

Multi-Modal Answer Validation for Knowledge-Based VQA

Jialin Wu¹, Jiasen Lu², Ashish Sabharwal², Roozbeh Mottaghi²

¹ The University of Texas at Austin

² Allen Institute for AI

jialinwu@utexas.edu, {jiasenl, ashishs, roozbehm}@allenai.org

Abstract

The problem of knowledge-based visual question answering involves answering questions that require external knowledge in addition to the content of the image. Such knowledge typically comes in various forms, including visual, textual, and commonsense knowledge. Using more knowledge sources increases the chance of retrieving more irrelevant or noisy facts, making it challenging to comprehend the facts and find the answer. To address this challenge, we propose Multi-modal Answer Validation using External knowledge (MAVEx), where the idea is to validate a set of promising answer candidates based on answer-specific knowledge retrieval. Instead of searching for the answer in a vast collection of often irrelevant facts as most existing approaches do, MAVEx aims to learn how to extract relevant knowledge from noisy sources, which knowledge source to trust for each answer candidate, and how to validate the candidate using that source. Our multi-modal setting is the first to leverage external visual knowledge (images searched using Google), in addition to textual knowledge in the form of Wikipedia sentences and ConceptNet concepts. Our experiments with OK-VQA, a challenging knowledge-based VQA dataset, demonstrate that MAVEx achieves new state-of-the-art results. Our code is available at <https://github.com/jialinwu17/MAVEX>

Introduction

Over the past few years, the domain of Visual Question Answering (VQA) has witnessed significant progress (Antol et al. 2015; Zhu et al. 2016; Hudson and Manning 2019). However, there is a recent trend towards knowledge-based VQA (Wang et al. 2017, 2018; Marino et al. 2019) which requires information beyond the content of the images. Besides visual recognition, the model needs to perform logical reasoning and incorporate external knowledge about the world to answer these challenging questions correctly. These knowledge facts can be obtained from various sources, such as image search engines, encyclopedia articles, and knowledge bases about common concepts and their relations.

Figure 1 illustrates a few visual questions and the knowledge from different external sources that helps answer them. Each question needs a different type of external knowledge. For example, to identify the movie that featured a man telling

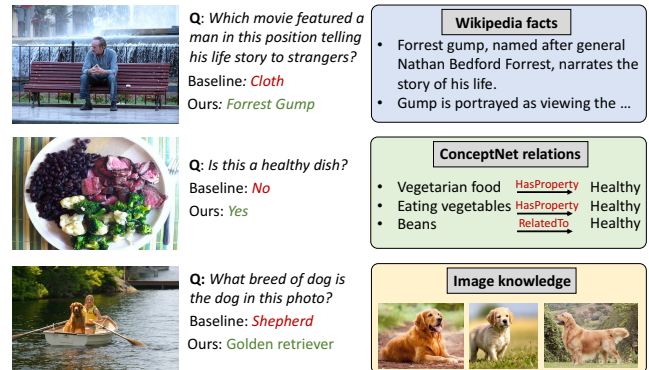


Figure 1: We address the problem of knowledge-based question answering. Retrieving relevant knowledge among diverse knowledge sources (visual knowledge, textual facts, concepts, etc.) is quite challenging. This paper aims to learn what knowledge source should be used for a particular question and how to validate a set of potential answer candidates using different knowledge sources.

his life story to strangers, we need to link the image content and question to some textual facts; Vegetarian food and eating vegetables are related to the concept of health; the retrieved images for a “golden retriever.” are visually similar to the dog in the question image. *The challenge is to retrieve and correctly incorporate such external knowledge effectively in an open domain question answering framework.*

We also witness a shift in knowledge-based VQA datasets from structured retrieved knowledge such as triplets and dense captions (Wang et al. 2017, 2018) to unstructured open knowledge (Marino et al. 2019). Most current knowledge-based VQA systems (Marino et al. 2019; Wang et al. 2018; Zhu et al. 2020; Marino et al. 2021) follow a two-stage framework, where a retriever first looks up knowledge relevant to the question and the image, and then a separate comprehension model predicts the answer.

