DUAL-MODALITY SEQ2SEQ NETWORK FOR AUDIO-VISUAL EVENT LOCALIZATION

Yan-Bo Lin, Yu-Jhe Li, and Yu-Chiang Frank Wang

Graduate Institute of Communication Engineering, National Taiwan University, Taiwan MOST Joint Research Center for AI Technology and All Vista Healthcare

ABSTRACT

Audio-visual event localization requires one to identify the event which is both visible and audible in a video (either at a frame or video level). To address this task, we propose a deep neural network named Audio-Visual sequenceto-sequence dual network (AVSDN). By jointly taking both audio and visual features at each time segment as inputs, our proposed model learns global and local event information in a sequence to sequence manner, which can be realized in either fully supervised or weakly supervised settings. Empirical results confirm that our proposed method performs favorably against recent deep learning approaches in both settings.

Index Terms— Audio-Video Features, Dual Modality, Event Localization, Deep Learning

1. INTRODUCTION

While a substantial amount of visual and auditory signals can be observed in real-world activities, several studies in neurobiology and human perception suggest that the perceptual benefits of integrating visual and auditory information are extensive. For computational models, they reflect in lip reading [1], where correlations between speech and lip movements provide a strong cue for linguistic understanding. In sound synthesis [2], where physical interactions with different types of material give rise to plausible sound patterns. In spite of these advanced models, they are still limited to some constrained domains. Hence, some works [3, 4, 5, 6] have begun to learn the correspondence between the visual scene and the sound, achieving the cross-modality scenario. However, such cross-modality learning methods typically assume that the audio and visual contents in a video are always presented and matched. Such assumptions might not be practical for analyzing/understanding unconstrained videos in the real world.

Audio-visual event localization addresses the task of matching both visible and audible components in a video for identifying the event of interest, as depicted in Fig. 1. Recently, some multi-modal deep networks for joint audio-visual representation have been studied [7, 8, 9]. Aside from works in representation learning, audio-visual cross-modal synthesis is studied in [10, 11], and associations between



Fig. 1. Illustration of audio-visual event localization. Note that green and red bounding boxes for video or audio data indicate the labels observed for single-modality data, while the ground truth labels are listed in the bottom row (i.e., non-background event labels are assigned only if cross-modality data exhibit same label information).

natural image scenes and accompanying free-form spoken audio captions are explored in [12]. However, how to handle learned audio and visual features for audio-video event jointly localization remains a challenging issue (i.e., hearing a dog barking and seeing a dog in the video simultaneously). To this end, [13] introduce such novel problem of audiovisual event localization in unconstrained videos. As shown in Fig. 1, they define an audio-visual event as an event that is both visible and audible in a video segment. In addition, they further collect an Audio-Visual Event (AVE) dataset to systemically investigate three temporal localization tasks: supervised and weakly-supervised audio-visual event localization, and cross-modality localization. Nevertheless, they only consider temporal relationship within video or audio data, and no cross-modality relationship is observed.

To address this challenging task, we propose an end-toend deep learning framework of Audio-Visual sequence-tosequence dual network (AVSDN). Based on sequence to sequence (seq2seq) [14] and autoencoder, our network architecture takes both audio and visual data at each time segment as inputs and exploits global and local event information in a seq2seq manner. More importantly, our model can be learned



Fig. 2. Overview of our Audio-Visual Sequence-to-sequence Dual Network (AVSDN). Our AVSDN is composed of three main components, which include encoder modules (in orange and blue) for learning visual and audio representations, a fusion network (in purple) for producing global video representation, and a decoder (in green) which jointly takes global and local features for event localization. Note that v and a denote visual and audio features, h and s are the hidden and cell states of LSTMs, and y indicates the event label.

in a fully or weakly supervised settings, i.e., ground truth event labels observed in the frame or video levels. While the technical details of our proposed networks will be presented in Sec. 2, contributions of our work are highlighted as follows:

- We propose a unique end-to-end trainable network for audio-visual event localization, which can be learned in fully or weakly supervised settings.
- Our proposed model jointly takes visual and audio data as inputs. By exploiting this cross-modality information across time, such encoded global features will be conditioned on the decoder for event localization.
- Experimental results demonstrate that our proposed model performs favorably against state-of-the-art approaches on the challenging AVE dataset.

