

Reasoning with Multi-Structure Commonsense Knowledge in Visual Dialog

Shunyu Zhang¹ Xiaoze Jiang¹ Zequn Yang¹ Tao Wan² Zengchang Qin^{1,*}

¹ Intelligent Computing and Machine Learning Lab, School of ASEE, Beihang University

² School of BSME, Beijing Advanced Innovation Center for Biomedical Engineering, Beihang University

{zhangshunyu, xzjiang, zqyang, taowan, *zcqin}@buaa.edu.cn

Abstract

Visual Dialog requires an agent to engage in a conversation with humans grounded in an image. Many studies on Visual Dialog focus on the understanding of the dialog history or the content of an image, while a considerable amount of commonsense-required questions are ignored. Handling these scenarios depends on logical reasoning that requires commonsense priors. How to capture relevant commonsense knowledge complementary to the history and the image remains a key challenge. In this paper, we propose a novel model by Reasoning with Multi-structure Commonsense Knowledge (RMK). In our model, the external knowledge is represented with sentence-level facts and graph-level facts, to properly suit the scenario of the composite of dialog history and image. On top of these multi-structure representations, our model can capture relevant knowledge and incorporate them into the vision and semantic fusion. Experimental results and analysis on VisDial v1.0 and VisDialCK datasets show that our proposed model effectively outperforms comparative methods.

1. Introduction

With the increasing interest in Visual Dialog task [7], which involves an agent to make a dialog conditional on an image, there exist loads of studies [5, 24, 42] concentrating on the reasoning of dialog history. Some recent works [1] showed that 10.96% of the questions in the validation set of the well-known dataset of *VisDial v1.0* [7] demand dialog history, while there are 10.62% of questions that require commonsense knowledge from their annotated data. However, there was little research studying the commonsense-required questions, compared to history-required ones. As shown in Figure 1, when answering “Where is the plane?”, without commonsense knowledge, the agent cannot easily figure out the place where the plane parks and only replies

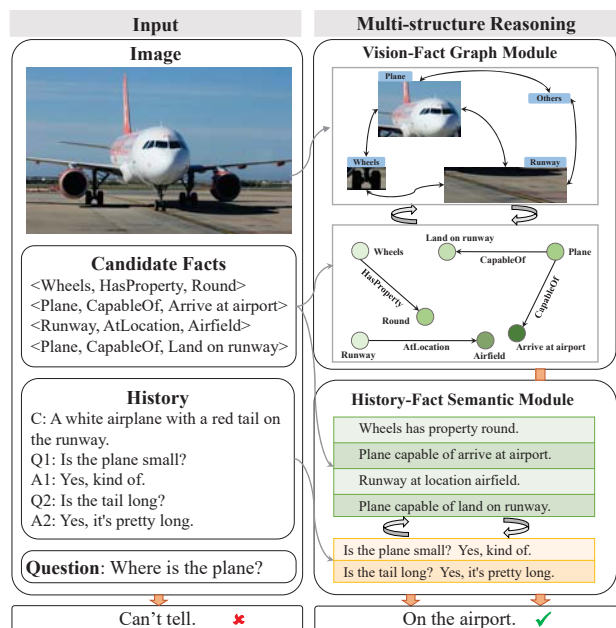


Figure 1. An illustration of RMK method. We represent Candidate Facts with multiple structures (graph-level and sentence-level) to reason with image and history for the optimal answer.

with the safe response “Can’t tell.”. Therefore, how to equip a visual dialog system with commonsense knowledge is unresolved and remains a challenge in the Vision and Language research.

There were quite a few attempts on knowledge-based visual question answering (KB-VQA) [20, 36]. The advanced solutions usually build a fact graph with filtered fact triplets and then reason on the graph to infer the best answer [22, 44]. However, Visual Dialog task requires an agent to comprehend the dialog history information additionally compared to the VQA tasks [2], so calls for more contextual logic. What’s more, graph-style knowledge has limited ability in capturing semantic-level information, since it pays

more attention to the relationship of the knowledge entities. Thus, the single-structure knowledge at semantic-level or graph-level may not satisfy the unique requirements of the visual dialog tasks.

To solve the above problems, we propose a novel multi-structure knowledge representations: i.e. *graph-level facts* and *sentence-level facts*, incorporating with two essential visual dialog components (i.e. image and dialog history). The graph-level facts are used to model relations in commonsense knowledge, and they can also complement the underlying visual relationship explicitly. Therefore we build a visual graph combined with graph-level facts, as shown in Fig.1. On the other side, the sentence-level facts tackle the knowledge semantics, it maps the knowledge in triplet to the text space. We equip them with sentence-level facts to better extract semantic features, for dialog history also contains semantic relations implicitly. Meanwhile, the advantage of this combination is that the image and dialog history is associated with homologous knowledge information, bridging the heterogeneous gap and complementary to each other.

As shown in Fig.2, our model consists of two modules: *Vision-Fact Graph Module*, *History-Fact Semantic Module*. Specifically, Vision-Fact Graph Module converts knowledge triplets to graph-level representation and further injects the commonsense knowledge into the graph-level vision bank. History-Fact Semantic Module involves sentence-level facts to the dialog history via cross-modal attention-based operations. Both two modules adopted three units, i.e. purification, injection, and aggregator to filter and incorporate relevant knowledge information. Finally, we adopt transformer-based multi-modal fusion and generate the response by the decoders.

Our contributions can be summarized as follows:

- We propose a novel method to represent commonsense knowledge in multi-structure: graph-level and sentence-level, to better suit the character of visual dialog and complement relevant information.
- Furthermore, we adopt a multi-structure reasoning network to encode vision-fact graph knowledge and history-fact semantic knowledge, to extract implicit dependence in different modalities. The principled ablation study and visualization show how different modules work in our model.
- We conduct comprehensive experiments on two datasets: *VisDial v1.0* [7] and *VisDialCK* [1]. Note that VisDiaCK (a validation subset of VisDial v1.0) is a collection of commonsense-required questions in Visual Dialog. The results demonstrate the superiority our model.

