

Audio Retrieval with Natural Language Queries

Andreea-Maria Oncescu^{1*} A. Sophia Koepke^{2*} João F. Henriques¹
 Zeynep Akata² Samuel Albanie¹

¹VGG, University of Oxford ²EML, University of Tübingen

<https://www.robots.ox.ac.uk/~vgg/research/audio-retrieval/>
 oncescu@robots.ox.ac.uk, a-sophia.koepke@uni-tuebingen.de

Abstract

We consider the task of retrieving audio using free-form natural language queries. To study this problem, which has received limited attention in the existing literature, we introduce challenging new benchmarks for text-based audio retrieval using text annotations sourced from the AUDIOCAPS and CLOTHO datasets. We then employ these benchmarks to establish baselines for cross-modal audio retrieval, where we demonstrate the benefits of pre-training on diverse audio tasks. We hope that our benchmarks will inspire further research into cross-modal text-based audio retrieval with free-form text queries.

1. Introduction

The large amounts of data being generated in recent years have prompted a real need to search evergrowing databases. A search query typically comprises natural language (text), which allows expressing virtually any concept. Spanning multiple modalities, different retrieval strategies were developed for content as diverse as text (including webpages and books), images [1], and videos [2, 3]. Surprisingly, while search engines currently exist for these modalities (e.g. Google, Flickr and YouTube, respectively), unstructured audio is not accessible in the same way. The aim of this paper is to address this gap.

It is important to distinguish content-based retrieval from that based on meta-data, such as the title of a video or song or an audio tag. Meta-data retrieval is feasible for manually-curated databases such as song or movie catalogs. However, content-based retrieval is more important in user-generated data, which often has little structure or meta-data. There are methods to search for audio which matches an audio query [4, 5], but satisfying the requirement to input an example audio query can be difficult for a human (e.g. making convincing frog sounds is difficult). We, on the other hand, propose a framework which enables the querying of a sound database using detailed free-form natural language descriptions of the desired sound (e.g. “A man talking as music is playing followed by a frog croaking.”). This enables the retrieval of audio data which matches the temporal sequence of events in the query instead of just a single class tag. Furthermore, natural language queries are a familiar user interface widely used in current search engines. Therefore, our proposed audio retrieval with free-form text queries could be a first step towards more natural and flexible audio-only search.

Text-based audio retrieval could also be beneficial for video retrieval. The majority of recent works that address the text-based video retrieval task focuses heavily on the visual and text domains [1, 2, 6, 7]. Since audio and visual information inherently have natural semantic alignment for a significant portion of video data, text-based audio retrieval could also be used for querying video databases by only considering the audio stream

of the video data. This would allow for video retrieval in the audio domain at reduced computational cost for cases in which audio and visual information correspond, with applications for low-power IoT devices, such as microphones in natural habitats, of particular interest for conservation and biology. Historical archives with extensive sound collections, such as the BBC Sound Effects Archive, would be easier to search, facilitating historical research and public access. Furthermore, text-based retrieval could enable appealing creative applications, such as automatically finding sounds which correspond to input text. This could be especially useful given the growing popularity of audio podcasts and audiobooks which are often supplemented with (background) sounds that match their content.

We propose two new benchmarks for text-based audio retrieval, based on the AUDIOCAPS [8] and CLOTHO [9] datasets. The former consists of a subset of 10-second audio clips from the AudioSet dataset [10] with additional human-written audio captions, while the latter contains audio captions for sounds sourced from the Freesound platform [11], varying between 15 and 30 seconds in duration. In contrast to sound event class labels, audio captions contain detailed information about the sounds. A user searching for a particular sound would usually describe the sound using text similar to an audio caption. AUDIOCAPS [8] and CLOTHO [9] allow to leverage the matching audio-caption pairs to train text-based audio retrieval frameworks. To establish baselines for this task, we adapt existing video retrieval frameworks for audio retrieval. We employ multiple pre-trained audio expert networks and show that using an ensemble of audio experts improves audio retrieval.

In summary, we make three contributions: (1) we introduce two new benchmarks for free-form text-based audio retrieval—to the best of our knowledge, these represent the first public benchmarks for this task; (2) we provide baseline performances with existing multi-modal video retrieval models that we adapt to text-based audio retrieval; (3) we demonstrate the benefits of combining multiple pre-training datasets.

2. Related Work

Our work relates to several themes in the literature: *sound event recognition*, *audio captioning*, *audio-based retrieval*, *text-based video retrieval* and *text-domain audio retrieval*. We discuss each of these next.