However, knowledge retrieved directly for the question and image is often noisy and not helpful in predicting the correct answer. For example, as shown in Figure 2, the sentences retrieved using only the words in questions and objects in images (top) or a wrong answer (middle) are hardly helpful to

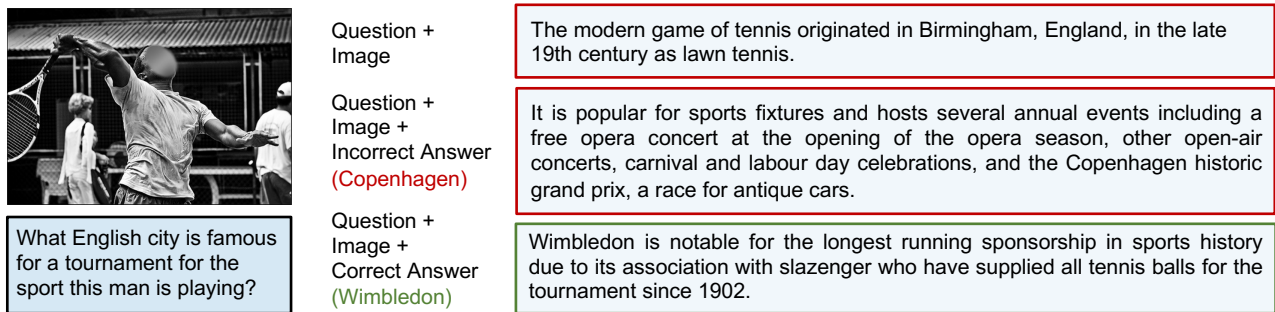


Figure 2: Examples of retrieved Wikipedia sentences using different sets of search words. The sentences retrieved using only the words in questions and objects in images (top) and the wrong answer (middle) are hardly helpful to answer the question. However, with the correct answer “Wimbledon” (bottom), the quality of the retrieved fact is significantly improved.

answer the question. This increases the burden on the answer predictor, leading to only marginal improvements from the use of retrieved knowledge (Marino et al. 2019). Interestingly, with the correct answer “Wimbledon” (bottom), the quality of the retrieved fact is significantly improved, making it suitable for answering the question. This observation motivates us to use retrieved knowledge for *answer validation* rather than for producing the answer.

To address this challenge, we propose a new system called MAVEx or **M**ulti-modal **A**nswer **V**alidation using **E**xternal knowledge. We use a three-stage framework. First, since state-of-the-art VQA models are surprisingly effective at generating a small set of promising answer candidates, we employ them for this purpose. Second, in the knowledge retrieval stage, we parse noun phrases from the question and answers, and generate queries for each phrase to query external resources. In order to better comprehend the retrieved knowledge, we embed it at multiple levels of granularity, from the basic query-level embedding to a noun-phrase-level embedding and finally to a question-level knowledge embedding. The goal of this multi-granular representation is to put more emphasis on queries that are important for each phrase, and on noun phrases that are critical for the question. Finally, in the validation stage, we predict how trustworthy each knowledge source is for the given question and answer candidate, and score the candidates accordingly.

We evaluate MAVEx on OK-VQA (Marino et al. 2019), the largest knowledge-based VQA dataset to date. Our approach achieves state-of-the-art results (score 40.3, ensemble score 41.4), demonstrating that answer-specific knowledge retrieval results in more informative supporting evidence and a more solid knowledge-based VQA system.

Our main contributions are: (a) We introduce a novel approach that uses answer candidates to guide knowledge retrieval among noisy facts; (b) We leverage multi-modal knowledge by retrieving from both visual and textual resources; and (c) We demonstrate that incorporating retrieved knowledge at multiple levels of granularity, based on the question and candidate answers, is an effective strategy.

Related Work

Visual Question Answering. Visual Question Answering (VQA) has made significant progress over the past few years

(Antol et al. 2015; Lu et al. 2016; Anderson et al. 2018; Kim, Jun, and Zhang 2018; Ben-Younes et al. 2017; Cadene et al. 2019). More recent VQA systems (Lu et al. 2019; Tan and Bansal 2019; Liu et al. 2019; Li et al. 2019; Yu et al. 2019; Li et al. 2020; Zhou et al. 2020; Chen et al. 2020; Lu et al. 2020) first extract visual features from a pre-trained object detector. Then they feed both visual and textual embeddings into a multi-modal transformer, which is pre-trained on auxiliary tasks using large-scale multi-modal datasets such as (Sharma et al. 2018; Hudson and Manning 2019; Kazemzadeh et al. 2014). These models achieve remarkable performance on the VQA (Antol et al. 2015) dataset; however, they can only reason based on the image content and do not have a mechanism to incorporate knowledge from external sources explicitly.

