

Do Audio-Language Models Understand Linguistic Variations?

Ramaneswaran Selvakumar^{*1}, Sonal Kumar^{*1}, Hemant Kumar Giri^{*2},
Nishit Anand¹, Ashish Seth¹, Sreyan Ghosh¹, Dinesh Manocha¹

¹University of Maryland, College Park, ²NVIDIA, Bangalore
{ramans, sonalkum, nishit, aseth125, sreyang, dmanocha}@umd.edu
hgiri@nvidia.com

Abstract

Open-vocabulary audio language models (ALMs), like Contrastive Language Audio Pretraining (CLAP), represent a promising new paradigm for audio-text retrieval using natural language queries. In this paper, for the first time, we perform controlled experiments on various benchmarks to show that existing ALMs struggle to generalize to linguistic variations in textual queries. To address this issue, we propose RobustCLAP, a novel and compute-efficient technique to learn audio-language representations agnostic to linguistic variations. Specifically, we reformulate the contrastive loss used in CLAP architectures by introducing a multi-view contrastive learning objective, where paraphrases are treated as different views of the same audio scene and use this for training. Our proposed approach improves the text-to-audio retrieval performance of CLAP by 0.8%-13% across benchmarks and enhances robustness to linguistic variation. We make our code publicly available ¹

1 Introduction

As user-generated audio content expands at an unprecedented pace, developing methods to index and search effectively across an ever-growing database becomes crucial. Open-vocabulary audio language models (ALMs) such as CLAP (Elizalde et al., 2023a,b) have emerged as a promising solution to this problem, achieving state-of-the-art (SOTA) results in text-based audio retrieval (Wu* et al., 2023). In a typical setting, a user would use a natural language query to describe an acoustic scene with various audio events and then use it to retrieve audio files that match the query. Natural language offers a powerful and intuitive interface for indexing and searching through audio

¹https://github.com/ramaneswaran/linguistic_robust_clap

^{*}These authors contributed equally to this work.

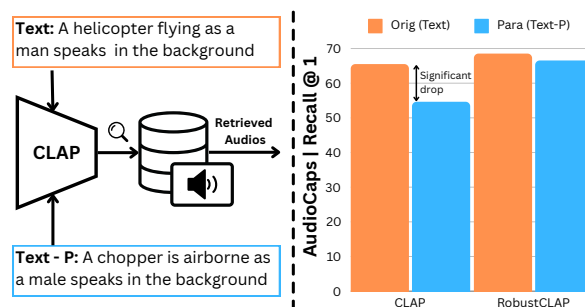


Figure 1: ALMs like CLAP struggle with linguistic variations in queries (**Text**), such as paraphrases (**Text-P**), resulting in a significant drop in retrieval performance. Our method, RobustCLAP, mitigates this issue while improving overall retrieval accuracy.

databases. It allows end-users to describe virtually any concept and provides the creative freedom to use linguistically diverse expressions to describe the scene. However, while humans naturally adapt to such linguistic variations, whether ALMs can generalize to these variations at test time remains to be determined. Our preliminary results suggest that the answer is **no**, and ALMs can observe up to a 16% drop in text-to-audio (T2A) retrieval performance on standard benchmarks with only slight changes in the wording of the text. This limitation can further lead to inconsistent retrieval results across natural language queries with the same intent (see Figure 1 for an example)

Main Contributions: To this end, in this paper, we present two novel contributions:

1. We present the first study to evaluate the robustness of ALMs for T2A retrieval. We construct five new benchmarks (synthetically with human-in-the-loop) to evaluate the performance of CLAP models in T2A retrieval across linguistically varied queries with similar intent. Our evaluation shows a consistent drop in retrieval recall scores (0.1% - 16%) across our benchmarks, highlighting the vul-

nerability to linguistic variation.

2. We propose RobustCLAP, a simple yet effective method to train CLAP-like ALMs that are robust to linguistic variations in input queries. We continually fine-tune pre-trained CLAP using a novel multi-view contrastive objective that gradually aligns the paraphrased captions with original captions and audio. By training on only a fraction of the original pre-training data, our method improves T2A retrieval performance on the original and paraphrased benchmarks by 0.8%-13%, demonstrating increased robustness to linguistic variation while maintaining computational and data efficiency.

2 Related Work

2.1 Retrieval With Query Variations

Although the effects of query variations for text (Zuccon et al., 2016; Voorhees and Harman, 1999) or image retrieval (Kim et al., 2024) have been explored before, there have been few attempts to address this issue in audio retrieval tasks such as Clotho (Drossos et al., 2019) and AudioCaps (Kim et al., 2019). Most prior efforts to improve audio language models (ALMs) have focused either on scaling the models (Wu* et al., 2023) or enhancing their reasoning capabilities (Ghosh et al., 2024b). However, as audio retrieval using ALMs is increasingly being used in tasks like audio captioning and question answering (Kong et al., 2024; Ghosh et al., 2024a), ensuring robustness to linguistic variation is critical to maintaining their effectiveness in real-world applications.

2.2 Synthetic Data For Retrieval

Synthetic data generation has been widely studied in text-based representation learning and information retrieval. InPars (Bonifacio et al., 2022), InParsv2 (Jeronymo et al., 2023) and Promptagator (Dai et al., 2022) generate synthetic queries from unlabelled documents for language encoder training. DINO (Schick and Schütze, 2021) generates synthetic textual similarity pairs for training cross-encoders while Gecko (Lee et al., 2024) extensively uses LLMs to generate synthetic queries and hard-negatives. LARMOR (Khramtsova et al., 2024) uses LLMs to generate synthetic data to adapt textual retrievers to a specific domain. On the other hand, synthetic data for improving

audio-language models (ALMs) is still under explored. Approaches like CompA (Ghosh et al., 2024b) pioneer the use of synthetic data to improve general and compositional representation of ALMs and train their models from scratch. In contrast, our approach adapts any off-the-shelf CLAP model and, with minimal additional training, enhances its robustness to linguistic variations while preserving its pre-trained knowledge and capabilities.