2. PROPOSED METHOD

2.1. Problem Formulation

Supervised event localization. In supervised settings, one needs to predict video events every second given audio and visual tracks. As mentioned in Section 1, non-background labels are determined only if visual and audio events are jointly observed. To this end, the segment-wised (every second) labels are given as $y_t = \left\{ y_t^k | y_t^k \in \{0,1\}, \sum_{k=1}^{C+1} y_t^k = 1 \right\}$, where *C* denotes total event categories. The number of total categories is C + 1 which includes one background label, and *t* is the time segment of one video. In supervised setting, segment-wise labels are observed during the training phase.

Weakly-supervised event localization. Compared with the aforementioned supervised task, we can only access the video-level event labels in a weakly-supervised manner. The video-level event labels are processed by averaging the segment event labels $Y = \frac{1}{T} \sum_{t=1}^{T} y_t$, where *T* is the length of a video. For the weakly-supervised task, we can lower the reliance on the well-annotated labels and evaluate the robustness of our proposed framework.

2.2. Audio-Visual Sequence-to-sequence Dual Network (AVSDN)

To address the audio-visual event localization problem in both supervised and weakly-supervised problems, we propose a novel framework named Audio-Visual sequence-to-sequence dual network (AVSDN). Based on seq2seq [14] and autoencoder, our network architecture takes both audio and visual data at each time segment as inputs and exploits global and local event information in a seq2seq manner.

As shown in Fig 2, the framework is composed of three components: encoder modules for learning visual and audio representations, fusion network for producing global video representation, and decoder to take global and local features for event localization. More details about the three components in AVSDN can be obtained as follows:

Encoder: learning global visual and audio representations. The encoder, which is the first part of our network AVSDN, is aimed to extract global visual and audio representations for fusion network.

Before learning the global features, we have to obtain the

segment visual and audio representations by utilizing CNNs. To better learn the visual and audio embedding features, our CNNs are learned from the large-scale dataset (ImageNet [15] and AudioSet [16]) which are highly shown useful for vision and audition tasks. To be specific, for visual frames we sample a frame and obtain the visual representation from pre-trained ResNet-152 [17], which has been trained on ImageNet. On the other hand, for one raw audio, we convert one segment to log mel spectrogram and extract an audio representation each 1s from VGGish [18] trained on AudioSet.

In order to further learn the global visual and audio representations, we now utilize the Long Short-Term Memory (LSTM) [19], which is known to exploit long-range temporal dependencies, to generate encoded temporal representation sequence. Generally, the most common implementation of vanilla LSTM [20] includes various gate mechanisms such as input gate, forget gate, output gate, memory state, and hidden state etc. The utilized LSTM unit in our proposed model is illustrated in Eq. (1).

$$f_t = \sigma_g(W_f x_t + U_f x_t + b_f)$$

$$i_t = \sigma_g(W_i x_t + U_i x_t + b_i)$$

$$o_t = \sigma_g(W_o x_t + U_o x_t + b_o) \qquad (1)$$

$$c_t = f_t \circ c_{t-1} + i_t \circ \sigma_c(W_c x_t + U_c h_{t-1} + b_c)$$

$$h_t = o_t \circ \sigma_h(c_t)$$

Each time step in Eq. (1) denotes subscript t, and $x_t \in \mathbb{R}^d$ denotes the given input of time step $t \, . \, f_t \in \mathbb{R}^h$, $i_t \in \mathbb{R}^h$ and $o_t \in \mathbb{R}^h$ are forget, input and output gate's activation vector respectively. $h_t \in \mathbb{R}^h$ and $c_t \in \mathbb{R}^h$ are hidden and cell state of the LSTM unit. When t = 0, $c_0 = 0$ and $h_0 = 0$ would be the initial values. $W \in \mathbb{R}^h \times d$, $U \in \mathbb{R}^h \times h$ and $b \in \mathbb{R}^h$ are weight matrices and bias vector which can be learned during train phase. Where d and h refer the number of input features and number of hidden units. Element-wise product is denoted by \circ . Activation function: σ_g is sigmoid function and σ_c is hyperbolic tangent function.

Eventually, all the audio and visual segments are the inputs of the two designed LSTM (audio and video separately). Hence, the last time step T of hidden and cell state can be generated as the global representations of audio and visual tracks.

Fusion network: learning video event representation.

