

# Multimodal Rationales for Explainable Visual Question Answering

Kun Li<sup>1</sup>, George Vosselman<sup>1</sup>, Michael Ying Yang<sup>2</sup>

<sup>1</sup>University of Twente, <sup>2</sup>University of Bath

<sup>1</sup>{k.li, george.vosselman}@utwente.nl, <sup>2</sup>myy35@bath.ac.uk

## Abstract

*Visual Question Answering (VQA) is a challenging task of predicting the answer to a question about the content of an image. Prior works directly evaluate the answering models by simply calculating the accuracy of predicted answers. However, the inner reasoning behind the predictions is disregarded in such a “black box” system, and we cannot ascertain the trustworthiness of the predictions. Even more concerning, in some cases, these models predict correct answers despite focusing on irrelevant visual regions or textual tokens. To develop an explainable and trustworthy answering system, we propose a novel model termed **MRVQA** (Multimodal Rationales for VQA), which provides visual and textual rationales to support its predicted answers. To measure the quality of generated rationales, a new metric **vtS** (visual-textual Similarity) score is introduced from both visual and textual perspectives. Considering the extra annotations distinct from standard VQA, MRVQA is trained and evaluated using samples synthesized from some existing datasets. Extensive experiments across three EVQA datasets demonstrate that MRVQA achieves new state-of-the-art results through additional rationale generation, enhancing the trustworthiness of the explainable VQA model. The code and the synthesized dataset are released under <https://github.com/lik1996/MRVQA2025>.*

## 1. Introduction

Visual Question Answering (VQA) takes as input an image and a natural language query and predicts the corresponding answer based on the understanding of the provided vision-language content. It is more challenging than other fundamental computer vision tasks (e.g., object detection [29]), since it requires fusion and reasoning cross different modalities. With the rapid developments of deep neural networks, recent approaches [2, 13, 20, 31] achieve promising performance on VQA. Normally, the answer accuracy [3] is employed in most VQA systems to measure the correctness of the predictions. However, users cannot know how reliable the predicted answers are, purely based on the accuracy

numbers. In contrast, explainable VQA (EVQA) tries to link the predicted answer to the given image-question pair by providing explanations in different formats (e.g., image regions of interest and textual explanations) to support the answering process. This explanatory mechanism enables VQA systems to reason for multimodal understanding.

Significant progress has been made in the field of “explainable artificial intelligence” (XAI) [40], uncovering the inner workings of AI algorithms and providing insights into how they arrive at their decisions. However, there remains a largely unexplored area in developing human-centric explanations for interactive vision-language tasks, such as VQA. Most of prevailing methods adopt single-track models to improve the explainability: either visual hints (termed “V-Exp”) or textual explanations (termed “T-Exp”), as illustrated in Fig. 1 (b). V-Exp methods [11, 34] predict answers and further extract the regions of interest by analyzing the heat maps or attention maps in the inner neural network layers. However, the provided information is usually coarse and inaccurate while they are restricted to a region level, which is not human-friendly for understanding the decision-making process. T-Exp methods [23, 46, 50] aim to describe the visual content of the image that helps explaining the answer in natural language. The explanation partially reproduces the reasoning process behind answering. However, it still lacks a rich comprehension of vision-language inputs when we aim to directly and visually identify the specific target (i.e., the zebra in the green bounding box from Fig. 1 (d)). Another drawback arises from the inconvenience of textual explanations in complex scenes (e.g., multiple zebras lying on the ground or the presence of horses), leading to longer and more detailed descriptions of the relevant objects. Based on the above concerns in mind, we observe that the current single-track EVQA methods cannot satisfy the credibility of an interactive vision-language system.

In fact, prior to this paper, a few works [8, 47] have made significant strides towards EVQA with compositional reasoning. However, these works essentially rely on various relationships among objects with external information (e.g., scene graphs [36]), resulting in a lack of visual recognition capabilities (as shown in Fig. 1 (c)). Specifically, the




Figure 1. Illustration of the differences between standard VQA, single-track EVQA, compositional EVQA and our multimodal EVQA. Our MRVQA answers the question and simultaneously predicts textual and visual rationales to support the answer.

textual explanations are either chosen from multiple choices or generated using functional programs from a predefined candidate pool. The shortcuts within the external information allow the model to directly utilize these relationships (*e.g.*, spatial relation and object category), which can harm the model’s general performance. In addition, the external information itself is expensive to collect. To our knowledge, few previous studies have explored free-form textual explanations alongside precise visual predictions.

Moreover, the assessment of generated explanation quality remains limited to basic text similarity measures. The existing methods [8, 23, 34] for EVQA employ widely-used metrics for text generation evaluation, such as BLEU [33], ROUGE [26], and CIDEr [44]. Normally, these metrics quantitatively compare the generated text to reference texts from various textual aspects (*e.g.*, n-gram overlap, semantic similarity, word order, and structure), while ignoring the role of visual content in cross-modality tasks. However, an ideal evaluation of multimodal explanations should consider both visual and linguistic perspectives.

Therefore, in this work, we attempt to address these problems with a multimodal EVQA model. Specifically, we propose a new model called **MRVQA** (Multimodal Rationale for Visual Question Answering) that not only predicts accurate answers but also delivers both textual and visual rationales. To highlight the differences between existing EVQA methods and our approach, an example is presented in Fig. 1. Compared to single-track ones, our approach can simultaneously provide multimodal rationales, offering more reliable and user-friendly explanations. For instance, the textual rationale “One zebra gazes in the opposite direction” and the visual rationale, represented by the green bounding box, complement each other to support the answer prediction “no”. Different from compositional ones, our approach generates free-form textual rationales and precise bounding boxes without relying on external information. From a data perspective, we synthesize a multimodal rationale dataset by augmenting existing datasets to alleviate the annotation effort required by EVQA. From a model perspective, we base our approach on a Transformer archi-

tecture [43], projecting features into the latent space of a large language model for textual predictions, and fusing features to extract bounding boxes by leveraging a pre-trained open-vocabulary detector. Additionally, we propose a loss function designed to promote the effective cross-alignment and balance among multitask predictions. From an evaluation perspective, a new metric is introduced, which incorporates both visual and textual views to assess generated rationales and advance EVQA research.

