

# Attention-Based Methods For Audio Question Answering

Parthasaarathy Sudarsanam, Tuomas Virtanen

*Audio Research Group, Tampere University, Tampere, Finland*

{parthasaarathy.ariyakulamsudarsanam, tuomas.virtanen}@tuni.fi

**Abstract**—Audio question answering (AQA) is the task of producing natural language answers when a system is provided with audio and natural language questions. In this paper, we propose neural network architectures based on self-attention and cross-attention for the AQA task. The self-attention layers extract powerful audio and textual representations. The cross-attention maps audio features that are relevant to the textual features to produce answers. All our models are trained on the recently proposed Clotho-AQA dataset for both binary yes/no questions and single-word answer questions. Our results clearly show improvement over the reference method reported in the original paper. On the yes/no binary classification task, our proposed model achieves an accuracy of 68.3% compared to 62.7% in the reference model. For the single-word answers multiclass classifier, our model produces a top-1 and top-5 accuracy of 57.9% and 99.8% compared to 54.2% and 93.7% in the reference model respectively. We further discuss some of the challenges in the Clotho-AQA dataset such as the presence of the same answer word in multiple tenses, singular and plural forms, and the presence of specific and generic answers to the same question. We address these issues and present a revised version of the dataset.

**Index Terms**—Audio question answering, attention mechanism, Clotho-AQA

## I. INTRODUCTION

Question answering (QA) is the task of producing natural language answers when posed with questions in natural language. Often, these questions are accompanied by a natural signal such as an image or audio and the questions posed are about the contents of these signals. If the auxiliary input is an image, the task is referred to as visual question answering (VQA) and if it is an audio signal, it is called audio question answering (AQA). Although the question answering framework is somewhat well-studied for image [1–5] and textual modalities [6–8], audio question answering is comparatively less explored. Audio question answering unlocks new possibilities in areas such as monitoring and surveillance, machine listening, human-technology interaction, acoustical scene understanding, etc.

One of the challenging aspects of any multimodal machine learning system is how the information from different modalities is fused to achieve a given task [9]. Traditionally, in question answering systems, the multimodal features are fused using point-wise multiplication [10] or they are concatenated [11] to generate a multimodal representation. This may not be an efficient strategy as these features are learned independently without any context from the other modality. Our hypothesis is that using an attention mechanism to learn a multimodal repre-

sentation helps the model to learn specific features in the audio representation that are closely related to the natural language words in the question and hence improve the performance of the system.

Recently, attention-based architectures [12] have achieved state-of-the-art performances in various tasks ranging from natural language processing [12], to image classification [13], sound event detection [14], sound event localization and detection [15] etc. The ability of transformers to learn powerful and meaningful representations is due to the self-attention and cross-attention layers. Traditional self-attention layers learn bidirectional temporal characteristics of their inputs efficiently. Cross-attention layers are useful to learn or combine features in multimodal translation tasks. For example, in visual question answering, [1, 16] used cross-attention layers to improve the ability of the model to find relevant visual cues depending on the question. However, the effect of attention layers for audio question answering task is unexplored.

In this work, we propose neural network architectures based on attention mechanisms and study their effectiveness for the audio question answering task. Our results show improvement over the reference method described for the Clotho-AQA dataset in [11].

The remainder of this paper is organized as follows. In Section II, our proposed method for the AQA task is explained. Then in Section III, the dataset, reference methods, and experimental setup are described in detail. Subsequently, in Section IV, the results of all our experiments are presented. Finally, in Section V, the conclusion of this work and possible future works are discussed.

## II. METHODS

### A. Proposed model

An AQA system processes an audio signal and an associated natural language question to produce a natural language answer. The answers can be produced either using a generative model or the answers can be chosen from a list of possible answers using a discriminative classification model.

