

MGA-VQA: MULTI-GRANULARITY ALIGNMENT FOR VISUAL QUESTION ANSWERING

Anonymous authors

Paper under double-blind review

ABSTRACT

Learning to answer visual questions is a challenging task since the multimodal inputs are within two feature spaces. Moreover, reasoning in visual question answering requires the model to understand both image and question, and align them in the same space, rather than simply memorize statistics about the question-answer pairs. Thus, it is essential to find component connections between different modalities and within each modality to achieve better attention. Previous works learned attention weights directly on the features. However, the improvement is limited since these two modality features are in two domains: image features are highly diverse, lacking structure and grammatical rules as language, and natural language features have a higher probability of missing detailed information. To better learn the attention between visual and text, we focus on how to construct input stratification and embed structural information to improve the alignment between different level components. We propose *Multi-Granularity Alignment architecture for Visual Question Answering task* (MGA-VQA), which learns intra- and inter-modality correlations by multi-granularity alignment, and outputs the final result by the decision fusion module. In contrast to previous works, our model splits alignment into different levels to achieve learning better correlations without needing additional data and annotations. The experiments on the VQA-v2 and GQA datasets demonstrate that our model significantly outperforms non-pretrained state-of-the-art methods on both datasets without extra pretraining data and annotations. Moreover, it even achieves better results over the pre-trained methods on GQA.

1 INTRODUCTION

Visual Question Answering (VQA) continues to be a topic of interest due to its broad range of applications, such as the early education system and visual chatbot. Multidisciplinary research on VQA has yielded innovations in multi-modality alignment (Ilievski & Feng, 2017), natural language understanding (Yi et al., 2018), image understanding (Goyal et al., 2017b), and even multimodal reasoning (Cadene et al., 2019).

Early methods addressed the multimodal problem by the simple concatenation of visual features obtained from Convolutional Neural Network (CNN) and natural language features obtained from Recurrent Neural Network (RNN) (Antol et al., 2015). Although they established milestones in this field and have provided insights into merging features,

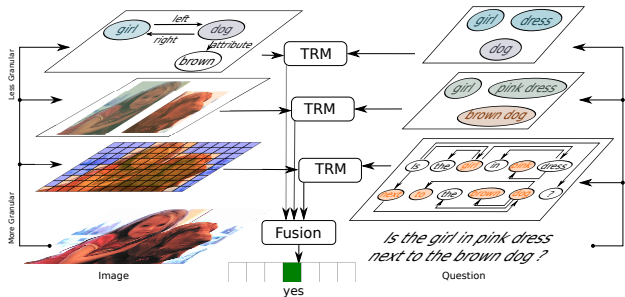


Figure 1: Multi-granularity alignment in Visual Question Answering. For each modality, the input is represented by different granularity levels, and Transformers (TRMs) are responsible for learning the correlations between levels.

such simple fusion methods do not offer good performance. This has motivated many additional related works that primarily focus on further processing the features before merging, i.e., embedding these features into a new space (Ren et al., 2015) or utilizing an attention mechanism to achieve better alignment (Nam et al., 2017). Another related research direction has involved the construction of graphs to represent image information (Hudson & Manning, 2019b). These explorations aimed to determine an effective way to achieve better alignment between multimodal features; however, the direct fusion of whole fine-grained image features and whole natural language sentences is complicated and lacks interpretability.

Taking Figure 1 as an example, while a human can easily answer the question, a model cannot directly achieve that if without further processing the raw visual and language features for a better alignment. However, if the multimodal components are correspondingly well-aligned, the model can work more straightforwardly to infer the answer. More specifically, a machine that attends to specific words or concepts in the questions and specific visual components in the image would arguably be more robust to linguistic variations, irrelevant to the question’s meaning and answer.

We tackle this problem by implementing granularity level alignment. Apart from only determining “where to look” for visual and textual attention, we try to achieve “looking correlation at the same level”. One solution for achieving efficient attention-based alignment is to implement a Transformer (TRM). Such architecture was first proposed as a machine translation model (Vaswani et al., 2017b), and later popularly implemented in multi-modality-based tasks due to its good alignment ability. However, many related works are pre-trained models, which require extra computational resources and extra human labor to collect additional data and its annotations (i.e., 3.3M images and captions for Conceptual Captions (Sharma et al., 2018)). To make it even worse, for some domains, i.e., medical, it is hard to obtain and annotate such a large-scale external dataset. Thus, unlike conventional Transformers, our focus is to learn multi-modality alignment without extra data more effectively. To achieve this, we embed graph-structured information into our model by involving the concept of lead graphs.

Our main contributions are as follows:

- We propose a novel multi-granularity alignment architecture that jointly learns intra- and inter-modality correlations at three different levels: concept-entity level, region-noun phrase level, and spatial-sentence level. Moreover, the results are then integrated with the decision fusion module for the final answer.
- We propose a co-attention mechanism that jointly performs question-guided visual attention and image-guided question attention and improves interpretability.
- The experiments are conducted on two challenging benchmark datasets, GQA and VQA-v2. They demonstrate our model’s effectiveness over the methods without extra pre-training data on both datasets. Moreover, our method even achieves better results over the pre-trained methods on GQA.