Sound event recognition. The task of sound event recognition consists of matching audio information with a text label. Research in this area has been driven by challenges, such as DCASE [12, 13], and by the collection of sound event datasets. These include TUT Acoustic scenes [14], CHIME-Home [15], ESC-50 [16], FSDKaggle [17], and AudioSet [10]. Of relevance to our approach, a number of prior works have employed deep learning for audio comprehension [18, 19, 20, 21, 22] on the AudioSet dataset. However, instead of considering the task

* Authors contributed equally.

of audio recognition, we focus on the task of retrieval with natural language queries.

Audio captioning. Audio captioning consists of generating a text description for a sound [23]. This requires a more detailed understanding of the sound than simply mapping the sound to a set of labels (sound event recognition). Recently, several audio captioning datasets have been introduced, such as CLOTHO [9] which was used in the DCASE automated audio captioning challenge 2020 [24], Audio Caption [25], and AUDIOCAPS [8]. Multiple works have addressed automatic audio captioning on the AUDIOCAPS dataset [26, 27, 28]. In this work, we use the AUDIOCAPS and CLOTHO datasets for cross-modal retrieval.

Audio-based retrieval. Multiple content-based audio retrieval frameworks, in particular query by example methods, leverage the similarity of sound features that represent different aspects of sounds (e.g. pitch, or loudness) [29, 30, 31, 5]. More recently, [4] use a siamese neural network framework to learn to encode semantically similar audio close together in the embedding space. [32] address multimedia event detection using only audio data, while [33] tackle near-duplicate video retrieval by audio retrieval. These are purely audio-based methods that are applied to video datasets, but without using visual information. [34] propose a two-step approach for video retrieval which uses audio (coarse) and visual (fine) information together.

Text-based video retrieval. More closely related to our work, a number of methods showed that embedding video and text jointly into a shared space (such that their similarity can be computed efficiently) is an effective approach [1, 3, 2, 6, 35, 7]. One particular trend has been to combine cues from several “experts”—pre-trained models that specialise in different tasks (such as object recognition, action classification etc.) to inform the joint embedding. In this work we propose to adapt two such methods: the Mixture of Embedded Experts method of [2] and the Collaborative Experts model of [6], by repurposing them for the task of audio retrieval (described in more detail in Sec. 3).

Text-domain audio retrieval. Methods that retrieve audio by matching associated text, such as meta-data or sound event labels, have the implicit assumption that the text is relevant [36]. [37] is an early work which links audio and text representations in hierarchical semantic and acoustic spaces. [38] propose a text-based sound retrieval framework which uses single-word audio tags as queries rather than caption-like natural language. The creative approach of [39] learns to align visual, audio, and text representations to enable cross-modal retrieval. Their framework is trained with captioned images and paired image-sound data (sourced from videos) and evaluated using the soundtrack of captioned videos. More recently, [40] use a siamese network to learn a shared latent text and sound space for cross-modal retrieval. While they use class labels as text labels, we study unconstrained text descriptions as queries.

3. Methods, Datasets and Benchmarks

In this section, we first formulate the problem of audio retrieval with natural language queries. Next, we describe two cross-modal embedding methods that we adapt for the task of audio retrieval. Finally, we describe the four datasets used in our experimental study, and the two benchmarks that we propose to evaluate performance on the audio retrieval task.

Problem formulation. Given a natural language query (i.e. a written description of an audio event to be retrieved) and a pool of audio samples, the objective of text-to-audio (abbreviated to $t2a$) retrieval is to rank the audio samples according to their

similarity to the query. We also consider the converse $a2t$ task, viz. retrieving text with audio queries.

Methods. To tackle the problem of audio retrieval, we propose to learn cross-modal embeddings. Specifically, given a collection of N audio samples with corresponding text descriptions, $\{(a_i, t_i) : i \in \{1, \dots, N\}\}$ we aim to learn embedding functions, ψ_a and ψ_t , that project each audio sample a_i and text sample t_i into a shared space, such that $\psi_a(a_i)$ and $\psi_t(t_i)$ are close when the text describes the audio, and far apart otherwise. Writing s_{ij} for the cosine similarity of the audio embedding $\psi_a(a_i)$ and the text embedding $\psi_t(t_j)$, we learn the embedding functions by minimising a contrastive ranking loss [41]:

$$\mathcal{L} = \frac{1}{B} \sum_{i=1, j \neq i}^B [m + s_{ij} - s_{ii}]_+ + [m + s_{ji} - s_{ii}]_+ \quad (1)$$

where B denotes the batch size, m the *margin* (set as a hyperparameter), and $[\cdot]_+ = \max(\cdot, 0)$ the hinge function.

