

# LaKo: Knowledge-driven Visual Question Answering via Late Knowledge-to-Text Injection

Zhuo Chen  
zhuo.chen@zju.edu.cn  
College of Computer Science and  
Technology, Zhejiang University

Yufeng Huang  
huangyufeng@zju.edu.cn  
School of Software Technology,  
Zhejiang University

Jiaoyan Chen  
jiaoyan.chen@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford

Yuxia Geng  
gengyx@zju.edu.cn  
College of Computer Science and  
Technology, Zhejiang University

Yin Fang  
fangyin@zju.edu.cn  
College of Computer Science and  
Technology, Zhejiang University

Jeff Z. Pan  
j.z.pan@ed.ac.uk  
School of Informatics, The University  
of Edinburgh

Ningyu Zhang  
zhangningyu@zju.edu.cn  
School of Software Technology,  
Zhejiang University

Wen Zhang\*  
zhang.wen@zju.edu.cn  
School of Software Technology,  
Zhejiang University

## ABSTRACT

Visual question answering (VQA) often requires an understanding of visual concepts and language semantics, which relies on external knowledge. Most existing methods exploit pre-trained language models or/and unstructured text, but the knowledge in these resources are often incomplete and noisy. Some other methods prefer to use knowledge graphs (KGs) which often have intensive structured knowledge, but the research is still quite preliminary. In this paper, we propose **LaKo**, a knowledge-driven VQA method via **Late Knowledge-to-text Injection**. To effectively incorporate an external KG, we transfer triples into textual format and propose a late injection mechanism for knowledge fusion. Finally we address VQA as a text generation task with an effective encoder-decoder paradigm, which achieves state-of-the-art results on OKVQA datasets.

## CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; **Knowledge representation and reasoning**; *Semantic networks*.

## KEYWORDS

Visual Question Answering; Knowledge Graph; Knowledge-to-Text; Late Knowledge Injection

### ACM Reference Format:

Zhuo Chen, Yufeng Huang, Jiaoyan Chen, Yuxia Geng, Yin Fang, Jeff Z. Pan, Ningyu Zhang, and Wen Zhang. 2022. LaKo: Knowledge-driven Visual Question Answering via Late Knowledge-to-Text Injection. In *Proceedings of The*

\*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

IJCKG'22, October 27–29, 2022, Virtual Event, China

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM. . . \$15.00

<https://doi.org/10.1145/nmnnnnn.nmnnnnn>

11th International Joint Conference on Knowledge Graphs (IJCKG'22). ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nmnnnnn.nmnnnnn>

## 1 INTRODUCTION

The task of Visual Question Answering (VQA) [2] is to answer natural language questions according to given images. Recently, some VQA methods [38, 48, 59] are developed to utilize the external knowledge for open-world scene understanding (a.k.a. knowledge-based VQA). According to how to incorporate knowledge, we divide the current works into two categories.

The first category is directly exploiting the knowledge in language model's parameters to answer questions [5, 40, 41, 44]. Specifically, inspired by the knowledge-based language model in NLP field, some methods [17, 37, 50] trying to inject the common-sense or factual knowledge as part of the model's parameter during model training. However, the knowledge in language model sometimes is insufficient for VQA scenario, and they are likely to fail when referring to **new knowledge that is out of origin training corpus**. More importantly, these encoder-based finetuning frameworks with a MLP attached behind the last layer limits the utilization of knowledge within the model itself [16].

Works from the second category are based on the knowledge retrieval strategy. We observe that those methods [25, 36, 42, 65] usually pass the vision-linguistic information through a search engine where the **network delay** might become a bottleneck. Others retrieve relevant corpus from encyclopedia articles, which leads to lots of **irrelevant information** and interferes with the model's judgment.

To address these challenges, in this paper, we propose **LaKo**, a knowledge driven VQA method via **Late Knowledge-to-text Injection**, to effectively incorporate the external information, as shown in Figure 1. Specifically, we develop a retriever-reader VQA architecture, with a knowledge retriever and a late injection mechanism between the knowledge and input corpus in reader. First of all, we construct a common-sense knowledge graph (KG) for VQA, and the retriever queries this KG according to the vision-language input to recall the target triples. Specifically, the modality is unified via the

transformation of images to captions and triples to sentences with Knowledge-to-Text strategy. Special prefixes are added to the front of them as knowledge guidance. Besides, we convert the VQA into a text generation task via an encoder-decoder paradigm, separate knowledge from the input corpus during encoding and integrate them at the stage of answer generation within decoder. We observe that our method further boost the performance compared to traditional pipeline and obtain state-of-the-art results. To sum up, the main contributions are summarized below:

- We propose a new KG retrieval paradigm for VQA together with a late knowledge injection strategy, which works without relying on annotated ground-truth knowledge.
- We improve and build a large-scale common-sense KG targeted at knowledge-based VQA, proving that high quality KG benefits VQA performance.
- Our method obtains state-of-the-art results on the OKVQA dataset, and verifies that using a high quality KG as the external knowledge is better than using unstructured text and pure language model parameters. Our code is available at <https://github.com/hackerchenzhuo/LaKo>.

## 2 RELATED WORK

### 2.1 Visual Question Answering (VQA)

Since being proposed by [2], extensive VQA methods [1, 11, 28, 55] have emerged to focus on applying multi-modal feature fusion between questions and images for answer decision. Recently, with the development of the pretraining technique, more systems [32, 33, 35, 53, 63, 64] begin to utilize multi-modal transformer architectures for the VQA task via vision-language pretraining (VLP) technique.

### 2.2 Knowledge-based VQA

The knowledge-based VQA [38, 48, 59] requires the model to acquire factual or common-sense knowledge outside the image for question answering. According to the way of knowledge incorporation, we divide the current works into two categories.

**Exploit Knowledge in Language Model.** Recent efforts in NLP fields [5, 40, 41, 44, 60] emphasize the relational world knowledge contained in huge pre-trained language models (PLMs). Inspired by these perspectives, many works [46, 56, 62] directly apply the PLMs into VQA tasks, where the PLM plays a role for question and image understanding (a.k.a. reader). They hold the view that the knowledge within the PLMs is sufficient to support knowledge-based multi-modal tasks despite no additional knowledge input, which makes the VQA become a machine reading comprehension (MRC) problem. Specifically, Salaberria et al. [46] convert an image into a caption, and then feed it together with question into the BERT [14] to predict the answer with an added classification head. Nonetheless, the encoder-based finetuning framework with a MLP attached behind the [CLS] position or pooling layer limits the utilization of the knowledge within the model itself [16]. Recent works [56, 62] endeavor to apply prompt for decoder-based model to solve the VQA problem under the few shot setting, but the max input length of model itself limits the knowledge utilization.