2 RELATED WORKS

2.1 VISUAL QUESTION ANSWERING

Visual Question Answering (VQA) is a popular joint visual language task that has been of increasing interest over the last few years. Its main task is to answer questions about a provided image. Similar to other multimodal tasks, VQA requires the model to understand both the image and natural language and to combine their features of the two modalities. However, there is a greater reliance on reasoning when answering visual questions. Implementing this reasoning ability is challenging, since the features of the two modalities lie in two very different domains: images lack the grammatical structure of the natural language, and it is high possible that the latter is biased in the sense that it makes it difficult for the model to output correct answers. These further motivate the construction of a model that can align these two features properly.

Early solutions employed convolutional neural networks (CNNs) and recurrent neural networks (RNNs) to embed the image and question, respectively, and extracted features are them directly fused (i.e., via concatenation or a shallow network) to obtain the result (Fukui et al., 2016; Zhou et al., 2015). In other works, one modality is further processed before fusion (Xiong et al., 2016),

or both are jointly embedded into a new space by additional networks (Ren et al., 2015). Yet other works proposed architectures that imply element-wise summation (Lu et al., 2016; Yang et al., 2016) or multiplication (Li & Jia, 2016; Nam et al., 2017) to achieve better fusion of multimodal features. The simple feature fusion methods established milestones in VQA task, providing insights into merging multi-modal features.

2.2 CO-ATTENTION MECHANISMS IN VQA

Many works focus on exploring image attention models for VQA (Zhu et al., 2016; Yang et al., 2016; Andreas et al., 2015). In the natural language processing domain, there are many related works on modeling language attention (Yin et al., 2016; Bahdanau et al., 2014; Rocktäschel et al., 2015). Some works learn textual attention for questions and visual attention for images simultaneously. (Lu et al., 2016) presents a hierarchical co-attention model for VQA that jointly reasons image and question attention. (Nam et al., 2017) is a multi-stage co-attention learning model that refines the attention based on the memory of the previous attention. (Yu et al., 2019) proposed a deep Modular Co-Attention Network that consists of Modular Co-Attention layers cascaded in depth. These works focus on the alignment between text features and fine-grained image features, where images lack language structure and grammatical rules, leading to difficulty in obtaining a good result. In addition, most of these works process questions in a simple manner and ignore the inner logical relations in the natural language domain. These issues become a bottleneck for understanding the relationships between multimodal features.

2.3 TRANSFORMERS IN THE MULTIMODAL TASK

The Transformer (Vaswani et al., 2017b) was initially proposed as a powerful machine translation model. However it gained substantial attention due to its ability to learn attention in all positions. In contrast to recurrent neural networks, Transformer expands the ability of the model to focus on the inner relations between sentences, producing a so-called “self-attention” property. (i.e., demonstrative term in the sentence) For example, when translating the sentence “The animal didn’t cross the street because it was too tired,” it would be helpful to know which word “it” refers to, as this would greatly improve the translation result. Due to the self-attention property, the Transformer architecture has been applied to other tasks beyond image captioning (Lu et al., 2019b) or visual question answering (Tan & Bansal, 2019); it also shows up in the works of vision-and-language navigation (Hao et al., 2020) and video understanding (Sun et al., 2019). Furthermore, implementation of the Transformer architecture aids the model in learning crossmodal representations from a concatenated sequence of visual region features and language token embeddings (Su et al., 2019; Li et al., 2020a). In addition, it learns joint representations that are appropriately contextualized in both modalities. However, these alignments are always achieved by pre-training on additional dataset (Sharma et al., 2018; Lin et al., 2014; Vicente et al., 2016).

3 APPROACH

We now introduce our Multi-Granularity Alignment Transformer for VQA (MGA-VQA). Our main idea of the model is to align multiple information levels accordingly between multimodal inputs and to integrate the information to obtain the final prediction. Figure 2 illustrates the architecture of our proposed model, which consists of three sets of alignments with different granularity levels. First, objects are detected from the input image, with their names, corresponding attributes, and relations. On the question side, noun phrases, entities, and sentence grammatical structure are detected. Then lead graphs is used to furtherly guide alignment learning, and they are constructed from the structural information extracted in the above steps, where the nodes in the graphs are regarded as the token features for the next steps. These features are the basic components of the following three levels of granularity alignment transformers (GA-TRMs): information of the concept level and entity level, information of the region level and noun phrase level, and information of the spatial level and sentence level. The lead graphs are then used to assist in co-attention learning. Finally, the outputs of the three GA-TRMs are used to predict the answer via the decision fusion module.

