



Vela: Scalable Embeddings with Voice Large Language Models for Multimodal Retrieval

Ruofan Hu¹, Yan Xia¹, Minjie Hong¹, Jieming Zhu², Bo Chen², Xiaoda Yang¹, Minghui Fang¹, Tao Jin¹

¹Zhejiang University, China
²Huawei Noah's Ark Lab, China
jint.zju@zju.edu.cn

1

Abstract

Multimodal large language models (MLLMs) have seen substantial progress in recent years. However, their ability to represent multimodal information in the acoustic domain remains underexplored. In this work, we introduce Vela, a novel framework designed to adapt MLLMs for the generation of universal multimodal embeddings. By leveraging MLLMs with specially crafted prompts and selected in-context learning examples, Vela effectively bridges the modality gap across various modalities. We then propose a single-modality training approach, where the model is trained exclusively on text pairs. Our experiments show that Vela outperforms traditional CLAP models in standard text-audio retrieval tasks. Furthermore, we introduce new benchmarks that expose CLAP models' limitations in handling long texts and complex retrieval tasks. In contrast, Vela, by harnessing the capabilities of MLLMs, demonstrates robust performance in these scenarios. Our code will soon be available.

Index Terms: information retrieval, signal representation, multimodal interaction

1. Introduction

With the rapid advancement of MLLMs, there is an increasing need for embedding models that facilitate cross-modal retrieval. In the acoustic domain, while Contrastive Language Audio Pre-training (CLAP) models [1, 2, 3, 4, 5] have demonstrated impressive performance in text-audio retrieval by aligning acoustic and linguistic representations through contrastive learning. However, several limitations still remain in these works.

One major challenge lies in the excessive dependence on contrastive learning, which limits their ability to fully capture the nuances of each modality [6]. Moreover, the text encoders used are optimized for short captions, restricting their capacity to handle long and complex queries commonly encountered in practical scenarios [7, 8]. Together, these constraints result in models that are confined to coarse-grained alignment at the caption level, and demand a substantial amount of paired audio-text data to facilitate the contrastive learning process.

Recent studies [9] highlight the importance of enhancing modality-specific understanding to improve generation performance. Based on these findings, we introduce Vela, a novel framework that employs MLLMs with enhanced modality comprehension for direct embedding generation. This approach offers several key advantages: First, unlike CLAP models, which depend on contrastive learning and are constrained by pairwise alignment, the multi-task paradigm allows MLLMs to




Figure 1: *Distribution of audio embeddings and text embeddings from the last token in MLLMs. We compare the results with the ordinary prompt and with our representation method using the crafted prompt and in context learning design.*

deepen their understanding of multimodal semantics via cross-task knowledge transfer. Second, MLLMs benefit from advanced text encoders and external knowledge, enabling them to handle more complex and extended contextual sequences beyond the input constraints of CLAP. Third, MLLMs transfer their grounding and understanding of knowledge directly to the retrieval process, eliminating the need for additional paired audio-text data, which is challenging to collect.

Technically, we achieve this mechanism by guiding MLLMs to autonomously compress modalities, exploiting their autoregressive nature to derive representations from the token. However, the inherent modality gap between text and audio processing in MLLMs can adversely affect the alignment, as shown in Figure 1(a). To address this, we develop an efficient prompt engineering and in-context learning design that unifies cross-modal representations into a shared space [10, 11] in Figure 1(b). Subsequently, we conduct training within this unified space under single-modality conditions. Our approach maintains frozen parameters across all MLLM components, while solely fine-tuning the language model through contrastive learning with positive-negative text pairs to strengthen its discriminative capabilities. Our contributions are:

- We propose a method for generating universal multimodal embeddings using MLLMs, harnessing prompt-based design and in-context learning to activate retrieval task capabilities that are not natively supported by the original model.
- Our framework outperforms CLAP-based methods that rely on text-audio pairs in retrieval tasks, even when trained on a

⁰Tao Jin is the corresponding author.