In this work, the AQA task is tackled as a classification problem. Our proposed model architecture is shown in Figure 1. It consists of two input branches, one for processing the audio features and the other for textual features. The audio branch takes in the mel-spectrogram of the audio signal with 128 mel bands and  $T$  time frames and uses a pre-trained OpenL3 [17] to extract audio features. The OpenL3 model is

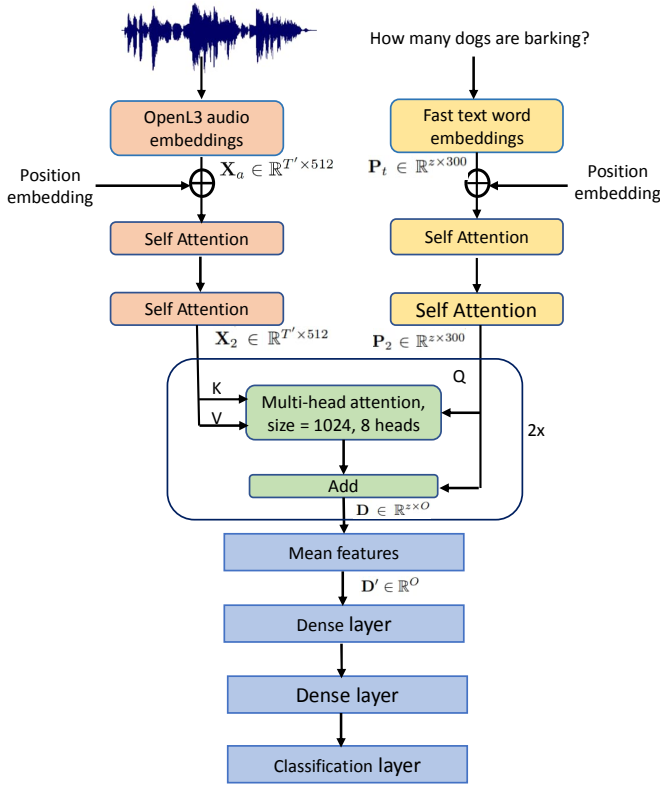


Fig. 1: Proposed attention model architecture

based on L<sup>3</sup>-Net [18] trained on videos from the Audioset [19] dataset for audio-visual correspondence task. The output from the pre-trained OpenL3 audio sub network is  $\mathbf{X}_a \in \mathbb{R}^{T' \times 512}$ , where  $T'$  is the number of output time frames in the OpenL3 model and 512 is the audio embedding size.

Similarly, the text branch produces the word vectors of the input textual question with  $z$  words using the pre-trained Fasttext [20]. The output representation from the Fasttext model is  $\mathbf{P}_t \in \mathbb{R}^{z \times 300}$ , where 300 is the dimension of the learned word vectors.

These extracted features are passed individually through a series of self-attention (SA) layers for both modalities. In the SA mechanism, each time step in the input feature attends to all other time steps to learn temporal relationships. Since all the time steps are fed simultaneously to the self-attention layer, the order of input features is not known. Hence, we add sinusoidal positional embeddings described in [12] to each of the time steps of the audio and textual features which aids the model in learning the relative positions of these features.

The SA layers calculate the dot-product attention of the input features with themselves. For any input  $\mathbf{H} \in \mathbb{R}^{t \times i}$ , where  $t$  is the number of time steps and  $i$  is the input dimension, the output of the SA layer is calculated as

$$\text{SA}(\mathbf{H}) = \text{softmax}(\mathbf{H}\mathbf{W}_q\mathbf{W}_k^T\mathbf{H}^T)\mathbf{H}\mathbf{W}_v \quad (1)$$

where,  $\mathbf{W}_q, \mathbf{W}_k \in \mathbb{R}^{i \times k}$  and  $\mathbf{W}_v \in \mathbb{R}^{i \times o}$  are learnable query, key and value matrices respectively. Here,  $k$  is the key dimension in the attention layer and  $o$  is the output dimension.

The softmax operation is performed over the rows. Note that in a self-attention layer, the query, key, and value are calculated from the same input.

