

Object-Aware Adaptive-Positivity Learning for Audio-Visual Question Answering

Zhangbin Li¹, Dan Guo^{1,2,3*}, Jinxing Zhou^{1*}, Jing Zhang¹, Meng Wang^{1,2}

¹School of Computer Science and Information Engineering, Hefei University of Technology

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³Anhui Zhonghuitong Technology Co., Ltd

lizhangbin.mail@gmail.com, guodan@hfut.edu.cn,

zhoujxfut@gmail.com, hfutzhangjing@gmail.com, wangmeng@hfut.edu.cn

Abstract

This paper focuses on the Audio-Visual Question Answering (AVQA) task that aims to answer questions derived from untrimmed audible videos. To generate accurate answers, an AVQA model is expected to find the most informative audio-visual clues relevant to the given questions. In this paper, we propose to explicitly consider fine-grained visual objects in video frames (object-level clues) and explore the multi-modal relations (*i.e.*, the object, audio, and question) in terms of feature interaction and model optimization. For the former, we present an end-to-end object-oriented network that adopts a question-conditioned clue discovery module to concentrate audio/visual modalities on respective keywords of the question and designs a modality-conditioned clue collection module to highlight closely associated audio segments or visual objects. For model optimization, we propose an object-aware adaptive-positivity learning strategy that selects the highly semantic-matched multi-modal pair as *positivity*. Specifically, we design two object-aware contrastive loss functions to identify the highly relevant question-object pairs and audio-object pairs, respectively. These selected pairs are constrained to have larger similarity values than the mismatched pairs. The positivity-selecting process is adaptive as the positivity pairs selected in each video frame may be different. These two object-aware objectives help the model understand *which objects are exactly relevant to the question and which are making sounds*. Extensive experiments on the MUSIC-AVQA dataset demonstrate the proposed method is effective in finding favorable audio-visual clues and also achieves new state-of-the-art question-answering performance. The code is available at <https://github.com/zhangbin-ai/APL>.

Introduction

Question-answering occupies a large part of human communication. To allow machines to learn to respond to questions posed by humans, Li *et al.* (Li *et al.* 2022) recently proposed the Audio-Visual Question Answering (AVQA) task, which requires answering questions considering both audio and visual information. Unlike other audio-visual tasks, such as the audio-visual segmentation (Zhou *et al.* 2022b, 2023b; Mao *et al.* 2023; Li *et al.* 2023b,a, 2021) and audio-visual event localization (Tian *et al.* 2018; Zhou *et al.* 2021, 2022a,

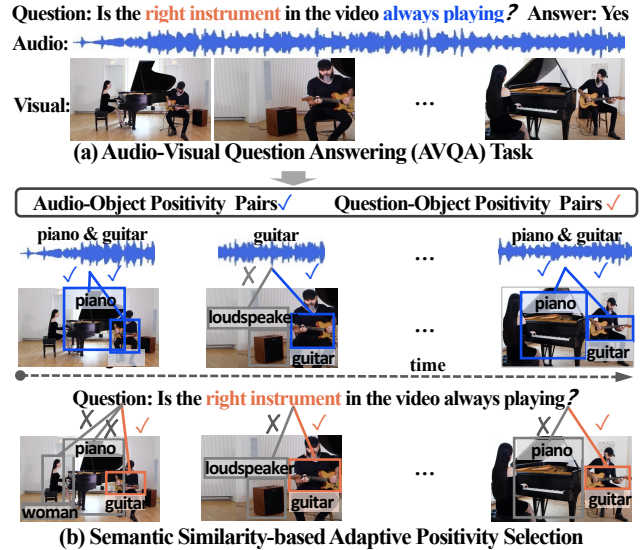


Figure 1: Illustration of the AVQA task and our adaptive-positivity selection strategy. (a) The AVQA task requires answering the question with effective audio and visual clues. (b) We explicitly explore the fine-grained visual *objects* and use them as the bridge to adaptively select the highly matched audio-object pairs (✓ in blue) and question-object pairs (✓ in orange) by measuring their semantic similarities. In this way, our method is object-aware, which not only knows which objects are making sounds but also identifies which audible objects are really relevant to the question.

2023a) independently explore the *spatial* and *temporal* relations between audio and visual modalities, the AVQA task requires comprehensive spatial-temporal reasoning of the audio-visual scenarios. As shown in Fig. 1 (a), to answer the question “*Is the right instrument in the video always playing?*”, an AVQA model should identify the object guitar at the right, *i.e.*, its visual appearance and *spatial* location, and listen to the audio of all the *temporal* segments.

Previous AVQA works (Li *et al.* 2022; Lao *et al.* 2023) usually adopt a two-stage strategy that first trains an audio-visual model to locate the sounding visual regions (from a spatial perspective) and then uses the question to high-

*Corresponding authors

light key temporal audio-visual segments (from a temporal perspective). Although this strategy is somewhat helpful for spatial-temporal reasoning in audio-visual scenarios, the sounding visual regions highlighted by the independent first stage may not relate to the question and the second stage can focus on the visual information related to the question only at the segment level. In this work, we simultaneously consider object-based spatial and temporal learning. This is inspired by following observations: 1) objects are fundamental units in human cognition (Spelke, Kinzler et al. 2007), 2) objects can be a bridge since a question usually focuses on specific objects, and the audio is also emitted by specific objects, and 3) it is easy to spatially differentiate objects in temporal visual segments and the objects convey clear semantics. Therefore, we propose to explicitly exploit the basic elements – objects, seeking fine-grained clues from the object level for the AVQA task. Although some video QA works (Dang et al. 2021; Xiao et al. 2022) also use visual objects, the audible objects are still in fancy and not explored in the studied AVQA task. In other words, we need to identify *which objects are corresponding to the audio* and *which are really relevant to the question* in each video frame.