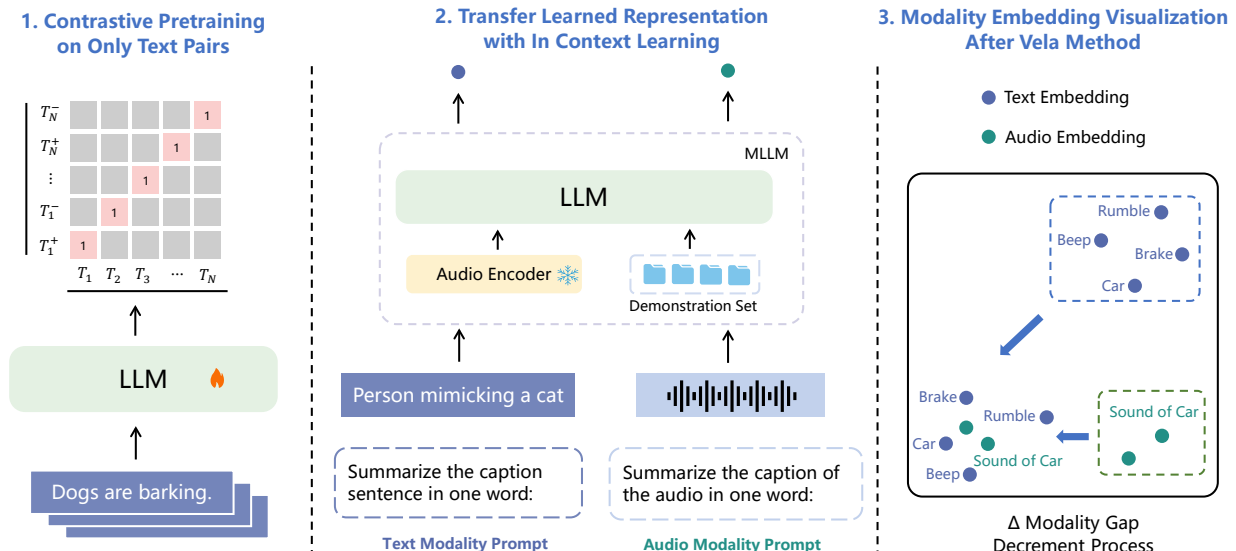


Figure 2: The whole framework of Vela. (1) In the **training** stage, by unifying multimodal representations into the same embedding space, Vela improves multimodal embeddings using **only contrastive learning on text pairs**. (2) In the **inference** stage, we freeze the modality encoder and projector in MLLM and replace the LLM with the pertaining one. With our representation method, embeddings are extracted. (3) Through the representation method and training of Vela, modality representations with similar semantics will aggregate together, which is further visualized using PCA.

limited number of text pairs in a single-modality setting.

- We introduce novel acoustic benchmarks for retrieval tasks, highlighting the limitations of traditional CLAP methods and showcasing the advantages of our Vela over them.

2. Method

2.1. Unifying Multimodal Embeddings

Inspired by previous work [11], we adopt the “**in one word**” method with MLLMs to unify multimodal embeddings. The core idea behind is to guide MLLMs in compressing the information, and then leveraging the embeddings of the last token as the representation. Through extensive experimentation, we find that using prompts like $\langle \text{text} \rangle$ Summarize the caption sentence in one word: and $\langle \text{audio} \rangle$ Summarize the caption of the audio in one word: achieves the best results. This design is driven by two key considerations. First, it guides autoregressive MLLMs to generate a one-word summary, with its hidden state reflecting the input’s probability distribution across the modalities. Second, explicitly incorporating terms like “caption” focuses the MLLM on the summarization task, thereby leveraging its multitasking capabilities to enhance the performance.

Additionally, previous studies have demonstrated that multimodal representations generated by the CLAP architecture exhibit a **modality gap** [10, 12, 13], which negatively affects the alignment in retrieval tasks. While our prompts effectively generate semantically compressed embeddings, the modality gap remains unresolved. This occurs because, during semantic compression, MLLMs tend to prioritize words with higher syntactic salience in their conditional probability distributions, rather than those semantically critical for modality alignment. For instance, in a sentence like “A person mimicking a cat’s meow”, the MLLM might emphasize “mimicking”, whereas the focus should clearly be on “meow”. This lexical salience bias ultimately perpetuates the modality gap despite effective informa-

tion compression.

