

ADVANCING NATURAL-LANGUAGE BASED AUDIO RETRIEVAL WITH PASST AND LARGE AUDIO-CAPTION DATA SETS

Paul Primus¹, Khaled Koutini^{1,2}, Gerhard Widmer^{1,2}

¹Institute of Computational Perception (CP-JKU)

²LIT Artificial Intelligence Lab

Johannes Kepler University, Austria

ABSTRACT

This work presents a text-to-audio-retrieval system based on pre-trained text and spectrogram transformers. Our method projects recordings and textual descriptions into a shared audio-caption space in which related examples from different modalities are close. Through a systematic analysis, we examine how each component of the system influences retrieval performance. As a result, we identify two key components that play a crucial role in driving performance: the self-attention-based audio encoder for audio embedding and the utilization of additional human-generated and synthetic data sets during pre-training. We further experimented with augmenting ClothoV2 captions with available keywords to increase their variety; however, this only led to marginal improvements. Our system ranked first in the 2023’s DCASE Challenge, and it outperforms the current state of the art on the ClothoV2 benchmark by 5.6 pp. mAP@10.

Index Terms— Natural-Language-Based Audio Retrieval, PaSST, ChatGPT

1. INTRODUCTION

Natural-language-based audio retrieval revolves around ranking audio recordings based on their relatedness to textual descriptions. Current state-of-the-art methods in this domain are based on the dual-encoder approach which converts both recordings and textual descriptions into high-level representations and then aligns them within a shared audio-caption space. The ranking of candidate audios is carried out by measuring their distance to the textual descriptions in the shared embedding space. The dual-encoder setup has been widely adopted in audio retrieval systems [1–4], because it permits fast ranking and the use of pre-trained audio and text embedding models. CNN architectures pre-trained on AudioSet [5] are the most common audio encoders and large transformer models, such as BERT [6] and RoBERTa [7], are the most popular text encoders. Recently, Mei et al. [8] managed to set a new state-of-the-art performance on ClothoV2 [9] by introducing WavCaps [8], a large dataset with synthetic captions.

In this work, we elaborate on our findings in the context of sub-task 6b of the 2023 DCASE challenge, which is concerned with natural-language-based audio retrieval. Our method is also based on the dual-encoder method but differs from previous methods in three central aspects: firstly, we use an audio spectrogram transformer called PaSST [10] instead of CNN14 [5] for audio embedding because PaSST outperforms CNN14 on AudioSet [11] and other relevant audio benchmarks [12]. PaSST employs Patchout [10] during training, which increases training speed and memory efficiency

while at the same time acting as a regularizer. Swapping CNN14 for PaSST improves the retrieval performance significantly, as we will show in Section 3.3. Secondly, we pre-train our models on AudioCaps [13] and WavCaps [8], two large audio captioning datasets with human-generated and synthetic captions, to alleviate the data scarcity problem. Using these large datasets to pre-train the models results in significantly better retrieval performance, as we will show in Section 3.1. And thirdly, we take advantage of the available metadata associated with each audio recording and augmented the captions during training by generating additional captions using the available keywords and the GPT3.5-turbo API, also known as ChatGPT. This augmentation method successfully reduces overfitting during fine-tuning but only results in a minor performance improvement (see Section 3.5). Our implementation and the keyword-augmented captions are available in our GitHub repository¹.

2. METHOD

Our model uses separate audio and caption embedding networks, denoted as $\phi_a(\cdot)$ and $\phi_c(\cdot)$, respectively, to embed pairs of spectrograms and descriptions $\{(a_i, c_i)\}_{i=1}^N$ into a shared D -dimensional space such that representations of matching audio-caption pairs are close. This behavior is achieved by contrastive training, which makes the embeddings of matching audio-caption pairs (a_i, c_i) more similar while pushing the representations of mismatching pairs $(a_i, c_{j;j \neq i})$ apart. The agreement between audio a_i and description c_j is estimated via the normalized dot product in the shared embedding space:

$$C_{ij} = \frac{\phi_a(a_i)^T \cdot \phi_c(c_j)}{\|\phi_a(a_i)\|^2 \|\phi_c(c_j)\|^2}$$