The **main contributions** of this paper are summarized as follows:

- We introduce **MRVQA**, a new EVQA model that facilitates the system to provide multimodal explanations for answer predictions through textual and visual rationale generation. Moreover, a new loss is introduced to balance multitask predictions and generate reasonable rationales.
- To assess the quality of generated rationales, we introduce a new metric called **vtS** for evaluating natural language from a visual perspective.
- We demonstrate the effectiveness of our model with extensive experiments under different settings. The proposed model achieves superior performance over existing methods on three EVQA datasets.

## 2. Related Work

### 2.1. Visual Question Answering

Visual Question Answering (VQA) has been defined as the task of providing an answer to a question about an image. As this task takes as input both images and questions, VQA models address the problem through diverse approaches in different stages, including feature extraction [2], alignment [28], fusion [5], and prediction [17]. With the rapid development of Transformers [43], both natural language processing and computer vision communities have witnessed remarkable advances in vision-language tasks [22, 37]. For instance, MCAN [48] introduced a deep modular co-attention network by leveraging intra-modal self-attentions and cross-modal guided attentions, which enhanced the feature representation and alignment. Some recent works [30, 51] investigate leveraging the vision-language pre-training from large-scale data to enhance the efficacy of VQA models and enable complex reasoning, pioneering the burgeoning trend in the vision-language community. These methods rely on extensive pre-training data and meticulous fine-tuning for downstream tasks, yet they still infer answers through a “black-box” system, which diminishes the explainability of results.

### 2.2. Explanations for VQA

Deep neural networks are typically “black-box” systems because their internal mechanics and decision-making processes are not easily interpretable or understandable. Some

pioneering works have emerged that bring explanations to VQA. Visualization is the first idea adopted to display the inherent details of models, aiding users in understanding the predictions. For instance, VQA-HAT [11] collected human attention maps of where humans choose to look so as to answer questions, visually providing coarse explanations for VQA. Textual explanations utilize natural language to articulate the reasoning behind the system’s predictions. VQA-X [34] combined textual explanations annotated by humans and visual pointing (similar to attention maps) to construct a multimodal reasoning dataset. Subsequently, Li *et al.* [23] further emphasized textual explanations by employing a semi-automatic construction method using caption annotations, creating a large-scale EVQA dataset (*i.e.*, VQA-E). VCR [49] transferred the generation task to a simpler classification one by building a multiple choice movie-scene dataset. GQA-REX [8] constructed a large-scale visually grounded textual explanation dataset by extending the previous GQA [16] dataset with an automatic scheme. However, the displayed decision-making process heavily relies on the pre-defined functional programs with external information (*e.g.*, scene graphs and correlations), which reduces its applicability in a diverse open world.

### 2.3. Generating Multimodal Explanations

The increasing interest in understanding the reasoning behind VQA has led to the development of various methods across both textual and visual perspectives [8, 49]. Park *et al.* [34] first adopted an LSTM [15] to generate textual explanations and subsequently displayed the internal visual focus through attention maps. Li *et al.* [23] employed two prediction heads to generate answers and textual explanations, respectively. Wu *et al.* [46] proposed to leverage the consistency between gradient-based visual explanations and answers to generate robust textual explanations. These earlier works employ encoders to separately extract intra-modality features, which are then simply fused being passed to the final prediction stage. Marasović *et al.* [32] incorporated pretrained language models with object recognition, visual semantics, and scene graphs to generate textual explanations. DMRFNet [50] leveraged pretrained semantic relation embeddings and adaptively fused spatial and semantic features through multi-graph reasoning and fusion layers. REX [8] enhanced cross-modality alignment by considering the semantic similarity between words and image regions. Subsequently, VCIN [47] extended relationships to include words, image regions, and explanation tokens through a gating transformer with variational causal inference. However, these methods restrict the relations with pre-defined compositions (*e.g.*, semantic relations, scene graphs and functional programs). In addition, visual explanations are usually generated through gradient-based segmentation proposals [46] or post-processed visual

regions [8, 47]. Different from the aforementioned methods, our proposed approach focuses on generating free-form textual explanations and object-level visual explanations.

## 3. Method

### 3.1. Problem Formulation

We first revisit the standard VQA setting. Given an image  $I$  and a natural language question  $Q$ , the desired output of a VQA model is an answer  $A$  (*i.e.*, a single word or a short phrase) that relates the queried image. To overcome the limitation in standard VQA, we expand the model’s predictions beyond a single answer to include multimodal rationales, offering explanations from both visual and textual perspectives. Regarding the textual part, we follow previous single-track EVQA methods [23, 50] and employ a similar text description to explain the related context for answering the question, which we refer to as the textual rationale  $TR$ . In contrast to previous methods [7, 34] that use region-level attention maps for visual explanations, we utilize precise bounding boxes of relevant objects as the visual rationale  $VR$ . To this end, given  $I$  and  $Q$ , the goal of our multimodal EVQA model is to predict multiple results: (1) an answer  $A$ ; (2) a textual rationale  $TR$ ; and (3) a visual rationale  $VR$ . It is simply formulated as follows,

$$\{A, TR, VR\} = P(I, Q; \theta), \quad (1)$$

where  $\theta$  represents the parameters of an EVQA model. It is particularly valuable for users who wish to explore more details in an interactive scenario. Moreover, these rationales enhance the model’s explainability and credibility.