To address this, we draw inspiration from **in-context learning** [14], which injects modality alignment priors into MLLMs. Specifically, our approach utilizes the `SoundBible.flac` dataset [3], carefully selected for its acoustically distinctive samples containing singular sound events. To construct high-quality teaching exemplars, we employ a two-stage refinement process: (1) GPT-4 [15] is used for the initial one-word summary generation, followed by (2) Gemini [16], which provides semantic fidelity scoring between the one-word summaries and the original captions, retaining the top 200 examples. In order to balance input length for autoregressive models with generation performance [17], we finalize the selection by keeping 65 text sentences and 35 audio clips as teaching exemplars.

2.2. Single Modality Training with Contrastive Learning

By unifying multimodal embeddings, we propose a novel single-modality training paradigm as illustrated in Figure 2. With the modality gap eliminated in the embeddings, this approach enables the transfer of single-modality representation capabilities to multimodal domains through contrastive learning on text pairs alone. To adapt to this modality-agnostic training framework, we streamline the architecture by freezing the modality encoder and projector, while training only the LLM component of the MLLM. We then utilize a simple NLI dataset [18], where each sentence pair (x_i, x_i^+, x_i^-) has a positive sample x_i^+ and a negative sample x_i^- for the input sentence x_i . We apply the same prompt and in-context learning examples for text modality to embed the sentence pairs into (h_i, h_i^+, h_i^-) . The training objective is as follows:

$$\mathcal{L} = -\log \frac{\exp(\cos(h_i, h_i^+)/\tau)}{\sum_{j=1}^N (\exp(\cos(h_i, h_j^+)/\tau) + \exp(\cos(h_i, h_j^-)/\tau))}$$

where τ is the temperature hyperparameter and N is the batch size in contrastive learning.

2.3. Complex and Longer Benchmark

The traditional CLAP framework is limited to coarse-grained, holistic representations and typically focuses on tasks involving short texts. In contrast, the Vela paradigm extends retrieval capabilities to long-text-to-audio tasks and complex retrieval tasks with designed instructions. To demonstrate this, we introduce two new benchmarks: Vela-long and Vela-conditional.

In **Vela-long**, we use an LLM [15] and Clotho [19] to expand the five captions into longer sentences while preserving their original meaning, increasing the average text length to **172.4** words—significantly more than the average of **11.32** words in the raw Clotho.

In **Vela-conditional**, we generate new samples by mixing diverse sounds [3, 20] and craft paired, narrative instructions that specify the target feature. For example, one new clip may include violin sounds, vehicle noises, and human speech, with the instruction guiding the model to ignore irrelevant elements and focus solely on the music. This methodology substitutes conventional comprehensive captions with succinct, directive instructions that emphasize partially features rather than overarching contextual elements. Each sound appears only once in the benchmark. We list all prompts in Table 1.

Table 1: *All the prompts we use in the benchmark.*

Sources	Prompts
Clotho	You will be given five short textual descriptions that correspond to a single clip. Your task is to merge these five descriptions into a single, coherent paragraph. The paragraph should be a reorganization and expansion of the original descriptions without adding any new information or altering the original meaning. Ensure the content is rich, but the meaning of each individual description must be preserved. The paragraph should flow naturally and read as a cohesive, well-structured piece of text. Make sure the final output is a smooth, connected paragraph that incorporates all five descriptions in a meaningful and coherent way.
FSD50K Wavcaps	You will be given one key label and three useless labels, you should design the prompt to instruct the system to search for the main label but ignore the useless labels.