For the audio features  $\mathbf{X}_a$  obtained from the OpenL3 model, the output of the  $n^{\text{th}}$  audio SA layer is  $\mathbf{X}_n \in \mathbb{R}^{T' \times N}$ , where  $N$  is the output attention size. Similarly, for the word vectors  $\mathbf{P}_t$  from the Fasttext model, the output of the  $n^{\text{th}}$  text SA layer is  $\mathbf{P}_n \in \mathbb{R}^{z \times M}$ , where  $M$  is the output attention size. In our experiments, the output attention size of the SA layers in the audio branch is fixed to 512, while that of the text branch is fixed to 300. We used two SA layers for both branches.

To perform a fusion of the audio and textual features, we use multi-head cross-attention (MHCA) layers to learn relevant multimodal features. Attention is applied from the textual features to the audio features to determine which audio features are important to each of the question words. Hence, the output of the final text self-attention layer  $\mathbf{P}_2 \in \mathbb{R}^{z \times 300}$  is used to compute the query and the output of the final audio self-attention layer  $\mathbf{X}_2 \in \mathbb{R}^{T' \times 512}$  is used to calculate the key and value inputs to the MHCA layers. The output of a cross-attention (CA) layer is computed as

$$\text{CA}(\mathbf{P}_2, \mathbf{X}_2) = \text{softmax}(\mathbf{P}_2\mathbf{W}_q\mathbf{W}_k^T\mathbf{X}_2^T)\mathbf{X}_2\mathbf{W}_v \quad (2)$$

where,  $\mathbf{W}_q \in \mathbb{R}^{300 \times 512}$ ,  $\mathbf{W}_k \in \mathbb{R}^{512 \times 512}$  and  $\mathbf{W}_v \in \mathbb{R}^{512 \times O}$  are learnable query, key and value matrices respectively and the softmax operation is performed over the rows. Here,  $O$  is the output dimension. For an MHCA layer with  $M$  attention heads, the outputs from all the heads are concatenated along the rows and  $\mathbf{W}_p \in \mathbb{R}^{MO \times O}$ , a learned projection matrix projects it into the desired output dimension. The output of the MHCA layer is given by

$$\text{MHCA}(\mathbf{P}_2, \mathbf{X}_2) = \text{Concat}_{m=1 \dots M}[\text{CA}_m(\mathbf{P}_2, \mathbf{X}_2)]\mathbf{W}_p \quad (3)$$

In all our experiments, the output attention size is fixed at 1024 with 8 attention heads. All the hyperparameters were tuned based on the model's performance on the validation data. The output of the MHCA layer is  $\mathbf{D} \in \mathbb{R}^{z \times O}$ . To obtain a fixed size representation, the mean is taken over the words axis of the output of the attention layer to produce  $\mathbf{D}' \in \mathbb{R}^O$ . This is then passed through two dense layers for combining the learned features and then to the classification layer.

We developed two classifiers using this architecture. A binary classifier for questions that have 'yes' or 'no' as answers and a multiclass classifier for questions that have other single-word answers in the Clotho-AQA dataset. The classification layer is a logistic regressor with one neuron for the binary classifier. In the case of the multiclass classifier, the final classification layer contains as many neurons as the number of unique answer words in the dataset followed by a softmax activation to predict the probabilities.

### III. EVALUATION

#### A. Dataset

We trained and evaluated our models on the recently proposed Clotho-AQA dataset [11]. The dataset contains 1991 audio files randomly selected from the Clotho dataset [21]. The Clotho dataset is an audio captioning dataset that contains 4981 audio files that are 15-30s in duration. It contains audio files of day-to-day sounds occurring in the environment such as water, nature, birds, noise, rain, city, wind, etc. In the Clotho-AQA dataset, for each of these audio files, there are four ‘yes’ or ‘no’ questions and two single-word answer questions. For each question, answers were collected from three independent annotators. Hence, each audio file is associated with 18 question-answer pairs. In the Clotho-AQA dataset, there are 828 unique single-word answers excluding ‘yes’ and ‘no’. The complete process of data collection, cleaning, and creating the splits is detailed in [11]. Henceforth, this dataset is referred to as Clotho-AQA\_v1 in this paper.