### 3.2. Multimodal EVQA Dataset Synthesis

Unfortunately, we have observed that the available EVQA datasets [8, 11, 23, 34, 49] lack either precise visual evidences or related textual descriptions. To avoid the labor-intensive process of annotating explanation-aware VQA examples, we propose a semi-automatic approach to synthesize examples by augmenting an existing EVQA dataset. We select VQA-E [23] for its comprehensive set of examples that include detailed textual explanations and, importantly, for its integration with COCO [27] that provides precise bounding-box annotations for object detection.

Specifically, we first extract all nouns from the question-answer-text triplets in the VQA-E dataset. These nouns usually identify the objects relevant to the explanation of answer predictions. However, analyzing all nouns is impractical due to the presence of over 10K categories. To address this, we conduct a frequency analysis of the nouns, selecting those that appear more than 20 times. This step yields 1,247 distinct nouns that can be matched with the visual labels. Unfortunately, COCO only includes annotations for

80 object categories. To address this, we implement a strategy that classifies candidate nouns into broader categories aligned with the existing COCO annotations. For instance, nouns like “man”, “woman”, “people”, and “girl” are categorized under the “Person” category in COCO. However, not all bounding boxes within a given category are relevant for explanations of a specific image-question pair. For example, although two sheep are present in an image, only one serves as the correct visual rationale for a question concerning the “right sheep”. Additionally, noun-based object extraction may miss key contextual elements. For example, in answering “What is the man doing?” with ground truth “surfing” and “a man on the wave”, only “man” yields a bounding box, while the full rationale should include both the man and the surfboard. To ensure the accuracy and relevance of visual rationales, two annotators, following detailed guidelines, manually refined the bounding boxes over two weeks by removing irrelevant objects and adding those essential to the explanation. This process results in the synthesis of a dataset (VQA-E-Syn). Overall, it comprises 33,726 question-answer pairs derived from 20,367 images, adhering to the original split in VQA-E. Each pair is accompanied by a textual rationale and an average of 2.8 bounding boxes for visual rationales.

### 3.3. Multimodal EVQA Network

The overall framework of the proposed MRVQA model is shown in Fig. 2. Here, the core components of our approach are introduced as follows.

**Input Representation.** Different from previous methods that use separate visual and linguistic feature extractors, we utilize a pre-trained vision-language model (*i.e.*, CLIP [39]) to simultaneously represent both the input image and question. The CLIP model, pre-trained on large-scale image-caption pairs and benefited from transferable representations, outperforms intra-encoders trained solely on in-domain data. With the pre-trained CLIP model, we generate robust visual and linguistic features for our EVQA task.

**Encoder-Decoder.** To effectively understand the queried focus from the visual and linguistic features, it is essential to exploit cross-modality fusion for VQA tasks. To enhance representations without introducing a new backbone or altering the model’s structure, we build on the popular VQA baseline MCAN [48], which cascades several Modular Co-Attention (MCA) layers. The MCA layer comprises self-attention units for intra-modal interaction and guided-attention units for cross-modal interaction.

**Projection Module.** To alleviate the mismatch between the pre-trained CLIP and our textual rationale generator (Large Language Models, LLMs described in the following part), we employ a Transformer-based [43] projection module, facilitating multimodal adaptation and integration. The self-attention mechanism enables the learning of global de-

pendencies between tokens, even within the context of long sentences. Inspired by Darcet et al. [10], we introduce two additional tokens as learnable constants for the CLIP features  $\underline{I}$  and  $\underline{Q}$ , respectively. These constant tokens act as registers, allowing the model to store and retrieve meaningful information efficiently. For the entire projection module, we utilize eight self-attention layers to process the dual-stream visual and linguistic features. The projected features  $I_{proj}$  and  $Q_{proj}$  are calculated as follows,

$$I_{proj} = SA(Cat(\underline{I}, C)), Q_{proj} = SA(Cat(\underline{Q}, C)), \quad (2)$$

$$SA(q, K, V) = Softmax\left(\frac{qK}{\sqrt{d}}\right)V, \quad (3)$$

where  $C$  represents the learnable constants,  $q$ ,  $K$ ,  $V$  and  $d$  denote query, Key, Value, and dimensionality of values, respectively. The projected features  $I_{proj}$  and  $Q_{proj}$  are fused by concatenation as  $f_{proj}$  before fed into the LLMs.

**Predictors.** For the answer prediction, we approach it as a classification problem and utilize a linear fusion module, followed by a sigmoid function, to process the cross-attended features from the decoder. For textual rationales, in contrast to earlier methods that rely on simple LSTMs [15], we utilize advanced LLMs to generate more robust and relevant textual rationales. However, given the complexity of handling multiple tasks in EVQA, the LLMs integrated into the model should facilitate computationally efficient fine-tuning for multimodal learning. To address this, we leverage the auto-regressive large language model GPT-2 [38]. For each token in the sequence, GPT-2, conditioned on the projected features  $f_{proj}$ , generates probability distributions over the entire vocabulary and predicts the next token in the sequence. For visual rationales, we generate precise object-level bounding boxes using Grounding-DINO [29], an open-vocabulary object detector based on text. In contrast to the original work, our approach inputs the detector with the image  $I$ , cross-attended features  $f_{cross}$ , and LLM-processed features  $f_{llm}$ . In addition, we set all detected bounding boxes as a single category. This setting addresses the challenge of aligning objects with the input question for discerning the category information from the image-question pair due to the lack of explicit context.