To achieve this goal, an AVQA model needs to first understand the question accurately and know what information to look for in audio and visual modalities. Next, the model is expected to find informative clues from audio and visual modalities to answer the question. This two-step decision-making process mimics how humans answer a question. As the example shown in Fig. 1 (a), given the question, we highlight the keywords “*right instrument*” and start to notice the corresponding *guitar* at the right position when watching the visual frames. Meanwhile, we listen to the entire audio to judge if the guitar is “*always playing*”. After going through the entire video, we pick out the audio-visual clues and find that they satisfy the question, giving the answer “*Yes*”. Motivated by this, we design a Question-conditioned Clue Discovery (QCD) module and a Modality-conditioned Clue Collection (MCC) module. QCD is used to seize the keywords in the question related to audio and visual modalities respectively, and MCC is used to enhance the key audio-visual clues, *e.g.*, the relevant objects and audio segments. The QCD and MCC modules are equipped with the Transformer (Vaswani et al. 2017) architecture to highlight key clues via cross-modal attention. These two modules form our backbone network for encoding the multi-modal relations among the audio, objects, and question.

More importantly, we propose an object-aware adaptive-positivity learning strategy to guide the network to enhance cross-modal features. This is inspired by the observation that the video frames contain multiple objects, whereas some of them may not be relevant to the question (especially the ones not mentioned in the question) or the audio (silent objects). As shown in Fig. 1 (a), for the second video segment, the question cares about the “*right instrument*” and only the *guitar* is making sound, thereby the detected “*loudspeaker*” is a distraction since it does not match the question and audio. We define the semantic-matched question-object pair or audio-object pair as *positivity*. These positivity pairs should be highlighted and distinguished from the mismatched pairs

for a better understanding of the audio-visual scenarios. We illustrate the positivity selection process in Fig. 1 (b). For each video segment, we calculate the similarity values between the features of the question-object pairs and the audio-object pairs. The question-object positivity pairs (✓ in orange) and audio-object positivity pairs (✓ in blue) can be adaptively distinguished if their similarity values surpass a pre-set threshold, while the remaining pairs are merged into a negativity set. For model optimization, our positivity learning strategy adopts two object-aware contrastive loss functions to respectively constrain the question-object pairs and audio-object pairs in the positivity set to be close, while the unmatched ones in the negativity set are far away. Note that this contrastive learning scheme can be applied in each video segment, the positivity pairs can be adaptively selected from both spatial and temporal perspectives. As the example shown in Fig. 1 (b), multiple objects are treated as positivities if they are relevant to the question/audio (from a spatial perspective at current frame) and simultaneously exist in the changeable positivity object sets along the timeline (from a temporal perspective in the video). As a result, the model is able to recognize the objects mostly related to the question and audio in each segment, achieving better audio-visual scene understanding.

We refer to the full methodology as Adaptive-Positivity Learning (APL) and evaluate it on the widely-used MUSIC-AVQA (Li et al. 2022) dataset and the experimental results demonstrate that our method is able to identify informative audio-visual clues for better question answering. In summary, we have the following contributions:

- We are the first to explore the fine-grained audible object clues in the AVQA task and present an object-oriented network that effectively encodes the multi-modal relations among the objects, audio, and question.
- We propose an object-aware adaptive-positivity learning strategy including two contrastive loss objectives that adaptively highlight the relevant question-object and audio-object pairs for better model optimization.
- Extensive experimental results show that our method is effective and achieves new state-of-the-art performance on the MUSIC-AVQA dataset.

Related Work

Question Answering in Different Modalities. Intelligent questioning and answering (QA) tasks originated from text-based QA (Mishra and Jain 2016; Song et al. 2022) and are increasingly incorporating more and more modality data such as text, images, video, and audio, resulting in Visual QA (VQA) (Ben-Younes et al. 2017; Liu et al. 2021; Zhou et al. 2023c), Video QA (Lei et al. 2018; Shen et al. 2023), Audio QA (Chuang et al. 2019; Li, Xu, and Hu 2023), and the recently emerged Audio-Visual QA (AVQA) (Yun et al. 2021; Li et al. 2022; Lao et al. 2023) tasks. Among these tasks, the AVQA task in this study is more challenging, as it requires synthesizing audiovisual information to answer the questions, involving more sophisticated cross-modal semantic understanding and spatio-temporal reasoning.