Knowledge-Based VQA Datasets. KB-VQA dataset (Wang et al. 2017) includes 2,402 questions generated by templates for 700 images. F-VQA (Wang et al. 2018) contains 5,826 questions, where each question-answer sample is annotated with a ground-truth fact triplet retrieved from the knowledge base. OK-VQA dataset (Marino et al. 2019) is a more recent open-domain dataset that covers a wide range of topics and includes 14,055 questions on 14,031 images. Our focus is on the OK-VQA dataset since it provides a larger scale dataset that requires open-domain knowledge. Due to the difficulty of collecting such datasets, knowledge-based VQA datasets are typically small compared to the traditional VQA datasets. The small scale of the datasets adds to the challenges of learning robust models.

Knowledge-Based VQA Models. Recent methods for knowledge-based VQA mainly follow two trends, template fitting, and learning-based approaches. For example, (Wang et al. 2017) fit the query to several pre-defined query templates and explicitly reason about the answer using the templates. The main limitation of the template fitting approaches is that the template is hand-designed, and it is hard to accommodate the rich knowledge required to answer the questions.

Learning-based approaches are proposed to fetch helpful facts and commonsense knowledge for better performance. Narasimhan and Schwing (2018) propose to retrieve relevant facts from a knowledge base. Wang et al. (2018) design a system to find the mappings from the question to a query triplet. GCN (Tompson et al. 2014) is applied on the fact graph in (Narasimhan, Lazebnik, and Schwing 2018)

where each node is a representation of an image-question-entity triplet. Li, Wang, and Zhu (2020) introduce a knowledge graph augmentation approach to retrieve context-aware knowledge subgraphs and then learn to aggregate the useful visual and question-relevant knowledge. Zhu et al. (2020) propose a modality-aware heterogeneous GCN capturing the most supporting evidence.

Most recent KB-VQA systems (Gardères et al. 2020; Marino et al. 2021; Shevchenko et al. 2021) utilize multi-modal transformers (Lu et al. 2019; Li et al. 2019) as base systems to incorporate the implicit knowledge they gathered from the large scale pre-training. In particular, (Gardères et al. 2020; Marino et al. 2021) combine the implicit knowledge with external symbolic knowledge and (Shevchenko et al. 2021) focus on injecting knowledge from the knowledge base into finetuning transformers. In contrast to these approaches, we formulate our problem as *answer validation*, where the idea is to learn to validate a set of potential answers using multi-modal noisy knowledge sources.

External Knowledge in Knowledge-Based VQA. To answer knowledge-based visual questions, most systems acquire external knowledge from various textual resources. For example, Wikipedia articles and ConceptNet concepts are frequently used as sources to provide factual and commonsense knowledge. There are two common approaches to utilize the knowledge. The first approach is parsing the knowledge in a symbolic format (Zhu et al. 2020; Narasimhan and Schwing 2018; Narasimhan, Lazebnik, and Schwing 2018; Li, Wang, and Zhu 2020; Marino et al. 2021) that usually consists of a collection of (subject, relation, object) triplets. Although each triplet presents explicit knowledge, many conditions and context are lost when producing these triplets. The missing information could prevent VQA systems from learning whether the knowledge is valid for the question and disambiguating the entities in the triplet. On the contrary, the second approach relies on simply using free-form knowledge (Wu et al. 2016; Marino et al. 2019; Qu et al. 2021) and use the raw text as input. While preserving most information, identifying helpful knowledge is quite challenging, especially in the multi-modal settings. (Wu et al. 2016; Marino et al. 2019) employ rule-based approaches that find knowledge relevant to one or a combination of objects detected in images or mentioned questions. (Qu et al. 2021) employ a dense retrieval approach that measures the relevance of the articles and the question-image pair in the feature space. However, being relevant to certain visual content cannot guarantee the helpfulness of the knowledge to predict the answer. To this end, we present answer-guided knowledge retrieval where the retrieved knowledge contains both the answer candidates and relevant visual content. Besides textual knowledge from Wikipedia and ConceptNet, we also explore external visual knowledge retrieved from the Google Images search engine.