Some other works inject knowledge during the model training, letting the common-sense knowledge become part of the reader's

parameters. For example, [50] employ an auxiliary training objective that encourages the entity representation to align with the corresponding graph embedding in a KG. ConceptBert [17] introduces a multi-modal representation which learns a joint Concept-Vision-Language embedding. KRISP [37] exploits the implicit reasoning of transformer models, integrates symbolic representations from a knowledge graph (KG), and combines them together through a relational graph convolutional network (RGCN) [47].

Although black-box models make knowledge an implicit representation, they are likely to fail when requiring implicit new knowledge that is out of the origin knowledge base (KB). In this study, we not only make full use of the implicit knowledge in the PLMs, but also selectively carry out additional KG corpus to supplement those ancillary explicit knowledge which may be ambiguous within the PLMs. Besides, we decouple the knowledge from models, hoping that the reader could focus on understanding the input auxiliary corpus. This avoids the interference of model semantic understanding during the knowledge injection process, and keeps our model being sensitive to the new knowledge.

**Knowledge Retrieval Strategy.** It is natural to think of adding a separate retrieval module (a.k.a. retriever) to recall the required explicit knowledge as external input of the downstream reader. In order to take advantage of the information on the Internet, [25, 36, 38, 42] pass the vision-linguistic information through a search engine (e.g., Google) to retrieve relevant corpus (e.g., sentences from Wikipedia articles or snippets in searching result) as weak positive knowledge samples, which are further passed to the reader module for knowledge incorporation. Within the above methods, Luo et al. [36] apply the previously retrieved snippets as a KB, and assign those snippets which contain the answer words as weak-supervised signals for retriever training. Besides, Wu et al. [61] leverage not only the structured knowledge, but also the image knowledge (from Google image search) to revise the answer. However, the network delay might become a bottleneck for all these policies when take the search engine as the retriever, and the unstructured knowledge probably leads to the decrease of knowledge density.

Considering that highly dense knowledge is stored in structured KG triples, [39, 59, 65] construct context-aware subgraph from a large scale KG (e.g., ConceptNet [52]) based on entity name matching or embedding similarity. But all of them preserve the original graph structure with a GNN-based model followed, which is deemed insufficient to exploit all useful evidence provided by external knowledge [5]. In this paper, we propose a vision-language KG retriever together with a Knowledge-to-Text transformation strategy. They unify the structured knowledge and visual data into a text modality to exploit the semantic understanding capability of PLMs, rather than relying on the message passing mechanism in GNNs.<sup>1</sup>

## 3 METHODOLOGY

A VQA task is to provide an answer *ans* given an image *v* paired with a question *q*. Following [2], there are a list of (usually ten) acceptable ground truth answers ( $GT_{ans}$ ) for each (*v*, *q*) pair.

<sup>1</sup>KAT [20], K-LITE [49] and TRiG [15] are the works in the same period as ours so far, which both take advantage of the knowledge bases and the richness of PLMs.




Figure 1: The overview model architecture of LaKo. Given a  $(v, q)$  pair, the generated background text and knowledge text retrieved by vision-language KG retriever are separately send to reader for late knowledge injection. The predicted *ans* tokens are decoded in turn.

### 3.1 Vision-Language KG Retriever

For the retrieval of factual triples from large scale KB, some works [30, 39] utilize all detected objects in an image as the reference for knowledge retrieval. However, this brings in a lot of irrelevant noise and makes the model easy to lose focus, especially when the number of appeared objects is not small. We observe that the image caption naturally contains a human-like attention mechanism on vision. So, instead of directly using vision modality data, our first step is to transform the visual content into the textual format. Given a  $(v, q)$  pair, we convert the input image  $v$  into corresponding image caption  $C(v)$  with a pre-trained model at first. Besides, we further apply optical character recognition (OCR) technique [58] for text extraction to improve the information integrity, and concatenate them with  $C(v)$  to get an image representation  $\tilde{v}$  with a text form:  $\tilde{v} = \text{Concat}[O(v); C(v)]$ , where  $O(\cdot)$  denotes OCR output. We note that the pre-trained captioning model and the OCR model could be regarded as the modules for caption feature extraction. Similar to ResNet [21] for image feature or BERT [14] for textual feature, these models generate the simple descriptions toward the image that are possibly related to ground truth background knowledge but not exactly equal to.

A KG may contain thousands of facts about a concept, but only several of them are relevant to the given  $(v, q)$ . Therefore, we reduce the scope of the KG through the establishment of a stem corpus in the VQA field, and make sure that all triples in the KG contain at least one stem within this corpus (see Section 4.2 for details). Then, we query the KG to get target triples based on a Knowledge-to-Text technique and a newly stem-based BM25 [45] algorithm.

**Knowledge-to-Text Transformation.** Knowledge facts are usually triples while the questions and answers are textual format. To realize the unification of three different modalities (i.e., vision, unstructured language, and structured knowledge) data, firstly, we translate the KG factual triples  $t_f$  into the sentence  $s_f$ . Specially, for

those relations with high frequency, we follow the template-based method in [5, 22] with pre-defined cloze templates and conduct manual calibration. For those long-tail or newly added relations, we apply BERT [14] tokenizer for coarse-grained word segmentation (e.g., “locatedin” is converted into “located” “##in”), and then generate cloze templates automatically after normalization. Despite many relations are readable for the PLMs even without the above preprocess, other matching-based retrieval approaches like BM25 will benefit greatly from those templates.

**Stem-based BM25.** The major discrepancy from the original BM25 is that our stem-based one defines the a word stem<sup>2</sup> as the smallest semantic unit rather than an entire word. Our motivation is to maximize the knowledge from the limited VQA and KG corpus via stem merging. Particularly, we remove these extra words with repeated stems in  $\tilde{v}$  before concatenating it with stem in  $q$  rather than de-duplicate on final  $S_{query}$ , since we want to maintain those important information emphasized by the  $q$ . Then we get a stem-based sequence  $S_{query}$  with  $s_1, s_2, \dots, s_t$ , and calculate the score for each factual triple sentence  $s_f$  via:

$$\text{Score}(S_{query}, s_f) = \sum_{i=1}^t w_i * R(s_i, s_f), \quad (1)$$

where  $w_i$  represents the significance of  $s_i$ :

$$w_i = IDF(s_i) = \log \frac{N - n(s_i) + 0.5}{n(s_i) + 0.5}, \quad (2)$$

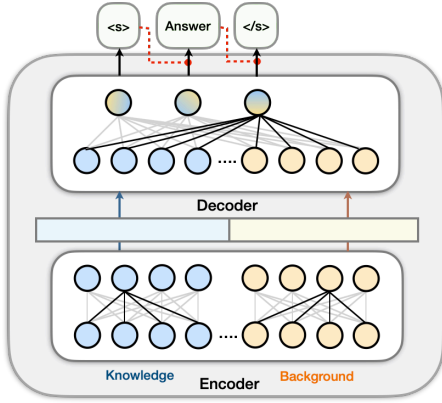
where  $N$  denotes the total number of  $s_f$  in KG and  $n(s_i)$  denotes the number of  $s_f$  containing the stem  $s_i$ . Hyperparameter 0.5 is mainly for smoothing computation and  $R(s_i, s_f)$  measures the semantics correlation between  $s_i$  and  $s_f$  [45].

