

AUDIOCARDS: STRUCTURED METADATA IMPROVES AUDIO LANGUAGE MODELS FOR SOUND DESIGN

Sripathi Sridhar^{1,2}, Prem Seetharaman², Oriol Nieto², Mark Cartwright¹, Justin Salamon²

¹New Jersey Institute of Technology ²Adobe Research

ABSTRACT

Sound designers search for sounds in large sound effects libraries using aspects such as sound class or visual context. However, the metadata needed for such search is often missing or incomplete, and requires significant manual effort to add. Existing solutions to automate this task by generating metadata, i.e. captioning, and search using learned embeddings, i.e. text-audio retrieval, are not trained on metadata with the structure and information pertinent to sound design. To this end we propose audiocards, structured metadata grounded in acoustic attributes and sonic descriptors, by exploiting the world knowledge of LLMs. We show that training on audiocards improves downstream text-audio retrieval, descriptive captioning, and metadata generation on professional sound effects libraries. Moreover, audiocards also improve performance on general audio captioning and retrieval over the baseline single-sentence captioning approach. We release a curated dataset of sound effects audiocards to invite further research in audio language modeling for sound design.

Index Terms— structured metadata, audio language models, sound design, audio captioning, text-audio retrieval

1. INTRODUCTION

Sound designers need to search in large sound libraries to find relevant sounds for their tasks. To achieve this, they typically search using aspects such as sound effect class, visual context (e.g. “car chase scene”), acoustic attributes, and causal relationships [1, 2]. However, existing professional sound libraries do not always have this information in their metadata, and significant human effort is required to provide this information, which is especially challenging for personal sound effects libraries recorded and/or collected by sound designers [3, 4].

Audio captioning models can predict descriptive information about audio [5, 6, 7], but they are generally trained on sentence-long captions from crowd-sourced workers [8], which do not necessarily contain the attributes or structure relevant to sound designers. While text-audio retrieval models can search for sounds without metadata using just a text query [9, 10], these embeddings are also trained on similar sentence captions which are not aligned with the shorter queries favored by users [1]. Furthermore, sound designers prefer to assess search results without having to listen to each file, which motivates the need for descriptive metadata [11].

To address these gaps, we propose structured metadata called *audiocards*. Audiocards are multi-field descriptions of audio files grounded in their acoustic attributes, sound classes, and usage context. Based on conversations with a professional sound designer, we include fields such as nouns, noun-verb pairs, and Universal Category System (UCS) [12] tags, as well as visual context and adjectives, which are aligned with sound designer needs [1, 2, 13]. To construct an audiocard, we few-shot prompt an off-the-shelf LLM

with metadata about an audio file, and a set of example handwritten cards, as seen in Figure 1. Generating captions after all other audiocard fields also leads to more descriptive captions by grounding in attributes relevant to sound design and inducing chain-of-thought like reasoning [14].

With audiocards, we can support both metadata generation and text-audio retrieval for sound design in multiple use cases. When text metadata is available, audiocards generated from LLMs can be used for keyword search on sound-design relevant fields. However, we can also train downstream captioning models on audiocards, enabling structured metadata and descriptive caption generation for sound libraries with incomplete or missing metadata. Finally, we can train text-audio embedding models on audio and audiocard field pairs to enable embedding-based audio retrieval aligned to the needs of sound designers. Overall, these approaches support sound design search scenarios such as using noun-verb pairs, e.g. “dog barking” [1] or an example visual scene [2] e.g. “haunted house”, and they can generate sound-design-relevant fields and descriptions that can be used to organize sound effects libraries and provide additional context during retrieval.

To validate the usefulness of audiocards we curate a human-verified subset of a professional sound effects library Audition SFx, called ASFx eval. We evaluate our captioning model on metadata generation and a descriptive captioning task on ASFx eval, as well as general audio captioning [8]. We evaluate our text-audio representation model on retrieval, both on a professional sound effects library and a general retrieval dataset [8]. We compare against state-of-the-art large audio language models (LALMs) and general contrastive text-audio embedding models on their respective tasks.