2. Related Work

Visual Dialog. For the visual dialog task [7], it aims to generate responses depending on an image, a caption, and the dialog history. LF [7], MN [7], CorefNMF [14] and CoAtt [8] utilize kinds of attention mechanisms as the backbone to locate the related visual objects. To solve the history-required problems such as visual co-reference, RVA [24] design recursive visual attention, inferring the co-reference through recursively inspecting the history dialog and improving the visual attention. Zheng *et al.* [42] propose an EM-style inference algorithm to obtain the latent relations among history dialogs. MCA [1] focuses on an iterative question-conditioned context-aware graph, including both fine-grained visual and history semantics. DualVD [12] constructs a scene graph to represent the image, which emphasizes the essential role of vision for the referred visual content may change remarkably. Another line of work targeted on response generation for visual dialog by carefully designed decoders. DMRM [6] adopts multi-step reasoning based on dual attention to iteratively update related visual objects for a more relevant response. DAM [13] designs an adaptive decoder with memory to store the state of dialog history and visual information. Recently, pre-trained models [21, 40] have also achieved impressive results in visual dialog. VisualBERT [21] and VDBERT [40] exploit large extra datasets to explore in visual dialog via pretraining language models.

Though these works have achieved great success in performance, the commonsense-required problems are ignored and it still has space to improve by considering commonsense knowledge.

Knowledge-based VQA. Visual question answering (VQA) [2] needs to give an accurate answer based on an image and a relevant question. Recently, there are many works proposed on knowledge-based VQA, including diverse benchmarks and systems. FVQA [36] is a fact-based VQA dataset that provides image-question-answer-supporting fact tuples. KBVQA [37] divides data into three categories in which it needs visual concept, basic common sense, or higher-level knowledge with explicit reasoning. KVQA [29] consists of questions requiring world knowledge of named entities in images. Furthermore, OK-VQA [20] covers 11 categories of knowledge, such as cooking and food, science and technology, plants and animals, etc.

Another line is the knowledge-based VQA models tapping into knowledge representations and reasoning strategies. Out of the Box [22] applies graph convolution networks to reason on the knowledge graph, whose nodes are attached by image and semantic embeddings. In addition, Mucko [44] reasons on visual, fact, and semantic graphs separately, and utilizes cross-modal networks to aggregate information together for knowledge reasoning. KRISP [19]

employs a BERT-pretrained model to better understand semantics and exploit implicit knowledge. MAVEx [38] votes among textual and visual knowledge from different sources. However, these works cannot apply to visual dialog directly, since visual dialog demands reasoning on both dialog history and image. Thus, how to design a knowledge fusion scheme adaptive to visual dialog appears particularly significant. Inspired by this, we design a multi-structure knowledge model to densely interact with both visual and dialog components in visual dialog.

Vision and Language Modeling. Approaches for multimodal vision and language tasks have explored diverse modeling strategies, such as GNN-based models (e.g. [12]) or transformer-based ones (e.g. [40]). Teney *et al.* [35] propose the first GNN-based VQA method, which builds a scene graph of the image and parses the sentence structure of the question. Li *et al.* [17] encodes each image into a graph and model inter-object relations via graph attention mechanism. Huang *et al.* [10] propose a novel dual-channel graph convolutional network to better integrate visual and textual information. GNN-based methods have also achieved impressive progress in visual dialog [5, 12], benefiting from the reasoning ability of graph network.

Over the past few years, multimodal transformers have made significant progress through pre-training on large-scale image and text pairs and then fine-tuning on downstream tasks. VisualBERT [21], Unicoder-VL [16] and VL-BERT [33] propose the single-stream architecture on both images and text. ViLBERT [18] and LXMERT [34] propose a two-stream architecture to process visual and textual information independently first and fused them later. CLIP [26] aligns visual and language representations by contrastive learning and achieves state-of-the-art results in image-text retrieval.

Different from these work that uses transformer or other methods separately, our model first infers on the multi-structure knowledge with GNN’s reasoning ability and then fuse different modalities via a transformer to better improve the interpretability and performance.

3. Methodology

The visual dialog tasks are as follows: given an image I and the dialog history $H = \{C, (Q_1, A_1), \dots, (Q_{t-1}, A_{t-1})\}$, where C is the image caption. The task is to infer the best answer to the current question Q_t by ranking a list of 100 candidate answers. Our work mainly focuses on the protocol of introducing external *commonsense knowledge* to enhance the visual dialog system to reason for better answers. Based on the characteristics of the image and the dialog history, we observe commonsense knowledge as two profiles: graph-level and sentence-level. On top of them, we incorporate them into the dialog system adaptively, and

we also visualize the reasoning clue in Fig.3.

3.1. Multi-structure Facts Representation

The image and dialog history are two key components in visual dialog. For the image, visual graph is widely adopted to handle the object relation [12], and the dialog history is indispensable for its contextual information [42]. Therefore, single-structure commonsense knowledge cannot meet the diverse information demand. To fit the characteristics of them in a visual dialog, we represent commonsense knowledge in two aspects: *sentence-level facts* and *graph-level facts*.

Sentence-level Facts. In open-domain conversational systems, the semantics shared with commonsense knowledge is vital for establishing effective interactions [43]. To capture the contextual semantics of fact triplets $\langle \text{subject}, \text{relation}, \text{object} \rangle$, we convert it to semantic domain as the fact description “subject relation object”. Then feed the description to an LSTM to get the sentence-level facts representation s_i^F .