Section 3.1 describes the extraction of the different granularity level features, construction of the lead graphs, and formation of the token features from the image and question. Section 3.2 explains in

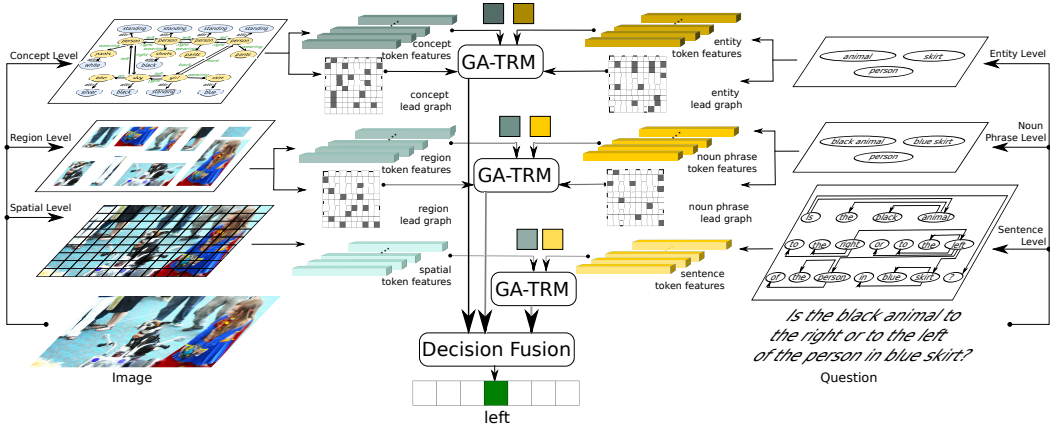


Figure 2: Overview of our approach. The following three granularity levels represent the image information: concept level, region level, and spatial level. The input questions are described by entity level, noun phrase level, and sentence level information. Later, token features and lead graphs are extracted from each level, then input into three granularity alignment Transformers (GA-TRMs) accordingly. (The lead graph in the figure is for demonstration, where its value is randomly set.) Finally, the outputs are combined by the decision fusion module to obtain the final prediction.

detail how the GA-TRMs use the token features and lead graphs. Section 3.3 illustrates the merging of the outputs from the three GA-TRMs by the fusion module.

3.1 GRANULARITY LEVELS IN VQA

In our model, three granularity levels are set for image and question, describing different levels of information. They are used to be aligned to the corresponding levels between modalities, i.e., concept-entity, region-noun phrase, and spatial-sentence.

3.1.1 GRANULARITY INFORMATION IN IMAGE

Given an input image (Img), three levels of features are extracted with different levels of granularity. For each granularity level, there are an associated set of token features and lead graph pairs. We first construct graph $\mathcal{G} = \{E, L\}$ representing the current granularity level information, where the tokens (E) are the entities in the graph, and the lead graph (L) is the connection pairs in the corresponding adjacency matrix. We now describe each level in detail as follows.

Concept Level

The concept level consists of the semantic features of objects, attributes, and relations between objects. We first extract this information from the image and build the corresponding graph \mathcal{G}_c , i.e., the top left section in Figure 1. To better input \mathcal{G}_c to the next stage, we first regard relations as extra nodes; then, we split this graph into node sequence (E_c) and pairs that represent node connections by index (L_c); in the example of Figure 1, $E_c = [girl, left, right, dog, brown]$ and $L_c = [(0, 1), (1, 3), (3, 2), (2, 1), (3, 4)]$, where the directional connections are consist of index pairs, describing both the “subject-predicate-object” relation and the “subject/object-attribute” information. Besides, we split the “subject-predicate-object” triple into “subject \rightarrow predicate” and “predicate \rightarrow object” pair so that it is easy to use index pair to describe the relation. Note that the token sequence feature $\mathbf{T}_c = \{t_{c_1}, t_{c_2}, \dots, t_{c_N}\}$ is computed from node sequence $E_c = \{e_{c_i}\}_{i=1}^N$ by GloVe embedding (Pennington et al., 2014) and the Multi-layer perceptron (MLP).

Region Level

The region level describes the middle-level visual features, which represent the visual region of object. Unlike the object features in the concept level, features in this level describe the object information visually instead of semantically. The token sequence features $\mathbf{T}_r = \{t_{r_i}\}_{i=1}^M$ are extracted by the Faster R-CNN method, and the relation pairs L_r are similar to L_c , where if there is a semantic relation between two objects at the concept level, there is a corresponding relation pair at the region level: for the previous example, $\mathbf{T}_r = [t_{girl}, t_{dog}]$ and $L_r = [(0, 1), (1, 0)]$.

Spatial Level

The spatial level describes the holistic but highest granularity visual features and provides detailed, supplementary information to the previous two levels, i.e., scene information. The token sequence features \mathbf{T}_{sp} are extracted from the backbone CNN, and L_{sp} is equal to the fully connected relations for all feature cells.