Our experiments show that audiocards improve audio-language modeling for sound design, specifically, that (1) we can generate audiocards from audio without accompanying metadata better than state-of-the-art LALMs; (2) training on audiocards leads to significantly better captions for professional sound effects libraries compared to both a baseline approach and state-of-the-art LALMs; and (3) training on audiocards improves embedding-based text-audio retrieval when using human-annotated captions as queries on a professional sound effects library.

2. AUDIOCARDS

We propose *audiocards*, structured descriptions of audio files with fields consisting of attributes contextually relevant to users’ audio understanding activity and goals. In this work, we focus on audiocards to aid in audio-language modeling for sound design.

2.1. Fields

We propose audiocard fields for sound design based on conversations with a professional sound designer. We opt for a noun-verb

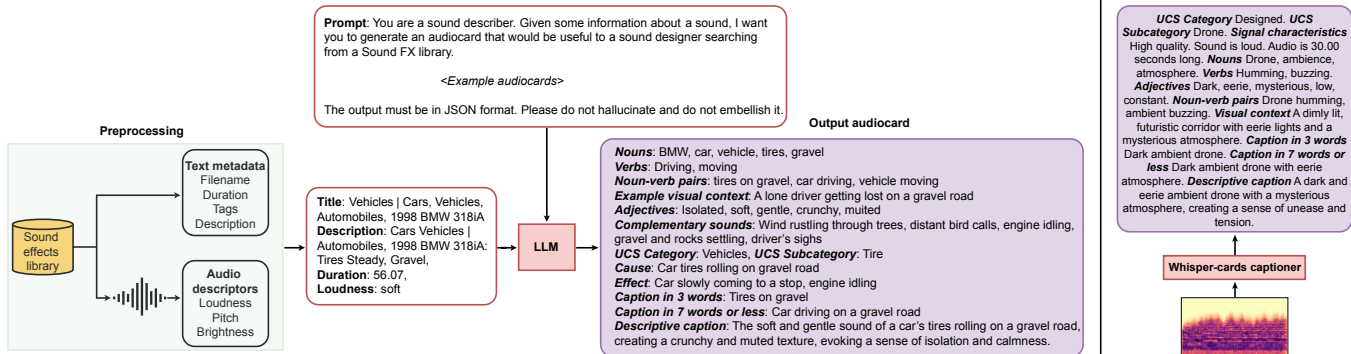


Fig. 1. Left: We propose *Audiocards*, structured metadata which describes an audio file with attributes relevant to sound designers. We prompt an LLM with the available text metadata and audio descriptors, and generate an audiocard, which can be used for text-based search and to train audio-language models. Right: Audiocard generated by our *Whisper-cards* audio captioneer from input audio without text metadata.

structure to describe the key sound-producing actors (what) and actions (how) in the audio file [15], through the fields *nouns*, *verbs* and *noun-verb pairs*. We also include the *Universal Category System (UCS) category* and *UCS subcategory* [12], which is a hierarchical ontology consisting of 82 categories and 753 sub-categories, commonly used to tag sound effects in professional sound libraries. This helps ground the audiocard vocabulary in what professional sound designers may use.

To better exploit the world knowledge of the LLM, we include a field called *example visual context* that acts as a bridge to potential visual scenes where this sound effect may be used. We provide additional auditory context using the fields of *complementary sounds*, which denotes potential co-occurring sounds, and *cause, effect* which aim to exploit a physical understanding of how the sound was created. We include an *adjectives* field that allows scope for sonic and emotional descriptive keywords. Conditioned on these fields, we then include *caption in 3 words*, *caption in 7 words or less*, and *descriptive caption*, to round out the card. A full example audiocard is depicted in Figure 1.