3 Methodology

3.1 Paraphrased Audio Text Retrieval Benchmark

To study the impact of linguistic variation in input queries, we introduce new benchmarks by carefully extending the following five audio-text retrieval benchmarks with their paraphrased captions: 1) AudioCaps (Kim et al., 2019) 2) Clotho (Drossos et al., 2019) 3) DCASE (Lagrange et al., 2022) 4) Audioset Strong Labels (Hershey et al., 2021) and 5) SoundDesc (Koepke et al., 2023).

To obtain the paraphrased captions, we generate new captions such that the vocabulary and the linguistic structure differ while preserving the key concepts and intent. This task requires linguistic expertise and a strong understanding of the concept behind real-world sounds. For instance, accurately differentiating between a bird’s "tweet" and a "chirp" involves recognizing subtle differences in tone and context, which are crucial for maintaining the accuracy and relevance of the paraphrases. On the other hand, Large Language Models (LLMs) have shown remarkable aptitude in natural language understanding and real-world common-sense knowledge. Consequently, we propose using LLMs to generate paraphrased captions in a two-step process: Step 1: We instruct the LLM to generate a paraphrase based on custom human-written ICL examples for each benchmark. Step 2: We instruct the LLM to carefully reason (Wei et al., 2023) whether the paraphrase is accurate and to correct it if required. We detail these steps and give examples below.

Paraphrase Generation: We instruct the LLM to generate an initial paraphrase (Text-P') of the original caption, such that we describe the acoustic events using varied vocabulary and sentence structures while preserving the original meaning. We give an example below:

Benchmark →	AudioCaps		Clotho		Audioset SL		SoundDesc		DCASE	
Model ↓	TEST	TEST-P	TEST	TEST-P	TEST	TEST-P	TEST	TEST-P	TEST	TEST-P
ML-ACT	35.53	34.87	27.54	23.90	21.52	17.91	08.72	06.06	10.12	08.77
MSCLAP-22	84.74	84.63	86.74	43.94	27.73	23.72	14.33	11.87	39.91	30.99
MSCLAP-23	80.77	77.63	51.14	42.12	55.12	39.15	38.27	24.89	47.84	39.21
CompA	97.17	96.23	51.28	42.49	43.03	40.24	33.32	23.56	49.54	39.51
LAION-CLAP	97.80	95.92	52.03	43.98	46.91	41.94	24.62	18.09	44.73	37.81
RobustCLAP	98.64	98.22	57.27	53.47	57.44	53.64	25.48	21.54	54.66	50.35

Table 1: Recall@10 scores (higher is better) for text-to-audio retrieval on the original test set (TEST) and paraphrased test set (TEST-P). All ALMs show a consistent, significant drop in performance on TEST-P. RobustCLAP not only improves overall retrieval performance on TEST but also mitigates the drop in TEST-P. The best scores for each benchmark are highlighted in bold.

Sample caption, paraphrase and corrected paraphrase

Text: A person talking which later imitates a couple of meow sounds.

Text-P’: An individual speaks, subsequently mimicking some cat cries.

Text-P: An individual speaks, subsequently mimicking some cat meows.

Paraphrase Correction: It is crucial that the paraphrased caption accurately conveys the nuances of the original acoustic events. To ensure this, we instruct the LLM to evaluate the paraphrase for both accuracy and specificity, making corrections where necessary. For example, in the paraphrase above, the LLM identified that "cat cries" typically implies a distressed or loud sound, which may not align with the softer or more playful tone often associated with "meows". As a result, the LLM corrects the paraphrase to use "meows" ensuring it better reflects the intended meaning.

For these tasks, we employ LLaMA-3-70B (AI@Meta, 2024) with in-context learning examples crafted by humans. Following insights from (Shen et al., 2022) we conducted a qualitative study to evaluate the quality of paraphrase generation and correction. **Paraphrase Quality:** Human evaluators rated 100 random paraphrases on a 1-5 Likert scale, with an average score of 4.89. **Paraphrase Correction:** For 50 paraphrases and their corrected versions, evaluators preferred the corrected captions 98% of the time. We refer readers to Appendix F, B.2 for additional details on the implementation and evaluation.

3.2 Improving CLAP With Paraphrases

To improve the robustness of audio retrieval to linguistic variation, we propose further training of a pre-trained CLAP model using paraphrases of the training data. Specifically, we reformulate the standard CLAP loss as a multi-view contrastive loss that uses two levels of paraphrases as two views to gradually align the text representations with their paraphrased counterparts. At the first level (T^{P_1}), only the linguistic structure is modified while maintaining the same vocabulary. At the second level (T^{P_2}), both the vocabulary and structure are altered. By presenting the model with progressively more complex paraphrases at each training step, we enable it to learn a more generalizable mapping between semantic content and its diverse linguistic expressions. This enhances the model’s robustness to linguistic variations in real-world queries.

A CLAP model takes in an input of an audio-text pair (A, T) and comprises i) audio-encoder $e_A = E(A)$ and (ii) text encoder $e_T = E(T)$. In this notation, we compute similarity score as:

$$S(T, I) = \exp\left(\frac{1}{\tau} \cdot \frac{e_T^\top e_A}{\|e_T\| \|e_A\|}\right), \quad (1)$$

where τ is a learned temperature parameter.

Contrastive Loss For a generated paraphrase $T_i^{P_k}, k \in \{1, 2\}$ produced from the text T_i corresponding to audio A_i , we compute the contrastive loss L^{P_k} as a combination of the following two losses:

$$L_{P_k}^T = \sum_i \left[-\log \left(\frac{S(T_i^{P_k}, T_i)}{\sum_j S(T_i^{P_k}, T_j)} \right) \right] \quad (2)$$

$$L_{P_k}^A = \sum_i \left[-\log \left(\frac{S(T_i^{P_k}, A_i)}{\sum_j S(T_i^{P_k}, A_j)} \right) \right] \quad (3)$$

Overall, the final loss is computed as follows:

$$L_{final} = L_{clap} + L^{P_1} + L^{P_2} \quad (4)$$

Here, L_{clap} is the original CLIP-loss (Radford et al., 2021) used to train the CLAP models. This is necessary to prevent the CLAP model from forgetting its knowledge acquired during pretraining.