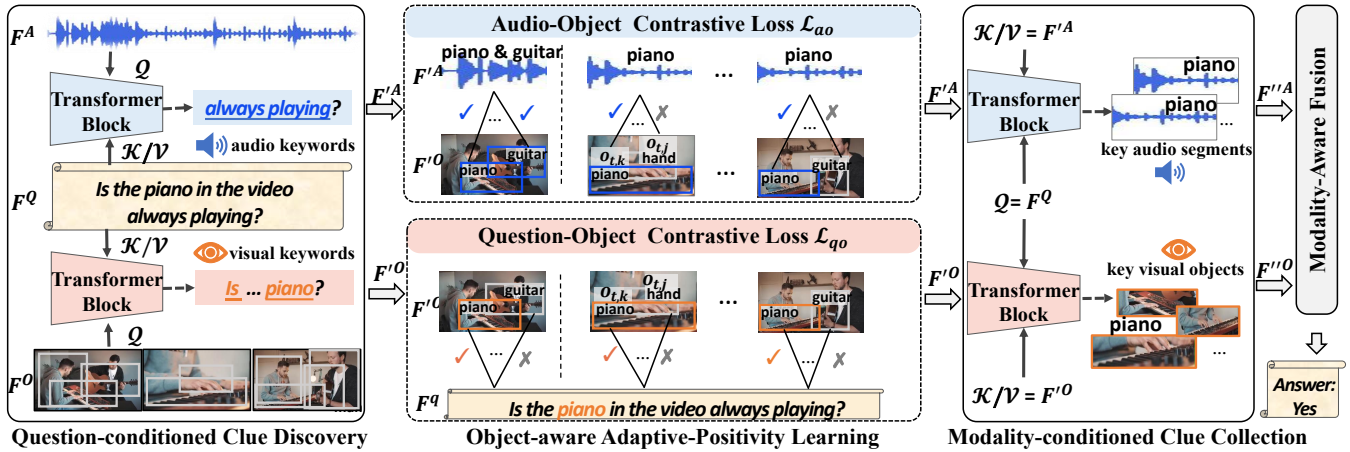


Figure 2: Method Overview. Given the extracted audio/object/question features ($F^{A/O/Q}$), we design two modules to build the backbone network for multimodal interaction. The *Question-conditioned Clue Discovery* (QCD) module individually employs different modalities, *i.e.*, audio and visual (in-frame objects), to highlight the responsive keywords of the question (Left). Then, the *Modality-conditioned Clue Collection* (MCC) module collects relevant audio and visual clues for answering the question (Right). Between these two modules, we design a novel *Object-aware Adaptive-Positivity Learning* strategy that contains two contrastive loss functions to adaptively recognize the mostly relevant question-object and audio-object positivity pairs. This object-aware design facilitates informative cross-modal feature encoding and better model optimization. The learned cross-modal features are finally input to a *Modality-aware Fusion* module to balance the audio-visual clues for answer prediction.

Audio-Visual Question Answering. In the AVQA task, existing works focus on the perception learning of cross-modalities. AVSD (Schwartz et al. 2019) incorporates a multi-modal cross-modal attention mechanism and an LSTM module for answer generation. Pano-AVQA (Yun et al. 2021) studies the audiovisual question-answering in panoramic video scenarios. They propose a transformer-based method that adopts multiple transformer encoders to integrate the cross-modal relations in audio, visual, and language (question). The questions in Pano-AVQA are generally simple and the videos are short (5 seconds). Later, Li et al. (Li et al. 2022) build a new MUSIC-AVQA dataset in long music scenarios and propose a two-stage method. A spatial grounding module is first used to highlight the visual regions corresponding to the audio. Then, the updated visual and audio features are sent to a temporal grounding module that uses the question to generate modality-related attention weights. Lao et al. (Lao et al. 2023) adopt the same framework but utilize the casual analysis to alleviate the shortcut bias problem. Unlike these works, our method explores fine-grained visual objects and activates the audio-matched and question-matched objects to learn multimodal relations. This helps the model to gather more precise audio-visual clues related to the question thus giving better answers.

Methodology

As our model is an object-aware solution, we focus on the exploitation of the object semantics. We first introduce the backbone network that incorporates the object, audio, and question for multimodal feature interaction. Then, we elaborate on the proposed object-aware adaptive-positivity learning strategy that regularizes the question-object and audio-

object semantic alignments for better feature encoding and model optimization.

Backbone Network

The main modules of the proposed network are shown in Fig. 2. Given an audible video, we first down-sample it to T segments and then extract its audio feature and object feature, denoted as $F^A \in \mathbb{R}^{T \times d}$ and $F^O \in \mathbb{R}^{T \times N \times d}$, respectively. N is the total number of detected objects in one video frame, and d is the feature dimension. For the raised question, we extract the word-level feature $F^Q \in \mathbb{R}^{L \times d}$ and the sentence-level feature $F^q \in \mathbb{R}^{1 \times d}$ using an LSTM, where L is the word length of the question. Then, the extracted multimodal features are processed by our multimodal interaction scheme that contains two modules to effectively discover informative audio-visual clues by encoding the cross-modal relations with the question. Specifically, the *Question-conditioned Clue Discovery* (QCD) module encourages the audio/visual modality to focus on relevant keywords in the question. Then, the *Modality-conditioned Clue Collection* (MCC) module enhances the features from informative audio segments and visual objects. After that, all the visual, audio, and textual features are incorporated to predict the answer in the *Modality-Aware Fusion* module. We elaborate on the details of each module below.

Question-conditioned Clue Discovery (QCD). To exactly answer a question, each modality should emphasize more on the relevant keywords in the question. As shown in Fig. 2, there is a question “*Is the piano in the video always playing?*”. The audio modality should focus on the words “*always playing*” in the question, and the visual modality should pay strong attention to the word “*piano*”. This

inspires us to design a question-conditioned clue discovery module. Specifically, we use the Transformer encoder (TFM) (Vaswani et al. 2017) to encode the cross-modal relations of the audio or visual objects to the question. The TFM mainly consists of a multi-head attention (MHA) and a feed-forward network (FF). Overall, the TFM can be defined as:

$$\text{TFM}(\mathcal{Q}, \mathcal{K}, \mathcal{V}) = \text{FF}(\text{MHA}(\mathcal{Q}, \mathcal{K}, \mathcal{V})), \quad (1)$$

where \mathcal{Q} , \mathcal{K} , and \mathcal{V} denotes the *query*, *key*, and *value*, respectively. Then, the QCD process is summarized as:

$$\mathbf{F}'^m = \text{TFM}(\mathbf{F}^m, \mathbf{F}^Q, \mathbf{F}^Q), \quad (2)$$

where $m \in \{O, A\}$ denotes the object feature and audio feature. Hence, we obtain the enhanced visual object feature $\mathbf{F}'^O \in \mathbb{R}^{T \cdot N \times d}$ and the audio feature $\mathbf{F}'^A \in \mathbb{R}^{T \times d}$.

