

BRIDGING HIGH-QUALITY AUDIO AND VIDEO VIA LANGUAGE FOR SOUND EFFECTS RETRIEVAL FROM VISUAL QUERIES

Julia Wilkins^{1,2} Justin Salamon² Magdalena Fuentes¹ Juan Pablo Bello¹ Oriol Nieto²

¹ Music and Audio Research Laboratory, New York University, NY, USA

² Adobe Research, San Francisco, CA, USA

jw3596@nyu.edu

<https://juliawilkins.github.io/sound-effects-retrieval-from-video/>

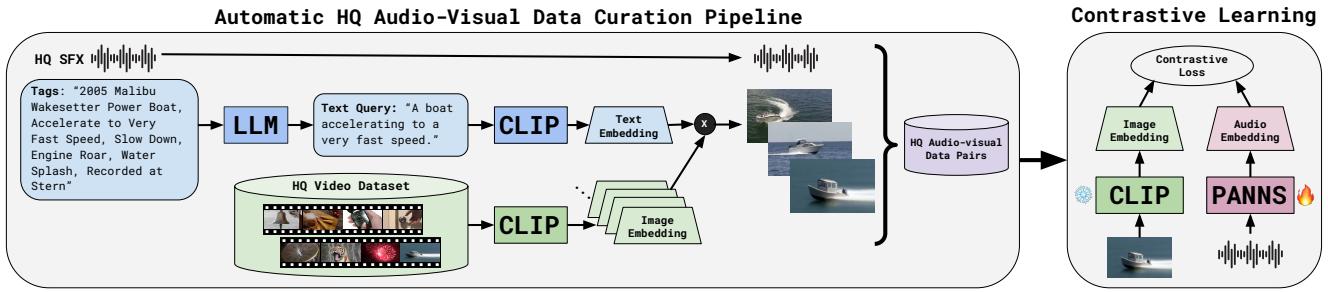


Figure 1: Our automatic audio-visual data curation pipeline (left). The new data are used to learn multimodal audio-visual embeddings (right).

ABSTRACT

Finding the right sound effects (SFX) to match moments in a video is a difficult and time-consuming task, and relies heavily on the quality and completeness of text metadata. Retrieving high-quality (HQ) SFX using a video frame directly as the query is an attractive alternative, removing the reliance on text metadata and providing a low barrier to entry for non-experts. Due to the lack of HQ audio-visual training data, previous work on audio-visual retrieval relies on YouTube (“in-the-wild”) videos of varied quality for training, where the audio is often noisy and the video of amateur quality. As such it is unclear whether these systems would generalize to the task of matching HQ audio to production-quality video. To address this, we propose a multimodal framework for recommending HQ SFX given a video frame by (1) leveraging large language models and foundational vision-language models to bridge HQ audio and video to create audio-visual pairs, resulting in a highly scalable automatic audio-visual data curation pipeline; and (2) using pre-trained audio and visual encoders to train a contrastive learning-based retrieval system. We show that our system, trained using our automatic data curation pipeline, significantly outperforms baselines trained on in-the-wild data on the task of HQ SFX retrieval for video. Furthermore, while the baselines fail to generalize to this task, our system generalizes well from clean to in-the-wild data, outperforming the baselines on a dataset of YouTube videos despite only being trained on the HQ audio-visual pairs. A user study confirms that people prefer SFX retrieved by our system over the baseline 67% of the time both for HQ and in-the-wild data. Finally, we present ablations to determine the impact of model and data pipeline design choices on downstream retrieval performance. Please visit our companion website to listen to and view our SFX retrieval results.

Index Terms— Multimodal machine learning, cross-modal retrieval, audio-visual representation learning, data augmentation

1. INTRODUCTION

Video creation and editing commonly entails matching sound effects (SFX) to video to create an engaging multimedia experience. Yet finding the right SFX to pair with moments in a video is challenging and often time-consuming. Regardless of skill level, this process typically relies heavily on the quality and completeness of text metadata associated with an SFX library (i.e., for a text-based search), which can vary significantly by library. Further, it is common to use sounds from one object to represent entirely different objects (e.g., the sound of shoes on glass shards to represent walking through snow), an art known as “Foley” [1, 2], meaning sonically relevant sounds may be missed if we rely on metadata for retrieval.