3. Experiment

3.1. Datasets

For the training data, we use NLI [18], which contains approximately 273k sentence pairs. For the test data, we conduct text-audio retrieval on two open-source datasets, AudioCaps [21] and Clotho [19]. To show the limitations of the CLAP structure in long-text and fine-grained retrieval tasks with instructions, we employ Vela-conditional, built on the FSD50 and Wavcaps [3, 20], with approximately 400 examples, and Vela-long, based on Clotho, which includes around 1000 examples.

3.2. Experimental Setup

For the Vela backbone, we use Qwen2-Audio-7B-Instruct [22], which is based on the Qwen2 [23] series and utilizes a frozen

Whisper-v3 [24] as the audio encoder. Due to hardware constraints, we trained the model using 8 NVIDIA 3090 GPUs, with a batch size of 168, over approximately 4 days. To optimize GPU memory usage, we adopted the QLoRA [25] paradigm, applying LoRA with $r = 32$, $\alpha = 8$, and dropout = 0.05, and added LoRA modules to the linear layers of all 4-bit quantized models.

4. Result

4.1. Text-Audio Retrieval

Text-audio retrieval involves searching for an audio clip or a caption in a database based on a query from another modality. In our experiments, we selected several strong baselines based on the CLAP paradigm, including Microsoft-CLAP [1], LAION-CLAP [2], WavCaps [3], AudioSetCaps [4] and Audio-ACD [5]. We select the **zero-shot version** when available.

We report Recall@K (R@K) for K=1, 5, 10 for both audio retrieval and text retrieval in the table. Compared to these strong baselines, Vela, as a universal multimodal embedding model, achieves competitive performance on both the AudioCaps and Clotho datasets. However, our model shows slightly lower performance than LAION-CLAP in terms of Recall@10 on Audiocaps. This can be attributed to the fact that AudioCaps is a subset of the LAION-CLAP training data, leading to a similar data distribution. Additionally, since Vela was trained solely on the text modality, its **zero-shot** nature contributes to a slight performance drop. This is further supported by the results on WavCaps, where the zero-shot version, despite being trained on a larger dataset, shows lower performance during testing than other CLAP models.

Moreover, MLLMs such as Qwen2-Audio are typically designed for Next-Token prediction tasks in audio understanding. We demonstrate that, by using task-specific prompts not included in the training data, these models can still effectively transfer single-modality representation capabilities to multimodal embeddings. This indicates that the model can leverage its understanding to boost its generative capabilities, aligning with the current trend of unified understanding and generation paradigm in multimodal models [9].

4.2. Retrieval on Lengthy Caption and Conditional Instruction

Most existing methods focus on short captions, which highlights a gap in cross-modal retrieval between audio and long texts (e.g., lengthy captions or documents) [8]. To explore this unexplored area, we used LLMs [15] to expand the captions in the Clotho evaluation set, significantly increasing their length while preserving semantic integrity. As shown in Table 3, we draw two main conclusions. First, as evidenced by the comparison with Table 2, most existing models experience performance degradation when confronted with long-text-to-audio retrieval tasks with an equal number of test samples from clotho. Microsoft-CLAP, in particular, **fails** to support such long-text retrieval tasks. This underscores the need to improve the text encoders, even when using well-established models like HTSAT-BERT. Additionally, since our paragraphs are composed of five useful captions, they inherently provide more descriptive information about the audio content. As a result, Vela, leveraging the text encoder of LLMs, achieves better retrieval performance compared to using shorter texts in Table 2. This further demonstrates the effectiveness of the MLLM framework in the retrieval domain.

Table 2: Comparison of audio and text retrieval performance on Audiocaps test set and Clotho evaluation datasets.