<sup>2</sup>We get word stem via Porter Stemmer algorithm <https://tartarus.org/martin/PorterStemmer>

Finally,  $s_f$  are retrieved according to their stem-based BM25 score. The Top-K  $s_f$  are concatenated to get the  $S_{fact}$  as external knowledge for each  $(v, q)$  pair, which contributes to the late knowledge injection within the reader.

### 3.2 Late Knowledge Injection

Recent VLP-based methods are substantially based on an encoder architecture with a MLP attached to the [CLS] position or the pooling layer, which limits the utilization of knowledge within the pre-trained model [16]. Inspired by previous works [7, 22] which explore the knowledge within PLMs, we unify all data into textual to fully exploit the semantic understanding capability of the text-only PLM. Specifically, we apply the encoder-decoder transformer architecture as the reader, following [12] to convert the knowledge-based VQA from a classification task into a text generation task. Differently, considering that the entities in the KG do not exist in isolation and a closed loop is formed among triples, we adapt the Fusion-in-Decoder (FiD) [24] into a new late injection paradigm to avoid the interference of vision-language information on knowledge self-integrate process. As the self-attention architecture of



**Figure 2: The self-attention architecture for late knowledge injection. The knowledge and background information are interact in encoder independently and fused in decoder to collectively predict the answer.**

Late Knowledge Injection shown in Figure 2, this policy can achieve the independent encoding within  $S_{fact}$  for knowledge aggregation, and assist on late joint decoding among  $(q, \tilde{v}, S_{fact})$  for knowledge searching. Besides, independent encoding also reduces the computation within the encoder during the self-attention process from  $O((N + M)^2)$  to  $O(N^2 + M^2)$  where  $N$  and  $M$  denote the length of the knowledge and background input, respectively.

In particular, we first add special prefixes `question:`, `context:` and `fact:` before the  $q$ ,  $\tilde{v}$  and  $S_{fact}$  as the knowledge guidance, which make up two texts with independent semantics:

- Background context: [question:  $q$ , context:  $\tilde{v}$ ]
- Knowledge context: [fact:  $S_{fact}$ ]

Then the encoder independently processes the background and knowledge context through  $N$  layers transformer [57]. The output hidden state from each layer of the encoder form a **global representation**  $X$  of dimension  $(L_b + L_k) * d$ , where  $L_k / L_b$  denote the

length of tokenized knowledge / background context, and  $d$  is the hidden state dimension of the model. As a regular autoregressive model, there are self-attention, cross-attention and feed-forward modules in each layer. For each head in one transformer layer of the decoder, the attention is defined as follow:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (3)$$

where  $K, Q, V$  denotes matrices of key, query, value for input tokens [57], respectively, and  $Q = W_q H$ ,  $K = W_k H$ ,  $V = W_v H$  denotes the output of previous (self) attention layer ( $H \in \mathbb{R}^d$ ). Specially, the cross-attention process in each layer of the decoder is the only part for message exchange between encoder and decoder, where  $K = W_k X$  and  $V = W_v X$ , which is the key of late knowledge injection.

Finally, the *ans* tokens is decoded one by one with the begin of the start token (e.g., `<s>`), and stop with the end token (e.g., `</s>`). Meanwhile, the whole encoder-decoder framework is optimized via minimizing the negative log-likelihood:

$$\mathcal{L}_\theta = - \sum_{j=1}^{|y|} \log P_\theta(y_j | y_{<j}, q, v, S_{fact}), \quad (4)$$

where  $y$  are tokenized from  $GT_{ans}$  for a given  $(v, q)$  pair.

## 4 EXPERIMENTS

### 4.1 Dataset

**VQA2.0** [2] is a large standard VQA dataset containing about 1.1 million open-ended questions with 204,721 images. Each question is associated with 10 different answers obtained by crowdsourcing. **OKVQA** [38] is a recent dataset where the visual content of an image is not sufficient to answer the question. There is not any exact ground truth common-sense fact triple for question support and all *ans* are annotated by volunteers. In addition, all the images are from COCO 2014 validation set.

We note that other datasets like the FVQA [59] is smaller and easier than OKVQA since it targets at a reasoning over a given KB rather than visual reasoning with the open knowledge. Thus we mainly focus on OKVQA for model validation.

### 4.2 Knowledge Graph Construction

Following [37], we consider taking the common-sense knowledge (e.g., what are paper made of) and scientific knowledge (e.g., what genus are cats) to construct a new KG toward the knowledge-based VQA. Differently, we exclude situational knowledge (e.g., where do cars tend to be located) outside since it may mislead the reader sometimes when the image scene is not typical (e.g., unseen scenarios). Several knowledge sources are incorporated (we manually add “\_” on relations just for easy reading):

- **ConceptNet** [52] is a semantic network which contains human common-sense knowledge about the world;
- **WebChild** [54] contains triples which connect nouns with adjectives via more fine-grained relations (e.g., “*has\_shape*”, “*faster*”);
- **DBpedia** [3] includes knowledge extracted from Wikipedia, which covers many fields and our daily life;

- **hasPart KB** [4] collects “*has\_part*” relationships between common objects such as  $\langle \text{dog}, \text{has\_part}, \text{whiskers} \rangle$  or scientific ones like  $\langle \text{molecules}, \text{has\_part}, \text{atoms} \rangle$ ;

Firstly, we collect triples from the above four knowledge sources to constitute the original KG (more than 900K triples). In particular, in WebChild we filter the first 100K triples according to the normalized triple confidence score. The relations from Dbpedia are mainly “*category*”, and about 50k “*has\_part*” triples are fetched from hasPart KB. Several relations with a large number but low potential contribution (e.g., “*Synonym*”, “*Antonym*”) are removed.