The Clotho-AQA\_v1 dataset has a few limitations in single-word answers due to crowdsourcing. The dataset contains issues relating to specificity, tense, singular and plural words, etc in the answers. For example, to questions like ‘What is making the chirping sound?’, some annotators provided ‘bird’ as the answer while some provided ‘seagull’ as a more specific answer. Although both these answers can be correct, they are considered different answer classes which creates confusion in the system. Tense issues generally occur when the same question is posed in different tenses for different audio files. For example, for questions like ‘What is the person doing?’ and ‘What does the person do?’ the answers are ‘running’ and ‘run’ respectively. An example of singular-plural answers is, for the question ‘Which object is making the metallic sound?’, some annotators answered with the singular form ‘key’ while some used the plural form ‘keys’. These are considered different answer classes in the Clotho-AQA dataset and thus affect the performance of the classifier.

Since all the answers in the dataset are single words, the AQA system is trained as a classifier where each unique answer word is denoted by a class index for the ground truth labels. Hence, the system does not learn any language modeling from the answer words. Therefore, we addressed these three issues by replacing specific answer words with their parent classes (for example, ‘seagull’ to ‘bird’), all plural words to singular, and all answer tenses to simple present. After this cleaning process, we ended up with 650 unique single-word answers compared to 828 in Clotho-AQA\_v1. This new version of the dataset is referred to as Clotho-AQA\_v2 in this paper. The distribution of unique answer words in Clotho-AQA\_v2 is shown in Figure 2.

Since each question is answered by three independent annotators, it is also important to analyze if the answers are the same or different from each other. For example in the Clotho-AQA\_v1 test split for single-word answers, out of 946 unique questions, 203 questions have unanimous answers provided by all the annotators and 381 questions have two out of the

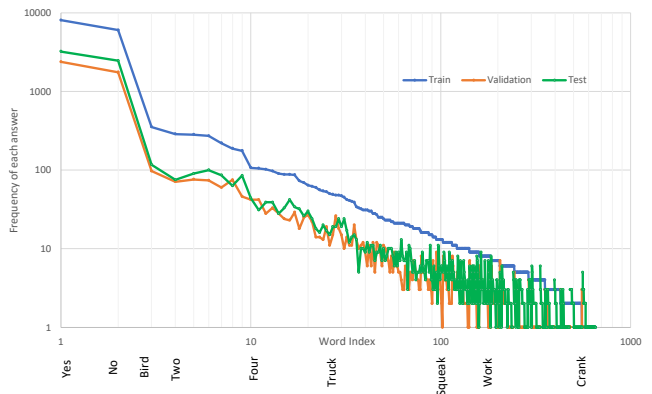


Fig. 2: Count of unique answers in each of the splits in Clotho-AQA\_v2

three annotators providing the same answer. This means that the maximum possible top-1 accuracy of a system without modeling the characteristics of the annotator will be 61%. In the Clotho-AQA\_v2, the maximum achievable accuracy of the multiclass classifier increased to 65%.

#### B. Reference methods

In order to study the effects of cross-attention and self-attention mechanisms on the AQA task, two reference architectures are used. As the first reference model, we used the architecture proposed in [11]. This model architecture does not have any attention blocks. Similarly to the proposed model, the mel-spectrogram of the audio input is processed using the pre-trained OpenL3 model and the pre-trained word vectors are obtained using the Fasttext model. Compared to the proposed model, where the pre-trained audio and textual features are passed separately through SA layers, in this model these features are passed separately through a series of Bi-LSTM layers to learn the temporal features and create fixed-size representations. The hidden state of the final time step of the last Bi-LSTM layer is used as the fixed-size representation. These fixed-size representations from the audio and text branches are concatenated to produce the multimodal representation. The multimodal features are then processed by the dense and classification layers similarly to the proposed model.