Multimodal machine learning methods have become popular across the visual, audio, and text domains [3, 4], due to their ability to learn from the correspondence between signals in different modalities, often via contrastive learning [5, 6, 7], and the growing availability of large multimedia datasets [8, 9, 10, 11, 12]. Previous approaches for audio-visual retrieval [13, 14, 15, 16, 17] rely on YouTube, aka “in-the-wild”, videos [11, 18], where the audio is often noisy and the video of amateur quality. These videos are not necessarily representative of the production-quality videos and associated audio that video creators often work with, so it is unclear whether systems trained on noisy video datasets can generalize to this type of high-quality (HQ) content. Datasets of HQ SFX moments in videos are lacking, though recent attempts have been made to address this by assuming sound event occurrences via video captions [9, 19] or even generation [20]. For music, unlike SFX, there are more HQ audio-visual data available (namely music video clips), which have been leveraged to train music retrieval systems for video [21, 22, 23]. To circumvent the need for HQ audio-visual pairs, Lin et al. [24] use CLIP [25] to match embeddings of video frames against embeddings of text metadata associated with audio

SFX, to recommend SFX for a given video.

In this work, we address the lack of HQ audio-visual data for SFX via a novel automatic data curation pipeline, and use it to train an audio-visual SFX retrieval system that generalizes both to production-grade and in-the-wild media content. Our contributions can be summarized as follows: (1) we propose an automatic data curation pipeline that bridges HQ audio and visual data via large language models and foundation vision-language models; (2) we show both quantitatively and qualitatively that by using this pipeline to train a contrastive audio-visual model, we are able to significantly outperform baselines on SFX retrieval given a visual input query; (3) we show that our model generalizes to noisy “in-the-wild” video despite being trained on the auto-curated HQ data; and (4) we present ablations to study the impact of key model and data pipeline design choices on downstream retrieval performance.

2. METHOD

2.1. Automatic Audio-visual Data Curation Pipeline

To generate HQ audio-visual training data, we propose an automatic data curation pipeline that leverages language and vision-language models. First, we use text metadata associated with each SFX to generate text embeddings via CLIP [25], and in parallel generate CLIP image embeddings of video frames. Next, we retrieve the most similar visual frames to each SFX-based text query via cosine similarity. This method is depicted in Figure 1 (left). Studies have shown that CLIP performs best at image retrieval when the input text is a clear natural language description [26, 27]. The text metadata in most HQ audio SFX libraries, including the one used for training in this work, does not fit this criterion. Rather, it is a list of tags (Figure 1, left) that may be over-specific or include audio-centric terms (e.g. “left to right sound”) that are not helpful for retrieving relevant visual frames. To alleviate this, we propose to use a large language model (LLM) in the data curation pipeline to convert the tag data into natural language sentences before computing their CLIP embeddings. In this way, we obtain a text query that is both in the right format (natural language) and level of specificity for successful text-image matching using CLIP. We use Bloom [28], an open-source LLM with 176B parameters and leverage prompt analogies [23] to improve the quality of the output sentence. We found that prompt analogies, in addition to producing overall better quality sentences, allowed us to control their level of specificity and filter out information from over-detailed tags. Later in our experiments, we perform an ablation to capture the quantitative effect of using the LLM-transformed SFX tags for text-image matching as opposed to using the tags directly.

Our approach takes inspiration from Lin et al. [24], with two key distinctions. First, we use an LLM to enhance the text metadata of SFX and optimize the image-text matching via CLIP. Most importantly, in [24] they use this image-text relationship directly as the retrieval system. In our method, we use this as a way to generate training data to then train an audio-visual model such that we can retrieve SFX using visual queries *without* the need for text. The entire pipeline is depicted in Figure 1.