Next, we collect all of the symbolic entities from the dataset, including the words on questions, answers, generated image caption, and OCR recognition. Based on the distribution of **word frequency** (from high to low), we remove common stop words and keep those that may have impact on knowledge representation (e.g. “*can*”). The remaining words constitute the **VQA corpus** with their stem representation and we only retain those triples in KG whose heads and tails both contain stems in VQA corpus.

In addition, we define those relations that occur more than 10,000 times in our KG (e.g. “*related\_to*”) as the frequent relations. For the triples that have identical subjects and objects (e.g.,  $\langle \text{person}, \text{related\_to}, \text{hand} \rangle$  and  $\langle \text{person}, \text{has\_part}, \text{hand} \rangle$ ), we remove those triples associated with frequent relations. For example, according to our statistics, 13836, 2584, 2533, 2391 triples are deleted, which separately contain the relation of “*related\_to*”, “*used\_for*”, “*at\_location*”, “*is\_a*”. Finally, we got a KG with 300,559 triples, 96191 entities and 2198 relations.

### 4.3 Metrics

**Acc.** For those classifier-based approaches, we apply the standard evaluation code<sup>3</sup>, which calculates the accuracy (Acc) metric recommended in the VQA challenge [2]:

$$\text{Acc}(ans) = \min\left(1, \frac{\#\{\text{human that said that ans}\}}{3}\right). \quad (5)$$

**EM.** For our text generation VQA framework, we use exact match (EM) when calculate the Acc, where the generated *ans* is compared to the ground truth answers ( $GT_{ans}$ ) after normalization. Nevertheless, the unfixed answer length makes the size of answer space indefinite (i.e., even much bigger than the PLM’s vocabulary size) for the autoregressive model. Likewise, the EM metric cannot make fair judgment sometimes which may result in some potential answers being left out (e.g.,  $GT_{ans}$  is “in oven” while *ans* is “oven”). Hence, we introduce two new metrics that are EM variants:

**Inc.** Inclusion-based Acc metric regards the *ans* as correct when it includes one answer in  $GT_{ans}$  or is included by one answer in  $GT_{ans}$  after normalization.

**Stem.** Stem-based Acc metric makes a small change toward **Inc**: making judgment according to whether *ans* and  $GT_{ans}$  have an intersection on stem (e.g., the stem of “*happy*” and “*happiness*” are both “*happi*”).

It is noteworthy that all the **EM**, **Inc** and **Stem** metrics will preferentially match the high score answer first. Meanwhile, the normalization procedure should strictly remove those stop words in both *ans* and  $GT_{ans}$  to avoid the disturbance.

<sup>3</sup><https://github.com/GT-Vision-Lab/VQA>

### 4.4 Training Details

During knowledge retrieval, we utilize a SOTA caption model VinVL [64] for generating image descriptions, and we finetune its pre-trained checkpoint with COCO 2014 training set to prevent the leakage of  $v$  in test data. The top-10 generated factual sentence  $s_f$  ( $K = 10$ ) are adopted to constitute the final  $S_{fact}$ . According to our statistics, the average word number in  $S_{fact}$  is about 57.

We respectively initialize the reader in LaKo-large/base with official pre-trained T5-large/base parameters. We utilize AdamW optimizer with initialized learning rate  $4e-5$  (with warm-up ratio 6%) and the model is trained for 20 epochs with max sequence length 130, early stop patience 5 (i.e., the training algorithm waits 5 epochs before early stop if no progress on the validation set). For the large version of LaKo, the mini-batch size is set as 8, which is 16 in the base version.

### 4.5 Overall Results

Table 1 summarizes the main result on the OKVQA dataset. Our LaKo follows the large version with late knowledge injection strategy and is re-pretrained using VQA2.0 training data unless otherwise specified. Specifically, we compared LaKo with recent LM-based (BERT [14], RoBERTa [34], GPT-3 [6]) or VLP-based ( VisualBERT [31], ViLBERT [35], LXMERT [53], MMF [51]) protocols for fairness. In addition to these methods, we also leverage some other strong baselines:

- **LXMERT** [53]. A SOTA two stream VLP model, where we extract image region features from a pre-trained Faster R-CNN [43] as visual input.
- **T5** [53]. A recent seq-to-seq LM with encoder-decoder paradigm. We simply take the question and caption as the input with special prefixes put ahead<sup>4</sup>.

We find that many previous works involve explicit knowledge retrieval from search engines (e.g., “Google Search” and “Google Image”), or store a huge number of unstructured encyclopedia texts as the background KB in advance (e.g., “Wikipedia”). However, in realistic application, the network bandwidth may become the bottleneck, while noise contained in unstructured text may limit the knowledge scale due to the limitation of max sequence length. Nevertheless, we compare LaKo with these approaches, observing that the performance of our model still exceeds the SOTA method (PICa-Base (C+T), with 175B parameters GPT-3 without multi-query ensemble) and further improves by 3.71%. We also concatenate the  $S_{fact}$  with  $q$  as input text to LXMERT, which achieves an improvement of 1.01% on Acc. Most importantly, since our results are based on EM, the actual acceptable answer accuracy may be higher than our statements. We believe that T5’s generalization on VQA tasks mainly comes from its large pre-trained Colossal Clean Crawled Corpus (C4) [53] dataset, and the basic encoder-decoder architecture makes it flexible when migrated to other tasks. Moreover, we get rid of the dependence on search engines and annotated ground-truth knowledge, which is easy for other researchers to follow.

We guess that two points are the key issues to the low performance of those traditional classifier-based approaches:

<sup>4</sup>Other encoder-decoder models such as BART [29] performs not as good as T5 as we have tested in VQA field. Thus we take T5 as the backbone of LaKo.