3.1.2 GRANULARITY INFORMATION IN QUESTION

Similar to *Img*, three levels of granularity are extracted from input question (Q). Most previous works only focused on the image features, ignoring inter-correlation within sentence. In contrast to them, our proposed method extracts structural and grammatical information, for better alignments.

Entity Level

The entity-level features represent individual objects in Q without attributes and help the model to achieve alignment in the abstract. The token features \mathbf{T}_e are processed in a similar manner as concept features for *Img*, and the corresponding lead graph pair L_e corresponds to the fully connected pair.

Noun Phrase Level

We filter the result from a constituency parser for our noun phrase level, discarding the determiners (e.g., 'a', 'the') and filtering out the words expressing positional relations (e.g., 'left', 'right') to save computational resources. The noun phrase level is constructed to align with the region-level features in *Img*, where the visual features contain attributions. Since most of the components are composed of multiple words, instead of merging them into a single token, we split them and process the GloVe features with the MLP, and obtain the token features as \mathbf{T}_{np} . In addition, the corresponding lead graph pair L_{np} corresponds to the fully connected pair.

Sentence Level

For sentence level, we process Q with the dependency parser and get the corresponding adjacency matrix (Dep_s) from the dependency graph. Since visual features are of a higher level and require less context aggregation than natural language (Lu et al., 2019a), sentence-level features need to be furtherly processed before alignment to better embed the structural information with the input. Instead of directly input the sentence token into the Transformer to fuse multi-modality features, we first use an extra Transformer module to process the sentence to get the context-aware features \mathbf{T}_s .

$$\mathbf{T}_s = Trm(\mathcal{MLP}(GloVe(Q)), Dep_s), \quad (1)$$

where $GloVe(\cdot)$ is the GloVe word embedding, $\mathcal{MLP}(\cdot)$ is the Multi-layer perceptron, and $Trm(t, g)$ is the Transformer module with input tokens t and attention mask m . Since the connection information is already embedded in \mathbf{T}_s , the lead graph pair for sentence level (L_s) consists of the fully connected pair. Details are shown in the Section 4.2.

3.2 MULTI-MODALITY GRANULARITY ALIGNMENT

In this section, we explain our design for using token features and lead graphs for alignment learning. For the lowest granularity level, the concept and entity token features from *Img* and Q respectively, are used as the token inputs of the GA-TRM, which obtains the most abstract information of both modalities. For the middle granularity level, the object region and noun phrase token features from *Img* and Q , respectively, are fed into the GA-TRM to learn the co-attention. For the highest granularity level, the spatial token features from *Img* and the sentence token features from Q are aligned, supporting the model with the most detailed information from the two modalities.

3.2.1 GA MULTI-HEAD ATTENTION

The Transformer architecture was initially proposed in (Vaswani et al., 2017a), in which it used stacked self-attention and pointwise, fully connected layers for both the encoder and decoder. The attention function can be described as mapping a query and a set

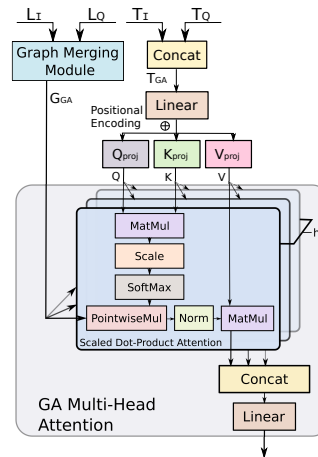


Figure 3: Granularity Alignment (GA) Attention.

of key-value pairs to an output. Our model follows the same architecture, as shown in detail in Figure 3. The token features from the image ($\mathbf{T}_I \in \{\mathbf{T}_c, \mathbf{T}_r, \mathbf{T}_{sp}\}$) and question ($\mathbf{T}_Q \in \{\mathbf{T}_e, \mathbf{T}_{np}, \mathbf{T}_s\}$) modalities are concatenated. After linear projection, they then employed with the learnable positional encoding (Vaswani et al., 2017a) to include both relative and absolute position information. For each token, a query vector (\mathbf{Q}), a key vector (\mathbf{K}), and a value vector (\mathbf{V}) are created, by multiplying the embeddings of the three matrices that are trained during the training process. Instead of utilizing a single attention module, we also linearly project \mathbf{Q} , \mathbf{K} , and \mathbf{V} h times with different, learned linear projections. Each of the sets of vectors is then input into the scaled dot-product attention, and pointwise multiplied with the lead graph (\mathbf{G}_{GA}) from the graph merging module:

$$\text{Att}_{GA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{norm}(\text{softmax}(\frac{\mathbf{QK}^T}{\sqrt{d_k}})) \circ \mathbf{G}_{GA} \mathbf{V}, \quad (2)$$

where d_k represents the dimensionality of the input, and $\text{norm}(\cdot)$ is the normalization over rows. Then, they are concatenated and once again projected, resulting in the final values.

3.2.2 GRAPH MERGING MODULE

Our graph merging module is designed to convert graph pairs (L) to merged lead graph (\mathbf{G}_{GA}).