After obtaining the audio and visual global representations, our goal is to convert these two representations into one video event representation. To perform such a fusion mechanism, our fusion network is designed to fuse cross-modality features which are built based on dual multimodal residual network (DMRN) fusion block [13]. As mentioned above, the last time step T of hidden and cell state from encoder can be given as the representations of audio and visual tracks. Following [13], with time step T, audio and visual hidden state (h_T^a, h_T^v) and cell state (c_T^a, c_T^v) can be fused with Eq.(2). After fusion hidden and cell state, these fused state will be the initial state of the decoder LSTM (one LSTM of right-half Fig.2).

$$h_T^{a'} = \sigma_c (h_T^a + \frac{1}{2} (g_\theta(h_T^a) + g_\theta(h_T^v)))$$

$$h_T^{v'} = \sigma_c (h_T^v + \frac{1}{2} (g_\theta(h_T^a) + g_\theta(h_T^v)))$$

$$h_T^f = h_T^{a'} + h_T^{v'}$$
(2)

where σ_c is denoted as hyperbolic tangent function and $g_{\theta}(.)$ is multilayer perceptrons (MLP) with parameters θ . The cell states can be fused like hidden states. The fused states, h_t^f and c_t^f , turn to be the initial states in our decoder. Because compared with vanilla LSTM, a representative initial state can benefit a LSTM for prediction [14]. Thus, we take the fused states for the initialization for the decoder.

Decoder: localization of video events using global and local cross-modality representations. Generally, our decoder is aimed to perform the supervised and weakly-supervised event localization. Thus, given both the fused global representations from the fused network and local features of audio and video, our decoder will generate the corresponding labels segment-wisely. The architecture of our decoder is a single LSTM. Different from each encoder, the inputs of the decoder are concatenated features which are global and local crossmodality representations. We concatenate a_t and v_t which are audio and visual segment features from pre-trained CNN. Our decoder is designed to not only learn spatial cross-modality representations but temporal ones. Especially weakly supervised setting, we can only access to the video-level labels in the training phase. All the individual predictions will be aggregated by average pooling in Eq.(3),

$$\hat{m} = avg(m_1, m_2, ..., m_T) = \frac{1}{T} \sum_{t=1}^T m_t,$$
 (3)

where $m_1, ..., m_T$ are the predictions from the last fully connected layer of our model. The average prediction \hat{m} over softmax function can be the probability distribution of the event category. For both the weakly-supervised and supervised setting, the predicted probability distribution can be optimized by video-level labels through binary cross-entropy.

3. EXPERIMENTS

3.1. Dataset

Following [13], we consider *Audio-Visual Event*(AVE) [13] dataset (a subset of Audioset [16]) for experiments. This AVE dataset includes 4143 videos with 28 categories and videos are labeled with audio-visual events every second. AVE dataset covers wide range domain events (e.g., Church bell, Dog barking, Truck, Bus, Clock, Violin, etc.).

Table 1. Comparisons with the state-of-the-art method of [13] in a **supervised** manner (all ground truth y_t are observed during training. The number in bold indicates the best result.

Method	Accuracy (%)	Remarks
AVEL [13]	59.5	audio only
	55.3	visual only (VGG16)
	71.4	audio+visual (VGG16)
	72.7	audio+visual w/ att (VGG16)
	65.0	visual only (ResNet-151)
	74.0	audio+visual (ResNet-151)
	74.7	audio+visual w/ att (ResNet-151)
AVSDN	75.4	audio+visual (ResNet-151)

Table 2. Comparisons with the state-of-the-art method of [13] in a weakly supervised manner (only ground truth Y is observed for training). The number in bold indicates the best result.

Method	Accuracy (%)	Remarks
AVEL[13]	53.4	audio only
	52.9	visual only (VGG16)
	63.7	audio+visual (VGG16)
	66.7	audio+visual w/ att (VGG16)
	63.4	visual only (ResNet-151)
	71.6	audio+visual (ResNet-151)
	73.3	audio+visual w/ att (ResNet-151)
AVSDN	74.2	audio+visual (ResNet-151)

3.2. Comparison results

For both fully supervised and weakly-supervised audio-visual event localization, we apply **frame-wise accuracy** as an evaluation metric in both supervised and weakly-supervised settings. We compute the percentage of correct matchings over all the testing frames as prediction accuracy to evaluate the performance of audio-visual event localization.