The proposed data pipeline is fully automatic and thus highly scalable, constrained only by the size of the SFX and video datasets. Additionally, each SFX can be matched against the top k relevant visual frames, serving as a form of data augmentation and producing a dataset of audio-visual pairs that is k times the size of the SFX dataset. While our focus is on HQ SFX and video in this work, the

pipeline is applicable to any combination of audio and visual data as long as the audio is associated with tag metadata.

Datasets: To create audio-visual training pairs using our pipeline, we draw data from high-quality SFX and video datasets. For SFX we use the ProSoundEffects (PSE) Core 3 Pro Plus¹ bundle. It contains 150k professional-grade audio from over 80 SFX categories, along with metadata including category labels and notably a detailed list of metadata tags per audio file. Some files contain multiple SFX events separated by silence, in which case we segmented them into a total of 336k individual SFX. For video, we use a large collection of over 1 million HQ stock videos² that are mostly 10 seconds or less in duration, and have associated text metadata of varied quality and level of detail. Since our focus is on SFX for notable visual events, we use the metadata to filter the SFX and video datasets to exclude categories outside the scope of our work. For audio, we removed human sounds (e.g., speech, crying) and music (e.g., saxophone, orchestra). For video, we aimed to keep visual scenes depicting sound-producing foreground objects and actions, removing largely ambient scenes (e.g., a video of an empty classroom is not of interest in our setup). After filtering the PSE SFX library and segmenting the audio files, we obtained 332k SFX with a median duration of 3 seconds. We obtained 115k videos after similar filtering. We sample the videos at 1FPS, and extract up to 10 random frames per video, yielding approximately 1.1 million video frames to use in our automatic data curation pipeline. We split the SFX into train, validation, and test splits, the latter two containing 4k SFX each with the rest used for training. The audio and visual data are both completely disjoint across the data splits.

2.2. Contrastive learning for SFX retrieval

We use our HQ audio-visual training pairs to train a multimodal model that maps audio and video frames into a shared embedding space (Figure 1, right). Once trained, the model can be used for retrieval by computing the embedding distance between a query frame and a dataset of SFX to rank the SFX by their similarity to the visual query. Similar to recent work including Wav2CLIP [13], AudioCLIP [14], and CLAP [15], we use contrastive learning with a standard contrastive loss [5] to train an audio encoder to project an audio signal into the CLIP image embedding space. We do so by keeping the original CLIP image encoder frozen during training, forcing the audio encoder to project audio into this pre-trained space. The choice of audio encoder is discussed in Section 3.4.

3. EXPERIMENTAL DESIGN

3.1. Evaluation Datasets

We evaluate our system on SFX retrieval from visual queries using three test datasets: two created using our automatic data curation pipeline (PSE-Test and ASFX-Test), and one “in-the-wild” dataset of YouTube videos (AVE-Test). **PSE-Test** is the test split described in Section 2, containing 4k unique audio-visual pairs spanning 40 SFX categories. It is the most “in-domain” of the three test sets, though note again that it is completely disjoint from the training data in terms of the audio and video files used to create the pairs. **ASFX-Test** was created similarly to PSE-Test, but using SFX from the publicly available Adobe Audition SFX Library³ paired with the

¹<https://www.prosoundeffects.com>.

²<https://stock.adobe.com/video>.

³<https://www.adobe.com/products/audition/offers/AdobeAuditionDLCSFX.html>.

same HQ video dataset used above. It contains 5k unique audio-visual pairs across 19 SFX categories, and is completely disjoint from PSE-Test and the training data. **AVE-Test** is a subset of the AVE dataset [10], itself a curated subset of AudioSet [18] containing videos with a high degree of audio-visual correspondence. We filter AVE to exclude videos containing human sounds (notably speech) or music, which as noted earlier are outside the scope of this work, leaving 2.4k videos across 17 categories.⁴ While our target application is matching HQ SFX to production-grade video, as represented by the first two datasets, AVE-Test allows us to evaluate our system's ability to generalize to "in-the-wild" noisy data.