4 Experimental Setup

Training Dataset: We train our model on a combination of AudioCaps (Kim et al., 2019) and Clotho (Drossos et al., 2019), which we augment with our two levels of paraphrased captions.

Evaluation Dataset: For T-A retrieval, we adopt the evaluation setup from previous work (Koepke et al., 2023) and employ AudioCaps, Clotho, Audioset SL (Hershey et al., 2021), SoundDesc (Koepke et al., 2023) and DCASE (Lagrange et al., 2022). We evaluate for Recall@10.

Baselines: For baselines we use ML-ACT (Mei et al., 2022), MSCLAP-22 (Elizalde et al., 2023a), MSCLAP-23 (Elizalde et al., 2023b), CompA (Ghosh et al., 2024b) and LAION-CLAP (Wu* et al., 2023). We use LAION-CLAP as the base model for RobustCLAP.

5 Results And Analysis

Quantitative Analysis: Table 1 shows that current ALMs struggle with linguistic variations, as evidenced by a significant drop (0.1%-16%) in recall scores for paraphrased captions compared to the original captions. In contrast, RobustCLAP not only 1) improves recall scores on the original benchmarks by 0.8% to 13% compared to its base model but also 2) mitigates the performance drop on the paraphrased benchmarks, improving scores by 2% to 12% compared to the respective second best-performing model. CompA and MSCLAP-23, being trained on SoundDesc, perform better on that dataset. However, they show a significant 10-14% drop on the paraphrased SoundDesc benchmark, illustrating that fine-tuning can worsen the issue. We evaluate RobustCLAP on zero-shot audio classification tasks using ESC-50 (Piczak, 2015) and FSD50K (Fonseca et al., 2022). CLAP gets a mAP@10 score of 94.25 and 52.20, while RobustCLAP gets 94.07 and 52.81, respectively,

on ESC-50 and FSD-50K. We observe a negligible drop in performance, which indicates that prior knowledge is retained. RobustCLAP also outperforms ALMs on paraphrased audio-to-text retrieval, these results are indicated in Table 7.

Qualitative Analysis: We conducted a qualitative experiment to assess how often CLAP retrieves incorrect audio compared to RobustCLAP. We sampled 100 instances where CLAP failed to retrieve the correct audio for a paraphrased query, while RobustCLAP succeeded. We then asked human evaluators to listen to the audio retrieved by CLAP and judge whether they matched the query. The results showed that in 97% of cases, the retrieved audios were indeed incorrect, while only 3% were correct. The latter result highlights a challenge inherent in retrieval benchmarks like AudioCaps and Clotho, where a small set of audio files may contain the same acoustic events, mainly when only one or two events are present. Moreover, we observed the following three common mistake patterns. First, CLAP often prioritizes sound events mentioned directly in the query, showing a spurious correlation to non-paraphrased sound events. Second, while the model captures the dominant context or setting of the scene, it frequently lacks precision in identifying all the sound events mentioned in the query. Finally, CLAP fails to recognize attributes or modifiers of a sound event.

Impact Of Sound Event And Attributes: In an acoustic scene, such as the "steady humming of an engine," the sound event refers to the sound and entity producing the sound (e.g., the "humming of an engine"); sound attributes describe its qualities (e.g., "steady"). We study how paraphrasing these elements affects retrieval performance by instructing the LLM to replace specifically the event and attributes with synonyms while maintaining the original linguistic structure. In Table 2 we observe that paraphrasing sound attributes leads to a 3.8% drop in Recall@1, while RobustCLAP significantly reduces this decline to just 0.4%. However, it is important to note that only 20% of the samples contain sound attributes, which limits the overall effect of this variation. Paraphrasing sound sources, on the other hand, has a much more significant impact, with recall dropping by as much as 15%. RobustCLAP mitigates this effect substantially, reducing the performance drop to 3%.

Dataset	Model	
	CLAP	RobustCLAP
AudioCaps	65.51	68.54
+ Sound attributes mod.	61.96	68.12
+ Sound events mod.	50.24	65.48

Table 2: We paraphrase sound attributes (row 1) and sound events (row 2), keeping the linguistic structure fixed, to study their impact on R@1 scores. Sound attributes contribute to the drop, while sound events have a greater impact. RobustCLAP mitigates the effects of these paraphrases.

6 Conclusion

This paper shows that current audio language models lack robustness to linguistic variation in natural language inputs. To demonstrate this phenomenon, we extend several audio-text retrieval benchmarks with paraphrased captions generated through a two-step LLM-based process. To address this issue, we propose a simple mitigation strategy, training CLAP models with a multi-view contrastive loss on a small set of paraphrased data. The resulting model, RobustCLAP, improves retrieval recall scores on the original benchmarks and their paraphrased versions while retaining its prior pre-trained knowledge. We hope our work fuels further studies into improving the robustness of audio-language models.

7 Acknowledgements

This project is supported in part by NSF#1910940.