We consider two frameworks for learning such embedding functions ψ_a and ψ_t : Mixture-of-Embedded Experts (MoEE) [2] and Collaborative-Experts (CE) [6]. MoEE, originally designed for video-text retrieval, constructs its video encoder from a collection of “experts” (features extracted from networks pre-trained for object recognition, action classification, sound classification, etc.) which are computed for each video, aggregated along their temporal dimension with NetVLAD [42], projected to a lower dimension via a self-gated linear map and then L2-normalised. The text encoder first embeds each word token with word2vec [43] and aggregates the results with NetVLAD [42]. The result is projected by a sequence of self-gated linear maps (one for each expert) into a shared embedding space with the outputs of the video encoder. Finally, a scalar-weighted sum of the embedded experts in each joint space is used to compute the overall cosine similarity between the video and text (see [2] for more details). CE adopts the same text encoder as MoEE and similarly makes use of multiple video experts. However, rather than projecting them directly into independent spaces against the embedded text, CE first applies a “collaborative gating” mechanism, which filters each expert with an element-wise attention mask that is generated with a small MLP that ingests all pairwise combinations of experts (see [6] for further details).

To adapt these frameworks for audio retrieval, we use the same text encoder structures ψ_t , and build audio encoders ψ_a by mimicking the structure of their video encoders, replacing the “video experts” with “audio experts” (described in Sec. 4).

Datasets. As the primary focus of our work, we study two *audio-centric datasets*—these are datasets which comprise audio streams (sometimes with accompanying visual streams) paired with natural language descriptions that focus explicitly on the content of the audio track. To explore differences between audio retrieval and video retrieval, we also consider two *visual-centric datasets*, that comprise audio and video streams paired with natural language which focus primarily (though not always exclusively) on the content of the video stream. Details of the four datasets we employ are given next.

1. **AUDIOCAPS [8] (audio-centric)** is a dataset of sounds with event descriptions that was introduced for the task of audio captioning, with sounds sourced from the AudioSet dataset [10]. Annotators were provided the audio tracks together with category hints (and with additional video hints if needed). We use a subset of the data, excluding a small number of samples for which either: (i) the YouTube-hosted source video is no longer

Dataset	Anno. Focus	Pool	Text \Rightarrow Audio		Audio \Rightarrow Text	
			R@1 \uparrow	R@10 \uparrow	R@1 \uparrow	R@10 \uparrow
AUDIOCAPS [8]	audio	816	18.0 \pm 0.2	62.0 \pm 0.5	21.0 \pm 0.8	62.7 \pm 1.6
CLOTHO [9]	audio	1045	4.0 \pm 0.2	25.4 \pm 0.5	4.8 \pm 0.4	25.8 \pm 1.7
ACTNETCAPS [46]	visual	4917	1.5 \pm 0.1	9.2 \pm 0.3	1.4 \pm 0.1	8.5 \pm 0.3
QUERYD [47]	visual	1954	3.7 \pm 0.2	17.3 \pm 0.6	3.8 \pm 0.2	16.8 \pm 0.2

Table 1: **Audio retrieval on audio-centric and visual-centric datasets.** Performance is strongest on the audio-centric AUDIOCAPS dataset and weakest on the visual-centric ACTIVITYNET-CAPTIONS (ACTNETCAPS) dataset.

available, (ii) the source video overlaps with the training partition of the VGGSound dataset [44]. Filtering to exclude samples affected by either issue leads to a dataset with 49,291 training, 428 validation and 816 test samples.¹

2. CLOTHO [9] (audio-centric) is a dataset of described sounds that was also introduced for the task of audio captioning, with sounds sourced from the Freesound platform [11]. During labelling, annotators only had access to the audio stream (i.e. no visual stream or meta tags) to avoid their reliance on contextual information for removing ambiguity that could not be resolved from the audio stream alone. The descriptions are filtered to exclude transcribed speech. The publicly available version of the dataset includes a dev set of 2893 audio samples and an evaluation set of 1045 audio samples. Every audio sample is accompanied by 5 written descriptions. We used a random split of the dev set into a training and validation set with 2,314 and 579 samples, respectively.

3. ACTIVITYNET-CAPTIONS [46] (visual-centric) consists of videos sourced from YouTube and annotated with dense event descriptions. It allocates 10,009 videos for training and 4,917 videos for testing (we use the public val_1 split provided by [46]). For this dataset, descriptions also tend to focus on the visual stream.

4. QUERYD [47] (visual-centric) is a dataset of described videos sourced from YouTube and the YouDescribe [48] platform. It is accompanied by *audio descriptions* that are provided with the explicit aim of conveying the video content to visually impaired users. Therefore, the provided descriptions focus heavily on the visual modality. We use the version of the dataset comprising trimmed videos with 9,114 training, 1,952 validation, and 1,954 test samples.