The similarity matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$ holds the agreement of matching pairs on the diagonal and the agreement of mismatching pairs off-diagonal. We train the system using the NT-Xent [14] loss, which is defined as the Cross-Entropy (CE) between the ground truth and the posterior over the text queries and the audio recordings; the ground truth is given by the identity matrix $\mathbf{I} \in \mathbb{R}^{N \times N}$:

$$\mathcal{L} = \frac{1}{2 \cdot N} \sum_{i=1}^N \text{CE}(\mathbf{C}_{i*}, \mathbf{I}_{i*}) + \text{CE}(\mathbf{C}_{*i}, \mathbf{I}_{*i})$$

¹https://github.com/OptimusPrimus/dcase2023_task6b

Instruction	GPT3.5-turbo’s Response
<p>I will give a description of an audio recording and tags associated with the audio recording. Generate 5 audio captions describing the sound event. Each audio caption should be one sentence with less than 20 words. Use grammatical subject-verb-object sentences. Do not include names of cities, countries, and persons. Do not include the time. Do not include recording device names. Do not write introductions or explanations. Do not use “heard” or “recorded”. Start each output sentence with its index. <i>‘A person is making noise by tapping their fingernails on a solid surface.’</i> [dull, metal]</p>	<ol style="list-style-type: none"> 1. The sound of nails tapping on a metal surface is heard. 2. Fingernails are producing a dull sound when tapped. 3. A person is tapping their nails on a solid, metal surface. 4. The noise of someone tapping their fingernails is audible. 5. Nails are rhythmically tapping on a hard metal object.

Table 1: Example query fed to GPT3.5-turbo to augment a ClothoV2 caption (in blue) with the available keywords (in orange; query inspired by WavCaps [8]). The response is a list of rephrased captions, some of which take the keywords into account (highlighted in orange).

2.1. Audio Embedding Models

We choose the Patchout faSt Spectrogram Transformer (PaSST) [10] to convert audio recordings into a compact, high-level vector representation because it achieves state-of-the-art results on multiple audio classification benchmarks [15] while keeping memory and computational complexity low compared to the vanilla audio spectrogram transformers [16]. PaSST uses ImageNet [17] pre-trained parameters from a vision transformer [18, 19] and fine-tunes them on AudioSet [11] for general-purpose audio tagging. The relatively low computational and memory footprint is achieved by dropping patches from the input sequence. This procedure, called Patchout [10], additionally regularizes the model during training. We conducted experiments with PaSST models that take audios of up to ten seconds in length as input and extract overlapping or non-overlapping patches of size 16×16 from the input spectrogram. Pre-trained PaSST models are available on GitHub². We additionally experimented with two convolutional neural networks pre-trained on AudioSet, namely CNN10 and CNN14 [5]. These models can handle inputs of arbitrary length, so we directly input up to 30 seconds long audio instead of cutting them into shorter segments. Table 2 gives an overview of all audio embedding models used in our experiments.

	patch stride	patch out	input length (s)	AS mAP	number params
CNN10	-	-	30	38.0	6.3M
CNN14	-	-	30	43.1	81.8M
PaSST-L	10 × 10	4; 50	10	45.9	41.8M
PaSST-N	16 × 16	2; 15	10	46.8	86.2M
PaSST-S	10 × 10	4; 50	10	48.6	86.2M
PaSST-S20	10 × 10	4; 80	20	47.4	86.2M

Table 2: Overview of the audio embedding models compared in our experiments. The first section shows CNNs from [5]. The second section summarizes PaSST variants [10].

2.2. Sentence Embedding Models

We compared five different sentence embedding models: bert-small, bert-base, bert-large, roberta-base, and roberta-large. All

²https://github.com/kkoutini/passt_hear21

models are bi-directional self-attention-based sentence encoders that underwent self-supervised pretraining on the BookCorpus [21] and WikiText datasets [22]. BERT- and RoBERTa-based models differ in the masking strategy used during training: the former was trained using 10 fixed masks for each sentence, while the latter used new, dynamically generated masks in each forward pass. For both models, we selected the output vector that corresponds to the class token as sentence embedding. The parameter counts for bert-small, bert-base, bert-large, roberta-base, and roberta-large are around 29, 110, 345, 123, and 354 million, respectively.