8 Limitations and Future Work

As part of future work, we would like to address the following limitations of RobustCLAP:

- We utilize LLM to generate the paraphrases for training and testing. Even though we use diverse human-written in-context examples and a correction mechanism, some paraphrases might not be exactly accurate due to hallucinations by the LLM.
- We have primarily experimented with diverse audio benchmarks, in future this work can be extended to related domains like music retrieval and speech retrieval.
- We use relatively shorter audio segments, in future this work can be extended to long audio.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. [Inpars: Data augmentation for information retrieval using large language models](#). *Preprint*, arXiv:2202.05144.
- Zhuyun Dai, Vincent Y. Zhao, Ji Ma, Yi Luan, Jianmo Ni, Jing Lu, Anton Bakalov, Kelvin Guu, Keith B. Hall, and Ming-Wei Chang. 2022. [Promptagator: Few-shot dense retrieval from 8 examples](#). *Preprint*, arXiv:2209.11755.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen. 2019. [Clotho: An audio captioning dataset](#). *Preprint*, arXiv:1910.09387.
- Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. 2023a. [Clap learning audio concepts from natural language supervision](#). In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Benjamin Elizalde, Soham Deshmukh, and Huaming Wang. 2023b. [Natural language supervision for general-purpose audio representations](#). *Preprint*, arXiv:2309.05767.
- Eduardo Fonseca, Xavier Favory, Jordi Pons, Frederic Font, and Xavier Serra. 2022. [Fsd50k: An open dataset of human-labeled sound events](#). *Preprint*, arXiv:2010.00475.
- Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. [Audio set: An ontology and human-labeled dataset for audio events](#). In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780.
- Sreyan Ghosh, Sonal Kumar, Chandra Kiran Reddy Evuru, Ramani Duraiswami, and Dinesh Manocha. 2024a. [Recap: Retrieval-augmented audio captioning](#). *Preprint*, arXiv:2309.09836.
- Sreyan Ghosh, Ashish Seth, Sonal Kumar, Utkarsh Tyagi, Chandra Kiran Reddy Evuru, Rameswarar S, S Sakshi, Oriol Nieto, Ramani Duraiswami, and Dinesh Manocha. 2024b. [Compa: Addressing the gap in compositional reasoning in audio-language models](#). In *The Twelfth International Conference on Learning Representations*.

- Shawn Hershey, Daniel P W Ellis, Eduardo Fonseca, Aren Jansen, Caroline Liu, R Channing Moore, and Manoj Plakal. 2021. [The benefit of temporally-strong labels in audio event classification](#). *Preprint*, arXiv:2105.07031.
- Vitor Jeronymo, Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, Roberto Lotufo, Jakub Zavrel, and Rodrigo Nogueira. 2023. [Inpars-v2: Large language models as efficient dataset generators for information retrieval](#). *Preprint*, arXiv:2301.01820.
- Ekaterina Khramtsova, Shengyao Zhuang, Mahsa Bakhtashmotlagh, and Guido Zuccon. 2024. [Leveraging llms for unsupervised dense retriever ranking](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2024, page 1307–1317. ACM.
- Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. 2019. [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)*, pages 119–132, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hyunjae Kim, Seunghyun Yoon, Trung Bui, Handong Zhao, Quan Tran, Franck Dernoncourt, and Jaewoo Kang. 2024. [Fine-tuning clip text encoders with two-step paraphrasing](#). *Preprint*, arXiv:2402.15120.
- A. Sophia Koepke, Andreea-Maria Oncescu, João F. Henriques, Zeynep Akata, and Samuel Albanie. 2023. [Audio retrieval with natural language queries: A benchmark study](#). *IEEE Transactions on Multimedia*, 25:2675–2685.
- Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, and Bryan Catanzaro. 2024. [Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities](#). *Preprint*, arXiv:2402.01831.
- Mathieu Lagrange, Annamaria Mesaros, Thomas Pellegrini, Romain Serizel, Gaël Richard, and Dan Stowell. 2022. *Proceedings of the 7th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE 2022)*. Tampere University, Nancy, France.
- Jinhyuk Lee, Zhuyun Dai, Xiaoqi Ren, Blair Chen, Daniel Cer, Jeremy R. Cole, Kai Hui, Michael Boratko, Rajvi Kapadia, Wen Ding, Yi Luan, Sai Meher Karthik Duddu, Gustavo Hernandez Abrego, Weiqiang Shi, Nithi Gupta, Aditya Kusupati, Praateek Jain, Siddhartha Reddy Jonnalagadda, Ming-Wei Chang, and Iftexhar Naim. 2024. [Gecko: Versatile text embeddings distilled from large language models](#). *Preprint*, arXiv:2403.20327.
- Xinhao Mei, Xubo Liu, Jianyuan Sun, Mark D. Plumbley, and Wenwu Wang. 2022. On metric learning for audio-text cross-modal retrieval. *arXiv preprint arXiv:2203.15537*.
- Karol J. Piczak. 2015. [Esc: Dataset for environmental sound classification](#). In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM ’15, page 1015–1018, New York, NY, USA. Association for Computing Machinery.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- Timo Schick and Hinrich Schütze. 2021. [Generating datasets with pretrained language models](#). *Preprint*, arXiv:2104.07540.
- Lingfeng Shen, Lemao Liu, Haiyun Jiang, and Shuming Shi. 2022. [On the evaluation metrics for paraphrase generation](#). *Preprint*, arXiv:2202.08479.
- E. Voorhees and D. Harman, editors. 1999. *The Eighth Text REtrieval Conference (TREC-8)*. National Institute of Standards and Technology (NIST). NIST Special Publication 500-246.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Yusong Wu*, Ke Chen*, Tianyu Zhang*, Yuchen Hui*, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. 2023. 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*.
- Guido Zuccon, Joao Palotti, and Allan Hanbury. 2016. [Query variations and their effect on comparing information retrieval systems](#). In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, CIKM ’16, page 691–700, New York, NY, USA. Association for Computing Machinery.

A Appendix

In the appendix we provide:

- Section B: Dataset Details
- Section C: Model Details
- Section D: Additional Results
- Section E: Additional Implementation Details
- Section F: Prompts Used

B Dataset Details

In this section we describe in detail the benchmarks that we used for evaluation. In Section B.1 we describe in detail the datasets that we used. Detailed information about samples present in these are given in Table 3. In Section F.1 we further detail our paraphrase generation and correction mechanism prompts.