Modality-conditioned Clue Collection (MCC). MCC aims to further collect all the useful audio-visual clues related to the question from each modality for effective answering. The implementation of MCC is similar to QCD but the question is used as the *query* in the multi-head attention. The calculation process is formulated as:

$$\mathbf{F}''^m = \text{TFM}(\mathbf{F}^Q, \mathbf{F}'^m, \mathbf{F}'^m), \quad (3)$$

where $m \in \{O, A\}$ denotes the object or audio modality. Until now, we update the question-related clues from each modality and obtain the corresponding cross-modal features, denoted as $\mathbf{F}''^O \in \mathbb{R}^{L \times d}$ and $\mathbf{F}''^A \in \mathbb{R}^{L \times d}$.

Modality-Aware Fusion. In general, both audio and visual modalities provide useful information for predicting answers. However, some types of questions may be more closely related to a particular modality. The AVQA model should be modality-aware when gathering audio-visual clues. We achieve this goal by generating a modality-aware weight vector $\beta \in \mathbb{R}^{1 \times 2}$ from the sentence-level question feature \mathbf{F}^q using a linear layer followed by a *softmax* activation. The object-related feature \mathbf{F}''^O and audio-related feature \mathbf{F}''^A are then transformed and weighted by β to obtain the final fused feature $\mathbf{f}_{out} \in \mathbb{R}^{1 \times d}$.

Answer prediction. The AVQA (Li et al. 2022) is performed in an open-ended manner, which aims to select the correct answer from a pre-defined answer vocabulary. To this end, we send the final output vector \mathbf{f}_{out} to a linear layer activated by *softmax* function for answer prediction. This generates an answer probability vector $\mathbf{p} \in \mathbb{R}^C$, where C is the number of candidate answers. During inference, the predicted answer is decided by $\hat{c} = \arg \max_c(\mathbf{p})$.

Object-aware Adaptive-Positivity Learning

Existing methods (Li et al. 2022; Lao et al. 2023) usually adopt the cross-entropy loss between the prediction and the ground truth answer to supervise model training. In this work, we emphasize that an AVQA model should understand the *semantic positivity* contained in the multimodalities for better question answering. On one hand, those objects relevant to the question should be highlighted. The video frames usually contain many objects whereas some of

them may be distractions as they do not match the question. On the other hand, those objects corresponding to the audio signal may also be considered. The audio-object semantic matched pairs would also help the model understand the audio-visual scene, *i.e.*, what the sounding objects are. We treat these semantically matched question-object and audio-object pairs as **positivity pairs**. Then, we design two multimodal positivity learning strategies, one for question-object pair matching and the other for audio-object pair matching.

Take the question-object positivity learning as an example, for the t -th segment, the sentence-level question feature and object feature are $\mathbf{q} = \mathbf{F}^q \in \mathbb{R}^{1 \times d}$ and $\mathbf{o}_t = \{\mathbf{o}_{t,i}\}_{i=1}^N = \mathbf{F}_t'^O \in \mathbb{R}^{N \times d}$, respectively. The relevance of the question-object pairs can be measured by the cosine similarity \mathbf{s}_t^{qo} ,

$$\mathbf{s}_t^{qo} = \text{softmax}\left(\frac{\mathbf{q}}{\|\mathbf{q}\|} \otimes \left(\frac{\mathbf{o}_t}{\|\mathbf{o}_t\|}\right)^\top\right), \quad (4)$$

where \otimes is the matrix multiplication, $t = 1, 2, \dots, T$, $\mathbf{s}_t^{qo} \in \mathbb{R}^{1 \times N}$ reflects the similarity of the question with all N detected objects at the t -th segment. A high similarity score of \mathbf{s}_t^{qo} , *e.g.*, $s_{t,i}^{qo}$ indicates that the i -th object is more likely to be relevant to the question.

We then use a simple thresholding strategy to select the highly relevant question-object pairs, forming the positivity set, denoted as \mathcal{P}_t^{qo} . While the remaining pairs are more likely to be irrelevant, forming the negativity set, denoted as \mathcal{N}_t^{qo} . This selection process can be summarized as,

$$\begin{cases} \mathcal{P}_t^{qo} = \{s_{t,i}^{qo} | s_{t,i}^{qo} > \varphi\}, \\ \mathcal{N}_t^{qo} = \{s_{t,i}^{qo} | s_{t,i}^{qo} \leq \varphi\}, \end{cases} \quad (5)$$

where φ is the threshold, $\{\cdot\}$ denotes the selected question-object pair set satisfying the condition. Assume there are K question-object positivity pairs in $\mathcal{P}_t^{qo} = \{s_{t,i,k}^{qo,+}\}_{k=1}^K$, $\mathcal{N}_t^{qo} = \{s_{t,i,j}^{qo,-}\}_{j=1}^{N-K}$ will contain $N - K$ pairs. N is the total number of detected objects and is fixed, whereas K is dynamic which means the positivity pair selection process is **adaptive**. This can be reflected in two aspects: 1) At the beginning of model training, a large portion of detected objects even irrelevant to the question may be regarded as positivity pairs. As the training proceeds, the model gradually acquires the ability to preserve mainly the connections between the most relevant question-object pairs. For each video frame, the model is able to adaptively recognize key objects at different *spatial* locations. 2) The positivity pair selection can be applied in all the T video segments. For each segment, it encourages the model to identify the question-related objects in the current timestamp. This is also an adaptive process along the *temporal* dimension. As the example shown in Fig. 2, our model adaptively selects the “piano” for all the video frames, while ignoring those objects that are irrelevant to the question, such as “hand” and “guitar”.

With the positivity set \mathcal{P}_t^{qo} and negativity set \mathcal{N}_t^{qo} , we then design a question-object positivity contrastive loss \mathcal{L}_{qo} to regularize question-object pairs in \mathcal{P}_t^{qo} to have large cross-modal similarities, while low values for pairs in \mathcal{N}_t^{qo} .