The second reference model examines the effect of the cross-attention mechanism on the fusion of multimodal features. MHCA layers are introduced after the Bi-LSTM layers of the first reference system for multimodal representation instead of feature concatenation. The output of the audio Bi-LSTM units is  $\mathbf{X}_{audio} \in \mathbb{R}^{T' \times 2h}$ , where  $h$  is the number of hidden units in the Bi-LSTM layer and  $T'$  is the number of output frames from the OpenL3 model. Similarly, the output of the textual branch Bi-LSTM is  $\mathbf{X}_{text} \in \mathbb{R}^{z \times 2h'}$ , where  $h'$  is the number of hidden units in the Bi-LSTM layer and  $z$  is the number of words in the natural language question. The cross-attention is calculated on these features as explained in Section II. Similarly to the proposed method, the mean is taken over

the words axis to produce a fixed-size representation and it is passed through dense and classification layers for predicting the answer class.

### C. Network training

All the models were trained and evaluated on both the Clotho-AQA\_v1 and Clotho-AQA\_v2 datasets. The data split for the binary classifier and the multiclass classifier is obtained by selecting ‘yes’ or ‘no’ answers and single-word answers respectively from the Clotho-AQA dataset splits as described in [11]. As a result, there are 1174, 344, and 473 audio files for training, validation, and test split respectively each associated with 12 yes/no and six single-word question-answer pairs.

The performance of the binary yes/no classifier is also analyzed on contradicting answers provided by different annotators to the same question similar to the approach proposed in [11]. In this regard, three cases are considered. In the first case, all the question-answer pairs are considered valid even if they contain contradicting answers. In the second case, only those question-answer pairs for which all three annotators have responded unanimously are considered valid. In the third case, a majority voting scheme is used, where for each question, the label is chosen as the answer provided by at least two out of the three annotators. These three cases are denoted as ‘Unfiltered data’, ‘Unanimous’, and ‘Majority votes’ respectively. Note that the binary classifier data set is the same in both Clotho-AQA\_v1 and Clotho-AQA\_v2.

All the models were trained for 100 epochs with cross-entropy loss and the model with the best validation score is used for evaluation on the test set.

## IV. RESULTS

The results of all the experiments on the Clotho-AQA data set for binary classification of ‘yes’ or ‘no’ answers are presented in Table I. In this table, ‘LSTM’ represents the first reference model that uses Bi-LSTM layers for feature extraction and feature concatenation for multimodal fusion. ‘LSTM-CA’ represents the model which uses Bi-LSTM layers for feature extraction and cross-attention layers for multimodal fusion. Finally, ‘SA-CA’ represents our proposed model which used self-attention layers for feature extraction and cross-attention layers for multimodal fusion.

Firstly, it is clear that using cross-attention layers significantly improves the performance compared to feature concatenation. This means that the cross-attention layer, helps the model to attend to audio features that are useful to answer the question. Secondly, using self-attention layers with positional embeddings to learn temporal relationships outperform Bi-LSTM layers.

The results of single-word answer multiclass classifier experiments on Clotho-AQA\_v1 and Clotho-AQA\_v2 are summarized in Table II and Table III respectively. Since the number of unique answer classes is large (828 in Clotho-AQA\_v1 and 650 in Clotho-AQA\_v2), top-5 and top-10 accuracy scores are also reported. These results indicate that the model is starting to learn relationships between the multimodal data.

Model	Unfiltered	Unanimous	Majority votes
LSTM	62.7	73.1	63.2
LSTM-CA	66.2	75.4	66.3
SA-CA	68.3	77.1	68.3
Oracle model	86.2	100	100

TABLE I: Accuracies (%) of binary ‘yes’ or ‘no’ classifier on Clotho-AQA data set.

Model	Top-1	Top-5	Top-10
LSTM	54.2	93.7	98.0
LSTM-CA	57.5	99.8	99.9
SA-CA	57.9	99.8	99.9
Oracle model	61	100	100

TABLE II: Accuracies (%) of single-word answers classifier on Clotho-AQA\_v1 data set.

Model	Top-1	Top-5	Top-10
LSTM	59.8	96.6	99.3
LSTM-CA	61.3	99.6	99.9
SA-CA	61.9	99.8	99.9
Oracle model	65	100	100

TABLE III: Accuracies (%) of single-word answers classifier on Clotho-AQA\_v2 data set.