**Table 1: Performance (%) on the OKVQA test-split. Since we only compared with those LM-based or VLP-based models, we do not highlight LM on “Knowledge Src”. The full names of these abbreviations are as follows: GS (Google Search), W (Wikipedia), KG (Knowledge Graph), GI (Google Image), Enc.Dec.(Encoder-Decoder), TG (Text Generation), IE (Information Extraction), CLS (Classification), Src (source), Acc (Accuracy), C (Caption), T (Tag).**

Method	Backbone	Architecture	Knowledge Src	Acc
ConceptBert [17]	BERT	Encoder (CLS)	KG	33.66
KRISP [37]	VisualBERT	Encoder (CLS)	W & KG	38.35
MAVE <sub>x</sub> [61]	ViLBERT	Encoder (CLS)	W & KG & GI	38.70
Caption-DPR + CReader [36]	LXMERT	Encoder (CLS)	GS	36.78
Caption-DPR + EReader [36]	RoBERTa	Encoder (IE)	GS	39.20
KGE Aligning [50]	LXMERT	Encoder (CLS)	KG	39.04
CBM + MMBERT [46]	BERT + MMF	Encoder (CLS)	-	39.20
PICa-Base (C+T) [62]	GPT-3	Decoder (TG)	GS	43.30
LXMERT [53]	-	Encoder (CLS)	-	36.91
LXMERT [53] + Knowledge	-	Encoder (CLS)	KG	37.92
T5[53] + Prefixes	-	Enc.Dec. (TG)	-	42.03
<b>LaKo</b>	T5	Enc.Dec. (TG)	KG	<b>47.01</b>

(i) In order to mitigate the long-tail problem in answers and the impact of disjoint answers between training & testing set (e.g., the total number of answers in VQA2.0 is 29,140 in our statistics, with just 15,259 intersections), they have to artificially prescribe the answer candidate set based on occurrence (Occ.) frequency and therefore determine the output dimension of the last MLP layer in classifier. It is a **trade-off between answer coverage and error rate**. This also inevitably affect the upper bound of model performance: the oracle Acc of VQA2.0 drops to 92.86% with Occ. 9, and OKVQA respectively drops to 72.08% (Occ. 10), 85.38% (Occ. 5), 91.44% (Occ. 3). (ii) The pluggable MLP layer attached behind the model, which is proved by many previous works like [16].

Besides, the phenomenon that recent SOTA works are rarely based on GNN supports the view that those GNN-based models may not fully exploit all evidence provided by external knowledge [5] and the knowledge within the PLM is essential. Since the single-modal data is much richer compared to limited image-text paired data, it makes the text-only PLM even stronger than VLP model on knowledge-based VQA. Furthermore, we hold the review that our model is not trained on natural **long language sequences** in the VQA field since the answer’s average length is 1.3/1.2 words in OKVQA/VQA2.0. According to our statistics, the length of the answers generated by LaKo is 1.23, which supports the idea that making VQA a simpler generative task is beneficial.

## 4.6 Ablation Study

**Effect of Knowledge and Finetuning.** In Table 2 we discuss the impact of factual knowledge and the pretraining on LaKo. “Finetune” refers to re-pretraining LaKo on VQA2.0 before finetuning on OKVQA, which is necessary since the VQA is a relatively unfamiliar logic for text-only PLMs. We observe that the knowledge retrieval and the re-pretraining in LaKo-large lead to 2.71% and 2.03% improvement, respectively, which is constant in LaKo-base. We own this to the fact that massive implicit knowledge is contained in PLMs’ parameter after the pretraining process on Internet corpus, and LaKo could further exploit those explicit external knowledge

to support more precise answer prediction. Besides, the result on Inc-based and Stem-based Acc usually can be 5%-7% higher than the original EM-based one, which shows the potential of our architecture toward multi-modal knowledge tasks.

**Table 2: Ablation study (%) for the effect of knowledge and finetuning. The full names of these abbreviations are w/o (with-out), KGR (Knowledge Graph Retrieval). “Finetune” refers to re-pretraining LaKo on VQA2.0 dataset first.**

Method	EM	Inc	Stem
LaKo-large	<b>47.01</b>	<b>53.09</b>	<b>53.97</b>
- w/o {KGR}	44.74	50.90	51.70
- w/o {Finetune}	44.06	49.48	50.25
- w/o {Finetune & KGR}	42.03	47.34	48.11
LaKo-base	42.21	48.17	49.06
- w/o {KGR}	40.27	45.91	46.85
- w/o {Finetune}	39.89	45.05	45.93
- w/o {Finetune & KGR}	38.71	43.80	44.32

**Table 3: Result (%) on VQA2.0 test-dev.**



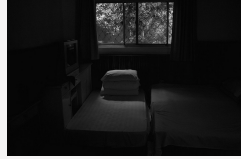

Method	Acc	Inc.
T5	66.38	68.85
ViLBERT [35]	67.90	-
LXMERT	69.15	-
LaKo	68.07	70.47
VinVL [64]	<b>75.95</b>	-

Furthermore, we also make experiments on the standard VQA2.0 dataset as shown in Table 3, where the results reflect that the knowledge is also effective for generic VQA (Acc is improved from 66.38% to 68.07%), and LaKo is comparable to part of VLP models like ViLBERT [35] and LXMERT [53]. However, it is not easy for us to surpass those SOTA VLP model without large-scale VLP model pretraining and fine-grained image feature providing.



**Table 4: The performance (%) discrepancy between late and early injection.**

Method	LaKo-large			LaKo-base		
	EM	Inc.	Stem	EM	Inc.	Stem
Early Injection	46.37	52.26	53.29	41.41	47.54	48.22
Late Injection	47.01 (↑ 0.64)	53.09 (↑ 0.83)	53.97 (↑ 0.68)	42.21 (↑ 0.80)	48.17 (↑ 0.63)	49.06 (↑ 0.84)

Question	Why is there a carriage in front of us?	How did the man extinguish one of the candles?	What is the pane of glass in the wall called?	Where might a person dress like this?
Image				
Caption	Horse drawn wagons with passengers pulled on path in forest ...	There is a man that is looking at a chocolate cake in the table with a candle in it ...	A very dark bedroom with unmade beds and a really big window ...	Person in business clothes sitting in a chair wearing a suit and tie with his hands crossed...
Retrieved Knowledge	<ul style="list-style-type: none"> <li>horse is located in harness pull carriage</li> <li>wagon is used for traveling</li> <li>horse is a large animal capable of carry person on it back</li> </ul>	<ul style="list-style-type: none"> <li>person is capable of blow out candle</li> <li>cake is related to with candle</li> <li>cake is related to blow candle</li> <li>...</li> </ul>	<ul style="list-style-type: none"> <li>pane is part of window.</li> <li>glass is related to window pane.</li> <li>pillows is located in bedroom.</li> <li>...</li> </ul>	<ul style="list-style-type: none"> <li>suit is related to business</li> <li>office is a where person do business</li> <li>tie is related to with shirt</li> <li>...</li> </ul>
GT Answer	travel: 1.0   tour: 0.6   trail: 0.6	blow: 1.0   blow on it: 0.6   blew it out: 0.6	window: 1.0	office: 1.0   wed: 0.6   interview: 0.6
Prediction	Lako: <b>travel</b> ✓ T5: <b>tourist</b> ✓ LXMERT: <b>annual</b> ✗	Lako: <b>blow</b> ✓ T5: <b>blow torch</b> ✓ LXMERT: <b>wave</b> ✗	Lako: <b>window</b> ✓ T5: <b>curtain</b> ✗ LXMERT: <b>sun</b> ✗	Lako: <b>office</b> ✓ T5: <b>business</b> ✗ LXMERT: <b>business</b> ✗

**Figure 3: We visualize some predictions and their corresponding retrieved facts  $s_f$ . Special, each question is associated with 10 different answers, and the score of  $GT_{ans}$  is calculated via  $\min(1, \frac{\#\{human\ that\ said\ that\ ans\}}{3})$ .**

**Benefits of Late Injection.** We studied the impact of late knowledge injection, which separate the background text and knowledge text on input and focuses on the knowledge aware answer generation within decoding stage via computing cross-attention between current generative tokens and these input corpora. According to the results shown in Table 4, we can see that the late injection process has a stable improvement (0.6%~ 0.9%) on answer prediction, which suggests that LaKo benefits from late knowledge injection rather than simply puts all these texts together as the input.