**Loss Functions.** A hybrid loss function is employed to refine the proposed model, with different loss components tailored for each type of prediction in our EVQA task. Consistent with previous methods [23], we employ the widely-used binary cross-entropy loss [42] as  $L_{ans}$  to supervise answer predictions. To train the model for generating textual rationales from the LLM-processed features  $f_{llm}$ , we utilize the simple yet effective cross-entropy loss as  $L_{tr}$  in our case. Given  $M$  questions,  $N$  candidate answers,  $M$  textual

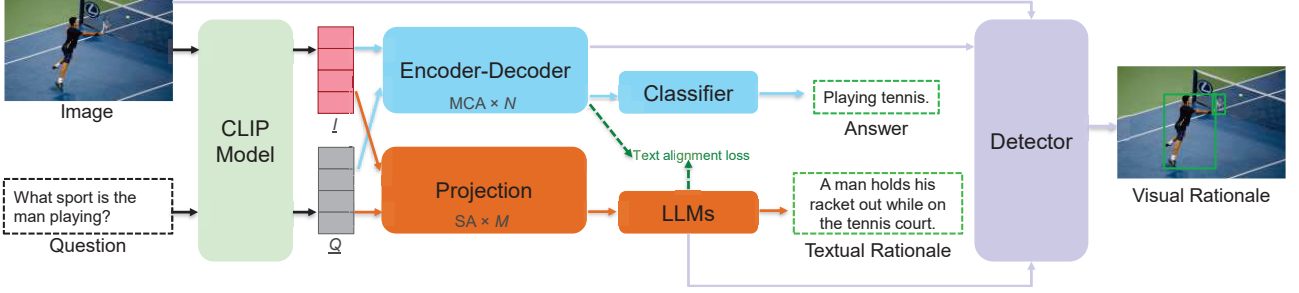


Figure 2. The framework of the proposed MRVQA model. The notations  $\underline{I}$  and  $\underline{Q}$  represent the pre-trained CLIP [39] features for the image and the question, respectively. The proposed text alignment loss helps to improve the coherence of text predictions.

rationales,  $L_{ans}$  and  $L_{tr}$  are calculated as follows,

$$L_{ans} = - \sum_i^M \sum_j^N s_{ij} \log(\hat{s}_{ij}) - (1 - s_{ij}) \log(1 - \hat{s}_{ij}), \quad (4)$$

$$L_{tr} = - \sum_i^M \sum_t^l \log p_{\theta}(w_t^i | f_{llm}^i, w_1^i, \dots, w_{t-1}^i), \quad (5)$$

where  $l$  and  $\theta$  represent the maximum length of tokens and the model’s trainable parameters, respectively. The sequenced tokens are denoted as  $w^i = w_1^i, \dots, w_l^i$ . For visual rationale generation, we directly adopt the existing loss function from Grounding-DINO [29] as  $L_{vr}$ , including L1 loss, GIoU loss, and contrastive loss. To train the entire model using multi-task learning, we take a linear combination of the different loss components.

Moreover, to improve the coherence between answers and their corresponding textual rationales, we introduce a new loss function termed **text alignment loss** (abbreviated as  $L_{ta}$ ). Specifically, the improvement of text coherence is achieved by the alignment between cross-attended features  $f_{cross}$  and LLM-processed features  $f_{llm}$ . We use the cosine similarity function to measure the feature similarity between these dual objectives from the textual side. The text alignment loss  $L_{ta}$  is obtained as follows,

$$L_{ta} = 1 - \sum_i^M \sum_t^l \cos(f_{llm}^i, f_{cross}^i), \quad (6)$$

where  $f'_{cross}$  and  $f'_{llm}$  represent the co-matched features in a common vector space through two MLP layers. Before incorporating it into the final loss function, we use a trade-off parameter  $\lambda$  to balance the basic and alignment components of the loss. The overall loss  $L$  of the model is defined as follows,

$$L = L_{ans} + L_{tr} + L_{vr} + \lambda L_{ta}, \quad (7)$$

### 3.4. Visual-Textual Similarity Score

We utilize the powerful pre-trained text embedding model, GTE [25], to map the generated textual rationales for similarity comparison to the ground truths. There are several

reasons for selecting this model: (1) The GTE model employs contrastive learning across a diverse range of datasets, which allows it to generalize effectively to various data; (2) It avoids the reliance on prompt formulation seen in LLMs; (3) It handles pairs of short and long texts more effectively without requiring word-level alignments, which is ideal for comparing our generated free-form texts; (4) It outperforms previous embedding methods in comprehensive benchmarks [25]. Specifically, we use the cosine function to assess textual similarity. To handle the negative values produced by the cosine function, we map the cosine similarity scores to a range from 0 to 1. The textual similarity ( $TS$ ) is calculated as follows,

$$TS = \frac{1 + \cos(\text{GTE}(T_{pred}), \text{GTE}(GT))}{2}, \quad (8)$$

where  $T_{pred}$  and  $GT$  represent the generated textual rationale and the ground truth, respectively. We assess the predicted bounding boxes by calculating the Average Precision [14] (AP) in comparison to ground truths.

To this end, we introduce a new metric **vtS** (visual-textual Similarity score), which combines visual evaluation and textual similarity to provide a balanced assessment of the generated rationale quality. The metric **vtS** is defined as follows,

$$vtS = \frac{2 \times TS \times AP}{TS + AP}, \quad (9)$$

## 4. Experiments

### 4.1. Experiment Setup

**Datasets.** We conducted primary experiments on the synthesized multimodal EVQA dataset (VQA-E-Syn) introduced in Section 3.2. In addition, we evaluated the model on two existing EVQA datasets: VQA-X [34] and GQA-REX [8], to access its adaptability and generalization capability under different settings.