Answer Validation. The idea of using answer candidates has been used in various question answering settings, including textual QA (Zhang, Vu, and Moschitti 2021), product QA (Zhang et al. 2020), commonsense VQA (Wu, Chen, and Mooney 2020), movie QA (Kim et al. 2019), allowing systems to perform a more in-depth examination of each answer candidate. We extend this idea to knowledge-based VQA,

where the system accesses answer-specific external knowledge to assess the correctness of each answer candidate.

The MAVEx Framework

We present the MAVEx framework, a three-stage scheme that first generates a set of promising answer candidates, retrieves knowledge guided by these answer candidates, and finally validates these answer candidates. Different from previous works (Gardères et al. 2020; Marino et al. 2021) that utilize textual knowledge, we propose to mine *multi-modal answer-specific* knowledge. In particular, we consider three knowledge sources: Wikipedia and ConceptNet for text, and Google for images. These provide factual, commonsense, and visual knowledge, respectively. For validation, we test each answer candidate using the retrieved multi-modal knowledge.

Answer Candidate Generation

In order to use answer candidates to inform knowledge retrieval, we use ViLBER-multi-task system (Lu et al. 2019), a state-of-the-art VQA model, to generate answer candidates. In particular, we finetune a ViLBER-multi-task model on the OK-VQA dataset that outputs a score for each answer collected from the training set. The highest-scoring answers are used as the candidates. Note that any VQA model or other approaches (for example, querying ontology knowledge bases) can be used for this purpose. However, as we will discuss in the experiments section, we found ViLBER to be particularly effective at generating a small set of promising candidates.

Answer Guided Knowledge Retrieval

Given a question q about an image I and a set of answer candidates A , we retrieve external knowledge supporting A in three main steps. Figure 3 shows the entire process for an example question and a candidate answer.

S1: Query Extraction. We first collect short phrases in q , each answer candidate in A , and concepts represented in I as a starting point for retrieving external information. This involves the following sub-steps:

Extract noun phrases from question and answers: We parse the question and the candidate answers using a constituency parser to obtain the parse tree. Then, we extract all the nouns on the leaves of the parse tree together with the words that describe the nouns and belong to one of the types from ADJP, ADVP, PP, SBAR, DT or JJ. We extract three kinds of noun phrases for modeling: (1) The target noun phrase that contains ‘wh’ or ‘how’ word (e.g. ‘*which movie*’), denoted by n_0^q . (2) Question noun phrases from the rest of the question (e.g., ‘*man in this position*’), denoted by $n_i^q, i \in \{1, \dots, N\}$. N is the number of noun phrases in the question. (3) Answer noun phrase for each answer a_i , denoted by n^{a_i} (e.g., ‘*Forrest Gump*’). These nouns help us link the mentioned objects to the images.

Link phrases to objects: As images usually contain plenty of question-irrelevant content, making the retrieval process hard, we propose narrowing the search to the objects referred to by the question or the answer candidates. In particular, we use a separate ViLBER-multi-task (Lu et al. 2020) model