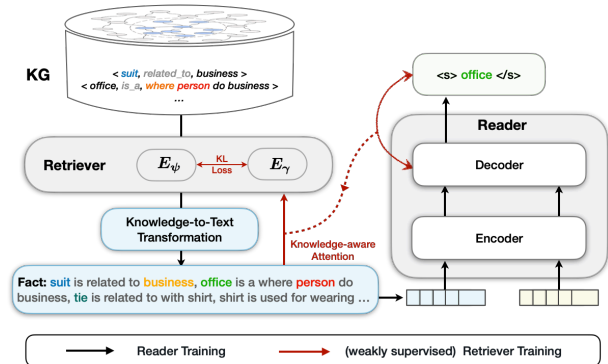
**Influence of KG quality.** We discuss the influence of KG quality on final performance of LaKo. As the result shown in Table 5, since the stem limitation of VQA corpus offers effective constraint for the knowledge range of the model, it mitigates the distraction from irrelevant information. Meanwhile, we observe that removing part of triples with redundant frequent relations also has a positive impact on the model.

**Table 5: “Stem Filter”: filtering out triples whose heads or tails do not contain stems in VQA corpus. “Freq. Rm”: selectively removing those triples which contain frequent relations (see Sec. 4.2 for details).**

KG	EM	Inc.	Stem
Final version	<b>47.01</b>	<b>53.09</b>	<b>53.97</b>
- w/o {Freq. Rm}	46.79	52.66	53.52
- w/o {Stem Filter & Freq. Rm}	46.26	52.17	52.91

### 4.7 Interpretability

To demonstrate the effectiveness of LaKo and that the knowledge injection is effective, we visualize some predictions and their corresponding retrieved facts  $s_f$ . As illustrated in Figure 3, LXMERT outputs wrong predictions in all these four cases, and sometimes outputs totally unrelated answers (e.g., “annual” in the first case). T5 outputs reasonable answers in the first and second cases, but makes mistakes in the third and fourth cases where there are distracting information within the caption or the query is ambiguous. In contrast, some retrieved  $s_f$  in LaKo are relevant to the ground truth answers (e.g., “office is where person do business” in the fourth case), which helps LaKo make correct prediction in those cases.



**Figure 4: The training process of weakly supervised differentiable retriever.**

## 4.8 Feasibility of Differentiable KG Retriever

The KG retriever plays an important role in our framework. So, is it feasible to train the retriever for better performance? Following [23], we leverage the cross attention between the token of the prediction output and input  $S_{fact}$  for retriever training, as shown in Figure 4. In particular, the minimum unit of attention score here is  $s_f$ . Therefore, no more than  $K$  (i.e., 10 in our work) scores are generated for each  $(v, q)$  pair, which are then utilized to training the retriever as the weak-supervised signals of corresponding  $s_f$ . We claim the retriever as a pseudo-siamese network [10] with  $E_Y$  for encoding  $s_f$  and  $E_\psi$  for encoding  $S_{query}$ , which only share the model architecture (BERT-base) rather than sharing parameters. It is optimized through minimizing the KL-divergence:

$$\mathcal{L}_{KL} = \sum_{f \in \mathcal{K}_{kg}} A_{q,f} \left( \log A_{q,f} - \log O(S_{query}, f) \right), \quad (6)$$


where  $\mathcal{K}_{kg}$  denotes the collection of those (top-k) retrieved textual triples via Knowledge-to-text transformation and

$$A_{q,f} = \frac{\exp(\text{Atten}_{q,f})}{\sum_{f' \in \mathcal{K}_{kg}} \exp(\text{Atten}_{q,f'})}, \quad (7)$$

$$O(S_{query}, f) = \frac{\exp(E_\psi(S_{query})^T E_Y(f))}{\sum_{f' \in \mathcal{K}_{kg}} \exp(E_\psi(S_{query})^T E_Y(f'))}. \quad (8)$$

For the aggregation of  $\text{Atten}_{q,f}$ , we apply several strategies:

- 1) taking the max or average or the top 1/2 attention scores over the input tokens corresponding to a  $s_f$ ;
- 2) taking the scores from last half layers or all the layers;
- 3) adding additional score bias (e.g., 1) to the  $s_f$  which contained answer stems.



**Figure 5: The results of recall (R) by different retrievers. Full names of these abbreviations: full/half (taking the scores from all/last half of the layers), max/mean/21mean (taking the max/average/the top 1/2 scores over the input tokens corresponding to a  $s_f$ ), pre (use the model finetuned at VQA2.0). Besides, DPR [27] refers to adding 1 additional score bias to the  $s_f$  which contains the answer stems.**

The motivation for choosing the top 1/2 attention scores is to extract the attention signals from the most valuable part of each  $s_f$  rather than a single token or simple averaging all the tokens;

Specifically, all the attention scores are computed toward the first output token of the decoder, and the scope of retrieval is within the TOP-500 facts retrieved via stem-based BM25 (we also strive to retrieve from the whole KG with 300K triples via faiss engine [26], only to get a poor result with Recall 23.37% on Top-100 facts). In order to measure the performance for retrieval, we introduce Inc-based Recall and define it as a **success recall when the retrieved knowledge sentence include the answer stem**. We train our retriever with different attention scores from both LaKo-large/base.

The Inc-based Recall rate for Top-K facts are shown in Figure 5. As we can see, the Recall rate by the stem-based BM25 is higher than the trained retrievers when  $K$  is smaller than 20, but becomes lower than some trained retrievers as  $K$  increases. Since we only utilize the Top-10  $s_f$  for vision-language reading, this result is not optimistic enough for iterative training on retriever, and we observe that the VQA Acc simultaneously drops with lower fact Recall rate. We guess that multiple different entities within the caption/question and the sparse embedding space impact the final similarity-based embedding recall, so we simply select the stem-based retrieval strategy in our primary experiment.