**Training and Inference Settings.** Under full supervision, we refined the model in an end-to-end manner, referred to as MRVQA-E. To adapt to various practical scenarios, we simplified the MRVQA model to predict the

required results under different conditions. For example, we developed a variant, MRVQA-C (combined), by removing the visual generator component, allowing us to test the model’s ability to generalize in the absence of visual rationale supervision. This variant utilized the generated textual rationales and input images to guide the detector in predicting bounding boxes during inference stage. In addition, we further simplified the model to focus exclusively on answer prediction or textual rationale prediction, developing variants called MRVQA-A and MRVQA-TR, respectively.

**Implementation Details.** We utilized the pre-trained CLIP model [39] with ViT-base [12] and Transformer [43] for input representation. The AdamW [18] optimizer was used with a learning rate of  $2 \times 10^{-5}$ . All the models were trained for 50 epochs with a batch size of 64 using two A40 GPUs, within a consistent PyTorch environment.

**Evaluation Metrics.** For answer evaluation, we maintained consistency with prior works by using accuracy as the metric. To evaluate generated textual rationales, we employed five mainstream metrics, namely BLEU-4 [33], METEOR [4], ROUGE [26], CIDEr [44] and SPICE [1] (B, M, R, C, S for their abbreviations, respectively). Additionally, the newly introduced metric vtS was used to capture the cross-modality similarity by leveraging the visual response.

## 4.2. Quantitative Results and Comparison

### 4.2.1. Results on the synthesized VQA-E-Syn

Table 1 presents a comparison with state-of-the-art EVQA methods trained and evaluated on the synthesized VQA-E-Syn dataset. Our proposed MRVQA-E model significantly enhanced the generation of textual rationales for multimodal EVQA, outperforming the previous methods. For example, our approach achieved a 2.41% improvement in CIDEr compared to VCIN [47], highlighting its superior fluency and coherence. To ensure a fair comparison with previous methods lacking visual prediction capabilities, we adopted a combined approach similar to our MRVQA-C variant. The results demonstrate that our approach achieved a 1.27% improvement over VCIN in vtS. Additionally, we also evaluated our MRVQA-TR and MRVQA-C models. Both variants demonstrated slightly superior performance compared to the previous methods. More importantly, the comparison between the end-to-end model and its variants highlights that incorporating visual rationale generation improves linguistic prediction performance.

Table 2 shows the results in terms of answer accuracy. The methods listed in the top section are trained solely for answer prediction, whereas those in the middle section have the capability to generate both answers and textual explanations. OFA [45], as a powerful unified sequence-to-sequence framework, facilitates generating multiple predictions through separate fine-tuning processes on task-specific datasets (e.g., VQA and image captioning). To adapt OFA

for our task, generating texts within a unified inference pass, we performed a single fine-tuning of the OFA-base model on the combined samples from VQA-E-Syn. Methods lacking answer prediction details or accuracy results [34, 35, 46] are excluded from this table. Our MRVQA-E model achieved an overall accuracy of 75.01% and an average accuracy of 72.46% on VQA-E-Syn, surpassing the previous methods. Notably, the model showed significant improvements in the *Number* and *Yes/No* question categories, with increases of 1.43% and 2.59%, respectively. Similarly, comparing MRVQA-E with its variants (last three rows) reveals that incorporating additional rationale supervision enhanced the performance of the answering process.

### 4.2.2. Results on VQA-X

Table 3 reports the results compared to previous EVQA methods on VQA-X [34]. As VQA-X only supports linguistic examples, to ensure a fair comparison, we conducted experiments with MRVQA-C for only answer and textual rationale generation. Our approach achieved a 78.8% overall accuracy, which is higher than the previous state-of-the-art VCIN by 1.1%. Also, the model showed significant improvements in BLEU-4 and CIDEr metrics, with increases of 0.7% and 4.7%, respectively, compared to VCIN.

### 4.2.3. Results on GQA-REX

Table 3 also shows the results on GQA-REX [8]. Note that the previous VCIN model predicts region-level patches, which relies on priors in visual encodings (e.g., patch location). Due to the single-category object detection used in our settings, we exclude this unfair comparison and only report the textual results. Our approach achieved an overall answer accuracy of 61.28%, representing a 0.67% improvement over the previous state-of-the-art VCIN. Meanwhile, our model maintained competitive performance even without external information, achieving 58.83% and 83.07% for BLEU-4 and ROUGE scores, respectively.

## 4.3. Ablation Studies

### 4.3.1. Ablation of the Projection Module

For comparison, we conducted an ablation study by replacing the projection module in MRVQA-E with a linear layer, as used in prior vision-language methods. The results in Table 4 show that our Transformer-based projection module significantly enhanced the performance of textual rationale generation compared to the linear layer across various metrics. For example, the incorporation of our module resulted in a 5.46% improvement in CIDEr score. Additionally, the enhancement in linguistic predictions contributed to a 2.02% increase in the multimodal rationale evaluation metric vtS, highlighting the efficacy of the module.

Table 1. Comparison results on VQA-E-Syn regarding textual rationale generation. All numbers are reported in percentage (%). The best results are **bold** while the second best are underlined.