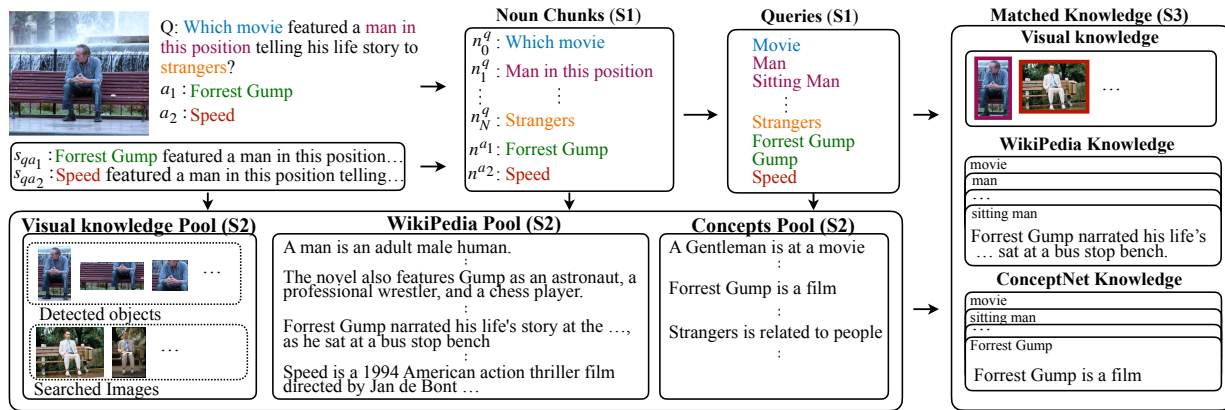


Figure 3: An example of the retrieval process for one question-answer pair. The numbers in parentheses denote the step number in Section . The noun phrase, its generated queries, and the matched visual knowledge are marked in the same color.

as the object linker, where it takes as inputs a set of detected objects and a noun phrase from the question, and outputs a linking score for each detected object to indicate how likely the noun phrase refers to the object. We approve the linking when the score is higher than 0.5 and extract the linked objects.

Generate search query set: We further generate a set of search queries to search the external knowledge base. For each noun phrase, we first extract the head of a phrase by finding the innermost NP from the dependency tree. Then, we obtain the visual attributes of the head of the noun phrase by using a pre-trained object-with-attribute detector (Anderson et al. 2018) for the corresponding linked objects. For example, the visually grounded queries for ‘man in this position’ are ‘man’ and ‘sitting man’ where *sitting* is inferred from visual attributes. We denote the set of queries as $r_i^n, i \in \{1, \dots, K\}$, where n is the corresponding noun phrase, K is the maximum number of queries per noun phrase.

S2: Answer Guided Knowledge Pool Construction. We now use the visually grounded queries from step S1 to construct knowledge pool as follows:

Conversion to a natural language statement: In order to use the answer candidate a to inform the retrieval step, we convert q and $a \in A$ into a natural language statement s_{qa} using a T5 model (Raffel et al. 2020) finetuned on the QA-NLI dataset (Demszky, Guu, and Liang 2018). Such conversion is effective as statements occur much more frequently than questions in textual knowledge sources (Khot, Sabharwal, and Clark 2017). These statements are later used to compute the relevance of the retrieved facts as described below.

Retrieval of textual facts and concepts: We search each query in the query set generated from the last sub-step in S1 in Wikipedia and ConceptNet. We compute the BERTScore (Zhang* et al. 2020) between each sentence from the retrieved article and each statement s_{qa} . For each statement, the top-15 sentences (according to the BERTScore) from each retrieved article are pushed to the sentence pool. Then, we decontextualize (Choi et al. 2021) each sentence in the Wikipedia pool for better knowledge quality.

Retrieval of visual knowledge: Pure textual knowledge is often insufficient due to two main reasons: (1) textual knowledge might be too generic and not specific to the question image, (2) it might be hard to describe some concepts using text, and an image might be more informative (e.g., the third question in Figure 1). Hence, visual knowledge can complement textual information, further enriching the external knowledge space. We consider both internal and external visual knowledge. For the given image, we utilize a MaskRCNN (He et al. 2017) object detector to detect common objects as internal knowledge. We use Google image search to retrieve the top-5 images using the statement s_{qa} as the query for each answer candidate a as the external visual knowledge.

S3: Matching Knowledge Pool to Queries. Instead of simply using each query’s top retrieved sentences as the query’s knowledge, we propose matching the sentences from the entire pool to each query. The intuition is that most queries cannot directly retrieve helpful facts; however, they can help retrieve important aspects that should be contained in the external knowledge.

Matching Textual Knowledge: For each query, the sentences from both Wikipedia and ConceptNet pool with a mean recall greater than 0.6 are considered the retrieved results. Mean recall is the average cosine similarity between the Glove embedding of the words in the query and their most similar word in the sentence. To ensure knowledge relevance, we remove sentences that are matched to only a single query. For each query r_i^n , according to the maximum BERTScore between the sentence and all of the statements S_q , we extract at most m sentences from both Wikipedia and ConceptNet pools, denoted by $W(r_i^n)$ and $C(r_i^n)$.

Matching Visual Knowledge: For each noun phrase in the question, we can directly use the results from the object linker defined in S1. Specifically, we find the top-3 referenced objects in the image for each question noun phrase, denoted $M(n)$.

For each answer noun phrase n^{a_i} , we use Google image search to retrieve the top-5 images, denoted $M(n^{a_i})$.