Method	audio retrieval						text retrieval					
	Clotho			Audiocaps			Clotho			Audiocaps		
	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
Contrastive Learning on audio-text pairs												
Wavcaps	15.8	39.1	50.2	35.5	68.1	78.5	17.1	39.2	51.2	37.1	70.2	81.7
Auto-ACD	15.5	40.4	48.9	33.1	67.8	83.6	14.3	34.2	47.6	32.5	65.8	82.8
AudioSetcaps	12.3	31.2	42.9	39.6	72.5	85.1	11.7	30.8	42.1	37.5	72.8	84.6
LAION CLAP	15.2	39.1	50.9	37.7	80.8	91.2	14.4	37.5	50.1	38.5	73.8	86.4
Microsoft CLAP	15.7	42.9	52.6	29.6	63.4	78.1	16.2	38.9	51.4	18.4	48.2	64.6
Contrastive Learning only on text pairs												
Vela	18.1	43.6	53.1	41.8	82.1	89.3	17.9	40.2	50.6	40.2	75.3	87.6

Table 3: Recall@K Comparison Across Methods on text retrieval under Vela-Long

Method	K=1	K=5	K=10
Wavcaps	10.1	26.1	40.0
Auto-ACD	11.4	27.2	41.9
AudioSetcaps	9.1	26.9	37.8
LAION CLAP	13.8	35.7	49.1
Microsoft CLAP	/	/	/
Vela	16.9	40.4	55.3

Additionally, we believe that future cross-modal retrieval will increasingly focus on conditional tasks with instructions instead of captions [26]. We also conducted experiments using Vela-conditional. Instead of relying on captions, we replaced them with instructions that guide the model to perform retrieval under more fine-grained conditions, such as: "Ignore irrelevant information as human talk and find samples with violin sounds." In contrast, traditional tasks would use a full caption sentence, such as: "In the midst of chaotic traffic noise, a man is speaking loudly while there is a soft violin playing." As shown in the Table 5, nearly all models exhibit a weaker understanding of instruction-based inputs, leading to a performance drop even with fewer samples. In contrast, Vela achieves higher performance due to its stronger understanding and reasoning capability enabled by MLLMs, contributing to the superior performance in our benchmark.

Table 4: Recall@K Comparison Across Methods on text retrieval under Vela-Conditional

Method	K=1	K=5	K=10
Wavcaps	7.1	19.4	29.7
Auto-ACD	7.2	18.5	24.3
AudioSetcaps	6.7	17.9	22.6
LAION CLAP	7.7	19.4	24.6
Microsoft CLAP	10.2	21.0	28.7
Vela	13.4	24.5	33.2

4.3. Ablation Study

Finally, to validate the effectiveness of our paradigm design, we conducted an ablation study, as shown in Table 5. We started with "in one word" prompt design and progressively incorporated in-context learning and contrastive learning separately into Vela. By utilizing these approaches, we demonstrated the gradual reduction of the modality gap [13], proving that MLLMs can map multimodal content into a unified space and support our premise of single-modality training. Additionally, our findings indicate that in-context learning plays a more significant role than contrastive learning. This is because MLLMs alone do not have the capacity to capture the semantic characteristics across different modalities that are crucial for retrieval tasks. Only by designing effective samples can they achieve unified mapping for paired modalities.

Table 5: Recall@K Comparison on Ablation Study. The Gap score reflects the difference between different modalities.

Configurations	K=1	K=10	Gap ↓
Prompt without "one word"	4.69	25.30	67.2
Prompt with "one word"	12.8	41.7	44.1
"One word" prompt with samples	13.6	47.1	28.9
"One word" prompt with training	7.04	34.3	40.6
With all proposals	18.1	53.1	21.9

5. Conclusion

In this work, we propose Vela, an MLLM-based universal multimodal framework capable of representing both acoustic and linguistic inputs. Technically, it employs a prompt-based representation design and an in-context learning approach to unify multimodal representations into a shared embedding space without fine-tuning. Through single-modality training on text pairs, Vela demonstrates strong performance on audio-language retrieval tasks. More importantly, by introducing our new benchmark, we highlight the limitations of existing CLAP-based models in long-text tasks and conditional retrieval tasks with instructions. We also conducted an ablation study to validate the effectiveness of our approach. We hope that our work can provide valuable insights for the development of unified multimodal representations in the future.

6. Acknowledgements

This work was supported by the “Pioneer” and “Leading Goose” R&D Program of Zhejiang under (Grant No. 2025C02110), Public Welfare Research Program of Ningbo under (Grant No. 2024S062), and Yongjiang Talent Project of Ningbo under (Grant No. 2024A-161-G).