Benchmark. To facilitate the study of the text-based audio retrieval task, we propose to re-purpose the two *audio-centric* datasets described above, AUDIOCAPS and CLOTHO, to provide benchmarks for text-based audio retrieval. This approach is inspired by precedents in the vision and language communities, where datasets, such as [49], that were originally introduced for the task of video captioning, have become popular benchmarks for text-based video retrieval [2, 35, 6, 7].

4. Experiments

In this section, we first compare text-based audio retrieval and audio-based text retrieval performance on audio-centric and visual-centric datasets. Next, we perform an ablation study on the contributions of different experts and present our baselines for the proposed AUDIOCAPS and CLOTHO benchmarks. Finally we perform experiments to assess the influence of pre-training and training dataset size, and give qualitative examples of retrieval results. Throughout the section, we use the standard retrieval metrics: recall at rank k (R@ k) which measures the percentage of targets retrieved within the top k ranked results (higher is better) and report the mean and standard deviation of

¹We make the sample list publicly available at the project page [45].

Expert	Text \Rightarrow Audio/Video		Audio/Video \Rightarrow Text	
	R@1 \uparrow	R@10 \uparrow	R@1 \uparrow	R@10 \uparrow
Visual experts only				
Scene	6.1 \pm 0.4	35.8 \pm 0.6	6.5 \pm 0.8	31.3 \pm 1.6
Inst	7.7 \pm 0.2	46.7 \pm 1.3	9.8 \pm 0.9	40.6 \pm 0.7
R2PID	8.2 \pm 0.5	44.7 \pm 0.9	10.3 \pm 0.4	41.8 \pm 3.1
Scene + Inst	8.7 \pm 0.5	47.4 \pm 0.5	10.6 \pm 0.6	41.4 \pm 1.5
Scene + R2PID	8.8 \pm 0.1	46.8 \pm 0.1	11.0 \pm 0.6	45.1 \pm 1.7
R2PID + Inst (CE-Visual)	10.1\pm0.2	49.6\pm1.1	12.1\pm0.4	46.1\pm1.3
Audio experts only				
VGGish	18.0 \pm 0.2	62.0 \pm 0.5	21.0 \pm 0.8	62.7 \pm 1.6
VGGSound	20.5 \pm 0.6	67.0 \pm 1.0	24.6 \pm 0.9	70.4 \pm 0.4
VGGish + VGGSound (CE-Audio)	23.1\pm0.8	70.7\pm0.7	25.1\pm0.9	73.2\pm1.6
Audio and visual experts				
CE-Visual + VGGish	23.9 \pm 0.7	74.4 \pm 0.2	29.0 \pm 2.0	77.2 \pm 1.9
CE-Visual + VGGSound	27.4 \pm 0.7	78.2 \pm 0.3	34.0 \pm 1.5	82.5 \pm 1.2
CE-Visual + CE-Audio	28.1\pm0.6	79.0\pm0.5	33.7\pm1.6	83.7\pm0.4

Table 2: **The influence of different experts on AUDIOCAPS.** A comparison of audio and visual experts (applied to the video from which the audio was sourced) using CE [6]. We observe that audio features are significantly more effective than visual features (which nevertheless provide some complementary signal as can be seen when jointly using audio and visual features).

three different randomly seeded runs.

Implementation details. To encode the audio signal, we use two pre-trained audio feature extractors: VGGish pre-trained for audio classification on YouTube8M [50], and ResNet18 pre-trained for audio classification on the VGGSound dataset [44].² All audio retrieval models were trained with a batch size of 128 using the contrastive ranking loss (Eqn. 1), with m set to 0.2. The learning rate was set to 0.01 with a weight decay of 0.001. Further details on optimisation, pre-trained experts (for audio and visual modalities) and hyperparameter choices are provided in the supplementary material.

Audio-centric vs. visual-centric queries. We first investigate audio retrieval with audio-centric and visual-centric queries. For this experiment, we use CE with a single expert (VGGish audio features). In Table 1, we observe that performance is strongest on the audio-centric AUDIOCAPS dataset and weakest overall on the ACTIVITYNET-CAPTIONS dataset. We note that the CLOTHO dataset is particularly challenging, with performance weaker (accounting for pool size) than the visual-centric QUERYD dataset. We hypothesise that this is for two reasons: (1) the significantly smaller training set size of CLOTHO, compared to all other datasets, (2) CLOTHO was constructed by selecting audio samples that were as diverse as possible, making it a potentially more difficult benchmark. We note, however, that the QUERYD experiments suggest that computationally efficient video retrieval using only the audio stream can still be obtained, although at a lower accuracy.