2.3. Shared Audio-Caption Space

The audio and text embeddings generated by the encoders are integrated into a shared audio-caption space by using a simple linear projection that maps the embedding models’ output to a size of 1024. Initial experiments suggested that using a non-linear projection (e.g., a multilayer neural network) does not significantly improve performance.

2.4. Datasets

Our final models were trained in two steps on multiple datasets. First, we performed pretraining on ClothoV2, AudioCaps, and WavCaps. The resulting models were then further finetuned on a custom, augmented version of ClothoV2 (called ClothoV2_GPT below), which also takes into account the available meta-data associated with each audio file.

2.4.1. ClothoV2

ClothoV2 [9] contains 10-30 second-long audio recordings and captions that are between 8 and 20 words long. The development set’s training, validation, and test split suggested by the organizers contains 3840, 1045, and 1045 recordings, respectively, and each recording is associated with five human-generated captions. The leaderboard evaluation split used for the final system ranking contains 1000 audio recordings and 1000 captions. We used the validation split to monitor the generalization performance and report the performance on the test split in Section 3.

2.4.2. AudioCaps

AudioCaps [13] contains 51, 308 audio recordings taken from AudioSet and one human-written caption for each of them. Each audio

	audio embedding	segment length (s)	overlap	text embedding	finetune	GPT-augment	mAP@10	R@1	R@5	R@10
	PaSST-N	10	✗	bert-small	✗	✗	32.98	21.45	48.71	62.05
	PaSST-N	10	✗	bert-base	✗	✗	35.22	23.07	51.48	65.36
	PaSST-N	10	✗	bert-large	✗	✗	35.78	23.75	52.27	65.57
	PaSST-N	10	✗	roberta-base	✗	✗	35.12	23.02	51.89	65.26
	PaSST-N	10	✗	roberta-large	✗	✗	36.65	24.26	53.89	66.87
	CNN10	30	✗	bert-base	✗	✗	23.72	14.18	36.59	49.21
	CNN14	30	✗	bert-base	✗	✗	28.06	17.86	40.82	54.56
	PaSST-L	10	✓	bert-base	✗	✗	33.47	21.67	49.24	63.16
	PaSST-N	10	✗	bert-base	✗	✗	35.22	23.07	51.48	65.36
	PaSST-S	10	✓	bert-base	✗	✗	32.83	20.90	48.82	62.60
	PaSST-N	10	✗	roberta-large	✓	✗	38.00	25.51	55.06	68.56
	PaSST-N	10	✗	roberta-large	✓	✓	38.56	26.07	55.27	69.30
[20]	CNN14	30	-	all-mpnet-base-v2	✗	✗	22.20	13.00	34.30	48.00
[8]	CNN14	30	-	bert-base	✓	✗	32.95	21.41	47.77	62.10

Table 3: Text-to-audio retrieval performance on the ClothoV2 test set for different combinations of language and audio embedding models (sections one and two, respectively). The impact of additional fine-tuning on ClothoV2 and ClothoV2GPT is shown in section three. Section four shows results from the DCASE baseline system [20] and the current state of the art [8] (values based on WavCaps’s GitHub repository).

recording has a duration of 10 seconds, and the captions are, on average, 9.8 words long. We concatenated the training, validation, and testing split of AudioCaps into one large dataset and used it for pretraining.

2.4.3. WavCaps

WavCaps [8] is a weakly-labeled audio-caption dataset that contains 403,050 audio recordings of varying length collected from FreeSound, BBC Sound Effects, SoundBible, and a strongly labeled subset of AudioSet [23]. Each audio file is associated with a synthetic audio caption that was created by instructing the GPT3.5-turbo model to extract relevant sound events from metadata and output a single-sentence description. The generated captions are, on average, 7.8 words long. The authors demonstrated the usefulness of these synthetic captions by successfully using the dataset for audio retrieval, audio captioning, and text-based sound generation.