7. References

- [1] B. Elizalde, S. Deshmukh, and H. Wang, “Natural language supervision for general-purpose audio representations,” 2023. [Online]. Available: <https://arxiv.org/abs/2309.05767>
- [2] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2023.
- [3] 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.
- [4] J. Bai, H. Liu, M. Wang, D. Shi, W. Wang, M. D. Plumbley, W.-S. Gan, and J. Chen, “Audiosetcaps: An enriched audio-caption dataset using automated generation pipeline with large audio and language models,” *arXiv preprint arXiv:2411.18953*, 2024.
- [5] L. Sun, X. Xu, M. Wu, and W. Xie, “Auto-acd: A large-scale dataset for audio-language representation learning,” in *ACM Multimedia 2024*.
- [6] Z. Xiao, M. Gong, P. Cascante-Bonilla, X. Zhang, J. Wu, and V. Ordonez, “Grounding language models for visual entity recognition,” in *European Conference on Computer Vision*. Springer, 2025, pp. 393–411.
- [7] L. Fan, D. Krishnan, P. Isola, D. Katabi, and Y. Tian, “Improving clip training with language rewrites,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 35 544–35 575, 2023.
- [8] Y. K. Jang, J. Kang, Y. J. Lee, and D. Kim, “Mate: Meet at the embedding—connecting images with long texts,” *arXiv preprint arXiv:2407.09541*, 2024.
- [9] S. Tong, D. Fan, J. Zhu, Y. Xiong, X. Chen, K. Sinha, M. Rabbat, Y. LeCun, S. Xie, and Z. Liu, “Metamorph: Multimodal understanding and generation via instruction tuning,” *arXiv preprint arXiv:2412.14164*, 2024.
- [10] A. Fahim, A. Murphy, and A. Fyshe, “Its not a modality gap: Characterizing and addressing the contrastive gap,” *arXiv preprint arXiv:2405.18570*, 2024.
- [11] T. Jiang, M. Song, Z. Zhang, H. Huang, W. Deng, F. Sun, Q. Zhang, D. Wang, and F. Zhuang, “E5-v: Universal embeddings with multimodal large language models,” *arXiv preprint arXiv:2407.12580*, 2024.
- [12] 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 *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [13] V. W. Liang, Y. Zhang, Y. Kwon, S. Yeung, and J. Y. Zou, “Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 17 612–17 625, 2022.
- [14] H. Tsukagoshi, R. Sasano, and K. Takeda, “Defsent: Sentence embeddings using definition sentences,” *arXiv preprint arXiv:2105.04339*, 2021.
- [15] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [16] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [17] B. Geng, Z. Huan, X. Zhang, Y. He, L. Zhang, F. Yuan, J. Zhou, and L. Mo, “Breaking the length barrier: Llm-enhanced ctr prediction in long textual user behaviors,” in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024, pp. 2311–2315.
- [18] T. Gao, X. Yao, and D. Chen, “Simcse: Simple contrastive learning of sentence embeddings,” *arXiv preprint arXiv:2104.08821*, 2021.
- [19] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: An audio captioning dataset,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 736–740.
- [20] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, “Fsd50k: an open dataset of human-labeled sound events,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 829–852, 2021.
- [21] 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)*, 2019, pp. 119–132.
- [22] Y. Chu, J. Xu, Q. Yang, H. Wei, X. Wei, Z. Guo, Y. Leng, Y. Lv, J. He, J. Lin *et al.*, “Qwen2-audio technical report,” *arXiv preprint arXiv:2407.10759*, 2024.
- [23] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, “Qwen technical report,” *arXiv preprint arXiv:2309.16609*, 2023.
- [24] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [25] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, “Qlora: Efficient finetuning of quantized llms,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [26] X. Li, J. Jin, Y. Zhou, Y. Zhang, P. Zhang, Y. Zhu, and Z. Dou, “From matching to generation: A survey on generative information retrieval,” *arXiv preprint arXiv:2404.14851*, 2024.