Method	Image Feature	Question Feature	Text Predictor	BLEU-4	METEOR	ROUGE	CIDEr	SPICE	vtS
PJ-X [34]	CNN	LSTM	LSTM	8.78	16.94	35.65	89.31	15.32	49.19
VQA-E [23]	CNN	GRU	LSTM	8.93	17.02	35.96	90.84	16.83	51.88
FME [46]	CNN	GRU	LSTM	12.81	21.26	37.69	89.16	18.77	53.82
CCM [35]	CNN	LSTM	LSTM	8.85	17.86	38.42	92.46	18.35	54.44
DMRFNet [50]	CNN	GRU	GRU	13.34	19.44	40.76	95.68	20.41	56.06
REX [8]	VilBert	VilBert	LSTM	14.71	21.35	41.86	100.75	21.08	56.38
OFA [45]	CNN	BPE	Seq2Seq	<u>16.38</u>	22.09	42.74	103.25	<u>22.65</u>	56.58
VCIN [47]	V-Bert	V-Bert	Trans	<u>15.77</u>	22.38	43.10	104.63	22.07	57.89
MRVQA-TR (ours)	CLIP	CLIP	GPT2	15.08	21.77	42.04	102.86	21.22	57.30
MRVQA-C (ours)	CLIP	CLIP	GPT2	15.84	<u>22.41</u>	<u>43.57</u>	<u>105.31</u>	22.19	<u>58.19</u>
MRVQA-E (ours)	CLIP	CLIP	GPT2	<b>16.97</b>	<b>23.02</b>	<b>44.28</b>	<b>107.04</b>	<b>23.68</b>	<b>59.16</b>

Table 2. Comparison results on VQA-E-Syn regarding answer accuracy. OA and AA represent overall and average accuracy, respectively.

Method	Number	Yes/No	Other	OA	AA
MUTAN [5]	49.23	65.52	60.27	61.07	58.34
BUTD [2]	52.68	64.07	65.72	64.13	60.83
MCAN [48]	59.19	71.70	72.37	71.09	67.76
LXMERT [41]	59.93	72.34	73.04	71.76	68.44
UNITER [9]	61.58	73.81	72.66	72.14	69.35
BLIP [20]	<u>63.68</u>	72.43	74.83	73.16	70.31
CMQEF [24]	60.02	71.63	72.20	71.03	67.95
VQA-E [23]	54.26	63.98	66.45	64.67	61.57
DMRFNet [50]	59.08	72.28	73.84	72.15	68.40
REX [8]	59.73	72.91	74.18	72.84	68.94
OFA-base [45]	60.48	73.26	75.37	73.49	69.70
VCIN [47]	61.59	73.47	73.10	72.43	69.39
MRVQA-A (ours)	57.02	72.80	73.41	71.89	67.75
MRVQA-C (ours)	63.20	<u>76.03</u>	<u>75.72</u>	<u>74.81</u>	<u>71.65</u>
MRVQA-E (ours)	<b>65.11</b>	<b>76.40</b>	<b>75.87</b>	<b>75.01</b>	<b>72.46</b>

Table 3. Comparison results on VQA-X [34] and GQA-REX [8] for answer and textual rationale evaluation.

Dataset	Method	B	M	R	C	S	OA
VQA-X	PJ-X [34]	19.5	18.2	43.7	71.3	15.1	-
	VQA-E [23]	20.7	18.6	44.2	75.6	15.6	70.2
	FME [46]	24.4	19.5	47.4	88.8	17.9	-
	CCM [35]	21.1	19.7	44.9	73.9	16.2	-
	DMRFNet [50]	20.5	19.9	41.3	74.5	17.6	72.6
	VCIN [47]	25.9	21.6	48.5	93.7	<b>19.4</b>	77.7
	MRVQA-C	<b>26.6</b>	<b>22.0</b>	<b>48.9</b>	<b>98.4</b>	19.2	<b>78.8</b>
GQA-REX	VQA-E [23]	42.56	34.51	73.59	358.20	40.39	57.24
	FME [46]	42.45	34.46	73.51	357.10	40.35	56.92
	REX [8]	54.59	39.22	78.56	464.20	46.80	57.77
	VCIN [47]	58.65	41.57	81.45	519.23	54.63	60.61
	MRVQA-C	<b>58.83</b>	<b>42.01</b>	<b>83.07</b>	<b>528.42</b>	<b>54.96</b>	<b>61.28</b>

Table 4. Ablation results of using different projection modules for multimodal rationale generation on VQA-E-Syn.

Module	B	M	R	C	S	vtS
Linear [19]	15.07	22.13	42.74	101.58	22.06	57.14
Ours-Transformer	<b>16.97</b>	<b>23.02</b>	<b>44.28</b>	<b>107.04</b>	<b>23.68</b>	<b>59.16</b>

Table 5. Ablation results of different strategies for text predictions on VQA-E-Syn.

Strategy	Data	B	S	OA	AA
Generator	{A + TR}	15.16	20.98	70.42	68.39
Classifier + Generator	{A, TR}	<b>15.84</b>	<b>22.19</b>	<b>74.81</b>	<b>71.65</b>

### 4.3.2. Ablation of Unified or Separate Text Predictor

Some approaches [21] use a unified strategy for all text-related downstream tasks, which often results in longer and more complex expressions. To compare it with our approach, we adapt the unified strategy by integrating answer prediction into the text generation, using samples in {answer, textual rationale} format (e.g., “the boy because only the boy can catch the frisbee”). Table 5 reports the performance of different text prediction strategies. While the performance of textual rationale generation showed a slight change, the overall accuracy significantly dropped by 4.39%, indicating a decrease in the robustness of an answering model. Thus, we employed dual-head predictors for linguistic outputs to address this issue in the EVQA task.

### 4.3.3. Ablation of the Text Generator

To evaluate LLMs’ effectiveness compared to commonly used text generation models, such as LSTMs and Transformer, we replaced this component in MRVQA-E with these alternative modules. The results are shown in Table 6. Comparing the first three rows, GPT-2 significantly outperformed standard LSTMs and yielded slightly better results than Transformer, even with frozen parameters. Subsequently, we allowed the parameters of GPT-2 to be trainable to boost generation performance. While this required additional effort to train the LLMs, it enhanced the quality of the generated textual rationales. For instance, it resulted in a 13.78% improvement in the CIDEr score. This advancement in linguistic understanding also positively impacted visual predictions, as evidenced by a 5.82% increase in the vtS score. We further replaced the text generator with the more advanced GPT-3 model [6], known for its superior performance in complex reasoning and contextual un-

Table 6. Ablation results of different text generators under either freezing or fine-tuning settings on VQA-E-Syn.