2.4.4. ClothoV2_GPT

Each audio recording in the ClothoV2 dataset is associated with additional metadata consisting of the file name, a list of keywords, a username, and a web URL. We took advantage of the additional information and used GPT3.5-turbo to augment the human-generated captions with the available keywords. To this end, we instructed the model to take the keywords into account and rephrase the available captions. We generated five new captions for each caption in the training set, resulting in 96,000 additional captions. Table 1 gives an example query and the resulting augmented captions: GPT3.5-turbo successfully combined the description and the keywords into five varying descriptions with similar content; four of these contain the provided keywords. We will refer to this augmented version of ClothoV2 as ClothoV2_GPT. The generated captions are available in our GitHub repository.

2.5. Preprocessing

To allow batched processing of recordings of varying lengths, we extracted random 30-second snippets from those audio recordings

that are longer than 30 seconds and zero-padded shorter recordings to the maximum duration in the current batch. The resulting waveforms were converted to 128-bin log-MEL spectrograms using a 1024-point FFT (32ms) and hop size of 320 (10ms). The spectrograms were centered and whitened with the approximate global mean and standard deviation before feeding them into the audio embedding model. The input sentences were pre-processed by transforming all characters to lowercase and removing punctuation. The resulting strings were tokenized with the WordPiece tokenizer, padded to the maximum sequence length in the current batch, and truncated if they were longer than 32 tokens.

2.6. Training

We pre-trained the models on AudioCaps, WavCaps, and the training set of ClothoV2. Both embedding models were jointly optimized using gradient descent with a batch size of 64. We used the Adam update rule [24] for 16 epochs, with one warmup epoch. Thereafter, the learning rate was reduced from 2×10^{-5} to 10^{-7} using a cosine schedule. The hyperparameters of the optimizer were set to PyTorch’s [25] defaults. We further used structured patchout as a regularizer during training with hyperparameters depending on the audio length and patch extraction (see Table 2). Finetuning on ClothoV2_GPT was done in a similar manner as pretraining but only for five epochs with a maximum learning rate of 8×10^{-6} . During the finetuning procedure, a caption was swapped with one of its five GPT-augmented versions with a probability of 0.3.

3. RESULTS

The performance of different audio and sentence embedding models is summarized in Table 3. The combination of PaSST-N and roberta-large, pre-training on all data sets, and fine-tuning with GPT augmentation outperforms the current state of the art [8] based on bert-base and CNN14 by 5.6 pp. mAP@10. In the following sections, we analyze our method in detail to identify each component’s impact on the overall performance.

3.1. Pre-training Data Sets

ClothoV2 is relatively small compared to captioning data sets in the image domain; to further enhance the performance, we additionally leveraged AudioCaps and WavCaps. In this section, we investigate the impact of the additional pre-training data sets on the final performance. To this end, we used PaSST-N and bert-base and trained them on different combinations of the three sets. We report the results in terms of mAP@10 on ClothoV2’s test set in Table 4. Pre-training on WavCaps or ClothoV2 results in similar performance of around 27 mAP@10; pretraining exclusively on AudioCaps is roughly 6 pp. worse. Surprisingly, adding AudioCaps to WavCaps did not further improve the result; however, adding ClothoV2 to WavCaps or AudioCaps yielded improvements of 6.8 and 9.6 pp., respectively. The best result overall was achieved by combining all three data sets.

AudioCaps	WavCaps	ClothoV2	mAP@10
✓	✗	✗	21.01
✗	✓	✗	27.62
✗	✗	✓	27.28
✓	✓	✗	27.13
✗	✓	✓	34.42
✓	✗	✓	30.64
✓	✓	✓	35.22

Table 4: Ablation study on the effect of pre-training data sets.