3.2. Evaluation Metrics

We use a given video frame as query to rank the audio from all test pairs based on their cosine similarity in the multimodal embedding space. We repeat this for all test pairs and compute standard retrieval metrics: Median Rank (MR) and Recall at 10 (R@10). A priori, we expect the metrics for matching the visual frame to the *exact* audio file in the audio-visual pair to be low, since there could be many other audio files containing SFX that are also relevant to that frame. For this reason, we also compute category-level metrics using the categories from the SFX metadata, where a retrieved SFX is considered a correct match for a frame if its from the same SFX category as the exact audio from the audio-visual pair. Since R@10 is a more lenient metric at the category-level, we instead focus on Precision at 10 (P@10), i.e., the fraction of top-10 results that are from the right SFX category. This metric is most representative of our target use case, i.e., presenting the user with a short list of SFX, ideally *all* of them from the right category given the visual query.

3.3. Baselines

We evaluate our system against two strong pre-trained baselines: Wav2CLIP [13] and AudioCLIP [14], both of which were trained to map audio input to the CLIP image embedding space, meaning they can be used for audio retrieval given a visual input in the same way as our system. **Wav2CLIP** uses a ResNet18 architecture for the audio encoder and was trained on 5-second chunks of video frames from VGGSound [11] sampled at a high frame rate and mean-pooled for segment-level CLIP embeddings. **AudioCLIP** uses an ESResNet [29] architecture initialized using ImageNet [30] pre-training for the audio encoder, and was trained on single random frames from AudioSet [18] videos. We use the AudioCLIP-Trimodal model in which the audio encoder was trained against CLIP image and CLIP text embeddings jointly. We use the original model weights provided by the authors for both baselines.

3.4. Model Architecture and Training

For our system, we experiment with two model architectures for the audio encoder: PANNS [31] and Wav2CLIP [13]. AudioCLIP is not included due to its lower performance as a baseline compared to Wav2CLIP, as we discuss in Section 4.1. PANNS is a 14-layer CNN architecture with 81M trainable parameters. It was originally trained for audio classification on AudioSet. Wav2CLIP's ResNet18 audio encoder architecture has 11.2M trainable parameters, and was originally trained on VGGSound. For each of the two architectures, we try both training the model from scratch and fine-tuning the original model weights using our auto-curated training set. We train for a maximum of 150 epochs using batch size of 64

⁴The list of SFX and videos used in evaluation is on our companion site.

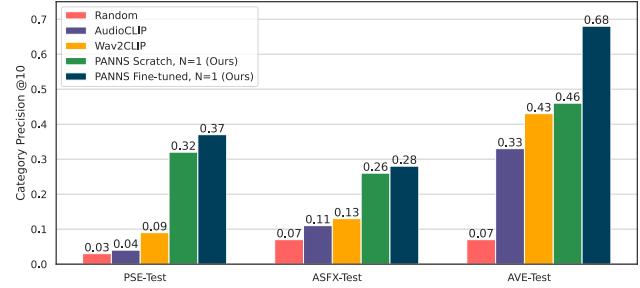


Figure 2: SFX retrieval results: P@10 for our best model (from scratch and fine-tuned) and the baselines on the three test sets.

and learning rate of 10^{-5} , and select the best epoch based on the Category Precision@10 metric on our validation split.