Ablation study. We next conduct an ablation study to investigate the effectiveness of different experts for audio retrieval on the AUDIOCAPS dataset. We present the results in Tab. 2, where we observe that audio experts significantly outperform visual experts (pre-trained models for visual tasks like scene classification, which we compute from the video from which the audio was sourced). We note that the combination of audio and visual experts performs strongest overall, suggesting that the audio-centric queries contain information that is more accessible from the visual modality. The strongest audio-only retrieval is achieved by combining VGGish and VGGSound features—we therefore adopt this setting for the remaining experiments.

Additionally, we experimented with using speech features (word2vec [43] encodings of speech-to-text transcriptions [51] of the audio stream). However, this did not improve the re-

²Since the AudioCaps test set is a subset of the AudioSet training set (unbalanced), we do not use audio experts pre-trained on AudioSet.


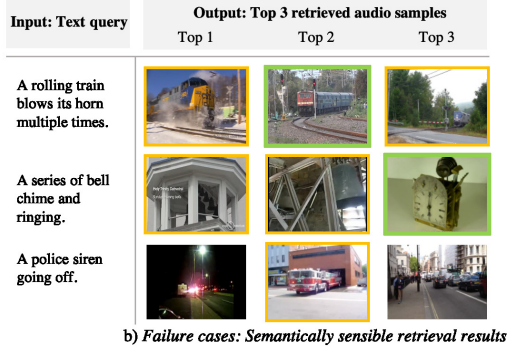



Figure 1: **Qualitative results.** Text-based audio retrieval results on AUDIOCAPS using CE with VGGish and VGGSound features. For an input text query, we visualise the top 3 retrieved audio samples using a video frame from the corresponding videos (the audio can be heard at the project webpage [45]). We mark the audio samples which correspond to the query with green boxes. Successful retrievals are shown in a), failures in b). Note, in particular, the examples in b), where the model’s top 1 retrieved audio is not the correct one, but the retrieved results nevertheless sound reasonable (visually convincing results are marked with yellow boxes).

Figure 2: **The influence of training data scale.** CE retrieval performance on AUDIOCAPS with different proportions of the available AUDIOCAPS training data.

Benchmark	Pre-training	Text \Rightarrow Audio		Audio \Rightarrow Text	
		R@1 \uparrow	R@10 \uparrow	R@1 \uparrow	R@10 \uparrow
AUDIOCAPS					
MoEE	None	22.5 \pm 0.3	69.5 \pm 0.9	25.1 \pm 0.8	72.9 \pm 1.2
CE	None	23.1 \pm 0.8	70.7 \pm 0.7	25.1 \pm 0.9	73.2 \pm 1.6
CLOTHO					
MoEE	None	6.0 \pm 0.1	32.3 \pm 0.3	7.2 \pm 0.5	33.2 \pm 1.1
CE	None	6.7 \pm 0.4	33.2 \pm 0.3	7.1 \pm 0.3	34.6 \pm 0.5
CLOTHO					
MoEE	AUDIOCAPS	8.6 \pm 0.4	39.3 \pm 0.7	10.0 \pm 0.3	40.1 \pm 1.3
CE	AUDIOCAPS	9.6 \pm 0.3	40.1 \pm 0.7	10.7 \pm 0.6	40.8 \pm 1.4

Table 3: **Audio retrieval benchmarks.** Text-audio and audio-text retrieval results for MoEE [2] and CE [6] with VGGish and VGGSound features on the proposed AUDIOCAPS and CLOTHO retrieval benchmarks. Pre-training on AUDIOCAPS improves the performance on CLOTHO.

trieval performance. Upon further investigation, we found that the audio captions and the spoken words in the AUDIOCAPS dataset do not have a significant overlap (corresponding to a METEOR [52] score of only 0.03).

Benchmark results. Incorporating the strongest combination of experts from the ablation study, we report our final baselines for text-audio and audio-text retrieval for two methods on the AUDIOCAPS and CLOTHO datasets in Tab. 3, where we observe that CE slightly outperforms MoEE. We also report the performance after pre-training models for retrieval on AUDIOCAPS and then fine-tuning on CLOTHO (lower part of Tab. 3). Here, we observe that pre-training on the audio-centric AUDIOCAPS brings a boost to both, CE and MoEE. Furthermore, we explored pre-training on CLOTHO and fine-tuning on the AUDIOCAPS dataset, but found negligible change in performance (likely due to the fact that the AUDIOCAPS training set is significantly larger than that of CLOTHO).