Method	A-QA			V-QA			AV-QA					All	
	CNT	Comp	Avg.	CNT	LOC	Avg.	Exist	LOC	CNT	Comp	Temp	Avg.	Avg.
AVSD (Schwartz et al. 2019)	72.47	62.46	68.78	66.00	74.53	70.31	80.77	64.03	57.93	62.85	61.07	65.44	67.32
Pano-AVQA (Yun et al. 2021)	75.71	65.99	72.13	70.51	75.76	73.16	82.09	65.38	61.30	63.67	62.04	66.97	69.53
STG (Li et al. 2022)	77.78	67.17	73.87	73.52	75.27	74.40	82.49	64.24	69.88	<u>64.67</u>	<u>65.82</u>	69.53	71.59
COCA (Lao et al. 2023)	79.94	67.68	75.42	75.10	75.43	75.23	83.50	66.63	69.72	64.12	65.57	69.96	72.33
APL-Faster R-CNN (ours)	<u>81.81</u>	<u>68.86</u>	<u>77.03</u>	<u>76.44</u>	<u>82.29</u>	<u>79.40</u>	82.29	67.28	<u>72.33</u>	63.49	65.33	<u>70.31</u>	<u>73.91</u>
APL-DETR (ours)	82.40	70.71	78.09	76.52	82.74	79.69	<u>82.99</u>	<u>66.68</u>	73.29	64.76	65.95	70.96	74.53

Table 1: Comparison with existing methods on the MUSIC-AVQA dataset. We evaluate the performance of each model on different question types and use the QA accuracy (%) as the metric. We implement our APL method with different object detectors, *e.g.*, the Faster R-CNN and DETR. The best and second-best results are bolded and underlined, respectively. The ‘CNT’, ‘Comp’, ‘LOC’, ‘Exist’, and ‘Temp’ are the abbreviations for question types ‘Counting’, ‘Comparative’, ‘Location’, ‘Existential’, and ‘Temporal’, respectively.

\mathcal{L}_{qo} can be computed by,

$$\begin{cases} \mathcal{L}_{qo} = \frac{1}{T} \sum_{t=1}^T l_t^{qo}, \\ l_t^{qo} = -\log \frac{\sum_{k=1}^K \exp(s_{t,k}^{qo,+}/\tau)}{\sum_{k=1}^K \exp(s_{t,k}^{qo,+}/\tau) + \sum_{j=1}^{N-K} \exp(s_{t,j}^{qo,-}/\tau)}, \end{cases} \quad (6)$$

where $s_{t,k}^{qo,+} \in \mathcal{P}_t^{qo}$, $s_{t,j}^{qo,-} \in \mathcal{N}_t^{qo}$, $s_{t,k}^{qo,+}$ denotes the similarity of the k -th question-object pair coming from the \mathcal{P}_t^{qo} set. τ is a hyper-parameter balancing the loss computation, and l_t^{qo} is the loss term at the t -th video segment.

Similarly, we can compute the audio-object positivity contrastive loss \mathcal{L}_{ao} that encourages the model to perceive the semantic-matched audio-object pairs. The final positivity learning loss \mathcal{L}_{pc} is the combination of \mathcal{L}_{qo} and \mathcal{L}_{ao} : $\mathcal{L}_{pc} = \mathcal{L}_{qo} + \mathcal{L}_{ao}$. \mathcal{L}_{pc} is then added to the basic cross-entropy loss \mathcal{L}_{ce} as final objective for model optimization:

$$\mathcal{L} = \mathcal{L}_{ce} + \lambda \mathcal{L}_{pc}, \quad (7)$$

where λ is a constant weight to balance the two loss items.

Experiment

Experimental Setup

Dataset. Experiments are conducted on the widely-used MUSIC-AVQA dataset (Li et al. 2022), which contains over 45K question-answer pairs distributed across 9,288 videos spanning more than 150 hours. The dataset encompasses 9 types of QA pairs related to diverse modalities, *i.e.*, the audio-visual (AV-), separate visual (V-), and separate audio (A-), depending on which modalities are used to discover question-related clues for answer prediction. The audio-visual QA pairs occupy a large portion of the entire dataset, and there are five audio-visual question types referring to *existential*, *counting*, *location*, *comparative*, and *temporal*. The dataset is split into the training, validation, and test sets, which comprise 32K, 4K, and 8K QA pairs, respectively.

Evaluation Metric. Following the protocol (Li et al. 2022; Lao et al. 2023), we test the model performance on each question type and use the average answer prediction accuracy *Avg.* as the metric, and it can be obtained by dividing the number of correct instances by the total number of questions of each particular type and then averaging.

Implementation Details. We down-sample each audible video into $T = 10$ non-overlapping segments, resulting in T audio segments and T frames per video. For the audio segments, we use the VGGish (Hershey et al. 2017) network pretrained on the AudioSet (Gemmeke et al. 2017) to extract the 128-D features. About the visual object features, we test two popular off-the-shelf object detectors, *i.e.*, Faster R-CNN (Ren et al. 2015) and DETR (Carion et al. 2020) for object detection and extracting corresponding features. We use independent linear layers to transform both object and audio features into the same 512 dimension. The detected object number N per frame is 36 for Faster R-CNN and 100 for DETR. Accordingly, we set φ in Eq. 5 to 0.028 and 0.011, respectively. During training, the parameter τ in Eq. 6 is set to 0.4 and λ in Eq. 7 is set to 0.3. The initial learning rate is set to $1.75e-4$ when using Faster R-CNN and $1e-4$ for DETR. The learning rate decreases by multiplying 0.1 every 8 epochs with the Adam optimizer. The batch size is set to 64 and we train the model for 20 epochs.