To improve the auditory grounding of these cards and mitigate hallucinations, we compute several descriptors from the audio signal. Specifically, we include duration, loudness, brightness, and pitch at the top of the audiocards. We pick these audio descriptors based on prior work that appended audio descriptors in captions to improve text-audio generation [16]. We include duration in seconds directly from the audio metadata. We compute loudness and bin them as *very soft* (-70 LKFS to -55 LKFS), *soft* (-55 LKFS to -40 LKFS), *loud* (-30 LKFS to -15 LKFS), and *very loud* (greater than -15 LKFS). We compute the spectral centroid as brightness, and classify values as *dull* (spectral centroid < 45) or *bright* (spectral centroid > 65) based on the distribution of values for our training data. We compute pitch using CREPE [17] on normalized and segmented audio, and classify pitch as *low* (pitch < 1.5 octaves) and *high* (pitch > 3.5 octaves). Values that fall outside these ranges are mapped to an empty string, allowing the model to learn from natural language audio descriptions with low ambiguity.

2.2. Training Data

We gather a large mix of proprietary, licensed sound effect datasets and publicly available CC-licensed general audio datasets, consisting of approximately 2M audio samples with corresponding meta-

2.3. Human verified sound effects evaluation dataset

Audio language models are typically evaluated on audio captioning datasets like Clotho [8]. However, these datasets may not be the best indicators of model performance in a sound design setting, as they do not reflect the data distribution of a sound effects database. Hence, we include Adobe Audition SFX, a publicly available library of 10K professionally recorded sound effects in our model evaluations¹.

Audition SFX primarily contains individually recorded sounds, and a small percentage of ambience recordings with overlapping sources. There are sound effects from 22 categories, such as Transportation and Production Elements, each with a descriptive title such as Drones/Drone Production Element Imaging Deep Synth Rumble Atonal 01.wav. We synthesized audiocards from these titles that we then used to evaluate our models. We extracted an evaluation subset for which we manually verified the generated audiocards accurately describe the audio file without incoherence or hallucination issues.

We built an annotation interface and annotated random examples on whether to keep them in the evaluation subset by listening to the audio file and examining the corresponding audiocard. We filtered out sound effects from the Noise, Tones, DTMF, DC, and Multichannel categories, which contain only test signals and pure tones. We uniformly sampled from each of the remaining categories. Once the annotator had annotated the example, we sampled the next example and repeated the process until we reached 500 approved files. We refer to this dataset of 500 approved files as ASFX eval and release audiocards for ASFX eval on our companion website².

2.4. Generation

We prompted an LLM to generate audiocards from the available file metadata. This metadata consists of fields which are always present, such as *filepath* and *duration*, and other fields which are sometimes present, such as *description*, *keywords*, *UCS category*, *UCS subcategory*, *extras*, and *tags*. In addition to a system prompt and warnings to stick to a JSON format, we included three handwritten audiocards in the prompt for in-context learning. We prompted the model to generate the audiocard fields in order—as the LLM is autoregressive, latter audiocard fields are conditioned on prior fields [14]. This enables us to generate audiocards grounded in audio descriptions and

¹<https://www.adobe.com/products/audition/offers/adobeauditionlcsfx.html>

²<https://sites.google.com/view/audiocards/>

to generate captions of different lengths that encapsulate all relevant information about the audio file.

As a baseline, we prompt the LLM to sequentially generate captions in odd lengths from two words to 15 words, so that the longer captions are conditioned on the shorter captions. We use a Pixtral-12B-2409 to generate both the audiocards and the baseline captions. We henceforth refer to these as Pixtral cards and Pixtral captions.

2.5. Improving UCS categories estimation with card classification

Sound effects libraries are often organized using UCS category—audio language models trained for sound design would thus benefit from learning semantic mappings between audio and UCS category. Since many examples in our data do not have human-labeled UCS tags, we aim to generate them to include in the audiocard and use in model training. However, we found in our initial experiments that LLMs used to generate audiocards could not reliably predict the UCS category and would often hallucinate, even when provided in-context examples and a list of possible UCS categories.

To this end, we trained a lightweight text classifier that takes the generated audiocards without UCS category as input, and predicts the UCS category and subcategory. We fine-tune a distilBERT-base-uncased [18] with a classifier head containing two linear layers on previously generated audiocards. We train on a portion of our data with human-annotated UCS tags as targets, achieving 96% accuracy on a held-out subset. We insert these pseudo-ground truth classifier-generated UCS predictions into the audiocards for examples without human annotated UCS tags during downstream model training.