Graph-level Facts. The graph structure has great capability in gripping the relation between the entities. Thus, we utilize the graph structure to further underline the relationship between each commonsense knowledge entity complementary to visual graph. In detail, the graph-level facts are denoted as $G^F = (E^F, R^F)$, in which the node is fact entity $e_i^F \in E^F$. To enhance the semantic information in the fact graph, the edge $r_{ij}^F \in R^F$ can be calculated as:

$$r_{ij}^F = \tanh(\mathbf{W}_r[r_{ij}^h, r_{ij}^d]) \quad (1)$$

where r_{ij}^d is Fact Description representation corresponding to entity e_i and e_j , r_{ij}^h is the embedding of relation in the triplet. “[\cdot , \cdot]” denotes concatenation, and \mathbf{W}_r (as well as $\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_n$ mentioned below) are learned parameters in linear layers.

To find the optimal supporting facts, we first retrieve relevant candidate facts from the knowledge base of facts [31], following a score based approach proposed in [22]. We compute the cosine similarity of the embeddings of every word in the fact with the words in the caption and the words of visual concepts detected in the image. Then we average these values to assign a similarity score to the fact. These facts are sorted based on the similarity and the highest scoring facts are retained.

3.2. Vision-Fact Graph Module

For the objects in the image lacking relation information [12], we combine the image with graph-level facts. As for the encoding strategy of image, we adopt the recent standard scheme [9], conducting a graph for the image. This module mainly contains three units to filter and select informative vision and fact information: *Vision-Fact Purifica-*

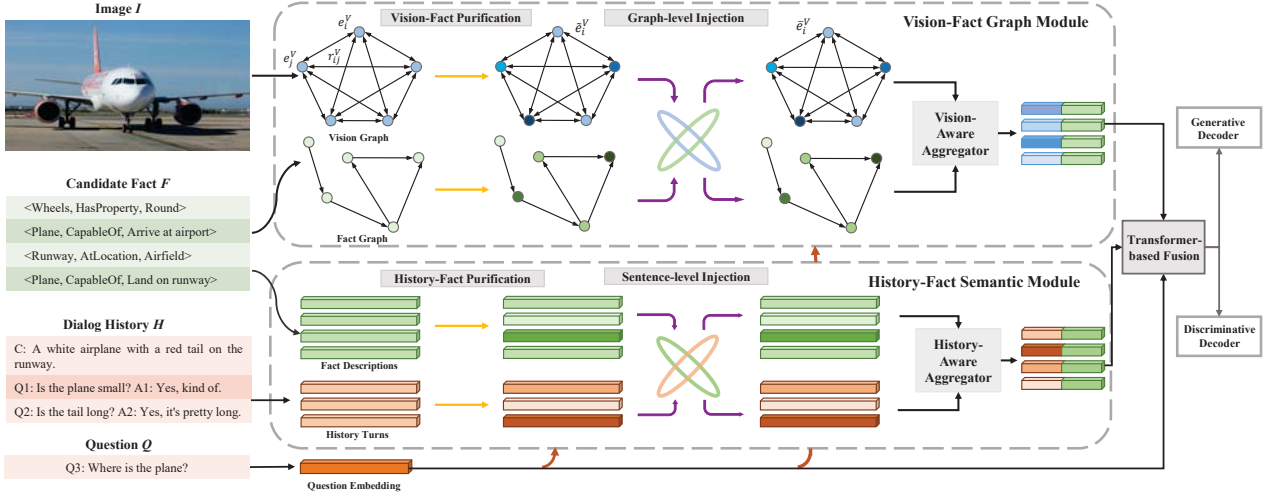


Figure 2. Overview structure of RMK. The model mainly contains two modules: Vision-Fact Graph Module and History-Fact Semantic Module, both of which contain three operators: Purification (yellow arrows), Injection (pink arrows) and Aggregator. And the orange arrows in the figure denote the question-guided way.

tion, Graph-Level Injection and Vision-Aware Aggregator, shown in Fig. 2.

Vision-Fact Purification. It aims to filter out less relevant information, for there may exist amounts of redundant information in the image and fact knowledge graph. In the visual feature graph $G^V = (E^V, R^V)$, the nodes $E^V = \{e_i^V\}^N$ are visual entity features extracted by a detector, where N is the number of detected objects. The edges $R^V = \{r_{ij}^V\}^{N \times N}$ are the visual relations between nodes provided by a visual relationship encoder [41]. The construction of the fact graph is described in Sec. 3.1. Then we adopted relation-aware GCN [12] methods to aggregate relation information among the entities in the vision graph and fact graph. And it results to purified vision feature \tilde{E}^V and fact feature \tilde{E}^F , respectively.

$$\begin{aligned}\tilde{E}^V &= GCN(E^V, R^V) \\ \tilde{E}^F &= GCN(E^F, R^F)\end{aligned}\quad (2)$$

Graph-Level Injection. The graph-level facts contain diverse knowledge, while the image may retain noisy entities that lack relevant information. The Graph-level Injection introduces external knowledge to help understand the visual information comprehensively, and also incorporates the visual knowledge into the facts graph to enhance the supported facts.

It strengthens the image information with commonsense knowledge, while further grasping the most relevant facts guided by vision, through cross-graph interaction. Specifically, to equip the image with useful facts, the graph message \tilde{v}_i^M is transferred from facts \tilde{v}_j^F to visual entity \tilde{v}_i^V between two graphs. The facts-injected image entity \tilde{v}_i^V is

generated as follows:

$$\begin{aligned}\gamma_{ij} &= softmax(\mathbf{W}_\gamma(\tanh(\mathbf{W}_1[Q_t, \tilde{e}_i^V, \tilde{e}_j^F]))) \\ \tilde{e}_i^M &= \sum_{j=1}^N \gamma_{ij} \tilde{e}_j^F \\ \tilde{e}_i^V &= \tanh(\mathbf{W}_2[\tilde{e}_i^V, \tilde{e}_i^M])\end{aligned}\quad (3)$$

Where Q_t is the question feature encoded by LSTM. We adopt additive attention [4] which is the concatenation followed by the weight matrix. The vision-injected facts entity \tilde{e}_i^F can be gained by swapping the position of \tilde{e}_j^F and \tilde{e}_i^V in the equations.