Generator	Mode	B	M	R	C	S	vtS
LSTM [15]	-	12.65	18.39	36.77	88.82	18.45	53.14
Transformer [43]	-	12.89	18.91	37.03	89.14	18.86	53.17
GPT-2 [38]	freezing	13.27	20.08	38.33	93.26	19.25	53.34
GPT-3 [6]	freezing	15.06	21.27	39.74	98.90	20.57	56.25
MRVQA-E	fine-tuning	<b>16.97</b>	<b>23.02</b>	<b>44.28</b>	<b>107.04</b>	<b>23.68</b>	<b>59.16</b>

Table 7. Ablation results of loss functions on VQA-E-Syn.

$Loss_A$	$Loss_{TR}$	$Loss_{VR}$	$Loss_{TA}$	B	S	vtS	OA
✓	✓			15.84	22.19	-	71.65
✓	✓	✓		15.13	22.34	59.01	72.28
✓	✓	✓	$\lambda = 0.1$	16.49	23.51	59.10	74.76
✓	✓	✓	$\lambda = 0.5$	<b>16.97</b>	<b>23.68</b>	<b>59.16</b>	<b>75.01</b>
✓	✓	✓	$\lambda = 1$	16.25	23.40	59.07	74.59

derstanding. However, GPT-3’s size (175 billion parameters compared to GPT-2’s 1.5 billion) poses significant challenges, and moreover, it is not available for fine-tuning. We utilized GPT-3 with frozen parameters for generating textual rationales. The results indicate that using more advanced LLMs does not significantly impact the performance of text generation for our EVQA task.

#### 4.3.4. Ablation of Loss Functions

To further evaluate the proposed loss function, we conducted ablation experiments on VQA-E-Syn to verify its effectiveness. The results are reported in Table 7. Comparing the first two rows, we can see that incorporating visual supervision significantly improved overall performance across all types of predictions, which is consistent with our earlier observations in Table 1 and Table 2. Comparing the last four rows, we observed a notable enhancement in performance due to the inclusion of our proposed text alignment loss. For example, all configurations utilizing  $Loss_{TA}$  showed at least a 1% improvement in SPICE scores over the others. In addition, we performed a series of experiments to determine the optimal value for the trade-off parameter  $\lambda$ . The results from the last three rows indicate that setting  $\lambda = 0.5$  yielded the best performance.

#### 4.4. Qualitative Analysis

Figure 3 displays three examples from the synthesized multimodal dataset VQA-E-Syn. Compared to the previous methods such as DMRFNet [50] and VCIN [47], our proposed approach demonstrates superior performance in both answering and rationale support. For example, in the first column, both DMRFNet and VCIN struggled to comprehend the query’s intent, focusing on the object “woman,” although VCIN ultimately predicted the correct answer “No”. In contrast, MRVQA effectively performed cross-modal understanding, accurately identifying the target “man” referring to the statue and correctly predicting the answer. The

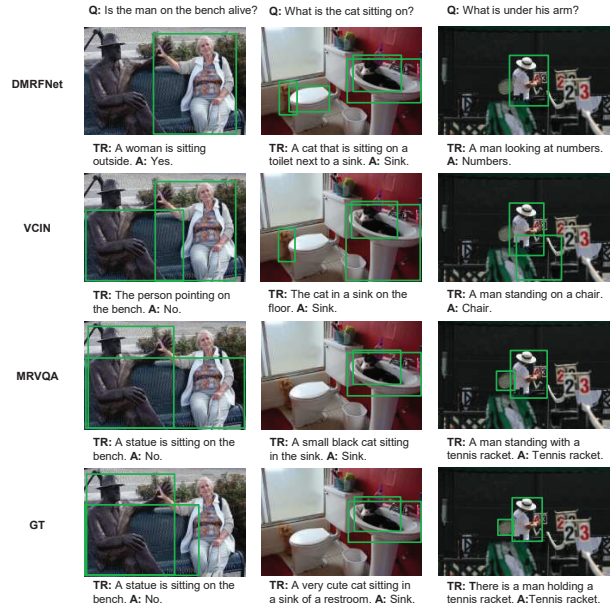


Figure 3. Qualitative comparison of results. The bounding boxes in green display the relevant visual rationales.

medium column illustrates our model’s ability to capture relevant objects in both visual and textual rationales. It successfully identified the action “sitting on” and located the target, while the other methods were misled by “the other cat” or “the toilet”. In the last example, despite the lack of category-level information, MRVQA accurately predicted the answer and delivered consistent rationales, as evidenced by the references to “the man” and “the tennis racket”.

## 5. Conclusion

In this paper, we propose a multimodal EVQA method that predicts not only answers but also textual and visual rationales to support the answers. To minimize the annotation effort, we synthesize a dataset by augmenting existing datasets through a semi-automatic matching engine. To solve the multimodal EVQA problem, we introduce MRVQA that generates textual rationales with large language models and visual rationales via an pre-trained detector. A loss function is introduced to ensure the consistency between the answer and textual rationale generation. Furthermore, we propose a metric for evaluating the quality of generated rationales from a visual perspective. The extensive experiments across three EVQA datasets demonstrate that MRVQA significantly outperforms existing state-of-the-art EVQA methods. We believe this work will inspire further research in the EVQA field.