3.5. Automatic Data Curation Design Choices

A key design choice in our automatic data curation pipeline is whether we allow the same video frame to be matched to multiple audio SFX (in the same data split), or whether we constrain the pairing. Since multiple SFX can end up with the same or similar natural language descriptions, if we pair them with the top matching video frame without any constraints, they may be paired to the same frame. This could lead to "visual hubs" that are potentially matched to hundreds of SFX. While this strategy ensures the best visual match for each SFX, it reduces the diversity of the dataset since we use fewer video frames overall. To address this, we can impose a matching constraint wherein each video frame can only be matched against N SFX after which we remove it from the pool of available video frames. In the extreme case, we can set $N = 1$, meaning each video frame can only be used in a single audio-visual pair, producing a dataset of completely unique audio-visual pairs. This has the effect of significantly increasing the diversity of the dataset, but at the cost of not always using the most closely matched video frame for each SFX. Having no constraint is equivalent to $N = \infty$, where a video frame can be matched to an unbounded number of SFX. In Section 4.3 we present an ablation study where we vary $N = \{1, 2, 5, 10, 100, \infty\}$ to generate variants of our training set and evaluate how they impact downstream retrieval performance.

3.6. User Study

Since our goal is to produce a system the users find helpful, we complement our quantitative evaluation with a user study in which we compare our best performing model against the Wav2CLIP baseline. The user was presented with two videos, each containing the same single video frame but different SFX, one with the top SFX retrieved by our model and one with the top SFX from the baseline. The user was asked to indicate "Which video has a sound that better matches the image?" We deliberately did not play the full video from which the frame was taken, since precise synchronization of the SFX to the video is outside the scope of this work. Using a pool containing a random sample of 400 video frames from PSE-Test and 400 frames from AVE-Test, each user was shown a random selection of 30 frames from the pool as queries. The ordering of the two videos in each comparison was also randomized. We gathered 28 participants, and each completed 30 comparisons, yielding 840 total comparisons, evenly split between PSE-Test and AVE-Test.

Audio Encoder	Training	Pairing	PSE-Test				ASFX-Test				AVE-Test			
			Exact		Category		Exact		Category		Exact		Category	
			MR	R@10	R@10	P@10	MR	R@10	R@10	P@10	MR	R@10	R@10	P@10
PANNS [31]	Scratch	∞	304	0.06	0.65	0.29	835	0.03	0.63	0.25	414	0.03	0.81	0.37
	Scratch	1	190	0.08	0.63	0.32	893	0.03	0.58	0.26	393	0.03	0.84	0.46
	Fine-tuned	∞	138	0.10	0.49	0.36	529	0.04	0.66	0.28	201	0.05	0.92	0.61
	Fine-tuned	1	122	0.11	0.66	0.37	597	0.03	0.61	0.28	193	0.05	0.95	0.68

Table 1: Ablation study over model design and data pipeline design choices.

4. RESULTS

4.1. Quantitative and Qualitative Results

Based on the PSE *validation* split, our best performing model uses the PANNS architecture for the audio encoder, starting from the pre-trained model weights and fine-tuning the encoder on our auto-curated training set which is generated with a pairing constraint of $N = 1$ (for model selection we only tried $N = \{1, \infty\}$). In Figure 2 we present the P@10 scores for this model, the PANNS architecture trained from scratch, the two baselines and a random baseline on the three test sets described above. Our best model significantly outperforms the baselines across all three test sets. The results indicate two important takeaways: first, that our model is superior at matching HQ SFX to production quality video, our target application, as indicated by the results on PSE-Test and ASFX-Test. Second, and more surprising, our model successfully generalizes to noisy in-the-wild data, represented by AVE-Test, which is particularly interesting for our model trained from scratch as it was trained with *only* HQ auto-curated audio-visual pairs. Conversely, the baselines, which were trained on large in-the-wild video collections, fail to generalize to the clean audio-visual data. This validates our proposed data curation pipeline as useful for training models that generalize to both clean HQ data and noisy in-the-wild data.

Turning to the qualitative results, the user study further validates our findings: participants preferred the SFX recommended by our model over the baseline 68.1% and 66.4% of the time for data from PSE-Test and AVE-Test, respectively. Two binomial tests, one per dataset, indicated that these proportions were significantly higher than the expected 50%, $p < 0.01$ (1-sided).