B.1 Benchmark Datasets

AudioCaps: The Audioscapes (Kim et al., 2019) dataset is a large-scale captioning dataset developed by google. It has 46K audio clips with 10 sec duration sourced from Audioset (Gemmeke et al., 2017) along with textual descriptions written by human annotators. They give detailed descriptions of the audio, highlighting specific sound events, their sources, and the context.

Clotho: The Clotho (Drossos et al., 2019) dataset is an audio captioning dataset with sound clips (15-30 seconds) sourced from Freesound. People have captioned them, describing environments, music, and activities. Along with AudioCaps it is one of the most widely used audio-retrieval benchmarks.

Audioset SL: Audioset SL (Strong Labels) (Hershey et al., 2021) is a significant component of Google’s Audioset project (Gemmeke et al., 2017), which involves annotating over 2 million 10-second audio clips from YouTube with specific labels. These labels include sounds like "dog barking," "car engine," or "crowd cheering." Although it does not provide full captions, the extensive sound event labeling in Audioset SL provides a rich source for generating artificial captions. We use these temporally strong audio event labels and instruct an LLM to generate a natural language audio caption.

SoundDesc: The SoundDesc (Koepke et al., 2023) is a dataset that provides detailed descriptions for diverse sound clips. It includes everyday sounds, natural environments, and specific events. Each clip is paired with a detailed description capturing the sound’s essence, source, and context.

DCASE: The Detection and Classification of Acoustic Scenes and Events (DCASE) (Lagrange et al., 2022) dataset is a comprehensive collection of audio recordings. It includes various environments like streets, parks, and indoor settings, each annotated with specific sound event or acoustic scene labels. This dataset is crucial for the DCASE community challenges, fostering advancements in the field. It’s essential for developing and evaluat-

ing models that recognize and classify sounds in complex environments.

Benchmark	# Audio Samples		# Captions	
	Train	Test	Train	Test
AudioCaps	49,275	958	49,275	4,790
Clotho	3,840	1,045	19,200	5,225
Audioset SL	N/A	1,471	N/A	1,471
SoundDesc	23,085	3,250	23,085	3,250
DCASE	N/A	997	N/A	997

Table 3: Overview of the datasets used, including the number of audio samples and captions available for both training and testing.

Score	Guideline
1	Completely different meanings with no semantic overlap.
2	Paraphrased caption shares some similar words but convey different overall meanings.
3	Common topic but differ in details or emphasis
4	Largely similar meanings with minor variation in detail
5	The core information being conveyed is same

Table 4: The likert scale guideline used for paraphrase quality assessment.

B.2 Benchmark Paraphrase Evaluation

We conduct a qualitative analysis to study both the final paraphrases as well as the performance of paraphrase correction. The volunteers for this study were computer science MS and PhD students.

Paraphrase Generation: In this experiment, we sample 100 random paraphrases and ask human paraphrases and ask human evaluators to listen to the audio and read the original caption and rate the paraphrase on a LIKERT scale of 1-5. Overall, we obtained an average score of 4.89 indicating that our pipeline of generation and subsequent correction if required is able to generate good paraphrases. The Likert scale guidelines are presented in Table 4.

Paraphrase Correction: Following (Piczak, 2015) we conduct an experiment to understand if users prefer the corrected paraphrases as opposed to original paraphrases. We sample 50 random paraphrases and their final corrected versions. We then ask human evaluators to choose one caption which describes the audio better. The corrected versions of the paraphrases were preferred 98% of the time.

C Model Details

C.1 Baseline Details

ML-ACT (Mei et al., 2022). This model uses a PANN model trained on Audioset (Gemmeke et al.,

2017) and a BERT (Devlin et al., 2019) model and employs the NT-Xent loss adapted from self-supervised learning.

LAION-CLAP (Wu* et al., 2023). This is a contrastive language-audio pretraining (CLAP) model from LAION-AI trained on LAION-Audio-630K (Wu* et al., 2023), a large collection of 633,526 audio-text pairs from different data sources. To improve the model’s ability to handle audio inputs of variable lengths and boost overall performance, it integrates a feature fusion mechanism and keyword-to-caption augmentation. This enables the model to effectively align and process both audio and text data for enhanced learning.

LAION-CLAP Music (Wu* et al., 2023). This is a music-specific version of the LAION-CLAP model. This version is trained both on audio and music, with the LAION-Audio-630K dataset contributing a major portion of its training data. The details of the music-text data being used for training are not specified.

MS-CLAP 22 (Elizalde et al., 2023a). This is a contrastive language-audio pretraining (CLAP) model from Microsoft. This version is trained on 128k audio and text pairs.

MS-CLAP 23 (Elizalde et al., 2023b). This is a follow-up to the MS-CLAP 22, from Microsoft. This version of CLAP uses two innovative encoders and is trained on massive 4.6M audio-text pairs. To learn audio representations, the authors trained an audio encoder on 22 audio tasks instead of the standard training of sound event classification. To learn language representations, they trained an autoregressive decoder-only model instead of the standard encoder-only models.

CompA (Ghosh et al., 2024b). This is a CLAP model that is trained specifically to enhance its compositional reasoning abilities. The authors introduce improvements to contrastive training by incorporating composition-aware hard negatives, allowing for more precise and focused training. Additionally, they propose a modular contrastive loss designed to help the model learn fine-grained compositional understanding.

D Additional Results

D.1 Performance On Zero-Shot Audio Classification

We evaluate CLAP and RobustCLAP on zero-shot audio classification task (ZSAC) on the ESC-50 (Piczak, 2015) and FSD-50K (Fonseca et al.,

Model	ESC-50	FSD-50K
CLAP	94.25	53.20
RobustCLAP	94.07	52.81

Table 5: Zero-shot audio classification results in terms of mAP@10. We observe there is negligible performance decrease for RobustCLAP compared to CLAP

2022) datasets. CLAP gets a mAP@10 score of 94.25 and 52.20, while RobustCLAP gets 94.07 and 52.81, respectively, on ESC-50 and FSD-50K. We observe negligible performance decreases, demonstrating that our approach does not lead to catastrophic forgetting of previously learned knowledge. Fine-tuning the CLAP model on AudioCaps and Clotho enables it to capture the descriptive features (of individual acoustic events), which are beneficial for audio retrieval based on rich natural language descriptions. However, it doesn’t necessarily help CLAP learn the discriminative features necessary for zero-shot audio classification.