Comparison with State-of-the-arts

We conduct the proposed APL model with the DETR object detector here. The experimental results are shown in Table 1. We report the model performance on different question types. **First**, our method significantly surpasses the previous SOTA method COCA (Lao et al. 2023) on audio question answering and visual question answering: the corresponding average accuracy values are improved by 2.67% and 4.46%, respectively. Questions under these two settings are simpler and only related to a single modality. The performance improvements indicate the proposed method has a better ability to perceive single-modal scenes. **Second**, our method is also superior in more challenging audio-visual question answering. Performance of the most competitive method COCA is close to the previous method STG (Li et al. 2022), while our method achieves 1% improvements. These results reveal that our approach is superior in various QA settings. These results verify the benefits of our object-aware method design.

Ablation Studies

Impacts of Object Detectors. To explore the impact of object detectors in our method, here we employ the pop-




Figure 3: Ablation study of the adaptive contrastive losses. \mathcal{L}_{ce} computes the cross-entropy loss of the predicted answer and ground truth. \mathcal{L}_{qo} and \mathcal{L}_{ao} are the proposed contrastive losses.

#	Module		Avg. (%)			
	QCD	MCC	A-QA	V-QA	AV-QA	All
1	-	-	44.01	18.41	24.08	26.09
2	✓	-	74.74	76.88	68.70	71.94
3	-	✓	73.68	77.50	68.21	71.64
4	✓	✓	78.09	79.69	70.96	74.53

Table 2: Ablation study of main network modules. We report the average accuracy of audio, visual, and audio-visual question answering. ‘-’ denotes that module is not used.

ular Faster R-CNN and DETR to implement the proposed method, respectively. As shown in the bottom part of Table 1, our method using the stronger DETR has the best performance, approaching 74.53% QA accuracy. Also, it is noteworthy that our method equipped with each object detector is superior to existing methods. These results demonstrate the generalization and robustness of our method. In the following experiments, unless otherwise specified, we adopt the DETR as the object detector.

Effects of Main Network Modules. Our network uses the question-conditioned clue discovery (QCD) module and modality-conditioned clue collection (MCC) module to encode question-related multi-modal features. We conduct an ablation study to explore the impacts of these modules. The experimental results are shown in Table 2. The first row (#1) denotes that the QCD and MCC are simply replaced with several fully connected layers to encode features. The model performance in this extreme case without QCD and MCC modules is very low. After adopting the QCD or MCC module independently, the model performance improved substantially to $\sim 72\%$ (#2 and #3 in Table). The model achieves the best performance when using both QCD and MCC modules (#4). These results directly show the effectiveness of each proposed module, which is helpful to encode question-related audio-visual features for answer prediction.

Effects of Adaptive Contrastive Losses \mathcal{L}_{qo} and \mathcal{L}_{ao} . For model optimization, in addition to the basic cross-entropy loss \mathcal{L}_{ce} , we propose two adaptive-positivity contrastive losses, *i.e.*, \mathcal{L}_{qo} and \mathcal{L}_{ao} , which encourage the model to be aware of the semantic matched question-object pairs and audio-object pairs, respectively. We use different loss functions to train the backbone network. As shown in Fig. 3, the model generally performs worst for the majority of the question types merely using \mathcal{L}_{ce} . The model performance is improved after combining the proposed \mathcal{L}_{qo} or \mathcal{L}_{ao} with \mathcal{L}_{ce} for training. In particular, \mathcal{L}_{qo} lifts the visual QA aver-

#	Threshold φ	Avg. (%)			
		A-QA	V-QA	AV-QA	All
1	$\varphi = 0$	77.52	78.52	69.64	73.39
2	$\varphi = 0.009$	77.34	79.22	69.74	73.60
3	$\varphi = 0.010$	77.96	80.26	70.09	74.18
4	$\varphi = \mathbf{0.011}$	78.09	79.69	70.96	74.53
5	$\varphi = 0.012$	77.72	79.52	70.62	74.24
6	$\varphi = 0.013$	77.65	79.85	70.04	73.98

Table 3: Impact of threshold φ used in positivity learning. This experiment uses the DETR that detects $N=100$ objects in each frame. We explore φ around $1/N = 0.010$, *i.e.*, the same attention to all objects.

age accuracy from 78.52% to 80.47%, and \mathcal{L}_{ao} increases the audio-visual QA average accuracy from 69.64% to 70.43%. The QA performance is further boosted when using both contrastive loss items, as shown in the right part of the figure. These results verify the benefits of the proposed contrastive learning that adaptively identifies the highly relevant question-object and audio-object pairs, facilitating the spatial-temporal reasoning of the audio-visual scenes.

Effects of Threshold φ in \mathcal{L}_{qo} and \mathcal{L}_{ao} . φ is the threshold for selecting highly relevant question-object and audio-object positivity pairs for contrastive learning (Eq. 5). Here we test several values of φ to explore its effect. As shown in Table 3, the QA accuracy values of most question types change as the φ varies. And the average performance across all question types approaches the highest when φ is set to 0.011. Note that we use the DETR object detector in this experiment and it extracts $N = 100$ objects for each video frame. The cross-modal similarity of each question-object (or audio-object) pair will be 0.01 if they share the same attention. $\varphi = 0.011$ is slightly larger than this value to ensure that fewer highly matched pairs are selected. When all the objects are selected as positivities ($\varphi = 0$), \mathcal{L}_{qo} (Eq. 6) and \mathcal{L}_{ao} are equal to 0, the model has sub-optimal performance.

Qualitative Results

In this subsection, we display some qualitative results to verify the effectiveness of the proposed APL method. The experiments are based on the Faster R-CNN object detector.