Vision-Aware Aggregation. After Graph-Level Injection, the entities in a graph are injected with local complementary information from the other. We then aggregate facts graph to global representation via attention mechanism, and further concatenate it with visual features. The aggregated vision-fact representation \tilde{I} can be gained by:

$$\begin{aligned}\delta_i &= softmax(\mathbf{W}_\delta(Q_t \circ (\mathbf{W}_3 \tilde{e}_i^F))) \\ \tilde{I} &= \mathbf{W}_v[\tilde{e}_j^V, \sum_{i=1}^N \delta_i \tilde{e}_i^F]\end{aligned}\quad (4)$$

3.3. History-Fact Semantic Module

Distinct from the image, the dialog history has different characteristics in manifestations. The contextual relation information is included in the sentences implicitly, and the graph-level facts have limited ability in handling the semantics among sentences. Thus, we further introduce the sentence-level facts, which are denoted as $\{s_i^F\}^K$, where

Table 1. Result on VisDial v1.0 val set using generative decoder.

Method	NDCG \uparrow	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	Mean \downarrow
MN [7]	51.86	47.99	38.18	57.54	64.32	18.60
CoAtt [39]	59.24	49.64	40.09	59.37	65.92	17.86
DMRM [6]	-	50.16	40.15	60.02	67.21	15.19
DAM [13]	60.93	50.51	40.53	60.84	67.94	16.65
KBGN [11]	60.42	50.05	40.40	60.11	66.82	17.54
GoG [5]	62.63	51.32	41.25	61.83	69.44	15.32
LTMI [23]	61.61	50.38	40.30	60.72	68.44	15.73
LTMI-RMK	63.57	51.76	41.56	62.16	69.83	15.05

K is the number of facts. The dialog history is denoted as $\{s_i^H\}^T$, where T is the rounds of history. We adopted similar methods in previous graph module, after minor modification, to filter and fuse them: *History-Fact Purification*, *Sentence-level Injection* and *History-Aware Aggregator*.

In this module, *History-Fact Purification* aims to evaluate the relevance of textual facts and history to the current question. Specifically, the sentence-level facts are purified by the guidance of question-aware attention.

$$\eta_i = \text{softmax}(\mathbf{W}_\eta(Q_t \circ \mathbf{W}_7 s_i^F))$$

$$\tilde{s}_i^H = \eta_i s_i^F \quad (5)$$

And the purified history features are gained in the same way.

As for *Sentence-level Injection* and *History-Aware Aggregator*, we similarly adopt the paradigm in Graph Module. And we computed Eq.3 and Eq.4 on the top of textual features, finally resulting to aggregated history-fact features \bar{H} . It can enrich dialog history and related facts with each other.

3.4. Multi-modal Fusion

After obtaining the fact-aware representations, we fuse the question representation Q_t , history-fact feature \bar{H} , vision-fact feature \bar{I} through a multi-modal fusion strategy. It can be any existing visual dialog model to learn the joint representation. In our experiments, we adopt a light-weight transformer-based method LTMI [23] to fuse them.

$$E = \mathcal{F}(Q_t, \bar{I}, \bar{H}) \quad (6)$$

Then the fused representation E is fed to the decoder to generate responses to the given question. As for the decoder, we follow the previous studies [7] to set discriminative and generative decoders and adopt multi-task learning [23] by minimizing the sum of the generative loss and the discriminative loss.

4. Experiments

4.1. Datasets

VisDial v1.0. For VisDial v1.0 dataset, the train, validation, and test splits contain 123k, 2k, and 8k dialogs, re-

Table 2. Results on VisDial v1.0 test-std set using discriminative decoder. Underline are the highest results except for pretraining-based models, which are trained with extra training data.

Method	NDCG \uparrow	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	Mean \downarrow
LF [7]	45.31	55.42	40.95	72.45	82.83	5.95
MN [7]	47.50	55.49	40.98	72.30	83.30	5.92
CoreFMN [14]	54.70	61.50	47.55	78.10	88.80	4.40
RvA [24]	55.59	63.03	49.03	80.40	89.83	4.18
DualVD [12]	56.32	63.23	49.25	80.23	89.70	4.11
CAG [9]	56.64	63.49	49.85	80.63	90.15	4.11
KBGN [11]	57.60	64.13	50.47	80.70	90.16	4.08
GoG [5]	60.38	63.13	49.88	79.65	89.05	4.39
VDBERT [40]	75.35	51.17	38.90	62.82	77.98	6.69
VisualBERT [21]	74.47	50.74	37.95	64.13	80.00	6.28
LTMI [23]	<u>60.92</u>	60.65	47.00	77.03	87.75	4.90
LTMI-RMK	58.48	64.14	50.58	80.72	90.28	4.14

Table 3. Results comparison on VisDialCK using discriminative decoder, where \dagger means re-implemented with the same settings as ours for fair comparison.