In this paper, we use different visual pre-trained embedding compared with Tian et al [13]. Visual pre-trained embedding in Tian et al is VGG16 [21]. Thus, we re-implement the model with ResNet-151 [17] visual pre-trained embedding and show each one modality results. In a fully supervised manner, all the frame-wise labels are used during training. In Table 1, our model has better results compared with state-of-the-art methods even if the model [13] is with crossmodality attention mechanism [22] which can find the audio location in the video scene [13, 4]; In a weakly supervised manner, Table 2 shows our model outperforms other methods as well. Compared with the supervised task, our model can tickle noise labels better. We give an example of our visualization comparison with the existing method in Fig. 3.

Table 3. Ablation studies on our network design, i.e., our decoder taking hidden and cell states of global visual/audio representations as conditioned LSTM inputs. Note that weakly supervised learning is considered in this table.

Method	Accuracy (%)
Visual input only	70.2
Audio input only	70.9
Audio + Visual (label guided)	72.6
Fusion (Ours)	74.2



Fig. 3. Visualization comparisons between AVEl and ours.

3.3. Ablation studies

In this section, we will discuss the results on the different initial state for decoder LSTM. In Table 3, first-two rows show the results which initial state is only from visual or audio content respectively. Further, we want to explore whether encoder LSTM can learn global event information or not. With the last hidden states from visual and audio modality individually, these hidden states are guided by video-level labels through a simple multilayer perceptron (MLP). The third row in Table 3 shows the result is improved compared with only one modality. However, extra loss functions for guiding the last hidden states are not needed. The last hidden states are well-learned during the training of our AVSDN.

4. CONCLUSION

In this work, we proposed Audio-Visual sequence-to-sequence dual network (AVSDN) for video event localization, which can be learned in fully or weakly supervised fashions. Our network takes both audio and visual local features, together with integrated global representation, to perform event localization in a sequence to sequence manner. From the experimental results, the use of our network and its design can be successfully verified.

Acknowledgements. This work is supported by the Ministry of Science and Technology of Taiwan under grant MOST 107-2634-F-002-010.

5. REFERENCES

- [1] Joon Son Chung, Andrew W Senior, Oriol Vinyals, and Andrew Zisserman, "Lip reading sentences in the wild.," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [2] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H Adelson, and William T Freeman, "Visually indicated sounds," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon, "Learning to localize sound source in visual scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Relja Arandjelović and Andrew Zisserman, "Objects that sound," in *Proceedings of the European Conference* on Computer Vision (ECCV), 2018.
- [5] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba, "The sound of pixels," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [6] Relja Arandjelovic and Andrew Zisserman, "Look, listen and learn," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [7] Di Hu, Xuelong Li, et al., "Temporal multimodal learning in audiovisual speech recognition," in *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [8] Xitong Yang, Palghat Ramesh, Radha Chitta, Sriganesh Madhvanath, Edgar A Bernal, and Jiebo Luo, "Deep multimodal representation learning from temporal data," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] Douwe Kiela, Edouard Grave, Armand Joulin, and Tomas Mikolov, "Efficient large-scale multi-modal classification," *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [10] Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, and Tamara L Berg, "Visual to sound: Generating natural sound for videos in the wild," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] Lele Chen, Sudhanshu Srivastava, Zhiyao Duan, and Chenliang Xu, "Deep cross-modal audio-visual generation," in *Proceedings of the on Thematic Workshops of* ACM Multimedia 2017. ACM, 2017, pp. 349–357.

- [12] David Harwath, Antonio Torralba, and James Glass, "Unsupervised learning of spoken language with visual context," in *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [13] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu, "Audio-visual event localization in unconstrained videos," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [14] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems* (*NIPS*), 2014.
- [15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2009.
- [16] Jort F. Gemmeke, Daniel P. W. Ellis, et al., "Audio set: An ontology and human-labeled dataset for audio events," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing* (ICASSP), 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [18] Shawn Hershey, Sourish Chaudhuri, et al., "Cnn architectures for large-scale audio classification," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [19] Sepp Hochreiter and Jrgen Schmidhuber, "Long shortterm memory," vol. 9, pp. 1735–80, 12 1997.
- [20] Klaus Greff, Rupesh Kumar Srivastava, Jan Koutnk, Bas R. Steunebrink, and Jrgen Schmidhuber, "Lstm: A search space odyssey.," *CoRR*, vol. abs/1503.04069, 2015.
- [21] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems (NIPS), 2017.