It is again evident from the results that the self-attention and cross-attention mechanism significantly improve the evaluation metrics in the case of multi-class classifier as well. It is also noticeable that the model performs better on the Clotho-AQA\_v2 dataset after resolving some issues present in the Clotho-AQA\_v1 dataset. The proposed multi-class classifier has also reached close to the maximum possible accuracy by an Oracle model with both Clotho-AQA\_v1 and Clotho-AQA\_v2 data sets.

## V. CONCLUSION

In this paper, we proposed self-attention and cross-attention based architectures for AQA task. We trained and evaluated our models on the recently proposed Clotho-AQA dataset referred to as Clotho-AQA\_v1 in this work. We also discussed some of the challenges of this dataset such as the presence of the same answer word in multiple tenses, the presence of specific and generic answers to the same question, and the presence of singular and plural forms of the same word. These challenges were addressed and a revised version of this dataset Clotho-AQA\_v2 was created. The results of our proposed attention models on both these datasets clearly prove that the cross-attention mechanism helps the model to learn better relationships between the input question and the audio compared to the reference methods. It is also evident that using self-attention layers with positional embeddings is powerful in learning useful audio and textual features compared to the Bi-LSTM layers used in the reference methods.

## ACKNOWLEDGMENT

The authors wish to acknowledge CSC-IT Center for Science, Finland, for the computational resources used in this research.

## REFERENCES

- [1] Z. Yang, X. He, J. Gao, L. Deng, and A. J. Smola, "Stacked Attention Networks for Image Question Answering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [2] J. Lu, D. Batra, D. Parikh, and S. Lee, "ViLBERT: Pre-training Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks," in *Advances in Neural Information Processing Systems*, 2019.
- [3] D. Teney, P. Anderson, X. He, and A. van den Hengel, "Tips and Tricks for Visual Question Answering: Learnings From the 2017 Challenge," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach, "Multimodal compact bilinear pooling for visual question answering and visual grounding," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016.
- [5] A. Agrawal, D. Batra, and D. Parikh, "Analyzing the Behavior of Visual Question Answering Models," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [7] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A Lite BERT for Self-supervised Learning of Language Representations," in *International Conference on Learning Representations*, 2020.
- [8] C. Xiong, V. Zhong, and R. Socher, "Dynamic coattention networks for question answering," in *International Conference on Learning Representations*, 2017.
- [9] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, February, 2019.
- [10] A. Agrawal *et al.*, "VQA: Visual question answering," *International Journal of Computer Vision*, vol. 123, pp. 4–31, 2015.
- [11] S. Lipping, P. Sudarsanam, K. Drossos, and T. Virtanen, "Clotho-AQA: A Crowdsourced Dataset for Audio Question Answering," in *30th European Signal Processing Conference (EUSIPCO)*, 2022.
- [12] A. Vaswani *et al.*, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [13] A. Dosovitskiy *et al.*, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in *International Conference on Learning Representations (ICLR)*, 2021.
- [14] K. Miyazaki, T. Komatsu, T. Hayashi, S. Watanabe, T. Toda, and K. Takeda, "Convolution-augmented transformer for semisupervised sound event detection," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, 2020.
- [15] P. Sudarsanam, A. Politis, and K. Drossos, "Assessment of Self-Attention on Learned Features For Sound Event Localization and Detection", in *Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2021.
- [16] K. J. Shih, S. Singh, and D. Hoiem, "Where To Look: Focus Regions for Visual Question Answering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] J. Cramer, H.-H. Wu, J. Salamon, and J. P. Bello, "Look, listen, and learn more: Design choices for deep audio embeddings," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [18] R. Arandjelovic and A. Zisserman, "Look, listen and learn," *IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, Oct. 2017.
- [19] J. F. Gemmeke *et al.*, "Audio set: An ontology and human-labeled dataset for audio events," in *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [20] T. Mikolov, E. Grave, P. Bojanowski, C. Puhrsch, and A. Joulin, "Advances in pre-training distributed word representations," in *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [21] K. Drossos, S. Lipping, and T. Virtanen, "Clotho: An audio captioning dataset," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.