From graph pairs (L_*) to lead graphs (\mathbf{G}_*)

The lead graphs from image (\mathbf{G}_I) and question (\mathbf{G}_Q) are binary graphs that are first constructed from the corresponding graph pairs from images ($L_I \in \{L_c, L_r, L_{sp}\}$) and questions ($L_Q \in \{L_e, L_{np}, L_s\}$). The dimension of the \mathbf{G}_I is $\|\mathbf{T}_I\| \times \|\mathbf{T}_I\|$, while that of \mathbf{G}_Q is $\|\mathbf{T}_Q\| \times \|\mathbf{T}_Q\|$. For each pair in L_I and L_Q , we assign the corresponding cell in the binary graph a value of 1, while the others are assigned a value of 0. e.g., $L_I = [(0, 1), (1, 3), (3, 2), (2, 1)]$, and $\|\mathbf{T}_I\| = 4$, then $\mathbf{G}_I = [[0, 1, 0, 0], [0, 0, 0, 1], [0, 1, 0, 0], [0, 0, 1, 0]]$, where the dimension of \mathbf{G}_I corresponds to $\|\mathbf{T}_I\|$, and the 2-D index of the nonzero value in \mathbf{G}_I corresponds to L_I .

From lead graphs (\mathbf{G}_*) to merged graphs (\mathbf{G}_{GA})

The merged lead graphs are a set of binary graphs of dimension $(\|\mathbf{T}_I\| + \|\mathbf{T}_Q\|) \times (\|\mathbf{T}_I\| + \|\mathbf{T}_Q\|)$. We set different lead graphs for different layers of encoders.

For the first layer, the lead graph is

$$\begin{bmatrix} \mathbf{0}_{\|\mathbf{T}_I\| \times \|\mathbf{T}_I\|} & \mathbf{0}_{\|\mathbf{T}_I\| \times \|\mathbf{T}_Q\|} \\ \mathbf{0}_{\|\mathbf{T}_Q\| \times \|\mathbf{T}_I\|} & \mathbf{1}_{\|\mathbf{T}_Q\| \times \|\mathbf{T}_Q\|} \end{bmatrix}. \quad (3)$$

This is to make the model learn the self-attention of the question, since the visual features are relatively high-level and require limited context aggregation with respect to words in a sentence, the latter of which needs further processing.

For the second layer, the lead graph is

$$\begin{bmatrix} \mathbf{0}_{\|\mathbf{T}_I\| \times \|\mathbf{T}_I\|} & \mathbf{0}_{\|\mathbf{T}_I\| \times \|\mathbf{T}_Q\|} \\ \mathbf{0}_{\|\mathbf{T}_Q\| \times \|\mathbf{T}_I\|} & \mathbf{1}_{\|\mathbf{T}_Q\| \times \|\mathbf{T}_Q\|} \end{bmatrix}. \quad (4)$$

This is to make the model learn the co-attention between the modalities.

For the third layer, the lead graph is

$$\begin{bmatrix} \mathbf{G}_I & \mathbf{0}_{\|\mathbf{T}_I\| \times \|\mathbf{T}_Q\|} \\ \mathbf{0}_{\|\mathbf{T}_Q\| \times \|\mathbf{T}_I\|} & \mathbf{G}_Q \end{bmatrix}. \quad (5)$$

which makes the encoder focus on the existing connectivity in the two modalities.

3.3 MULTI-GRANULARITY DECISION FUSION

The outputs of each level alignment are \mathbf{H}_{ce} , \mathbf{H}_{rn} and \mathbf{H}_{ss} , which embed the alignments of concept-entirety, region-noun phrase and spatial-sentence, respectively. We define the linear multimodal fusion function as follows:

$$\mathbf{H}_{GA} = \mathbf{W}_{GA}^T [\mathbf{W}_{ce}^T \text{LayerNorm}(\mathbf{H}_{ce}); \mathbf{W}_{rn}^T \text{LayerNorm}(\mathbf{H}_{rn}); \mathbf{W}_{ss}^T \text{LayerNorm}(\mathbf{H}_{ss})], \quad (6)$$

where $[\cdot; \cdot; \cdot]$ is the concatenation operation on vectors, \mathbf{W}_{ce} , \mathbf{W}_{rn} , \mathbf{W}_{ss} and \mathbf{W}_{GA} are linear projection matrices, and $\mathcal{L}_{ayerNorm}(\cdot)$ is used to stabilize the training (Ba et al., 2016).

Different from sequence-to-sequence learning tasks, our model uses the Transformer module to perform classification. We individually compute the cross-entropy loss from the three alignment streams. An early fusion strategy is used to make each alignment stream learn the attention. We define the loss as follows:

$$L = \mathcal{L}_{CE}(f_{ce}, a) + \mathcal{L}_{CE}(f_{rn}, a) + \mathcal{L}_{CE}(f_{ss}, a) + \mathcal{L}_{CE}(f_{GA}, a), \quad (7)$$

where f_{ce} , f_{rn} , f_{ss} , and f_{GA} represent the logits for the above three streams and their fusion, respectively, and a is the answer to the question.