3. MODELS

3.1. Audio captioning

To evaluate the usefulness of audiocards for audio captioning, we trained audio captioners to generate audiocards, called Whisper-Cards, and captions, called Whisper-Baseline, from an audio file. Whisper-Cards first predicts the signal characteristics to provide auditory grounding similar to audiocard generation in Section 2.4. Specifically, we train the model to predict the audio quality, loudness, and duration, e.g., *Signal characteristics – High quality. Sound is loud. Audio is 2.98 seconds long.* Then, the captioner predicts the rest of the audiocard fields, e.g., *Nouns – Electricity, arc, spark. Verbs – , etc. up to the Descriptive caption field.* See Figure 1 for an example audiocard. We include additional examples of audiocards generated by the audio captioners on our companion website.

We trained Whisper-medium-v3 [19], a transformer encoder-decoder sequence-sequence model pre-trained on 680k hours of weakly labeled speech data, using audio from our dataset paired with Pixtral cards or Pixtral captions for Whisper-Cards and Whisper-Baseline respectively. We downsampled all our audio files to 16kHz to match Whisper’s input requirements. We finetuned on Pixtral cards or Pixtral captions for 100k iterations using batch size 16. We used a cosine decay schedule with linear warmup for 4000 steps, and an AdamW optimizer with learning rate $1e-5$.

We evaluate our captioners using metrics commonly used in the DCASE automated audio captioning challenge – SPIDeR [20] and FENSE [21]. SPIDeR is a combination of SPICE [22] and CIDEr [23] that compares the semantic graphs of the predicted and reference captions. FENSE measures similarity using SentenceBERT [24] and also accounts for fluency errors.

3.2. Text-audio contrastive representation learning

We also evaluated the utility of audiocards for text-audio sound effects retrieval by training text-audio contrastive models. We trained a Cards-CLAP model on (audio, Pixtral audiocard field(s)) pairs by randomly selecting fields in the card during training. Through initial experiments we found the best performing model to be the one trained on a random selection from [“nouns”, “noun-verb pairs”, [“nouns”, “noun-verb pairs”], [“noun-verb pairs”, “UCS Category”], “caption in 3 words”, “caption in 7 words or less”, “descriptive caption”]. When an option with multiple fields is picked, we concatenate the two fields together. We also trained a baseline model, called Captions-CLAP, on (audio, Pixtral caption) pairs by randomly selecting different caption lengths during training.

Our model architecture is based on the LAION-CLAP framework [9], with a RoBERTa [25] text encoder and a HTSAT [26] audio encoder. We trained using an overall batch size 1280 for 40k steps with a peak learning rate of $1e-4$, using a linear warmup of 3500 steps and a cosine decay schedule.

To evaluate text-audio retrieval, we computed audio embeddings for all examples in the dataset. Then, given a text query, we computed the cosine similarity between its embedding and all audio embeddings, picking the highest similarity top K as the retrieval results.

4. EXPERIMENTS

4.1. Can audio captioning models trained on audiocards generate structured metadata?

We benchmark audiocard field generation as a structured metadata generation task for sound design using Whisper-Cards (Section 3.1), which generates audiocards from an audio file without additional prompts. As a baseline we prompt Audio Flamingo 3 (AF3) [27], a state-of-the-art LALM, to generate audiocard fields in the same order as Whisper-Cards, given an audio file and audiocard field descriptions. We use ASFx eval Pixtral audiocards generated from available text metadata (Section 2.4) as the ground truth. We evaluate individually on the audiocard fields: *nouns*, *verbs*, *adjectives*, *noun-verb pairs*, and *descriptive caption*. We also evaluate UCS category prediction as a multi-class classification task, using pseudo-ground truth from the UCS category predictor (Section 2.5) as targets.

We see in Figure 2 that Whisper-Cards outperforms AF3 across all audiocard fields. Moreover, AF3 which is not trained using UCS tags, has multi-class UCS category prediction accuracy 0% while Whisper-Cards achieves 31%. While there is room for improvement, these results highlight the utility of audiocard training for sound effects metadata prediction. We also note that standard audio captioning models like Whisper-Baseline cannot generate audiocard fields as they are only trained on single sentence captions.