## 5 DISCUSSION AND FUTURE WORK

- Given a  $(v, q)$  pair, sometimes there are several correct answers in real world, but the generation-based approach only gets one answer at a time instead of sorting those potential ones. It is an interesting direction to consider generating multiple answers with particular policy, where the trie-based search [13] strategy could be considered for generative answers' domain constrain over the candidates.
- KG driven zero-shot problem [8, 9, 18, 19] on VQA also deserves a deeper research to further discuss the trade-off between PLMs and KGs, which requires the model to have better generalization ability toward the real world scenarios.
- we believe that the differentiable KG retriever in VQA field would be practicable in the future with better knowledge representation learning methods and high-quality knowledge annotations published.

## 6 CONCLUSION

In this paper, we propose LaKo, a knowledge-driven VQA method via late knowledge-to-text injection, to effectively incorporate both the knowledge from the KG and the PLM itself. Specifically, we address the VQA as a text generation task with an effective encoder-decoder paradigm under vision-language retriever-reader architecture, which achieves state-of-the-art result on standard knowledge-based VQA dataset OKVQA. Besides, we also pay attention to the KG construction, observing that KG with higher quality contributes to better performance of LaKo. This could be an exploration direction for future works. More importantly, we get rid of the dependence on annotated ground-truth knowledge and search engines, which is easy for other researchers to follow.

## 7 ACKNOWLEDGEMENTS

This work is funded by NSFCU19B2027/91846204 and the EPSRC project ConCur (EP/V050869/1).



## REFERENCES

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*. Computer Vision Foundation / IEEE Computer Society, 6077–6086.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2425–2433.
- [3] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. In *ISWC/ASWC (Lecture Notes in Computer Science, Vol. 4825)*. Springer, 722–735.
- [4] Sumithra Bhakthavatsalam, Kyle Richardson, Niket Tandon, and Peter Clark. 2020. Do Dogs have Whiskers? A New Knowledge Base of hasPart Relations. *CoRR* abs/2006.07510 (2020).
- [5] Ning Bian, Xianpei Han, Bo Chen, and Le Sun. 2021. Benchmarking Knowledge-Enhanced Commonsense Question Answering via Knowledge-to-Text Transformation. In *AAAI*. AAAI Press, 12574–12582.
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *NeurIPS*.
- [7] Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021. Knowledgeable or Educated Guess? Revisiting Language Models as Knowledge Bases. In *ACL/IJCNLP (1)*. Association for Computational Linguistics, 1860–1874.
- [8] Jiaoyan Chen, Yuxia Geng, Zhuo Chen, Ian Horrocks, Jeff Z. Pan, and Huajun Chen. 2021. Knowledge-aware Zero-Shot Learning: Survey and Perspective. In *IJCAL*. ijcai.org, 4366–4373.
- [9] Jiaoyan Chen, Yuxia Geng, Zhuo Chen, Jeff Z. Pan, Yuan He, Wen Zhang, Ian Horrocks, and Huajun Chen. 2021. Low-resource Learning with Knowledge Graphs: A Comprehensive Survey. *CoRR* abs/2112.10006 (2021).
- [10] Xinlei Chen and Kaiming He. 2021. Exploring Simple Siamese Representation Learning. In *CVPR*. Computer Vision Foundation / IEEE, 15750–15758.
- [11] Zhuo Chen, Jiaoyan Chen, Yuxia Geng, Jeff Z. Pan, Zonggang Yuan, and Huajun Chen. 2021. Zero-Shot Visual Question Answering Using Knowledge Graph. In *ISWC (Lecture Notes in Computer Science, Vol. 12922)*. Springer, 146–162.
- [12] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. Unifying Vision-and-Language Tasks via Text Generation. In *ICML (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 1931–1942.
- [13] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to algorithms*. MIT press.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*. Association for Computational Linguistics, 4171–4186.
- [15] Feng Gao, Qing Ping, Govind Thattai, Aishwarya N. Reganti, Ying Nian Wu, and Prem Natarajan. 2022. A Thousand Words Are Worth More Than a Picture: Natural Language-Centric Outside-Knowledge Visual Question Answering. *CoRR* abs/2201.05299 (2022).
- [16] Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making Pre-trained Language Models Better Few-shot Learners. In *ACL/IJCNLP (1)*. Association for Computational Linguistics, 3816–3830.
- [17] François Gardères, Maryam Ziaefard, Baptiste Abeloos, and Freddy Lécué. 2020. ConceptBert: Concept-Aware Representation for Visual Question Answering. In *EMNLP (Findings) (Findings of ACL, Vol. EMNLP 2020)*. Association for Computational Linguistics, 489–498.
- [18] Yuxia Geng, Jiaoyan Chen, Zhuo Chen, Jeff Z. Pan, Zhiquan Ye, Zonggang Yuan, Yantao Jia, and Huajun Chen. 2021. OntoZSL: Ontology-enhanced Zero-shot Learning. In *WWW*. ACM / IW3C2, 3325–3336.
- [19] Yuxia Geng, Jiaoyan Chen, Zhuo Chen, Jeff Z. Pan, Zonggang Yuan, and Huajun Chen. 2021. K-ZSL: Resources for Knowledge-driven Zero-shot Learning. *CoRR* abs/2106.15047 (2021).
- [20] Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander Hauptmann, Yonatan Bisk, and Jianfeng Gao. 2022. KAT: A Knowledge Augmented Transformer for Vision-and-Language. In *NAACL-HLT*. Association for Computational Linguistics, 956–968.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. IEEE Computer Society, 770–778.
- [22] Benjamin Heinzerling and Kentaro Inui. 2021. Language Models as Knowledge Bases: On Entity Representations, Storage Capacity, and Paraphrased Queries. In *EACL*. Association for Computational Linguistics, 1772–1791.
- [23] Gautier Izacard and Edouard Grave. 2021. Distilling Knowledge from Reader to Retriever for Question Answering. In *ICLR*.
- [24] Gautier Izacard and Edouard Grave. 2021. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In *EACL*. Association for Computational Linguistics, 874–880.
- [25] Aman Jain, Mayank Kothiyari, Vishwajeet Kumar, Preethi Jyothi, Ganesh Ramakrishnan, and Soumen Chakrabarti. 2021. Select, Substitute, Search: A New Benchmark for Knowledge-Augmented Visual Question Answering. In *SIGIR*. ACM, 2491–2498.
- [26] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.
- [27] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *EMNLP (1)*. Association for Computational Linguistics, 6769–6781.
- [28] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear Attention Networks. In *NeurIPS*. 1571–1581.
- [29] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*. Association for Computational Linguistics, 7871–7880.
- [30] Guohao Li, Xin Wang, and Wenwu Zhu. 2020. Boosting Visual Question Answering with Context-aware Knowledge Aggregation. In *ACM Multimedia*. ACM, 1227–1235.
- [31] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A Simple and Performant Baseline for Vision and Language. *CoRR* abs/1908.03557 (2019).
- [32] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. 2021. UNIMO: Towards Unified-Modal Understanding and Generation via Cross-Modal Contrastive Learning. In *ACL/IJCNLP (1)*. Association for Computational Linguistics, 2592–2607.
- [33] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020. Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. In *ECCV (30) (Lecture Notes in Computer Science, Vol. 12375)*. Springer, 121–137.
- [34] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR* abs/1907.11692 (2019).
- [35] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*. 13–23.
- [36] Man Luo, Yankai Zeng, Pratyay Banerjee, and Chitta Baral. 2021. Weakly-Supervised Visual-Retriever-Reader for Knowledge-based Question Answering. *EMNLP (2021)*.
- [37] Kenneth Marino, Xinlei Chen, Devi Parikh, Abhinav Gupta, and Marcus Rohrbach. 2021. KRISP: Integrating Implicit and Symbolic Knowledge for Open-Domain Knowledge-Based VQA. In *CVPR*. Computer Vision Foundation / IEEE, 14111–14121.
- [38] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge. In *CVPR*. Computer Vision Foundation / IEEE, 3195–3204.
- [39] Medhini Narasimhan, Svetlana Lazebnik, and Alexander G. Schwing. 2018. Out of the Box: Reasoning with Graph Convolution Nets for Factual Visual Question Answering. In *NeurIPS*. 2659–2670.
- [40] Fabio Petroni, Patrick S. H. Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. How Context Affects Language Models’ Factual Predictions. In *AKBC*.
- [41] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language Models as Knowledge Bases?. In *EMNLP/IJCNLP (1)*. Association for Computational Linguistics, 2463–2473.
- [42] Chen Qu, Hamed Zamani, Liu Yang, W. Bruce Croft, and Erik G. Learned-Miller. 2021. Passage Retrieval for Outside-Knowledge Visual Question Answering. In *SIGIR*. ACM, 1753–1757.
- [43] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2017. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 39, 6 (2017), 1137–1149.
- [44] Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How Much Knowledge Can You Pack Into the Parameters of a Language Model?. In *EMNLP (1)*. Association for Computational Linguistics, 5418–5426.
- [45] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.* 3, 4 (2009), 333–389.
- [46] Ander Salaberria, Gorka Azkune, Oier Lopez de Lacalle, Aitor Soroa, and Eneko Agirre. 2021. Image Captioning for Effective Use of Language Models in Knowledge-Based Visual Question Answering. *CoRR* abs/2109.08029 (2021).
- [47] Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling Relational Data with Graph Convolutional Networks. In *ESWC (Lecture Notes in Computer Science, Vol. 10843)*. Springer, 593–607.