4 EXPERIMENTAL SETUP AND RESULTS

4.1 DATASETS

[VQA-v2] The VQA (Goyal et al., 2017a) dataset is a real image dataset that contains over 204k images from COCO, 614k free-form natural language questions, and over 6 million free-form answers. For each question, 10 answers were gathered for robust inter-human variability. To be consistent with ‘human accuracy,’ the accuracy metric is $acc(ans) = \min\{\frac{\# ans is chosen}{3}\}$.

[GQA] The GQA (Hudson & Manning, 2019a) dataset consists of 20 million compositional questions involving a diverse set of reasoning skills and 1.5 million questions with closely controlled answer distributions. Compared with other real-image VQA datasets, the GQA dataset contains fewer language biases, involves more reasoning, and focuses on large vocabulary questions. We evaluate the reasoning ability of our MGA-VQA with the test split and conduct ablation studies to validate the effectiveness of each module in our model.

4.2 IMPLEMENTATION DETAILS

Our model setting is based on (Vaswani et al., 2017b). The encoder and decoder are separately composed of a stack of 3 identical layers, while the decoder is also composed of a stack of 3 identical layers. For the multi-head attention, we set 8 heads to achieve co-attention learning. The model is trained with distributed training in PyTorch, with 4 GeForce RTX 3090 GPUs. The learning rate is set to 10^{-4} with Adam optimizer, and batch size is set to 256. We merge same relation tokens to reduce computational load and update the lead graph accordingly; however, we do not change object category tokens. The [SEP] (special end-of-sequence) token is inserted after token features from image modality and is included in the corresponding dimension. The visual features are extracted from BUA (Anderson et al., 2018), and the scene graph is built in a similar manner as (Kim et al., 2020). The spatial level features are obtained from Resnet-101 backbone (He et al., 2016).

4.3 MODEL TRAINING DETAILS

Our model is designed to only use visual question answer annotation pairs to learn finer grained latent features. However, existing pre-training models rely on the additional labeled data describing finer-grained granularity (e.g., captioning) to explicitly learn finer-grained features, which require more computational resources and extra supervised data. Thus, our experiments are focused on training our model from scratch (a.k.a., non-pretraining). Nevertheless, our method not only achieves significant performance gain over existing non-pretraining methods, but also achieves competitive performance with some of pretraining methods. Next, we will show you detailed results.

4.4 QUANTITATIVE RESULTS ON GQA

[Overall performance] We evaluate our MGA-VQA model on the GQA dataset with the top-1 accuracy, as shown in Table 1. The methods that need the extra dataset (Sharma et al., 2018; Lin et al., 2014; Vicente et al., 2016) to pre-train the Transformer model are marked in the ‘‘Extra Data’’ column of the table. Among the evaluated methods, BUA (Anderson et al., 2018) is an attention-based model that enables calculation of attention at the region level rather than by using a uniform grid of equally sized image regions. LCGN (Hu et al., 2019), NSM (Hudson & Manning, 2019b),

Method			Performance		
No.	Name	Extra Data	Open	Binary	Overall
1	BUA (Anderson et al., 2018)	✗	34.83	66.64	49.74
2	LCGN (Hu et al., 2019)	✗	-	-	56.10
3	NSM (Hudson & Manning, 2019b)	✗	49.25	78.94	63.17
4	LXMERT (Tan & Bansal, 2019)	✓	-	-	60.33
5	LRTA (Liang et al., 2020)	✗	-	-	54.48
6	Oscar (Li et al., 2020b)	✓	-	-	61.62
7	VinVL (Zhang et al., 2021)	✓	-	-	65.05
8	Human (Hudson & Manning, 2019a)	✗	87.40	91.20	89.30
9	Ours	✗	54.29	78.25	65.51

Table 1: Overall accuracy on the GQA dataset.

and LRTA (Liang et al., 2020) mainly focus on solving complicated visual questions by first constructing graphs that represent the underlying semantics. LXMERT (Tan & Bansal, 2019), Oscar (Li et al., 2020b), and VinVL (Zhang et al., 2021) are Transformer-based models that solve the visual language problem by pre-training the model to align visual concepts and corresponding concepts in the text modality. The table shows that our model outperforms the state-of-the-art methods, even those needing pre-training.