4.2. Are captions generated in audiocards better than baseline and LALM captions?

To illustrate the quality of audiocard captions, we trained Whisper-Cards on just the last field from Pixtral audiocards, i.e., the descriptive caption, and Whisper-Baseline on 15 word Pixtral captions (Section 3.1). We evaluate using ASFx eval Pixtral cards and Clotho captions. In initial experiments, we found a domain gap between the training and evaluation data when models were trained on just Pixtral audiocards for our internal dataset. To rectify this, we added the Clotho development set to the training data. When training to generate the entire audiocard, we treat the Clotho caption as a *descriptive caption* and ignore the other audiocard fields that don’t exist in the ground truth captions in evaluation. We also compare against recent




Fig. 2. Audiocard field metadata prediction evaluated on ASFx eval. We use Audio Flamingo 3 (AF3) [27], a state-of-the-art large audio language model, in think mode to utilize its reasoning capabilities. Whisper-Cards, trained to generate audiocards on our data, achieves superior performance on all audiocard fields.

Table 1. Audio captioning results on the Clotho and ASFx eval. Whisper-Baseline is trained on Pixtral captions, while Whisper-Cards is trained on either just the descriptive caption from Pixtral audiocards or the full audiocard. * Results from the original paper. Bold and underlined values are the best and second best models.

| Model | Clotho | | ASFx eval | |
|-------------------------------|--------------|--------------|--------------|--------------|
| | SPIDER | FENSE | SPIDER | FENSE |
| Whisper-Baseline | 21.05 | 47.61 | 7.98 | 49.78 |
| Whisper-Cards (card caption) | 22.18 | 48.48 | 19.36 | 53.40 |
| Whisper-Cards (full card) | 22.07 | 48.67 | <u>18.61</u> | <u>51.78</u> |
| WavCaps CNN14 [6] | <u>31.0*</u> | – | 3.00 | 32.20 |
| WavCaps HTSAT [6] | 29.7* | – | 3.98 | 35.27 |
| GAMA [28] | 5.34 | 43.21 | 9.30 | 45.70 |
| Audio Flamingo 2 [29] | 7.25 | 33.71 | 1.68 | 30.08 |
| Audio Flamingo 3 [27] | 31.88 | 52.11 | 5.27 | 32.10 |
| Audio Flamingo 3 (think) [27] | 13.22 | <u>50.19</u> | 9.61 | 42.61 |

state-of-the-art LALMs, Audio Flamingo 3 (AF3) [27] and 2 (AF2) [29], GAMA [28], as well as WavCaps with CNN and HTSAT encoders [6]. We tried several recommended prompts for AF2, AF3 and GAMA, and report metrics from the best prompt.

We report Clotho and ASFx eval metrics in Table 1. We find that training on Pixtral audiocards improves performance over Pixtral captions on both datasets, suggesting that audiocards contain more descriptive captions. Our models underperform AF3 and WavCaps on Clotho, likely due to the distribution mismatch between Clotho and the rest of our training data. On ASFx eval however, we find that training on audiocard captions is significantly more useful than training on baseline captions, and that Whisper-Cards significantly outperforms state-of-the-art LALMs. This suggests audiocards can significantly improve audio captioning performance on sound design specific datasets.

4.3. Does generating audiocards lead to better captions?

We also investigate the effect of training on full audiocards on downstream captioning performance. We train Whisper-Cards to generate the entire audiocard as a series of the fields followed by their contents, given an input audio file. We report metrics for Whisper-Cards models trained on the full audiocard in Table 1. Across datasets, we find that training on the full audiocard also improves performance over Whisper-Baseline. However, we do not see a clear improve-

Table 2. Zero-shot text-audio retrieval results for models trained on audiocards (Cards-CLAP) and captions (Captions-CLAP), evaluated on Clotho and our internal dataset (ID). Audiocards outperform baseline captions across datasets, and significantly outperform comparable models on ID. * Model is trained on Clotho, result from the original paper. Bold models are best for that dataset.