3.2. Text Embedding Models

We assumed that larger sentence embedding models would lead to better retrieval performance. To test this hypothesis, we experiment with three variants of BERT and two variants of RoBERTa. The results are summarized in the first section of Table 3. Larger BERT sentence encoders indeed performed better (compare bert-small, bert-base, and bert-large), and a similar trend can be observed for RoBERTa (compare roberta-base and roberta-large). The best overall results were achieved by utilizing roberta-large.

3.3. Audio Embedding Models

We likewise assumed that using a self-attention-based architecture would lead to further improvements. To test this assumption, we compared two convolutional architectures (CNN10 and CNN14) to three recent spectrogram transformers (PaSST-L, PaSST-N, and PaSST-S); section two of Table 3 summarizes the results. Scaling up the audio embedding model from CNN10 to CNN14 yielded an improvement of 4.3 pp. mAP@10. Switching from CNN14 to PaSST-N further improved the mAP@10 by 7.1 pp. PaSST-S, which extracts overlapping spectrogram patches and performs better on AudioSet, surprisingly did not further improve the retrieval performance over PaSST-N. This inconsistency could be due to a suboptimal patchout configuration.

3.4. Audio Context Length

PaSST uses a learnable positional encoding with a fixed length of 10 seconds and consequently cannot handle longer audio segments. To deal with the up to 30-second long audio recordings in ClothoV2, we cut longer waveforms into shorter segments, embedded each segment separately, and averaged the resulting embeddings over

time to obtain a single vector representation. To investigate the impact of the segment length, we conducted experiments by splitting the recordings into 2, 5, 10, and 15 seconds long snippets. We used PaSST-S20 for those experiments, an architecture similar to PaSST-S, but with a positional encoding for audios of up to 20 seconds in length. The results are given in Figure 1. While a longer context is advantageous performance-wise, it also comes at the price of higher computational cost (which grows quadratically with the input size). We find that PaSST’s default maximum input length of ten seconds strikes a good balance.

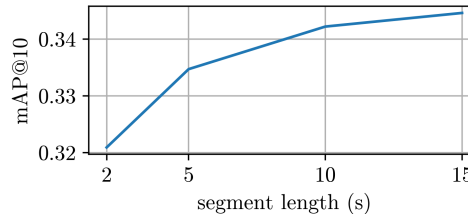


Figure 1: Performance of PaSST-S20 for different audio lengths.

3.5. Fine-Tuning

We further investigated the impact of additional finetuning on the target data sets. To this end, we continued training (as described in Section 2.6) on the training split of ClothoV2 with and without GPT augmentation. The results are given in the third section of Table 3. Finetuning on ClothoV2 without GPT augmentation further improved the pre-trained model by 1.3 pp. mAP@10. When finetuned with GPT-Augmentation, overfitting was slightly reduced, and the model improved by 1.9 pp. mAP@10. A similar advantage for the GPT-augmented data set can be observed for the hidden test set of the 2023’s DCASE Challenge.

4. DISCUSSION & CONCLUSION

This work presented a dual-encoder system for automatic audio retrieval, achieving state-of-the-art results on the ClothoV2 benchmark. The results of our experiments attribute the considerable performance gains to two factors: firstly, the additional data sets with human-generated and synthetic captions, and secondly, the audio spectrogram transformer, which scaled better with the additional data compared to convolutional neural networks. Augmentation of the captions with the additional keywords reduced overfitting during finetuning; however, it did not significantly improve retrieval performance. One possible explanation for this is that each recording in the ClothoV2 training set is associated with five different captions, which are likely to contain the most relevant keywords already; adding further captions increases the variety only marginally.

5. ACKNOWLEDGMENT

The LIT AI Lab is financed by the Federal State of Upper Austria. The computational results presented in this work have been partially achieved using the Vienna Scientific Cluster (VSC).