D.2 Error Analysis

We conduct a manual study of CLAP and RobustCLAP model performance. We sample 100 instances, where for a given paraphrased query, CLAP is not able to correctly retrieve audio whereas RobustCLAP is able to retrieve the audio correctly. We asked human evaluators to listen to the retrieved audio and score whether the audio retrieved by CLAP was correct. The main findings are

- In 97% of the cases, the audio retrieved by CLAP were actually wrong (We highlight some common mistake patterns later in our discussion)
- In 3% of the cases, the audio retrieved was correct according to the given query. This is a challenge inherent in retrieval benchmarks like AudioCaps, Clotho, where a small number of audio files might contain the exact same acoustic events, especially when only one or two events are present.

Overall, we were able to verify that CLAP was indeed retrieving the incorrect audio files, whereas RobustCLAP was able to retrieve the correct audio. We noticed some common mistake patterns that we highlight below.

1) Spurious correlation to non-paraphrased sound events: CLAP tends to prioritize sound events that are directly mentioned in the query without any paraphrasing. In this case, audios which are retrieved may contain an exact sound event such as “wind noise” or “background music” but overall have a completely different meaning compared to the given query. In the examples below the events “background music” and “gurgling and bubbling noises” are spuriously correlated during retrieval

Error Example

Paraphrased Query: A man’s voice is heard alongside background music and TV noise, then interrupted by kids’ giggles and chatter.

Retrieved Audio Description: A kid is speaking while rattling and tapping sounds are heard amidst the background music, with occasional breathing sounds and mechanisms in the background.

Paraphrased Query: Continuous music is accompanied by two instances of gurgling and bubbling noises.

Retrieved Audio Description: Water is poured, splashing and splattering, followed by gurgling and bubbling sounds, with a person breathing in the background towards the end.

2) Captures the dominant context but lacks precision: In this case the model understand the dominant context or the setting of the scene, but fails to precisely capture all the sound events in the query.

Error Example

Paraphrased Query: In an urban environment, a man talks as machines and vehicles hum in the background, punctuated by a final thud.

Retrieved Audio Description: A man is speaking amidst urban traffic noise, accompanied by birds chirping and wind blowing.

Explanation: CLAP is able to capture the context of a man speaking in urban setting, but does not capture the vehicle hum, but includes bird and wind sound.

3) Does not capture sound attributes: In this

case, the CLAP model fails to accurately capture the attributes that act as modifiers to a sound event.

Error Example

Paraphrased Query: A serene ambiance is created by an orchestra of bird melodies, punctuated by turkey calls and faint vehicle hums.

Retrieved Audio Description: A bird is singing along with occasional squawks amidst a constant vehicle noise.

Explanation: While CLAP model is able to capture most of the events, listening to the audio shows that a faint vehicle noise (which is in the background and muted) is a big contrast from a constant vehicle noise (which is in the foreground and loud)

D.3 Statistical Significance Test

We use a bootstrapping method to collect recall metrics for both CLAP and RobustCLAP. This involves repeatedly sampling with replacement from the test set and then computing the recall for each resampled set. These sets of recall values are used to perform a t-test, and we conclude that the improvement of RobustCLAP over CLAP is statistically significant.

E Additional Implementation Details

Model	# Params	Link
ML-ACT	140M	https://github.com/XinhaoMei/audio-text_retrieval
MSCLAP22	196M	https://github.com/microsoft/CLAP
MSCLAP23	159M	https://github.com/microsoft/CLAP
CompA	300M	https://github.com/Sreyan88/CompA
LAION-CLAP	158M	https://github.com/LAION-AI/CLAP/

Table 6: ALMs used in our project and their size (in millions of parameters). We use official implementations of these models

E.1 Model Parameters

The ALMs that consists of an audio-encoder and a BERT like text encoder. Typically these models are under 300M parameters, refer to Table 6 for more details. We use a Llama3-70B which consists of 70B parameters to generate paraphrases for training and validation.

E.2 Compute Infrastructure

RobustCLAP is trained on four NVIDIA A100 GPUs and takes around 2 hours to converge. Infer-

ence only requires 1 A100 GPU. To perform inference on the LLama3-70B model we use 4 NVIDIA A100 GPUs.

E.3 Implementation Software And Packages:

For all the ALMs that we implement we use their original GitHub repository. We provide links to these in Table 6. We build RobustCLAP on top of LAION-CLAP repository and use their base models. To perform accelerated inference on Llama3-70B we use vllm²

E.4 Potential Risks:

Our approach involves using an LLM to generate paraphrases for training and evaluation. While LLMs can sometimes hallucinate or produce incorrect or toxic outputs, we mitigated these risks through a qualitative analysis of the generated paraphrases. In our analysis, we observed no toxic outputs, and the paraphrases were of consistently high quality.

F Prompt Details

²<https://github.com/vllm-project/vllm>

F.1 Paraphrase Generation And Correction Prompts

Paraphrase Generation Prompt

<s>[INST] I will provide you with an audio caption of an audio. Paraphrase the caption while accurately describing the nuances and technical terms. Here are some input-output examples:

Input Caption: Gunfire, followed by a click and shattering glass.

Paraphrase Caption: Shots ring out, then a click and glass breaks into fragments.

Input Caption: Pots clatter as water flows from a turned-on faucet.

Paraphrase Caption: Utensils clatter while liquid streams from an open tap.

Input Caption: A man and woman laugh, followed by a man shouting and a woman joining in with childlike giggles.

Paraphrase Caption: A couple chuckles, then a male yells, and a female responds with youthful giggles.

Input Caption: A woman delivers a formal address.

Paraphrase Caption: A female presents an official speech.