4.2. Model Ablation

In Table 1 we present an ablation over model design and data pipeline design choices. We conducted the same ablation using the Wav2CLIP audio encoder, but since results were consistently better for the PANNS encoder, we only present the latter. From initial experiments, we hypothesize that performance improves using PANNS vs. Wav2CLIP due to both our HQ audio-visual data curation pipeline and increased model capacity (81M vs. 11.2M parameters). Fine-tuning a pre-trained PANNS works better than training from scratch, indicating we can benefit from transfer learning from AudioSet classification to our audio-visual matching task. Finally, when only comparing $N = 1$ to ∞ for the pairing constraint in the automatic data curation pipeline, the former yields better results, especially for AVE, indicating it is preferable to have a more diverse training set, even at the cost of not pairing every SFX to its top-matching video frame.

4.3. Automatic Data Curation Design Choice Ablations

To analyze the effect of the LLM in our data curation pipeline, we create a version of training data that uses the SFX metadata tags

directly to retrieve visual frames via CLIP (as opposed to using natural language sentences generated from the tags using an LLM). We train our best performing architecture on this data and evaluate on the same test sets detailed previously. This version of the model performs similarly on PSE-Test and ASFX-Test but about 4pp worse on AVE-Test for Category Precision@10. This shows that using the LLM in our pipeline leads to overall better quantitative performance. Qualitatively, we observed that the usage of the LLM creates more realistic, action-focused audio-visual pairs than when tags are used directly.

To further explore the impact of the data pairing constraint on downstream retrieval performance, we train our best performing model design on six variants of the auto-curated training data. Each training set variant is generated with a different pairing constraint N ranging from 1 to ∞ (no constraint). The results in Table 2 show that for the scenario of matching HQ SFX to production-grade video, represented by PSE-Test and ASFX-Test, we can actually relax the pairing constraint, with optimal results for N between 5-10, suggesting we can sacrifice some diversity in the training data for the sake of matching more SFX to their top ranked video frame. Conversely, if our goal is to generalize to in-the-wild noisy data, represented by AVE-Test, it is preferable to maintain a strict pairing constraint, indicating diversity is more important for generalizing to these data than maximizing pairing quality.

Pairing Limit	PSE-Val	PSE-Test	ASFX-Test	AVE-Test
1	0.37	0.37	0.28	0.68
2	0.39	0.38	0.29	0.68
5	0.40	0.40	0.29	0.65
10	0.39	0.39	0.30	0.64
100	0.36	0.37	0.30	0.62
∞	0.35	0.36	0.28	0.61

Table 2: Ablation over the pairing constraint in the automatic data-curation pipeline, using PANNS fine-tuned architecture. P@10.

5. CONCLUSION

In this work we proposed a scalable automatic data curation pipeline for creating HQ audio-visual data, by leveraging powerful language and vision-language models. We showed that using the resulting data to train a contrastive multimodal model significantly outperforms baselines on the task of SFX retrieval given an input visual query, both quantitatively and qualitatively. Furthermore, we showed that our method generalizes not only to unseen data from the same setup, but also to “in-the-wild” noisy videos whereas the baselines, which are trained on noisy videos, fail to generalize to the HQ data. For future work we plan to extend our system to utilize video clips as queries instead of single video frames, which will enable us to leverage the motion cues in the video in addition to the semantic cues contained in a static frame for SFX retrieval.