Visualization of the Multi-modal Attention Weights. Our network uses the QCD and the MCC modules (Fig. 2) for multi-modal interaction. To explore whether these two modules actually learn the question-related audio-visual clues, we visualize the attention weights in the multi-head atten-

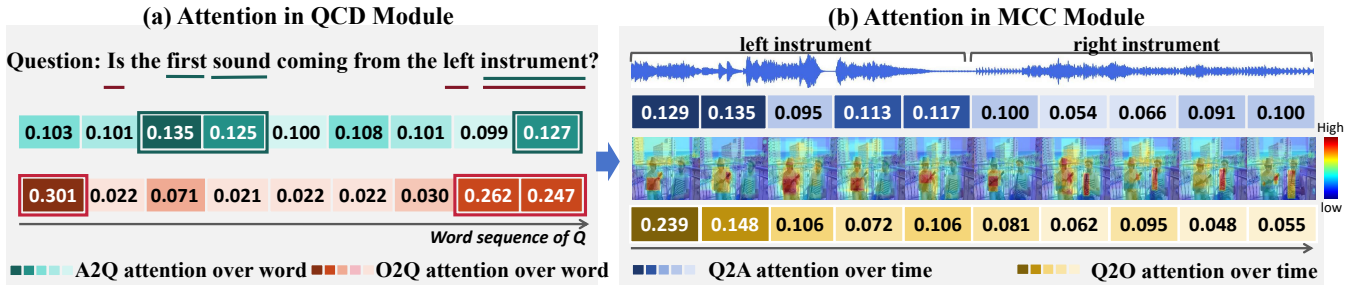


Figure 4: Multi-modal attention visualization during answer reasoning process. A2Q/O2Q from the QCD module is the abbreviation of audio/object to question, and Q2A/Q2O from the MCC module is the abbreviation of question to audio/object.

Q: How many types of musical instruments sound in the video? (Answer of $w/\mathcal{L}_{ao}\&\mathcal{L}_{qo}$: **Two** ✓) vs. (Answer of $w/o \mathcal{L}_{ao}\&\mathcal{L}_{qo}$: **One** X)




Figure 5: The ablation visualization of adaptive positivity learning with \mathcal{L}_{ao} and \mathcal{L}_{qo} . Two instruments make sounds in this example (a). With \mathcal{L}_{ao} and \mathcal{L}_{qo} , these objects relevant to the audio and question are adaptively selected (b, d); Otherwise, the model highlights many irrelevant objects (c, e).

tion (Eq. 1) in these modules. For the QCD module, we display the attention of each modality to each word of the question. As shown in Fig. 4 (a), the audio modality has higher weights to these audio-related words in the question, e.g., “first sound”. Also, the visual modality focuses more on those words relevant to the spatial location of visual objects, such as “left instrument”. With QCD module, each modality successfully seizes the relevant keywords in the question. Then in the subsequent MCC module, as shown in Fig. 4 (b), the left instrument emits sound at the first five video segments, and the attention of the question to audio/visual successfully moves to these temporal segments. Also, the heat maps in the video frames reflecting the question-object attention clue verify that the left instrument highly relevant to the question is highlighted in these earlier segments. These results show that our model learns to understand what visual objects or sound clues the question cares about, and the model further collects highly relevant clues from each modality, ultimately predicting the correct answer “Yes”.

Visualization of the Object-aware Adaptive Positivity Learning. To enhance the interpretability, we showcase an example to see what objects are adaptively selected with the objectives \mathcal{L}_{qo} and \mathcal{L}_{ao} . In Fig. 5 (a), we observe two men playing musical instruments (*clarinet* and *guitar*) in the scene, while some *saxophones* hanging on the wall are silent. In Fig. 5 (b)-(e), we visualize the bounding boxes of the selected objects based on the similarity of audio-object or question-object pairs. Bright brown color indicates higher similarity values. Fig. 5 (b) demonstrates the results with loss \mathcal{L}_{ao} , which adaptively selects the *clarinet* and *guitar* as positivity objects. However, in the absence of \mathcal{L}_{ao} con-

straint, the focus shifts to *saxophone* hanging on the wall, as depicted in Fig. 5 (c). This indicates that \mathcal{L}_{ao} helps to effectively locate sound-related objects. Similar phenomena can be seen in Fig. 5 (d). With loss \mathcal{L}_{qo} , the model successfully highlights the question-related objects, such as various musical instruments. While without \mathcal{L}_{qo} , in Fig. 5 (e), the model pays large attention to more objects irrelevant to the question, such as the *man* and *wall decorations*. With both \mathcal{L}_{ao} and \mathcal{L}_{qo} , our model is able to identify the key objects not only related to the audio but also relevant to the question, thereby it can provide the correct answer “Two”.

Conclusion

In this paper, we propose to explore fine-grained object clues for the challenging audio-visual question answering task. We present a multi-modal interaction network to encode the relations among the visual objects, audio, and question. The core processing units are the question-conditioned clue discovery module and the modality-conditioned clue collection module. The former module makes the audio/visual modality aware of the relevant keywords mentioned in the question, and the latter further highlights the informative clues contained in the audio/visual modality. In addition, we also propose an object-aware adaptive-positivity learning strategy. It uses two contrastive loss functions to encourage the model to adaptively recognize the highly relevant question-object pairs and audio-object pairs. Quantitative and qualitative experimental results verify that our method is able to seize the key clues in the question, audio, and visual objects, making it superior for audio-visual question answering.

Acknowledgements

We would like to express our sincere gratitude to Dr. Liang Zheng for his invaluable comments and insightful suggestions, which helped us to improve this paper significantly. This work was supported by the National Key R&D Program of China (2022YFB4500600), the National Natural Science Foundation of China (62272144, 72188101, 61725203, 62020106007, and U20A20183), and the Major Project of Anhui Province (202203a05020011).