| Model | Training regime | Dataset | R@10 | CatP@10 |
|----------------|-------------------|---------|---------------|--------------|
| Captions-CLAP | Baseline captions | ID | 73.45 | 77.66 |
| Cards-CLAP | Audiocard fields | ID | 75.40 | 78.73 |
| LAION-CLAP [9] | – | ID | 24.85 | 47.10 |
| Captions-CLAP | Baseline captions | Clotho | 50.12 | 35.00 |
| Cards-CLAP | Audiocard fields | Clotho | 52.44 | 35.26 |
| LAION-CLAP [9] | – | Clotho | 55.40* | – |

ment over the Whisper-Cards model trained only on the descriptive caption. This is likely because the model can focus modeling capacity on just the descriptive caption, which is what the metrics are computed on. However, Whisper-Cards can successfully generate full audiocards while still outperforming the baseline, and is suitable for structured metadata generation to organize a sound effects library from just the raw audio files.

4.4. Does training text-audio models on audiocards improve downstream retrieval performance?

We evaluate retrieval results using Recall@10 on Clotho and an internal professional sound effects dataset (ID) with around 6000 samples held out during training. To quantify semantic relevance in the search results, we compute the UCS category-level P@10 as the percentage of examples in the top 10 results with the expected UCS category [30]. For our internal dataset, we have ground truth UCS category annotations. For Clotho, we use the text classifier described in Section 2.5 to generate UCS category annotations. Note that neither Cards-CLAP nor Captions-CLAP are trained on Clotho or ID.

We compare the baseline Captions-CLAP and the proposed Cards-CLAP models in Table 2. We find that training with audiocards leads to better downstream text-audio retrieval performance across datasets, improving recall by 4.5% on Clotho and 2.6% on ID. We also see an improvement in category precision, suggesting that the retrieved results are more semantically aligned with the queries. We also compare against LAION-CLAP, a popular text-audio retrieval model with a similar model architecture to Cards-CLAP [9]. We find that ID is challenging for zero-shot retrieval, and that Cards-CLAP achieves comparable zero-shot performance with LAION-CLAP which is trained on Clotho.

5. CONCLUSION

We introduced audiocards, a novel structured metadata for sound design that are multi-field descriptions grounded in acoustic descriptors. Audiocards are more expressive than conventional single-sentence captions, enabling models to learn richer semantic associations between audio and text. Through extensive evaluations, we demonstrate that training downstream automated audio captioners and text-audio contrastive models on audiocards leads to consistent improvements over baseline captioning approaches across retrieval and captioning benchmarks.

Beyond performance gains, our work underscores the importance of adapting audio understanding models to domain-specific requirements. To this end, we release a curated dataset of audiocards, and benchmark a novel metadata generation task on audiocard fields. Through this work, we aim to catalyze research that bridges general-purpose audio modeling with the workflows of sound designers.

6. ACKNOWLEDGMENT

We thank Adolfo Hernandez Santisteban for insightful discussions during this project. This work was partially supported by NSF awards 2504642 and 2504643.