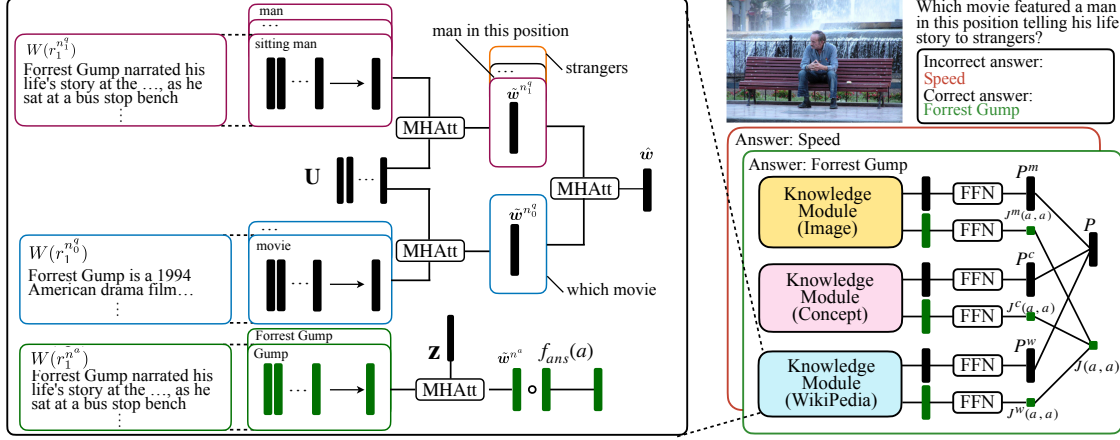


Figure 4: Model overview for validating two candidate answers. We explore three sources of external knowledge, *i.e.* Wikipedia, ConceptNet, and Google Images presented by the three parallel knowledge embedding modules. The black blocks denote features shared by all answer candidates, and the green blocks denote answer-specific features. Different colors denotes the features for different noun phrases and their queries.

Answer Candidate Validation

The answer validation module takes as input an answer candidate a_i and the supporting knowledge, and outputs a scalar score indicating how well the knowledge supports a_i . As we will discuss, in order to better aggregate the knowledge, we first compute the knowledge embedding for each query. Then, we compute an embedding for each noun phrase that aggregates the embedding for the queries generated from the noun phrase¹. Finally, the embedding for the entire question aggregates the embedding computed for all noun phrases.

We build MAVEx on top of the ViLBERT system. Given a question q and an image I , ViLBERT provides textual features $U \in \mathbb{R}^{|q| \times d}$, visual features $V \in \mathbb{R}^{|V| \times d}$ from the last layer, where $|q|$ is the number of tokens in q , d is the feature dimension, $|V|$ is the number of objects in the image plus one for the representation for the entire image, and a joint visual-textual representation $z \in \mathbb{R}^d$. For each sentence in the retrieved textual knowledge $W(r_i^n)$ and $C(r_i^n)$, we use TinyBERT (T-BERT) model (Turc et al. 2019) to extract the corresponding features. We further average the sentence features for each query r_i^n , resulting w_i^n and c_i^n .

For each image in the retrieved visual knowledge $M(n_a)$, we use MaskRCNN (He et al. 2017) to extract object features. Then, we average the object features of visual detection results as the image features and denote them as m_i^n . Note that we directly use the object features for the linked objects. Figure 4 shows the overview of the model.

Multi-Granular Knowledge Embedding Module. In order to better aggregate the retrieved knowledge, we employ a multi-granular knowledge embedding module that learns to recognize the critical queries for each noun phrase, and then the critical noun phrases for answering the question.

Note that our knowledge embedding module is identical for each knowledge source but with different learnable parameters. We only show knowledge from Wikipedia for brevity. Given the knowledge embeddings w_i^n for each query r_i^n in the question, we compute the knowledge embedding \tilde{w}^n for each noun phrase in question as follows:

$$\tilde{w}^n = \text{MHAtt}(u^n, \{w_i^n\}_{i \in \{1, \dots, K\}}, \{w_i^n\}_{i \in \{1, \dots, K\}}), \quad (1)$$

where $\text{MHAtt}(\text{query}, \text{key}, \text{value})$ is the multi-head attention operator. u^n is the attentive pooled (Lee et al. 2017) ViLBERT features according to the span $\{s, e\}$ of the phrase n . We use u^n as the query in MHAtt module to aggregate the retrieved knowledge, where the corresponding knowledge embeddings $\{w_i^n\}_{i \in \{1, \dots, K\}}$ serve as key and value.