Input Caption: High-pitched snoring echoes repeatedly.

Paraphrase Caption: Sharp snores resound over and over.

Here is the Input Caption: Constant rattling noise and sharp vibrations [/INST]

Prompt Paraphrase Correction

<s>[INST] I will provide you with an audio caption of an audio and its paraphrase. I want you to tell me if the caption is accurately paraphrased especially check if the paraphrased sound events convey the same nuance. Suggest if correction is required and provide corrected paraphrase by give your reasoning. Here are some input-output examples:

Input Caption: :A man talking as metal clanks together followed by footsteps on grass while insects buzz in the background.

Paraphrase Caption: A male speaks as metallic objects collide, succeeded by the sound of steps on a lawn amidst a gentle humming of bugs.

Correction: foo

Corrected Paraphrase Caption: A male speaks as metallic objects clatter, succeeded by the sound of steps on a lawn amidst a gentle humming of bugs

Reasoning: The term "collide" broadly implies contact but lacks the specific metallic sound detail conveyed by "clank." Using "metallic objects chime" or "metallic clatter" would better capture the resonant sound characteristic of metal without reusing the original word.

Input Caption: Men speak as someone snores.

Paraphrase Caption: Males converse amidst a person's heavy breathing.

Correction: foo

Corrected Paraphrase Caption: Males converse amidst a person's disruptive nasal noises.

Reasoning: "Heavy breathing" generally suggests deep breaths and lacks the unique, disruptive nature associated with snoring. A phrase like "disruptive nasal noises" more accurately conveys the irritating and unmistakable sounds of snoring,

highlighting its potential to interrupt or disturb. This emphasizes not only the sound but also the common reaction to it.

Input Caption: An ambulance travels with the siren blaring loudly and moves through traffic.

Paraphrase Caption: A rescue vehicle speeds along with its alarm wailing at full volume and navigates through congested roads.

Correction: bar

Corrected Paraphrase Caption: Not required

Reasoning: This is accurate.

Input Caption: An idle vehicle engine running.

Paraphrase Caption: A stationary car motor hums continuously.

Correction: bar.

Corrected Paraphrase Caption: Not required.

Reasoning: This is accurate.

Input Caption: A toy helicopter flying followed by wind blowing into a microphone.

Paraphrase Caption: A miniature aircraft whirs as it moves through the air, then a gust of air hits the recording device.

Correction: foo

Corrected Paraphrase Caption: A miniature aircraft whirs as it moves through the air, followed by wind rushing continuously against the recording device.

Reasoning: The phrase "wind blowing into a microphone" suggests a continuous or ambient wind noise, which is not precisely captured by "a gust of air hits the recording device." To better reflect the ongoing nature of the sound, the paraphrase could use "as wind rushes against the recording device" or as 'wind continuously interacts with the recording device.'

Input Caption: A man and a woman talking as paper crinkles.

Paraphrase Caption: A male and female converse amidst the rustling of documents.

Correction: bar

Corrected Paraphrase Caption: Not required

Reasoning: This is accurate.

Input Caption: White noise and then birds chirping.

Paraphrase Caption: A gentle hum precedes the sweet sounds of avian creatures.

Correction: foo

Corrected Paraphrase Caption: A continuous static hum precedes the crisp chirping of birds.

Reasoning: The term 'gentle hum' suggests a softer, more subdued sound compared to 'white noise,' which generally implies a more consistent, static-like background noise. To maintain the specific quality of 'white noise,' a more precise description like 'continuous static hum' could be used instead of 'gentle hum.' Additionally, 'the sweet sounds of avian creatures' does not capture the distinctive, rhythmic chirping of birds. A term like 'crisp chirping' would more accurately reflect the clear, melodic nature of bird calls.

Input Caption: Music is playing.

Paraphrase Caption: A melody fills the air.

Correction: [/INST]

F.2 Paraphrase Samples

AudioCaps

TEXT: People are talking while a motor vehicle engine is revving.

TEXT-P: A group of individuals engage in conversation amidst a car engine's loud, rapid revving.

TEXT: A lady laughing while a baby cries, then the lady speaks and a couple men also talk as well

TEXT-P: A female bursts into laughter as an infant wails, then she utters words and a pair of males join in the conversation too.

TEXT: Clicks followed by gunshots and breathing then some speaking

TEXT-P: Series of clicks precede gunfire, labored breathing, and subsequent conversation.

TEXT: Metal clanking followed by steam hissing as a truck engine is running then accelerating

TEXT-P: Clattering metal sounds precede a continuous hissing of steam as a lorry's motor hums and gains speed.

TEXT: A goat bleating with people speaking

TEXT-P: A goat lets out a loud, nasal cry while individuals converse.

Clotho

TEXT: Water goes down a drain pipe while water is dripping.

TEXT-P: Liquid flows down a drainage tube as droplets fall.

TEXT: The ripping of paper occurs at evenly spaced intervals.

TEXT-P: The tearing of a document happens at regular time gaps.

TEXT: Metal sliding together such as swords or knives.

TEXT-P: Metallic blades scraping against each other, similar to clashing swords.

TEXT: Someone walking slowly, their feet are crunching leaves.

TEXT-P: A person strolls at a slow pace, their footsteps crushing foliage.

TEXT: A man and woman are talking among themselves while others chat in the background.

TEXT-P: A gentleman and lady converse privately amidst murmurs of surrounding discussions.

Audioset SL

TEXT: A camera shutter is snapped twice during an ongoing music session.

TEXT-P: A camera shutter clicks twice, punctuating the ongoing musical performance.

TEXT: A vehicle is moving through an urban area filled with traffic noise, accompanied by a rooster's crowing and various bird vocalizations.

TEXT-P: A car navigates through a bustling cityscape with constant traffic din, interspersed with a rooster's loud, shrill crowing and varied bird vocalizations.

TEXT: Music plays while occasional mechanisms and impact sounds are heard,

including thuds and a ticking sound, with additional sound effects.