Method	NDCG \uparrow	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	Mean \downarrow
LF \dagger [7]	53.46	55.53	41.32	76.95	87.04	4.61
MN \dagger [7]	55.06	56.18	41.47	77.32	87.45	4.36
DualVD [12]	55.48	58.77	42.55	81.01	88.30	3.93
LTMI [23]	58.74	58.12	43.78	80.27	88.23	4.02
LTMI-RMK	60.94	65.78	54.92	81.76	90.23	3.91

spectively. In “train” and “val”, each image is accompanied by a 10-round dialogue, while in “test”, each image is followed by random rounds of question-answer pairs and an ongoing question for answer prediction. The training split is composed of 123k images and each dialog consists of 10-round QA pairs for each image. The following metrics are adopted: mean reciprocal rank (MRR), recall@ k ($k=1, 5, 10$), mean rank (Mean), and normalized discounted cumulative gain (NDCG). A lower value for Mean and higher for other metrics are desired. Note that we train the model on the VisDial v1.0 training set, and evaluate the model on the VisDial v1.0 val, test, and VisDialCK.

VisDialCK. For the purpose of verifying the effectiveness of RMK on commonsense-required questions in visual dialog, we also conduct evaluations on a commonsense-required dataset called *VisDialCK*. It is first proposed by [1], in which they conducted crowd-sourcing on VisDial v1.0 val to annotate the dialog into different categories, among which commonsense-required and history-required are the most two except for normal VQA kind (don’t need history and commonsense). However, they only focus on history-required ones. So we further collect commonsense-required ones from their raw data to form VisDialCK, a subset of VisDial v1.0 val, which contains 940 history-required dialog rounds. It can properly reflect the model’s capability to deal with the knowledge-required dialogs.

4.2. Implementation Details

To build the vocabulary, we retain words in the dataset with word frequency greater than 5. Each word in the dialog is embedded into a 300-dim vector with the GloVe embedding initialization [25]. The maximum sentence length of the dialog history and the current question are set to 20. The hidden state size of Transformer blocks is all set to 512. We adopt Adam optimizer with the initial learning rate of $4e-3$ and final learning rate of $5e-5$ via cosine annealing strategy with 16 epochs. The mini-batch size is 15 and the dropout [32] ratio is 0.5. The model is trained with a multi-class N-pair loss. We choose the widely adopted ConceptNet as the external commonsense knowledge source [31]. Following [3], we use bottom-up features of 36 proposals from images using a Faster-RCNN [27] pre-trained on Visual Genome [15] to get a bag of object-level 2048-d image representations. For the results on test set, we only report results for our best performing models as the number of allowed submissions to the challenge is limited.

4.3. Comparison Results

Baselines. In our experiment, the compared methods mainly include: (1) Fusion-based and Attention-based models: LF [7], MN [7], CorefNMN [14], RvA [24], DMRM [6], DAM [13]. (2) The pretraining model: VDBERT [40] and VisualBERT [21]. (3) Graph-based models: DualVD [12], FGA [28], CAG [9], KBGN [11]. These methods are our mainly compared baselines.

Generative Results. First, we compare the performance of generative results of different models. As shown in Table 1, our method outperforms all the compared methods with large margins on the val v1.0 split. Comparing with the results of LTMI [23] without commonsense knowledge, our model improves NDCG for 62.63 to 63.57 (+1.96), MRR from 50.38 to 51.76 (+1.38), R@1 from 40.30 to 41.56 (+1.26), Mean from 15.73 to 15.05 (+0.68) and more than 1% on other metrics. Notice that GoG [5] additionally parses the words relations in a question and builds a more complex graph-over-graph network. Our RMK validates that when incorporating commonsense knowledge, it improves significantly and outperforms other compared models on all metrics. It proves that RMK can improve the performance of visual dialog models by introducing explicit knowledge reasoning, which also illustrates that commonsense knowledge is helpful for visual dialog.

Discriminative Results. We also compare discriminative results in Table 2. Our method improves a lot compared to LTMI on the test-std v1.0 split, which is about +3% on MRR, R@1, R@5, and R@10. Compared to previous non-pretrained models, our method also achieves significant improvement on most metrics, which proves that our method is effective and beneficial. The performance of our

Table 4. Ablation study of model design on VisDial val v1.0. F in the second block is short for facts.

Model	MRR↑	R@1↑	R@5↑	R@10↑	Mean↓	NDCG↑
LTMI	62.32	48.94	78.65	87.88	4.86	62.72
LTMI-RMK	65.08	51.78	81.62	90.48	3.98	60.68
w/o All F	63.92	50.32	80.13	89.27	4.43	58.37
w/o Sentence F	64.97	51.26	81.32	90.12	4.09	59.63
w/o Graph F	64.48	50.86	80.82	89.74	4.24	59.21
w/o Purification	64.84	50.92	80.73	90.23	4.13	59.12
w/o Injection	63.92	51.13	80.78	90.07	4.42	58.72
w/o Aggregator	64.16	50.75	80.91	89.83	4.20	58.63

Table 5. Ablation study on different number of commonsense fact candidates on VisDial val v1.0.

# facts	MRR↑	R@1↑	R@5↑	R@10↑	Mean↓	NDCG↑
top 50	64.04	50.78	80.83	89.37	4.17	58.63
top 100	65.08	51.78	81.62	90.48	3.98	60.68
top 150	64.43	51.20	81.23	89.86	4.01	59.32
top 200	64.65	51.32	81.17	90.23	4.05	58.95

model even exceeds the performance of VDBERT [40] on all the metrics except NDCG. Notice that the pretrain-based model(VDBERT and VisualBERT) works for they use a lot of extra train data except for VisDial train set. These observations show that RMK can assist in the improvement of visual dialog tasks. The reason why our method is effective is that we incorporate multi-structure of commonsense knowledge through our designed network.

Results on VisDialCK. To certify whether our model can deal with the commonsense-required questions successfully, we compare RMK with previous models on VisDialCK [1]. As shown in Table 3, RMK outperforms them on all metrics. Our model substantially improves a lot on LTMI, on MRR and R@1 by about +8%, and on NDCG and R@10 by +2%, which proves that the model can help with the questions that require commonsense. It verifies that the traditional methods can not answer the questions that require commonsense knowledge well. And the significant improvement also indicates that our method can indeed assist in handling the commonsense-required questions.