6. REFERENCES

- [1] V. T. Ament, *The Foley grail: The art of performing sound for film, games, and animation*. Routledge, 2021.
- [2] A. De Götzen, E. Sikström, F. Grani, and S. Serafin, “Real, foley or synthetic? an evaluation of everyday walking sounds,” *Proceedings of Sound and Music Computing Conference (SMC)*, 2013.
- [3] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, “Multi-modal deep learning,” in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.
- [4] D. Ramachandram and G. W. Taylor, “Deep multimodal learning: A survey on recent advances and trends,” *IEEE signal processing magazine*, vol. 34, no. 6, pp. 96–108, 2017.
- [5] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [6] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [7] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon, “A survey on contrastive self-supervised learning,” *Technologies*, vol. 9, no. 1, 2021. [Online]. Available: <https://www.mdpi.com/2227-7080/9/1/2>
- [8] K. Grauman et al., “Ego4d: Around the World in 3,000 Hours of Ego-centric Video,” in *IEEE/CVF Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [9] R. Hebbar, D. Bose, K. Somanandpalli, V. Vijai, and S. Narayanan, “A dataset for audio-visual sound event detection in movies,” *arXiv preprint arXiv:2302.07315*, 2023.
- [10] Y. Tian, J. Shi, B. Li, Z. Duan, and C. Xu, “Audio-visual event localization in unconstrained videos,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 247–263.
- [11] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, “Vggssound: A large-scale audio-visual dataset,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 721–725.
- [12] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi, et al., “Musiclm: Generating music from text,” *arXiv preprint arXiv:2301.11325*, 2023.
- [13] H.-H. Wu, P. Seetharaman, K. Kumar, and J. P. Bello, “Wav2clip: Learning robust audio representations from clip,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4563–4567.
- [14] A. Guzhov, F. Raue, J. Hees, and A. Dengel, “Audioclip: Extending clip to image, text and audio,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 976–980.
- [15] 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,” *arXiv preprint arXiv:2211.06687*, 2022.
- [16] R. Zellers, J. Lu, X. Lu, Y. Yu, Y. Zhao, M. Salehi, A. Kusupati, J. Hessel, A. Farhadi, and Y. Choi, “Merlot reserve: Neural script knowledge through vision and language and sound,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 16375–16387.
- [17] S. Krishnamurthy, “Learning self-supervised audio-visual representations for sound recommendations,” in *Advances in Visual Computing: 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part II*. Springer, 2021, pp. 124–138.
- [18] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [19] A. Nagrani, P. H. Seo, B. Seybold, A. Hauth, S. Manen, C. Sun, and C. Schmid, “Learning audio-video modalities from image captions,” in *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV*. Springer, 2022, pp. 407–426.
- [20] T. Hayes, S. Zhang, X. Yin, G. Pang, S. Sheng, H. Yang, S. Ge, Q. Hu, and D. Parikh, “Mugen: A playground for video-audio-text multimodal understanding and generation,” in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 431–449.
- [21] L. Prétet, G. Richard, and G. Peeters, “Cross-modal music-video recommendation: A study of design choices,” in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–9.
- [22] D. Surís, C. Vondrick, B. Russell, and J. Salamon, “It’s time for artistic correspondence in music and video,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 10564–10574.
- [23] D. McKee, J. Salamon, J. Sivic, and B. Russell, “Language-guided music recommendation for video via prompt analogies,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 14784–14793.
- [24] D. C.-E. Lin, A. Germanidis, C. Valenzuela, Y. Shi, and N. Martelaro, “Soundify: Matching sound effects to video,” in *5th NeurIPS Workshop on Machine Learning for Creativity and Design*, 2021.
- [25] 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.
- [26] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, “Learning to prompt for vision-language models,” *International Journal of Computer Vision*, vol. 130, no. 9, pp. 2337–2348, 2022.
- [27] G. Yong, K. Jeon, D. Gil, and G. Lee, “Prompt engineering for zero-shot and few-shot defect detection and classification using a visual-language pretrained model,” *Computer-Aided Civil and Infrastructure Engineering*, 2022.
- [28] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al., “Bloom: A 176b-parameter open-access multilingual language model,” *arXiv preprint arXiv:2211.05100*, 2022.
- [29] A. Guzhov, F. Raue, J. Hees, and A. Dengel, “Esresnet: Environmental sound classification based on visual domain models,” in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 4933–4940.
- [30] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition (CVPR)*, 2009, pp. 248–255.
- [31] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumley, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.