References

- Ben-Younes, H.; Cadene, R.; Cord, M.; and Thome, N. 2017. Mutan: Multimodal tucker fusion for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2612–2620.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 213–229.
- Chuang, Y.-S.; Liu, C.-L.; Lee, H.-Y.; and Lee, L.-s. 2019. Speechbert: An audio-and-text jointly learned language model for end-to-end spoken question answering. *arXiv preprint arXiv:1910.11559*.
- Dang, L. H.; Le, T. M.; Le, V.; and Tran, T. 2021. Hierarchical Object-oriented Spatio-Temporal Reasoning for Video Question Answering. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, 636–642.
- Gemmeke, J. F.; Ellis, D. P.; Freedman, D.; Jansen, A.; Lawrence, W.; Moore, R. C.; Plakal, M.; and Ritter, M. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 776–780.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; et al. 2017. CNN architectures for large-scale audio classification. In *IEEE International conference on Acoustics, Speech and Signal Processing (ICASSP)*, 131–135.
- Lao, M.; Pu, N.; Liu, Y.; He, K.; Bakker, E. M.; and Lew, M. S. 2023. COCA: Collaborative CAusal Regularization for Audio-Visual Question Answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 12995–13003.
- Lei, J.; Yu, L.; Bansal, M.; and Berg, T. L. 2018. Tvqa: Localized, compositional video question answering. *arXiv preprint arXiv:1809.01696*.
- Li, G.; Wei, Y.; Tian, Y.; Xu, C.; Wen, J.-R.; and Hu, D. 2022. Learning to Answer Questions in Dynamic Audio-Visual Scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19108–19118.
- Li, G.; Xu, Y.; and Hu, D. 2023. Multi-Scale Attention for Audio Question Answering. *arXiv preprint arXiv:2305.17993*.
- Li, K.; Guo, D.; Wang, M.; et al. 2021. Proposal-free video grounding with contextual pyramid network. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, 1902–1910.
- Li, K.; Guo, D.; Wang, M.; et al. 2023a. ViGT: proposal-free video grounding with a learnable token in the transformer. *Science China Information Sciences (SCIS)*, 66(10): 202102.
- Li, K.; Li, J.; Guo, D.; Yang, X.; and Wang, M. 2023b. Transformer-based Visual Grounding with Cross-modality Interaction. *ACM Transactions on Multimedia Computing, Communications and Applications (TOMCCAP)*, 19(6): 1–19.
- Liu, Y.; Zhang, X.; Zhang, Q.; Li, C.; Huang, F.; Tang, X.; and Li, Z. 2021. Dual self-attention with co-attention networks for visual question answering. *Pattern Recognition*, 117: 107956.
- Mao, Y.; Zhang, J.; Xiang, M.; Zhong, Y.; and Dai, Y. 2023. Multimodal variational auto-encoder based audio-visual segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 954–965.
- Mishra, A.; and Jain, S. K. 2016. A survey on question answering systems with classification. *Journal of King Saud University-Computer and Information Sciences*, 28(3): 345–361.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1–14.
- Schwartz, I.; Schwing, A. G.; Hazan, T.; et al. 2019. A simple baseline for audio-visual scene-aware dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12548–12558.
- Shen, X.; Li, D.; Zhou, J.; Qin, Z.; He, B.; Han, X.; Li, A.; Dai, Y.; Kong, L.; Wang, M.; et al. 2023. Fine-grained Audible Video Description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10585–10596.
- Song, P.; Guo, D.; Zhou, J.; Xu, M.; and Wang, M. 2022. Memorial gan with joint semantic optimization for unpaired image captioning. *IEEE Transactions on Cybernetics*.
- Spelke, E. S.; Kinzler, K. D.; et al. 2007. Core knowledge. *Developmental Science*, 10(1): 89–96.
- Tian, Y.; Shi, J.; Li, B.; Duan, Z.; and Xu, C. 2018. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 247–263.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1–11.
- Xiao, J.; Yao, A.; Liu, Z.; Li, Y.; Ji, W.; and Chua, T.-S. 2022. Video as conditional graph hierarchy for multi-granular question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2804–2812.

- Yun, H.; Yu, Y.; Yang, W.; Lee, K.; and Kim, G. 2021. Pano-AVQA: Grounded Audio-Visual Question Answering on 360° Videos. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2011–2021.
- Zhou, J.; Guo, D.; Wang, M.; et al. 2022a. Contrastive positive sample propagation along the audio-visual event line. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Zhou, J.; Guo, D.; Zhong, Y.; and Wang, M. 2023a. Improving audio-visual video parsing with pseudo visual labels. *arXiv preprint arXiv:2303.02344*.
- Zhou, J.; Shen, X.; Wang, J.; Zhang, J.; Sun, W.; Zhang, J.; Birchfield, S.; Guo, D.; Kong, L.; Wang, M.; et al. 2023b. Audio-Visual Segmentation with Semantics. *arXiv preprint arXiv:2301.13190*.
- Zhou, J.; Wang, J.; Zhang, J.; Sun, W.; Zhang, J.; Birchfield, S.; Guo, D.; Kong, L.; Wang, M.; and Zhong, Y. 2022b. Audio-Visual Segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 386–403. Springer.
- Zhou, J.; Zheng, L.; Zhong, Y.; Hao, S.; and Wang, M. 2021. Positive sample propagation along the audio-visual event line. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8436–8444.
- Zhou, S.; Guo, D.; Li, J.; Yang, X.; and Wang, M. 2023c. Exploring Sparse Spatial Relation in Graph Inference for Text-Based VQA. *IEEE Transactions on Image Processing (TIP)*.