4.4. Ablation Study

In Table 4, we first remove the different levels of facts to validate the effect of multi-structure knowledge. The results in the second block show both the sentence-level and graph-level facts are crucial for visual dialog, and combining them can achieve better results. In the second block, we investigate the importance of different operations in our model. w/o Purification removes the purification stage in both Vision-Fact Graph Module and History-Fact Semantic Module and others as the same. Without any of these three stages, the performance consistently drops, which validates the effectiveness of these adaptive strategies.

As shown in Table 5, we vary the number of retrieved

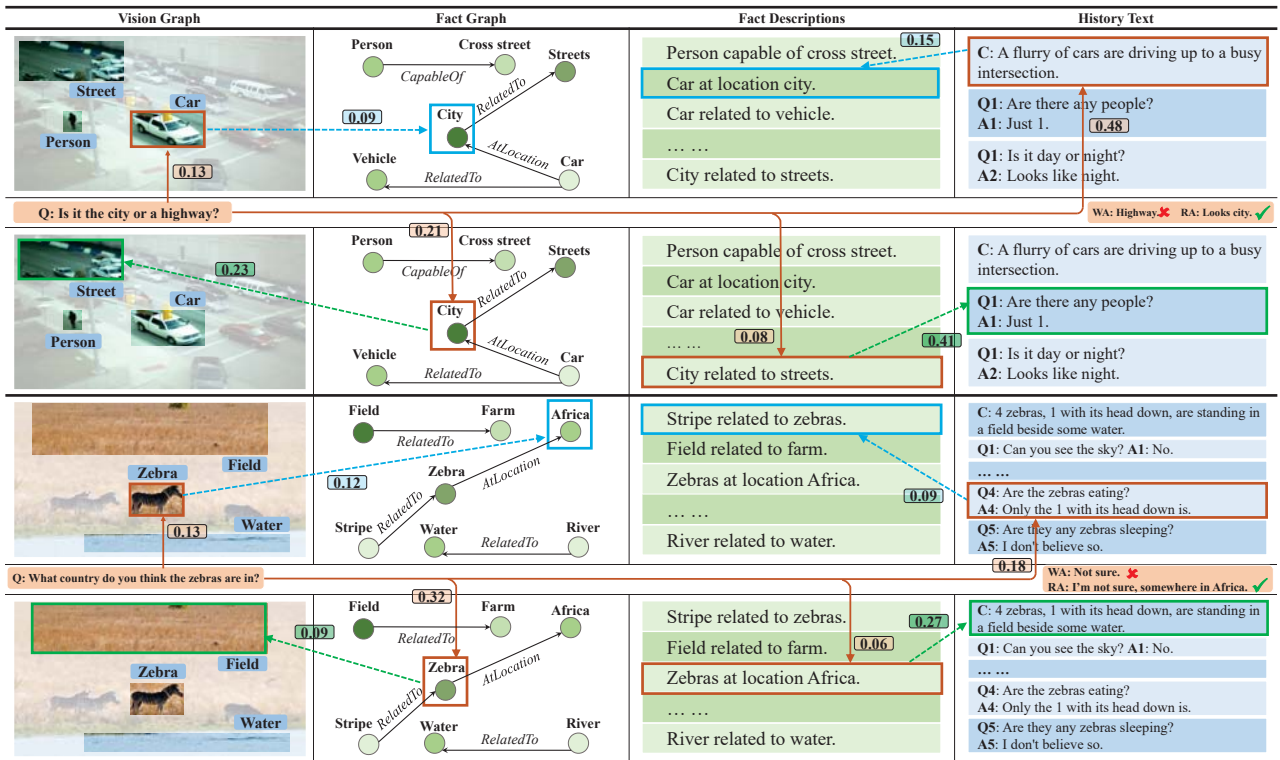


Figure 3. Qualitative results from our RMK. The WA means the wrong answer predicted by *LTMi*, while RA means the right answer by *LTMi-RMK*. The decimals on the arrows from question *Q* to other modalities indicate the normalized question-guided attention. The arrows between facts and visual graph or history text denote the cross-modal interaction weight, displaying the complementary information.

Table 6. Human evaluation on 100 sampled responses on VisDial val v1.0. M1: percentage of responses pass the Turing Test. M2: percentage of responses evaluated better than or equal to human responses.

	LTMi [23]	RMK
Method 1 (M1)	54	65
Method 1 (M2)	62	68

candidate facts for the model, in which top-k are ranked by the weighted score of fact confidence and visual object confidence. We achieve the best downstream metrics with the top 100 candidate facts (adopted by us). Fewer facts may not include the required facts for the questions, while too many facts may introduce much noise into the model.

4.5. Human Study

As shown in Table 6, we conduct human study to further demonstrate the effectiveness of our proposed RMK model. Our model achieves the highest scores both on the metrics M1 and M2 compared with LTMi model. These results show that our model can generate a contextually coherent response, which is more in line with human commonsense.

4.6. Qualitative Results

To figure out how the RMK model works, we visualize the reasoning paths on top of the multi-structure commonsense knowledge with vision and history information. Figure 3 shows two examples, in which the first one comes from VisDialCK and the second comes from VisDial val set. There are two reasoning clues for answering the question: one is reasoning through vision or history to support facts (the row above questions in Fig.3), and the other reasons from question directly to facts incorporated with vision or history information (the row below questions).