Qualitative results. The qualitative results in Figure 1 show examples in which the CE model with VGGish and VGGSound expert modalities (CE-Audio) is used to retrieve audio with natural language queries. The retrieved results mostly contain audio that is semantically similar to the input text queries. Observed failure cases arise from audio samples sounding very similar to one another despite being semantically distinct (e.g. the siren of a fire engine sounds very similar to a police siren).

Influence of scale. Finally, we present experiments using dif-

ferent proportions of the AUDIOCAPS dataset for training CE-Audio in Fig. 2. We observe that as more training data becomes available, retrieval performance increases monotonically. We also observe that there is still clear room for improvement in terms of retrieval results simply by collecting additional training data, motivating further dataset construction work to support future research on this important task.

5. Conclusion

We introduced a new benchmark for natural language based audio retrieval, and provided baseline results by adapting strong multi-modal video retrieval methods. Our results show that these methods are relatively well-suited for the audio retrieval task, however there is room for improvement, as expected for an under-explored problem. We hope that our proposed benchmarks will facilitate the development of future audio search engines, and make this large fraction of the world’s produced media available for public use.

Acknowledgements. AMO was supported by an EPSRC DTA Studentship. ASK and ZA were supported by the ERC (853489 - DEXIM) and by the DFG (2064/1 – Project number 390727645). JFH is supported by the Royal Academy of Engineering (RF\201819\18\163). SA was supported by EPSRC EP/T028572/1 Visual AI.

6. References

- [1] J. Dong *et al.*, “Word2visualvec: Image and video to sentence matching by visual feature prediction,” *arXiv:1604.06838*, 2016.
- [2] A. Miech *et al.*, “Learning a text-video embedding from incomplete and heterogeneous data,” *arXiv:1804.02516*, 2018.
- [3] N. C. Mithun *et al.*, “Learning joint embedding with multimodal cues for cross-modal video-text retrieval,” in *Proc. ACM ICMR*, 2018.
- [4] P. Manocha *et al.*, “Content-based representations of audio using siamese neural networks,” in *Proc. ICASSP*, 2018.
- [5] I. Lallemand *et al.*, “Content-based retrieval of environmental sounds by multiresolution analysis,” in *Proc. SMC*, 2012.
- [6] Y. Liu *et al.*, “Use what you have: Video retrieval using representations from collaborative experts,” in *Proc. BMVC*, 2019.