7. REFERENCES

- [1] B. Weck and F. Font, “The language of sound search: Examining user queries in audio search engines,” in *Proc. 9th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, Tokyo, Japan, Oct. 2024.
- [2] D. Sonnenschein, *Sound design*. Michael Wiese Productions Studio City, CA, 2001.
- [3] “Frank Bry Special: Methods of a Metadata Madman [Part 1].” [Online]. Available: https://designingsound.org/2010/12/16/frank-bry-special-methods-of-a-metadata-madman-part-1/?utm_source=chatgpt.com
- [4] D. Forshee, “Top 4 Tips for Organizing Your Sound Effects Library.” [Online]. Available: <https://blog.prosoundeffects.com/top-4-tips-for-organizing-your-sound-effects-library>
- [5] S.-L. Wu, X. Chang, G. Wichern, J.-W. Jung, F. Germain, J. Le Roux, and S. Watanabe, “Improving Audio Captioning Models with Fine-Grained Audio Features, Text Embedding Supervision, and LLM Mix-Up Augmentation,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2024, pp. 316–320, iSSN: 2379-190X.
- [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*, vol. 32, pp. 3339–3354, 2024.
- [7] G. Zhu, J. Daresky, and Z. Duan, “Cacophony: An Improved Contrastive Audio-Text Model,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 4867–4879, 2024.
- [8] 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)*, May 2020, pp. 736–740, iSSN: 2379-190X.
- [9] 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 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [10] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, “Clap learning audio concepts from natural language supervision,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [11] H. Liu, T. Deacon, W. Wang, M. Paradis, and M. D. Plumbley, “Exploring the user experience of AI-assisted sound searching systems for creative workflows,” in *Proc. 159th Audio Engineering Society Convention*, 2025.
- [12] “Universal Category System.” [Online]. Available: <https://universalcategorysystem.com/>
- [13] F. Riege, “TTA GUIDES: Tagging for Sound Design | Metadata for Production Music,” Apr. 2022. [Online]. Available: <https://tagteamanalysis.com/2022/04/tta-guides-tagging-for-sound-design/>
- [14] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [15] G. Wijngaard, E. Formisano, B. L. Giordano, and M. Dumontier, “Aces: Evaluating automated audio captioning models on the semantics of sounds,” in *2023 31st European Signal Processing Conference (EUSIPCO)*, 2023, pp. 770–774.
- [16] S. Kumar, P. Seetharaman, J. Salamon, D. Manocha, and O. Nieto, “SILA: Signal-to-Language Augmentation for Enhanced Control in Text-to-Audio Generation,” Dec. 2024, arXiv:2412.09789 [cs].
- [17] J. W. Kim, J. Salamon, P. Li, and J. P. Bello, “Crepe: A Convolutional Representation for Pitch Estimation,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 161–165, iSSN: 2379-190X.
- [18] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter,” in *Proc. 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing, NeurIPS*, 2019.
- [19] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, “Robust Speech Recognition via Large-Scale Weak Supervision,” in *Proceedings of the 40th International Conference on Machine Learning*. PMLR, Jul. 2023, pp. 28 492–28 518, iSSN: 2640-3498.
- [20] S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. Murphy, “Improved Image Captioning via Policy Gradient optimization of SPIDEr,” in *2017 IEEE International Conference on Computer Vision (ICCV)*. Venice: IEEE, Oct. 2017, pp. 873–881.
- [21] Z. Zhou, Z. Zhang, X. Xu, Z. Xie, M. Wu, and K. Q. Zhu, “Can audio captions be evaluated with image caption metrics?” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 981–985.
- [22] P. Anderson, B. Fernando, M. Johnson, and S. Gould, “SPICE: Semantic Propositional Image Caption Evaluation,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 382–398.
- [23] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, “Cider: Consensus-based image description evaluation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4566–4575.
- [24] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 3982–3992.
- [25] 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 *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022, pp. 646–650, iSSN: 2379-190X.
- [26] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A robustly optimized BERT pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [27] A. Goel, S. Ghosh, J. Kim, S. Kumar, Z. Kong, S.-g. Lee, C.-H. H. Yang, R. Duraiswami, D. Manocha, R. Valle, and B. Catanzaro, “Audio flamingo 3: Advancing audio intelligence with fully open large audio language models,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 38, 2025.
- [28] S. Ghosh, S. Kumar, A. Seth, C. K. R. Evuru, U. Tyagi, S. Sakshi, O. Nieto, R. Duraiswami, and D. Manocha, “GAMA: A large audio-language model with advanced audio understanding and complex reasoning abilities,” in *Proc. Conf. Empirical Methods in Natural Language Processing (EMNLP)*, 2024, pp. 6288–6313.
- [29] S. Ghosh, Z. Kong, S. Kumar, S. Sakshi, J. Kim, W. Ping, R. Valle, D. Manocha, and B. Catanzaro, “Audio flamingo 2: An audio-language model with long-audio understanding and expert reasoning abilities,” in *Proc. 42nd Int. Conf. Machine Learning (ICML)*, 2025, pp. 19 358–19 405.
- [30] J. Wilkins, J. Salamon, M. Fuentes, J. P. Bello, and O. Nieto, “Bridging high-quality audio and video via language for sound effects retrieval from visual queries,” in *2023 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2023, pp. 1–5.