- [48] Sanket Shah, Anand Mishra, Naganand Yadati, and Partha Pratim Talukdar. 2019. KVQA: Knowledge-Aware Visual Question Answering. In *AAAI*. AAAI Press, 8876–8884.
- [49] Sheng Shen, Chunyuan Li, Xiaowei Hu, Yujia Xie, Jianwei Yang, Pengchuan Zhang, Anna Rohrbach, Zhe Gan, Lijuan Wang, Lu Yuan, Ce Liu, Kurt Keutzer, Trevor Darrell, and Jianfeng Gao. 2022. K-LITE: Learning Transferable Visual Models with External Knowledge. *CoRR* abs/2204.09222 (2022).
- [50] Violetta Shevchenko, Damien Teney, Anthony R. Dick, and Anton van den Hengel. 2021. Reasoning over Vision and Language: Exploring the Benefits of Supplemental Knowledge. *CoRR* abs/2101.06013 (2021).
- [51] Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2020. Mmf: A multimodal framework for vision and language research.
- [52] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *AAAI*. AAAI Press, 4444–4451.
- [53] Hao Tan and Mohit Bansal. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *EMNLP/IJCNLP (1)*. Association for Computational Linguistics, 5099–5110.
- [54] Niket Tandon, Gerard de Melo, and Gerhard Weikum. 2017. WebChild 2.0 : Fine-Grained Commonsense Knowledge Distillation. In *ACL (System Demonstrations)*. Association for Computational Linguistics, 115–120.
- [55] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. 2018. Tips and Tricks for Visual Question Answering: Learnings From the 2017 Challenge. In *CVPR*. Computer Vision Foundation / IEEE Computer Society, 4223–4232.
- [56] Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. 2021. Multimodal Few-Shot Learning with Frozen Language Models, Vol. 34. 200–212.
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*, 5998–6008.
- [58] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge J. Belongie. 2016. COCO-Text: Dataset and Benchmark for Text Detection and Recognition in Natural Images. *CoRR* abs/1601.07140 (2016).
- [59] Peng Wang, Qi Wu, Chunhua Shen, Anthony R. Dick, and Anton van den Hengel. 2018. FVQA: Fact-Based Visual Question Answering. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 10 (2018), 2413–2427.
- [60] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. OFA: Unifying Architectures, Tasks, and Modalities Through a Simple Sequence-to-Sequence Learning Framework. In *ICML (Proceedings of Machine Learning Research, Vol. 162)*. PMLR, 23318–23340.
- [61] Jialin Wu, Jiasen Lu, Ashish Sabharwal, and Roozbeh Mottaghi. 2022. Multi-Modal Answer Validation for Knowledge-Based VQA. In *AAAI*. AAAI Press.
- [62] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2021. An Empirical Study of GPT-3 for Few-Shot Knowledge-Based VQA. *CoRR* abs/2109.05014 (2021).
- [63] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. ERNIE-ViL: Knowledge Enhanced Vision-Language Representations through Scene Graphs. In *AAAI*. AAAI Press, 3208–3216.
- [64] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. VinVL: Revisiting Visual Representations in Vision-Language Models. In *CVPR*. Computer Vision Foundation / IEEE, 5579–5588.
- [65] Zihao Zhu, Jing Yu, Yujing Wang, Yajing Sun, Yue Hu, and Qi Wu. 2020. Mucko: Multi-Layer Cross-Modal Knowledge Reasoning for Fact-based Visual Question Answering. In *IJCAI*. ijcai.org, 1097–1103.