- [7] V. Gabeur *et al.*, “Multi-modal transformer for video retrieval,” in *Proc. ECCV*, 2020.
- [8] C. D. Kim *et al.*, “Audiocaps: Generating captions for audios in the wild,” in *Proc. NACCL*, 2019.
- [9] K. Drossos *et al.*, “Clotho: An audio captioning dataset,” in *Proc. ICASSP*, 2020.
- [10] J. F. Gemmeke *et al.*, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. ICASSP*, 2017.
- [11] F. Font *et al.*, “Freesound technical demo,” in *Proc. ACM Multimedia*, 2013.
- [12] D. Stowell *et al.*, “Detection and classification of acoustic scenes and events,” *IEEE Transactions on Multimedia*, 2015.
- [13] A. Mesaros *et al.*, “DCASE 2017 challenge setup: Tasks, datasets and baseline system,” in *DCASE 2017*, 2017.
- [14] A. Mesaros *et al.*, “TUT Acoustic scenes 2017, Dev. dataset,” Mar 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.400515>
- [15] P. Foster *et al.*, “Chime-home: A dataset for sound source recognition in a domestic environment,” in *IEEE WASPAA*, 2015.
- [16] K. J. Piczak, “ESC: Dataset for environmental sound classification,” in *Proc. ACM Multimedia*, 2015.
- [17] E. Fonseca *et al.*, “Audio tagging with noisy labels and minimal supervision,” *arXiv:1906.02975*, 2019.
- [18] Q. Kong *et al.*, “Audio set classification with attention model: A probabilistic perspective,” in *Proc. ICASSP*, 2018.
- [19] C. Yu *et al.*, “Multi-level attention model for weakly supervised audio classification,” *arXiv:1803.02353*, 2018.
- [20] Q. Kong *et al.*, “Weakly labelled audioset tagging with attention neural networks,” *IEEE/ACM TASLP*, 2019.
- [21] L. Ford *et al.*, “A deep residual network for large-scale acoustic scene analysis,” in *INTERSPEECH*, 2019.
- [22] Q. Kong *et al.*, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM TASLP*, 2020.
- [23] K. Drossos *et al.*, “Automated audio captioning with recurrent neural networks,” in *IEEE WASPAA*, 2017.
- [24] “DCASE2020 challenge task 6: Automated audio captioning,” 2020. [Online]. Available: <http://dcase.community/challenge2020/task-automatic-audio-captioning>
- [25] M. Wu *et al.*, “Audio caption: Listen and tell,” in *Proc. ICASSP*, 2019.
- [26] Y. Koizumi *et al.*, “Audio captioning using pre-trained large-scale language model guided by audio-based similar caption retrieval,” *arXiv:2012.07331*, 2020.
- [27] X. Xu *et al.*, “Investigating local and global information for automated audio captioning with transfer learning,” *arXiv:2102.11457*, 2021.
- [28] A. Ö. Eren and M. Sert, “Audio captioning based on combined audio and semantic embeddings,” in *IEEE ISM*, 2020.
- [29] J. T. Foote, “Content-based retrieval of music and audio,” in *Multimedia Storage and Archiving Systems II*. International Society for Optics and Photonics, 1997.
- [30] E. Wold *et al.*, “Content-based classification, search, and retrieval of audio,” *IEEE Multimedia*, 1996.
- [31] M. Helén and T. Virtanen, “Query by example of audio signals using euclidean distance between gaussian mixture models,” in *Proc. ICASSP*, 2007.
- [32] Q. Jin *et al.*, “Event-based video retrieval using audio,” in *Proc. ISCA*, 2012.
- [33] P. Avgoustinakis *et al.*, “Audio-based near-duplicate video retrieval with audio similarity learning,” *arXiv:2010.08737*, 2020.
- [34] S. Hou and S. Zhou, “Audio-visual-based query by example video retrieval,” *Mathematical Problems in Engineering*, 2013.
- [35] M. Wray *et al.*, “Fine-grained action retrieval through multiple parts-of-speech embeddings,” in *Proc. ICCV*, 2019.
- [36] B. Elizalde *et al.*, “Nels-never-ending learner of sounds,” *arXiv:1801.05544*, 2018.
- [37] M. Slaney, “Semantic-audio retrieval,” in *Proc. ICASSP*, 2002.
- [38] G. Chechik *et al.*, “Large-scale content-based audio retrieval from text queries,” in *Proc. ACM ICMIR*, 2008.
- [39] Y. Aytar *et al.*, “See, hear, and read: Deep aligned representations,” *arXiv:1706.00932*, 2017.
- [40] B. Elizalde *et al.*, “Cross modal audio search and retrieval with joint embeddings based on text and audio,” in *Proc. ICASSP*, 2019.
- [41] R. Socher *et al.*, “Grounded compositional semantics for finding and describing images with sentences,” *Transactions of the Association for Computational Linguistics*, 2014.
- [42] R. Arandjelovic *et al.*, “Netvlad: Cnn architecture for weakly supervised place recognition,” in *Proc. CVPR*, 2016.
- [43] T. Mikolov *et al.*, “Efficient estimation of word representations in vector space,” *arXiv:1301.3781*, 2013.
- [44] H. Chen *et al.*, “Vggsound: A large-scale audio-visual dataset,” in *Proc. ICASSP*, 2020.
- [45] “Project page.” [Online]. Available: <https://www.robots.ox.ac.uk/~vgg/research/audio-retrieval/>
- [46] R. Krishna *et al.*, “Dense-captioning events in videos,” in *Proc. ICCV*, 2017.
- [47] A.-M. Oncescu *et al.*, “QuerYD: A video dataset with high-quality textual and audio narrations,” in *Proc. ICASSP*, 2021.
- [48] Video Description Research and Development Center, “Youdescribe,” 2013. [Online]. Available: <http://youdescribe.ski.org>
- [49] J. Xu *et al.*, “Msr-vtt: A large video description dataset for bridging video and language,” in *Proc. CVPR*, 2016.
- [50] S. Hershey *et al.*, “CNN architectures for large-scale audio classification,” in *Proc. ICASSP*, 2016.
- [51] Google, “Speech-to-Text API,” <https://cloud.google.com/speech-to-text>.
- [52] S. Banerjee and A. Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. ACL Workshop*, 2005.
- [53] S. Abu-El-Haija *et al.*, “Youtube-8m: A large-scale video classification benchmark,” *arXiv preprint arXiv:1609.08675*, 2016.
- [54] K. He *et al.*, “Deep residual learning for image recognition,” *arXiv preprint arXiv:1512.03385*, 2015.
- [55] R. Xu *et al.*, “Jointly modeling deep video and compositional text to bridge vision and language in a unified framework,” in *AAAI*, 2015.
- [56] D. Mahajan *et al.*, “Exploring the limits of weakly supervised pre-training,” in *Proc. ECCV*, 2018.
- [57] J. Deng *et al.*, “Imagenet: A large-scale hierarchical image database,” in *Proc. CVPR*, 2009.
- [58] G. Huang *et al.*, “Densely connected convolutional networks,” in *Proc. CVPR*, 2017.
- [59] B. Zhou *et al.*, “Places: A 10 million image database for scene recognition,” *IEEE PAMI*, 2017.
- [60] D. Tran *et al.*, “A closer look at spatiotemporal convolutions for action recognition,” in *Proc. CVPR*, 2018.
- [61] D. Ghadiyaram *et al.*, “Large-scale weakly-supervised pre-training for video action recognition,” in *Proc. CVPR*, 2019.
- [62] M. R. Zhang *et al.*, “Lookahead optimizer: k steps forward, 1 step back,” in *NeurIPS*, 2019.
- [63] L. Liu *et al.*, “On the variance of the adaptive learning rate and beyond,” *arXiv preprint arXiv:1908.03265*, 2019.
- [64] “Ranger optimiser.” [Online]. Available: <https://github.com/lessw2020/Ranger-Deep-Learning-Optimizer>