[Ablation study] To validate each component of our model, we conducted experiments for an ab-

No.	Single Modality	Multi-Modality	Low Level GA	Middle Level GA	High Level GA	Lead Graph	Accuracy
19	✗	✓	✓	✗	✗	✓	47.83
20	✗	✓	✗	✓	✗	✓	54.72
21	✗	✓	✗	✗	✓	✓	61.65
22	✓	✗	-	-	-	✓	50.22
23	✗	✓	✓	✓	✓	✗	61.28
9	✗	✓	✓	✓	✓	✓	65.51

Table 2: Ablation study on the GQA dataset.

lation study, and the results are shown in Table 2. The columns “Single Modality” and “Multi-Modality” represent whether features from both the image and the question are used for the final answer prediction. The single modality setting uses only the question as the input, which tests whether the model’s good performance comes from the language bias of the dataset. Since there are no image features under the single modality setting, the multi-granularity alignment is not applicable. For the multi-modality setting, three experiments were conducted in order to verify each level’s granularity alignment. To guarantee fair comparisons, we always set three different initialized Transformer modules for each single-level granularity alignment, preventing the model scale from affecting the experimental result.

Experiments 19, 20, and 21 show the performance of using a single alignment, validating the need for multiple granularity level alignments. In addition, more abstract information leads to a higher overall accuracy, which may be due to achieving a better and easier alignment. However, when there is inaccurate information in the high-level granularity alignment, the lower-level granularity alignment will help provide some of the details. More examples of this are shown in Section 4.6. Experiment 22 is used to test the validity of the lead graph, and the results indicate that with such guidance, the overall performance is improved. To further validate our model’s reasoning ability, we train and evaluate our model with the ground-truth scene graphs and achieve 92.64% result. Since there is no annotated scene information of the testing split, we randomly divided the validation split into two, one for validation and the other for testing.

4.5 QUANTITATIVE RESULTS ON VQA

We also evaluate our model on VQA-v2 dataset for both test-dev and test-std splits, and the results are summarized in Table 3. In addition to the methods tested with GQA, we also test the following methods. Counter (Zhang et al., 2018) aims to solve object counting problem. MUTAN (Ben-Younes et al., 2017), BLOCK (Ben-Younes et al., 2019) and MUREL (Cadene et al., 2019) mainly focus on different multimodal fusions. VL-BERT (Su et al., 2019) and ViLBERT (Lu et al., 2019b) are Transformer-based models that require extra pre-training with extra, large-scale training data.

Method			Performance							
No.	Name	Extra Data	Yes/No	Number	Other	Test-dev	Yes/No	Number	Other	Test-std
10	MUTAN (Ben-Younes et al., 2017)	X	85.14	39.81	58.52	67.42	-	-	-	67.36
11	Counter (Zhang et al., 2018)	X	83.14	51.62	58.97	68.09	83.56	51.39	59.11	68.41
12	BLOCK (Ben-Younes et al., 2019)	X	83.60	47.33	58.51	67.58	83.98	46.77	58.79	67.92
13	MuRel (Cadene et al., 2019)	X	84.77	49.84	57.85	68.03	-	-	-	68.41
14	ViLBERT (Lu et al., 2019b)	✓	-	-	-	70.55	-	-	-	70.92
15	VL-BERT (Su et al., 2019)	X	-	-	-	69.58	-	-	-	-
16	VL-BERT (Su et al., 2019)	✓	-	-	-	71.79	-	-	-	72.22
17	Oscar (Li et al., 2020b)	✓	-	-	-	73.61	-	-	-	73.82
18	Ours	X	86.09	54.66	59.80	70.05	86.38	54.20	60.18	70.39

Table 3: Overall accuracy on the VQA-v2 dataset.

4.6 QUALITATIVE RESULTS

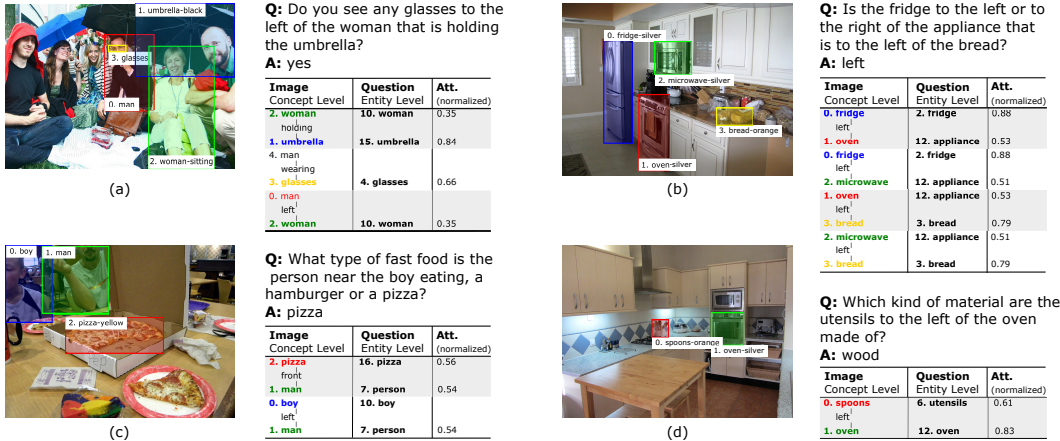


Figure 4: Visualization of our MGA-VQA multi-modality alignment. The attention value is the normalized learned attention of the last layer. The bounding boxes with IDs in the images correspond to the text with the same color in the table.