Similarly, for each answer a^2 , we compute the knowledge embedding w^a using a MHAtt module over the knowledge features $w_i^{n_a}$ as follows:

$$w^a = \text{MHAtt}(z, \{w_i^{n_a}\}_{i \in \{1, \dots, K\}}, \{w_i^{n_a}\}_{i \in \{1, \dots, K\}}), \quad (2)$$

where the joint visual-textual embeddings z from ViLBERT system is used as the keys.

Then, another MHAtt module is used to gather the knowledge from each noun phrase $n \in \{n_1^q, \dots, n_N^q\}$. Specifically, given the knowledge embedding for each noun phrase, the knowledge embeddings w is computed as follows:

$$\hat{w} = \text{MHAtt}(\tilde{w}^{n_q}, \{\tilde{w}^{n_i^q}\}_{i \in \{1, \dots, N\}}, \{\tilde{w}^{n_i^q}\}_{i \in \{1, \dots, N\}}) \quad (3)$$

Answer Prediction and Validation Module Given the knowledge embedding $k \in \{\hat{w}, \hat{c}, \hat{m}\}$ from each one of the three knowledge sources, MAVEx predicts the answers' probability as $P^k = \text{FFN}(k + z)$, where FFN denotes a feed-forward layer. The final prediction P is the answer that has

¹Recall that our queries r_i^n are created based on noun phrase n .

²For simplicity we omit the subscript index of the answer in this section when there is only one answer involved in the current step.

the maximum confidence over the three knowledge sources for each answer, *i.e.* $P = \max_k \{P^k\}$.

The validation module takes as inputs the answer candidate a and the knowledge features $k^{a'} \in \{w^{a'}, c^{a'}, m^{a'}\}$ from the three sources to learn how well the knowledge supports the answer candidate. We first embed the answer candidate using the summation of the BERT features of the corresponding statement and the glove features of the answer itself, *i.e.* $f_{ans}(a) = (\text{BERT}(s_{qa}) + \text{glove}(a))$. Then, the validation score $J(a, a')$ for answer candidate a using the knowledge retrieved for a' (a different candidate) is computed as $J^k(a, a') = \text{FFN}(f_{ans}(a) \circ k^{a'})$, where the \circ means element-wise multiplication. The final validation score is the maximum validation confidence over the three knowledge sources, *i.e.* $J(a, a') = \max_k \{J^k(a, a')\}$.

Consistency Criteria. The intuition behind our consistency criteria is that for the correct answer a , the knowledge retrieved for a from the most confident source (the one with the highest supportiveness score J for a) should support a more than it supports other answer candidates, and it should also support a more than knowledge retrieved for other answer candidates. Specifically, we approve the answer validation score $J(a, a)$ only if it is higher than the scores computed using this knowledge for all other answers as well as the score for a when using knowledge retrieved for other answers. We also eliminate the case where the top-1 prediction from P is not in the answer candidate set. Mathematically, the consistency criteria checks that $J(a, a) > J(a', a)$ and $J(a, a) > J(a, a')$ for all $a' \neq a$. If the above condition is not met, we output the answer with the maximum VQA prediction score $P(a)$; otherwise, we output the answer with the maximum VQA-weighted validation score $J(a, a)P(a)$.

Training and Implementation Details

Implementation. We implemented our approach on top of ViLBERT-multi-task (Lu et al. 2019), which utilizes a Mask-RCNN head (He et al. 2017) in conjunction with a ResNet-152 base network (He et al. 2016) as the object detection module. Convolutional features for at most 100 objects are then extracted for each image as the visual features, *i.e.* a 2,048-dimensional vector for each object. We used the constituent parser from AllenNLP to extract the nouns phrases in the question. For linking the mentioned objects, we adopt a separate ViLBERT-multi-task system. For converting the question and answer, we finetuned a T5-base model (Raffel et al. 2020) on the QA-NLI dataset (Demszky, Guu, and Liang 2018) for 4 epochs. We detected 100 objects using Mask-RCNN to encode the retrieved Google images. For question embedding, following (Devlin et al. 2019), we use a BERT tokenizer on the question and use the first 23 tokens as the question tokens. We encode at most 4 sentences per query, 3 queries per noun phrase. The number of hidden units in the multi-head attention modules is set to 512. We use Pytorch 1.4 on a single TITAN V GPU with 12M memory for each run, and it generally costs 22 hours to train a single model.