6. REFERENCES

- [1] H. Xie, S. Lipping, and T. Virtanen, “Language-based audio retrieval task in DCASE 2022 challenge,” in *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, November 2022.
- [2] X. Xu, Z. Xie, M. Wu, and K. Yu, “The SJTU system for DCASE2022 challenge task 6: Audio captioning with audio-text retrieval pre-training,” DCASE2022 Challenge, Tech. Rep., 2022.
- [3] X. Mei, X. Liu, H. Liu, J. Sun, M. D. Plumbley, and W. Wang, “Language-based audio retrieval with pre-trained models,” DCASE2022 Challenge, Tech. Rep., July 2022.
- [4] T. Pellegrini, “IRIT-UPS DCASE 2022 language-based audio retrieval system,” DCASE2022 Challenge, Tech. Rep., July 2022.
- [5] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “PANNs: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE ACM Trans. Audio Speech Lang. Process.*, 2020.
- [6] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proc. of the North American Ch. of the Ass. for Computational Linguistics: Human Language Technologies, NAACL-HLT*, 2019.
- [7] 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,” *CoRR*, vol. abs/1907.11692, 2019.
- [8] 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,” *CoRR*, vol. abs/2303.17395, 2023.
- [9] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: an Audio Captioning Dataset,” in *Proc. IEEE Int. Conf. Acoustic, Speech and Signal Process., ICASSP*, 2020.
- [10] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, “Efficient training of audio transformers with patchout,” in *23rd Annual Conf. of the Int. Speech Communication Association, Interspeech*, 2022.
- [11] J. F. Gemmeke, D. P. W. 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 *Proc. IEEE Int. Conf. Acoustic, Speech and Signal Process., ICASSP*, 2017.
- [12] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally, M. Henry, N. Pinto, C. Noufi, C. Clough, D. Herremans, E. Fonseca, J. H. Engel, J. Salamon, P. Esling, P. Manocha, S. Watanabe, Z. Jin, and Y. Bisk, “HEAR: holistic evaluation of audio representations,” in *NeurIPS 2021 Competitions and Demonstrations Track, NeurIPS*, 2021.
- [13] C. D. Kim, B. Kim, H. Lee, and G. Kim, “AudioCaps: Generating captions for audios in the wild,” in *Proc. of the North American Ch. of the Ass. for Computational Linguistics: Human Language Technologies, NAACL-HLT*, 2019.
- [14] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proc. of the 37th Int. Conf. on Machine Learning, ICML*, 2020.
- [15] K. Koutini, S. Masoudian, F. Schmid, H. Eghbal-zadeh, J. Schlüter, and G. Widmer, “Learning General Audio Representations With Large-Scale Training of Patchout Audio Transformers,” *HEAR: Holistic Evaluation of Audio Representations*, vol. 166, 2022.
- [16] Y. Gong, Y. Chung, and J. R. Glass, “AST: audio spectrogram transformer,” in *22nd Annual Conf. of the Int. Speech Communication Association, Interspeech*, 2021.
- [17] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *2009 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition CVPR*, 2009.
- [18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” *CoRR*, vol. abs/2010.11929, 2020.
- [19] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, “Training data-efficient image transformers & distillation through attention,” in *Proc. of the 38th Int. Conf. on Machine Learning, ICML*, 2021.
- [20] “Language-Based Audio Retrieval, Task description,” <https://dcase.community/challenge2023/task-language-based-audio-retrieval#description>, accessed: 2023-05-21.
- [21] Y. Zhu, R. Kiros, R. S. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler, “Aligning books and movies: Towards story-like visual explanations by watching movies and reading books,” in *IEEE Int. Conf. on Computer Vision, ICCV*, 2015.
- [22] S. Merity, C. Xiong, J. Bradbury, and R. Socher, “Pointer sentinel mixture models,” in *5th Int. Conf. on Learning Representations, ICLR*, 2017.
- [23] S. Hershey, D. P. W. Ellis, E. Fonseca, A. Jansen, C. Liu, R. C. Moore, and M. Plakal, “The benefit of temporally-strong labels in audio event classification,” in *Proc. IEEE Int. Conf. Acoustic, Speech and Signal Process., ICASSP*, 2021.
- [24] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd Int. Conf. on Learning Representations, ICLR*, 2015.
- [25] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An imperative style, high-performance deep learning library,” in *Annual Conf. on Neural Information Processing Systems, NEURIPS*, 2019.