To better evaluate our model, we visualize some examples of our multimodal alignment in Figure 4. The table shows the alignment between the concept-level components in the image and the entity-level components in the question. The corresponding object regions are highlighted in the image. The results highlight the promising performance of our method; our model turns out to be able to find correlated elements for both modalities, and even at the image concept level, if the attributes are not accurately detected, with the assistance of the low-level alignment, the model obtains the correct final results (i.e., in Figure 4(d), the attribute of the spoon is not detected as wood.)

5 DISCUSSION AND CONCLUSION

In this work, we propose a novel architecture named MGA-VQA. In contrast to previous works that align visual and question features at a single level, our proposed Transformer-based model achieves multi-granularity alignment and jointly learns the intra- and inter-modality correlations. In addition, we construct a decision fusion module to merge the outputs from different granularity Transformer modules. In experiments, our model achieves outperforming results in GQA dataset and decent results in VQA-v2 dataset. Furthermore, we conduct ablation studies to quantify the role of each component in our model.

Research on graph-based VQA remains ongoing. One direction is involving pre-training mechanisms into current multiple granularity alignment and building a better decision fusion module. The other direction is finding a strategy to better overcome misdetection or incorrect detections at a single granularity level. Our MGA-VQA is an attempt to explore this issue. However, we expect that additional, related research will be conducted in the future.

REFERENCES

- 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*, pp. 6077–6086, 2018.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Deep compositional question answering with neural module networks. corr abs/1511.02799 (2015). *arXiv preprint arXiv:1511.02799*, 2015.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, 2015.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- 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*, pp. 2612–2620, 2017.
- Hedi Ben-Younes, Remi Cadene, Nicolas Thome, and Matthieu Cord. Block: Bilinear superdiagonal fusion for visual question answering and visual relationship detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 8102–8109, 2019.
- Remi Cadene, Hedi Ben-Younes, Matthieu Cord, and Nicolas Thome. Murel: Multimodal relational reasoning for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1989–1998, 2019.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017a.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6904–6913, 2017b.
- Weituo Hao, Chunyuan Li, Xiujuan Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13137–13146, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Ronghang Hu, Anna Rohrbach, Trevor Darrell, and Kate Saenko. Language-conditioned graph networks for relational reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10294–10303, 2019.
- 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*, pp. 6700–6709, 2019a.
- Drew A Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. *arXiv preprint arXiv:1907.03950*, 2019b.

- Ilija Ilievski and Jiashi Feng. Multimodal learning and reasoning for visual question answering. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 551–562, 2017.
- Eun-Sol Kim, Woo Young Kang, Kyoung-Woon On, Yu-Jung Heo, and Byoung-Tak Zhang. Hypergraph attention networks for multimodal learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14581–14590, 2020.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 11336–11344, 2020a.
- Ruiyu Li and Jiaya Jia. Visual question answering with question representation update (qru). *Advances in Neural Information Processing Systems*, 29:4655–4663, 2016.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pp. 121–137. Springer, 2020b.
- Weixin Liang, Feiyang Niu, Aishwarya Reganti, Govind Thattai, and Gokhan Tur. Lrta: A transparent neural-symbolic reasoning framework with modular supervision for visual question answering. *arXiv preprint arXiv:2011.10731*, 2020.
- 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*, pp. 740–755. Springer, 2014.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. *Advances in neural information processing systems*, 29:289–297, 2016.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL <https://proceedings.neurips.cc/paper/2019/file/c74d97b01eae257e44aa9d5bade97baf-Paper.pdf>.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*, 2019b.
- Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 299–307, 2017.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.
- Mengye Ren, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering. *Advances in neural information processing systems*, 28:2953–2961, 2015.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. Reasoning about entailment with neural attention. *arXiv preprint arXiv:1509.06664*, 2015.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2556–2565, 2018.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7464–7473, 2019.

- Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017a. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- 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*, pp. 5998–6008, 2017b.
- Tomás F Yago Vicente, Le Hou, Chen-Ping Yu, Minh Hoai, and Dimitris Samaras. Large-scale training of shadow detectors with noisily-annotated shadow examples. In *European Conference on Computer Vision*, pp. 816–832. Springer, 2016.
- Caiming Xiong, Stephen Merity, and Richard Socher. Dynamic memory networks for visual and textual question answering. In *International conference on machine learning*, pp. 2397–2406. PMLR, 2016.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 21–29, 2016.
- Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Joshua B Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *arXiv preprint arXiv:1810.02338*, 2018.
- Wenpeng Yin, Hinrich Schütze, Bing Xiang, and Bowen Zhou. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics*, 4:259–272, 2016.
- 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*, pp. 6281–6290, 2019.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5579–5588, 2021.
- Yan Zhang, Jonathon Hare, and Adam Prügel-Bennett. Learning to count objects in natural images for visual question answering. *arXiv preprint arXiv:1802.05766*, 2018.
- Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*, 2015.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4995–5004, 2016.