Take the first example for detailed analysis. When answering the given question “*Is it the city or a highway?*”, to determine what is the image about, the model focuses on the main object *Car* which is directed to *City* in Fact Graph. Similarly, reasoning from question through caption *C* in history also leads to “*Car at location City*” in Fact Descriptions. Moreover, as seen in the blocks below the question, the model can link the question directly to the relevant fact entity *City* and fact description “*City related to streets*”. Finally, our model generates a more reliable answer “*Looks city*” rather than “*Highway*”, which is more in line with commonsense compared to the one without facts







	<p>Caption C: Traffic light in front a tree a statue of a lion is behind the tree, a large building is in the background.</p> <p>History Q1: Are there a lot of trees? A1: I see 1 tree and some bushes Q2: Are there people? A2: No Q3: Is this in black and white? A3: It's in color?</p> <p>Question Q4: Is this an urban setting?</p> <p>Answer A4: Yes, it is. ✓</p>		<p>Caption C: A group of dogs follow cows through a fence.</p> <p>History Q1: Is the picture in color? A1: Yes Q2: Is the picture time stamped? A2: No Q3: How many dogs are in the picture? A3: 3 Q4: How many cows in the picture? A4: A lot Q5: is there a gate on the fence A: fence Q6: Are they people in picture? A6: No</p> <p>Question Q7: Is this on a farm?</p> <p>Answer A7: Likely. ✓</p>
	<p>Caption C: The television is located beside the blue front door.</p> <p>History Q1: Is this in a living room? A1: Yes Q2: Is the tv on? A2: No ... Q5: Is the tv on a table? A5: Yes Q6: Are there any pictures? A6: No Q7: Can you see the wall? A7: Yes</p> <p>Question Q8: Are there any windows?</p> <p>Answer A8: Yes, but covered. ✓</p>		<p>Caption C: Many people lay on towels on a sandy beach.</p> <p>History Q1: Is the picture in color? A1: Yes Q2: How many people are there? A2: 7 ... Q6: How old do they look? A6: 20's Q7: What are they doing? A7: Lying down on toes Q8: Can you see the water? A8: No</p> <p>Question Q9: Is it sunny?</p> <p>Answer A9: Yes. ✓</p>
	<p>Caption C: Several signs are on 1 post, describing the intersection of east 42nd st and lexington ave ... at certain times.</p> <p>History Q1: Can you see roads? A1: No Q2: What color are the signs? A2: Blue with white writing Q3: How many signs are there? A3: 5 ... Q5: Can you see any buildings? A5: Yes Q6: Can you tell if it is day or night? A6: Day</p> <p>Question Q7: When can you not turn left onto Lexington?</p> <p>Answer A7: Not sure. ✗</p>		<p>Caption C: Sheep stand on a hill with a tree and houses in the background.</p> <p>History Q1: How many houses can you see? A1: 2 Q2: Is the hill grassy? A2: Yes, very Q3: Is it day time? A3: Yes Q4: Can you see the sky? A4: Yes Q5: Is the sky clear? A5: Yes</p> <p>Question Q6: How many sheep can you see?</p> <p>Answer A6: Two. ✗</p>

Figure 4. More qualitative examples of our model. We show the caption, dialog history, question, image, and the answers generated by our proposed RMK model.

knowledge. Similar observation exists in the second example. Faced with the difficult question of where the zebras are, RMK points the relevance of Africa in the facts and then chooses the optimal answer. With the commonsense knowledge, it generates a more informative answer “*somewhere in Africa*” instead of a safe response “*Not sure*”. It illustrates that our multi-structure knowledge reasoning architecture can not only extract the required information from the facts, but also capture the underlying dependence from vision and history.

In addition, we supply more qualitative examples from our model as shown in Figure 4. In the first four examples, our model can handle the diverse kinds of questions in visual dialog. The last two examples are the failure cases for our model. The second last one needs looking into the text on the image while our model not. For the last example, there are actually three sheep in the image, but the answer is “Two”. It shows that our model cannot well handle the

question related to the text on the image (may need OCR as in TextVQA [30]) and the complicated counting problem, which also remain open questions in multimodal systems.

5. Conclusion

In this paper, we introduce a novel model RMK for reasoning with commonsense knowledge in visual dialog. To properly suit the characteristics of dialog history and image in the task, we first represent commonsense knowledge at multi-structure level: sentence-level facts and graph-level facts. Then it captures and fuses relevant knowledge into visual dialog system, complementing with the visual graph and the history sentences. Experimental results on two datasets illustrate the superiority of our proposed model, and show the significant increase with external knowledge for VisDial task. The work will inspire research on visual dialog involving knowledge-based reasoning.