**Acknowledgement** Kun Li has been supported by China Scholarship Council. Michael Ying Yang has been supported by EU HORIZON-CL42023-HUMAN-01-CNECT XTREME (grant no.101136006).

## References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Proceedings of the European Conference on Computer Vision*, pages 382–398, 2016. 6
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6077–6086, 2018. 1, 2, 7
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015. 1
- [4] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, 2005. 6
- [5] Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. Mutan: Multimodal tucker fusion for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2612–2620, 2017. 2, 7
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020. 7, 8
- [7] Arjun Chandrasekaran, Viraj Prabhu, Deshraj Yadav, Prithvijit Chattopadhyay, and Devi Parikh. Do explanations make vqa models more predictable to a human? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1036–1042, 2018. 3
- [8] Shi Chen and Qi Zhao. Rex: Reasoning-aware and grounded explanation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15586–15595, 2022. 1, 2, 3, 5, 6, 7
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European Conference on Computer Vision*, pages 104–120, 2020. 7
- [10] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *International Conference on Learning Representations*, 2024. 4
- [11] Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Computer Vision and Image Understanding*, 163: 90–100, 2017. 1, 3
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 6
- [13] Radhika Dua, Sai Srinivas Kancheti, and Vineeth N Balasubramanian. Beyond vqa: Generating multi-word answers and rationales to visual questions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 1623–1632, 2021. 1
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 2010. 5
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 3, 4, 8
- [16] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6700–6709, 2019. 3
- [17] Kushal Kafle and Christopher Kanan. Answer-type prediction for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4976–4984, 2016. 2
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [19] Chunyuan Li. Large multimodal models: Notes on cvpr 2023 tutorial. *arXiv preprint arXiv:2306.14895*, 2023. 7
- [20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900, 2022. 1, 7
- [21] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, pages 19730–19742, 2023. 7
- [22] Kun Li, George Vosselman, and Michael Ying Yang. Hrvqa: A visual question answering benchmark for high-resolution aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 214:65–81, 2024. 2
- [23] Qing Li, Qingyi Tao, Shafiq Joty, Jianfei Cai, and Jiebo Luo. Vqa-e: Explaining, elaborating, and enhancing your answers for visual questions. In *Proceedings of the European Conference on Computer Vision*, pages 552–567, 2018. 1, 2, 3, 4, 7
- [24] Shengdong Li, Chen Gong, Yuqing Zhu, Chuanwen Luo, Yi Hong, and Xueqiang Lv. Context-aware multi-level question embedding fusion for visual question answering. *Information Fusion*, 102:102000, 2024. 7
- [25] Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*, 2023. 5
- [26] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 2, 6

- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755, 2014. 3
- [28] Fenglin Liu, Yuanxin Liu, Xuancheng Ren, Xiaodong He, and Xu Sun. Aligning visual regions and textual concepts for semantic-grounded image representations. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [29] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 1, 4, 5
- [30] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [31] Aihua Mao, Zhi Yang, Ken Lin, Jun Xuan, and Yong-Jin Liu. Positional attention guided transformer-like architecture for visual question answering. *IEEE Transactions on Multimedia*, 25:6997–7009, 2022. 1
- [32] Ana Marasović, Chandra Bhagavatula, Jae Sung Park, Ronan Le Bras, Noah A Smith, and Yejin Choi. Natural language rationales with full-stack visual reasoning: From pixels to semantic frames to commonsense graphs. In *Findings of the Association for Computational Linguistics: EMNLP*, pages 2810–2829, 2020. 3
- [33] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 2, 6
- [34] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8779–8788, 2018. 1, 2, 3, 5, 6, 7
- [35] Badri Patro, Shivansh Patel, and Vinay Namboodiri. Robust explanations for visual question answering. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1577–1586, 2020. 6, 7
- [36] Tianwen Qian, Jingjing Chen, Shaoxiang Chen, Bo Wu, and Yu-Gang Jiang. Scene graph refinement network for visual question answering. *IEEE Transactions on Multimedia*, 25: 3950–3961, 2022. 1
- [37] Tianwen Qian, Ran Cui, Jingjing Chen, Pai Peng, Xiaowei Guo, and Yu-Gang Jiang. Locate before answering: Answer guided question localization for video question answering. *IEEE Transactions on Multimedia*, 26:4554–4563, 2023. 2
- [38] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 4, 8
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763, 2021. 4, 5, 6
- [40] Wojciech Samek and Klaus-Robert Müller. Towards explainable artificial intelligence. *Explainable AI: interpreting, explaining and visualizing deep learning*, pages 5–22, 2019. 1
- [41] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 5100–5111, 2019. 7
- [42] Damien Teney, Peter Anderson, Xiaodong He, and Anton 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*, pages 4223–4232, 2018. 4
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, page 5998–6008, 2017. 2, 4, 6, 8
- [44] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4566–4575, 2015. 2, 6
- [45] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340, 2022. 6, 7
- [46] Jialin Wu and Raymond Mooney. Faithful multimodal explanation for visual question answering. In *Proceedings of the ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 103–112, 2019. 1, 3, 6, 7
- [47] Dizhan Xue, Shengsheng Qian, and Changsheng Xu. Variational causal inference network for explanatory visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2515–2525, 2023. 1, 3, 6, 7, 8
- [48] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6281–6290, 2019. 2, 4, 7
- [49] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6720–6731, 2019. 3
- [50] Weifeng Zhang, Jing Yu, Wenhong Zhao, and Chuan Ran. Dmrfnet: deep multimodal reasoning and fusion for visual question answering and explanation generation. *Information Fusion*, 72:70–79, 2021. 1, 3, 7, 8

- [51] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 13041–13049, 2020. [2](#)