TEXT-P: Music plays alongside intermittent mechanical noises, occasional thuds, and a steady ticking, accompanied by additional sound effects.

TEXT: A vehicle is accelerating in the midst of a noisy crowd and hubbub with people talking in the background.

TEXT-P: Amidst a chaotic and loud crowd with murmurs of conversation, a vehicle rapidly gains speed.

TEXT: A man speaks in a small room filled with mechanisms, where rodents are scurrying around.

TEXT-P: A male voice is audible amidst machinery sounds and rodents scurrying around in a confined space.

SoundDesc

TEXT: A monkey makes close-up snake alarm calls with birds in the background.

TEXT-P: A monkey's loud, close-up warning cries mix with bird sounds.

TEXT: Two seals challenge each other with close-up calls and snorts, accompanied by surf.

TEXT-P: Two seals get up close and personal, growling and snorting at each other.

TEXT: Chaffinches, crossbills, and great tits sing amidst the rustling of trees in high wind.

TEXT-P: Birds like chaffinches and crossbills belt out their tunes as the trees creak in the gusty breeze.

TEXT: A vintage car approaches, stops, and switches off.

TEXT-P: Old-school wheels roll up, come to a stop, and kill the engine.

TEXT: A lesser black-backed gull vocalizes closely, then attacks a juvenile, amidst herring gulls.

TEXT-P: A lesser black-backed gull squawks loudly, then swoops in on a young bird, surrounded by herring gulls.

DCASE

TEXT: A continuous chirp while birds chatter quietly in the background and then a meow from a cat.

TEXT-P: Birds chat softly in the background as a steady chirp flows, interrupted by a cat's meow.

TEXT: A truck drives by while a woman speaks in the background.

TEXT-P: A woman chats away as a truck zooms past in the distance.

TEXT: A train is coming closer and closer, then passes.

TEXT-P: A locomotive approaches, getting louder, then zooms by.

TEXT: Continuous 8-bit arcade game sounds that are building in pitch.

TEXT-P: Retro arcade sounds amp up, getting higher pitched

TEXT: A group of girls laughing harder and louder as time goes by.

TEXT-P: Girls' giggles escalate to uncontrollable laughter over time.

Retrieval Type →		Text-to-Audio Retrieval				Audio-to-Text Retrieval			
Benchmark	Model	R@1 ↑		R@10 ↑		R@1 ↑		R@10 ↑	
		TEST	TEST-P	TEST	TEST-P	TEST	TEST-P	TEST	TEST-P
AudioCaps	ML-ACT	08.36	07.92	35.53	34.87	07.97	06.42	29.17	26.14
	MSCLAP-22	39.18	36.99	84.74	84.63	33.33	16.50	79.41	59.35
	MSCLAP-23	37.30	24.24	80.77	77.63	28.42	22.57	77.84	68.44
	CompA	67.81	58.72	97.17	96.23	46.02	31.17	88.59	78.97
	LAION-CLAP	65.51	54.64	97.80	95.92	43.36	32.60	89.86	80.14
	RobustCLAP	68.54	66.35	98.64	98.22	45.76	40.96	89.34	86.31
Clotho	ML-ACT	12.87	11.42	27.54	23.90	13.20	12.03	52.71	48.87
	MSCLAP-22	36.19	29.78	86.74	43.94	19.76	12.24	51.89	45.93
	MSCLAP-23	37.03	30.47	51.14	42.12	22.87	16.26	61.53	51.19
	CompA	36.39	29.11	51.28	42.49	17.14	11.97	53.44	44.34
	LAION-CLAP	36.75	32.54	52.03	43.98	37.03	30.72	81.91	74.83
	RobustCLAP	39.43	38.66	57.27	53.48	39.43	37.32	82.49	82.30
Audioset SL	ML-ACT	04.31	04.01	21.52	17.91	05.54	03.77	22.02	18.91
	MSCLAP-22	06.45	04.74	27.73	23.72	07.00	05.57	30.38	26.03
	MSCLAP-23	21.02	16.85	55.12	39.15	15.43	13.46	51.66	46.29
	CompA	11.82	10.19	43.03	40.24	15.70	13.18	53.36	43.77
	LAION-CLAP	14.41	11.62	46.91	41.94	16.52	11.90	52.75	43.71
	RobustCLAP	21.82	19.10	57.44	53.64	15.84	14.41	50.37	47.99
SoundDesc	ML-ACT	01.10	00.65	08.72	06.06	00.74	00.60	08.96	07.32
	MSCLAP-22	02.33	01.96	14.33	11.80	01.84	01.44	09.72	09.63
	MSCLAP-23	09.75	05.53	38.27	24.89	06.58	05.72	26.36	25.60
	CompA	06.80	04.03	33.32	23.56	04.21	03.32	20.86	17.26
	LAION-CLAP	05.82	03.17	24.62	18.09	03.23	02.34	17.75	13.63
	RobustCLAP	05.45	05.02	25.48	21.54	03.78	02.95	19.08	16.92
DCASE	ML-ACT	01.47	01.12	10.12	08.77	02.93	02.87	13.50	11.63
	MSCLAP-22	09.82	07.02	39.91	30.99	10.53	05.71	39.71	27.08
	MSCLAP-23	13.84	10.43	47.84	39.21	15.64	11.73	49.24	40.92
	CompA	14.84	10.61	49.54	39.51	14.44	08.92	48.54	35.10
	LAION-CLAP	13.34	11.23	44.73	37.81	17.25	10.93	54.86	43.53
	RobustCLAP	17.45	15.95	54.66	50.35	14.84	13.14	48.65	45.94

Table 7: Text-to-audio and audio-to-text on the original test set (TEST) and paraphrased test set (TEST-P). All ALMs show a consistent, significant drop in performance on TEST-P. RobustCLAP not only improves overall retrieval performance on TEST but also mitigates the drop in TEST-P. The best scores for each benchmark are highlighted in **bold**.