References

- [1] Shubham Agarwal, Trung Bui, Joon-Young Lee, et al. History for visual dialog: Do we really need it? In *ACL*, pages 8182–8197, 2020. [1](#), [2](#), [5](#), [6](#)
- [2] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, et al. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. [1](#), [2](#)
- [3] 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 *IEEE Conference on Computer Vision and Pattern Recognition Artificial Intelligence*, pages 6077–6086, 2018. [6](#)
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. [4](#)
- [5] Feilong Chen, Xiuyi Chen, Fandong Meng, et al. GoG: Relation-aware graph-over-graph network for visual dialog. In *Findings of ACL*, pages 230–243, 2021. [1](#), [3](#), [5](#), [6](#)
- [6] Feilong Chen, Fandong Meng, Jiaming Xu, et al. Dmrm: A dual-channel multi-hop reasoning model for visual dialog. In *AAAI*, pages 7504–7511, 2020. [2](#), [5](#), [6](#)
- [7] Abhishek Das, Satwik Kottur, Khushi Gupta, et al. Visual dialog. In *Proceedings of the IEEE international conference on computer vision*, pages 326–335, 2017. [1](#), [2](#), [5](#), [6](#)
- [8] François Gardères, Maryam Ziaefard, Baptiste Abeloos, and Freddy Lecue. Conceptbert: Concept-aware representation for visual question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 489–498, 2020. [2](#)
- [9] Dan Guo, Hui Wang, Hanwang Zhang, , et al. Iterative context-aware graph inference for visual dialog. In *Proceedings of the IEEE international conference on computer vision*, pages 10055–10064, 2020. [3](#), [5](#), [6](#)
- [10] Qingbao Huang, Jielong Wei, Yi Cai, Changmeng Zheng, Junying Chen, Ho-fung Leung, and Qing Li. Aligned dual channel graph convolutional network for visual question answering. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 7166–7176, 2020. [3](#)
- [11] Xiaoze Jiang, Siyi Du, Zengchang Qin, Yajing Sun, and Jing Yu. Kbn: Knowledge-bridge graph network for adaptive vision-text reasoning in visual dialogue. In *ACM MM*, page 1265–1273, 2020. [5](#), [6](#)
- [12] Xiaoze Jiang, Jing Yu, Zengchang Qin, , et al. Dualvd: An adaptive dual encoding model for deep visual understanding in visual dialogue. In *AAAI*, pages 11125–11132, 2020. [2](#), [3](#), [4](#), [5](#), [6](#)
- [13] Xiaoze Jiang, Jing Yu, Yajing Sun, et al. Dam: Deliberation, abandon and memory networks for generating detailed and non-repetitive responses in visual dialogue. In *IJCAI*, pages 687–693, 2020. [2](#), [5](#), [6](#)
- [14] Satwik Kottur, José MF Moura, Devi Parikh, et al. Visual coreference resolution in visual dialog using neural module networks. In *ECCV*, pages 153–169, 2018. [2](#), [5](#), [6](#)
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. [6](#)
- [16] 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, pages 11336–11344, 2020. [3](#)
- [17] Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. Relation-aware graph attention network for visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10313–10322, 2019. [3](#)
- [18] 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. [3](#)
- [19] Kenneth Marino, Xinlei Chen, Devi Parikh, Abhinav Gupta, and Marcus Rohrbach. Krisp: Integrating implicit and symbolic knowledge for open-domain knowledge-based vqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14111–14121, 2021. [2](#)
- [20] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, et al. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE international conference on computer vision*, pages 3195–3204, 2019. [1](#), [2](#)
- [21] Vishvak Murahari, Dhruv Batra, Devi Parikh, et al. Large-scale pretraining for visual dialog: A simple state-of-the-art baseline. In *ECCV*, pages 336–352, 2020. [2](#), [3](#), [5](#), [6](#)
- [22] Medhini Narasimhan, Svetlana Lazebnik, , and Alexander G Schwing. Out of the box: Reasoning with graph convolution nets for factual visual question answering. In *Advances in neural information processing systems*, pages 2654–2665, 2018. [1](#), [2](#), [3](#)
- [23] Van-Quang Nguyen, Masanori Suganuma, and Takayuki Okatani. Efficient attention mechanism for visual dialog that can handle all the interactions between multiple inputs. In *ECCV*, pages 223–240, 2020. [5](#), [6](#), [7](#)
- [24] Yulei Niu, Hanwang Zhang, Manli Zhang, et al. Recursive visual attention in visual dialog. In *Proceedings of the IEEE international conference on computer vision*, pages 6679–6688, 2019. [1](#), [2](#), [5](#), [6](#)
- [25] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014. [6](#)
- [26] 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. PMLR, 2021. [3](#)
- [27] Shaoqing Ren, Ross Girshick, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2017. [6](#)

- [28] Idan Schwartz, Seunghak Yu, Tamir Hazan, et al. Factor graph attention. In *Proceedings of the IEEE international conference on computer vision*, pages 2039–2048, 2019. 6
- [29] Sanket Shah, Anand Mishra, Naganand Yadati, and Partha Pratim Talukdar. Kvqa: Knowledge-aware visual question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8876–8884, 2019. 2
- [30] Amanpreet Singh, Vivek Natarjan, Meet Shah, Yu Jiang, Xinlei Chen, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8317–8326, 2019. 8
- [31] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. In *AAAI*, pages 4444–4451, 2017. 3, 6
- [32] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. 6
- [33] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2019. 3
- [34] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, 2019. 3
- [35] 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. 3
- [36] Peng Wang, Qi Wu, Chunhua Shen, , et al. Fvqa: Fact-based visual question answering. *TPAMI*, 40(10):2413–2427, 2017. 1, 2
- [37] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Henge. Explicit knowledge-based reasoning for visual question answering. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1290–1296, 2017. 2
- [38] Jialin Wu, Jiasen Lu, Ashish Sabharwal, and Roozbeh Motlaghi. Multi-modal answer validation for knowledge-based vqa. *arXiv preprint arXiv:2103.12248*, 2021. 3
- [39] Qi Wu, Peng Wang, Chunhua Shen, et al. Are you talking to me? reasoned visual dialog generation through adversarial learning. In *Proceedings of the IEEE international conference on computer vision*, pages 6106–6115, 2018. 5
- [40] Wang Yue, Joty Shafiq, R. Lyu Michael, , et al. Vd-bert: A unified vision and dialog transformer with bert. In *EMNLP*, pages 3325–3338, 2020. 2, 3, 5, 6
- [41] Ji Zhang, Yannis Kalantidis, Marcus Rohrbach, et al. Large-scale visual relationship understanding. In *AAAI*, pages 9185–9194, 2019. 4
- [42] Zilong Zheng, Wenguan Wang, Siyuan Qi, et al. Reasoning visual dialogs with structural and partial observations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6669–6678, 2019. 1, 2, 3
- [43] Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, pages 4623–4629, 2018. 3
- [44] Zihao Zhu, Jing Yu, Yujing Wang, , et al. Mucko: Multi-layer cross-modal knowledge reasoning for fact-based visual question answering. In *IJCAI*, pages 1097–1103, 2020. 1, 2