for zero-shot chain-of-thought," in *Findings of the Association for Computational Linguistics ACL 2024*, pp. 9191–9200, 2024.
- [10] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "PANNs: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [11] 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 *International Conference on Acoustics, Speech and Signal Processing*, pp. 646–650, 2022.
- [12] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, "BEATs: audio pre-training with acoustic tokenizers," in *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [13] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, "CLAP: Learning audio concepts from natural language supervision," in *International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2023.
- [14] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *Transactions on Machine Learning Research*, 2023.
- [15] X. Liu, Q. Huang, X. Mei, H. Liu, Q. Kong, J. Sun, S. Li, T. Ko, Y. Zhang, L. H. Tang, *et al.*, "Visually-aware audio captioning with adaptive audio-visual attention," in *Interspeech*, pp. 2838–842, 2023.
- [16] J. Kim, J. Shin, and J. Kim, "AVCap: Leveraging audio-visual features as text tokens for captioning," in *Interspeech*, 2024.
- [17] Y. Lin, J. Zhang, Z. Huang, J. Liu, Z. Wen, and X. Peng, "Multi-granularity correspondence learning from long-term noisy videos," in *The Twelfth International Conference on Learning Representations*, 2024.
- [18] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Proceedings of the 26th International Conference on Neural Information Processing Systems*, p. 2292–2300, 2013.
- [19] C. D. Kim, B. Kim, H. Lee, and G. Kim, "Audiocaps: Generating captions for audios in the wild," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 119–132, 2019.
- [20] H. Dinkel, Y. Wang, Z. Yan, J. Zhang, and Y. Wang, "Ced: Consistent ensemble distillation for audio tagging," in *International Conference on Acoustics, Speech and Signal Processing*, pp. 291–295, 2024.
- [21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763, 2021.
- [22] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [23] 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 *The Tenth International Conference on Learning Representations*, 2022.
- [24] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: a method for automatic evaluation of machine translation," in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, p. 311–318, 2002.
- [25] C.-Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in *Text Summarization Branches Out*, pp. 74–81, 2004.
- [26] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pp. 65–72, 2005.
- [27] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4566–4575, 2015.
- [28] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "Spice: Semantic propositional image caption evaluation," in *Proceedings of the European Conference on Computer Vision*, pp. 382–398, 2016.
- [29] S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. Murphy, "Improved image captioning via policy gradient optimization of spider," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 873–881, 2017.