A. Supplementary Material

We give details on the pre-trained experts used, and on optimisation and hyperparameter choices below.

A.1. Pre-trained experts

This section contains information about the audio and visual expert features that were obtained from pre-trained networks.

A.1.1. Audio experts

We use two different audio experts, which we abbreviate as VGGish and VGGSound.

VGGish. These audio features are obtained with a VGGish model [50], trained for audio classification on the YouTube-8M dataset [53]. To produce the input for this model, the audio stream of each video is re-sampled to a 16kHz mono signal, converted to an STFT with a window size of 25ms and a hop size of 10ms with a Hann window, then mapped to a 64 bin log mel spectrogram. Finally, the features are parsed into non-overlapping 0.96s collections of frames (each collection comprises 96 frames, each of 10ms duration), which is mapped to a 128-dimensional feature vector.

VGGSound. These features are extracted using a ResNet-18 model [54] that has been pre-trained on the VGGSound dataset (model H) [44]. We modify the last average pooling layer to aggregate along the frequency dimension, but keep the full temporal dimension. This results in features of dimension $t \times 512$, where t denotes the number of time steps.

A.1.2. Visual experts

We use three different visual experts, abbreviated as Inst, Scene, and R2P1D.

Inst. These features are extracted using a ResNeXt-101 model [55] that has been pre-trained on Instagram hashtags [56] and finetuned on ImageNet [57] for the task of image classification. Features are extracted from frames extracted at 25 fps, where each frame is resized to 224×224 pixels. Embeddings are 2048-dimensional.

Scene. Scene features are extracted from 224×224 pixel centre crops with a DenseNet-161 model [58] pre-trained on Places365 [59]. Embeddings are 2208-dimensional.

R2P1D. Features are extracted with a 34-layer R(2+1)D model [60] trained on IG-65M [61] which processes clips of 8 consecutive 112×112 pixel frames, extracted at 30 fps. Embeddings are 512-dimensional.

A.2. Optimisation details and hyperparameter choices

A.2.1. Optimiser

We used the Lookahead solver [62] in combination with RAdam [63] (implementation by [64]) with an initial learning rate of 0.01 and weight decay of 0.001. We use a learning rate decay for each parameter group with a factor of 0.95 every epoch. Most models were trained for a maximum of 20 epochs whilst the Activity-Net model was trained for 40 epochs. We chose the models that gave the best performance on the geometric mean of R@1, R@5, and R@10.

A.2.2. NetVLAD hyperparameters

For text, we used 20 VLAD clusters and one ghost cluster. For the VGGish and VGGSound features we used 16 VLAD clus-

ters.

A.2.3. Zero-padding hyperparameters

In the main paper, we used zero-padding of variable length experts to a common shape $z = (z_t, z_a, z_v)$. For AUDIOCAPS we have used $z_t = 20$ for text tokens, $z_a = 29$ for VGGish feature tokens, and $z_v = 29$ for the VGGSound features. For CLOTHO we have used zero-padding with $z_t = 21$ for text tokens, $z_a = 31$ for VGGish feature tokens, and $z_v = 95$ for the VGGSound features.

For QuerYD we have used $z_t = 70$ for text tokens and $z_a = 500$ for VGGish feature tokens. Lastly, for Activity-Net we used $z_t = 20$ for text tokens and $z_a = 29$ for VGGish feature tokens.