Training. The OK-VQA test images are a subset of COCO

validation images that are used to pre-train most transformer-based vision and language models (Lu et al. 2019; Tan and Bansal 2019; Li et al. 2019). Although the test questions never appear in the pre-training process, other questions on the test images may help the system understand the image better, leading to higher performance. Besides, there is also data contamination from extra object annotations from Visual Genome (VG) dataset, which also contains some OK-VQA test images. As the VG dataset is used to pre-train the object detector, those test images can access the ground truth object annotations. Therefore, we carefully remove all OK-VQA test images from the pre-training and re-train the ViLBERT-multi-task model and the object detector from scratch using the default configurations.

We finetune the ViLBERT-multi-task model on OK-VQA using the default configuration for 150 epochs for answer candidate generation. Binary cross-entropy loss and VQA soft score are employed to optimize the system. OK-VQA provides five annotations for each question. Soft scores are 0, 0.6, and 1 corresponding to 0, 1, more than 1 matching answer annotations. We use the finetuned model to extract the top 5 answers for each question in the training and test set. We follow the default settings of ViLBERT and apply the BertAdam optimizer (Devlin et al. 2019) with a linear warmup learning rate.

For the training of the answer validation module, we optimize the validation score $J(a, a')$ using the loss in Eq. 4 for the three knowledge sources, where $s(a)$ denotes the VQA soft scores for answer a . We also add the standard VQA losses on the predictions from the three external sources. We train the system for 75 epochs using a learning rate of $2e-5$ for the ViLBERT parameters and $5e-5$ for the additional parameters introduced in the validation module. We freeze the first 10 layers of the ViLBERT base network. We use \mathcal{L}_{bce} to denote binary cross-entropy loss.

$$\begin{aligned} \mathcal{L}_{\text{MAVEX}} = & \mathcal{L}_{bce} \left(\max_{\substack{a \\ \text{s.t. } a \neq a'} J(a, a'), \mathbf{0} \right) \\ & + \mathcal{L}_{bce} \left(\max_{\substack{a'} \\ \text{s.t. } a \neq a'} J(a, a'), \mathbf{0} \right) \\ & + \mathcal{L}_{bce} \left(J(a, a), s(a) \right) \end{aligned} \quad (4)$$

Experiments

We evaluate our framework on the OK-VQA dataset. We first briefly describe the dataset and then present our results, comparing with the current state-of-the-art systems.

OK-VQA dataset (Marino et al. 2019) is the largest knowledge-based VQA dataset at present. The questions are crowdsourced from Amazon Mechanical Turkers, leading to two main advantages: (1) the questions indeed require outside knowledge beyond images; (2) there are no existing knowledge bases that cover all the questions, thus requiring systems to explore open-domain resources. The dataset contains 14,031 images and 14,055 questions covering a variety of knowledge categories. The metric is the VQA soft score.

Method	Knowledge Resources	Performance
ArticleNet (AN) (Marino et al. 2019)	Wikipedia	5.3
Q-only (Marino et al. 2019)	—	14.9
MLP (Marino et al. 2019)	—	20.7
BAN (Kim, Jun, and Zhang 2018)	—	25.2
+ AN (Marino et al. 2019)	Wikipedia	25.6
+ KG-AUG (Li, Wang, and Zhu 2020)	Wikipedia + ConceptNet	26.7
MUTAN (Ben-Younes et al. 2017)	—	26.4
+ AN (Marino et al. 2019)	Wikipedia	27.8
Mucko (Zhu et al. 2020)	Dense Caption	29.2
ConceptBert (Gardères et al. 2020)	ConceptNet	33.7
KRISP (Marino et al. 2021)	Wikipedia + ConceptNet	38.9*
RVL [†] (Shevchenko et al. 2021)	Wikipedia + ConceptNet	39.0 [†]
MAVEx (ours)	Wikipedia + ConceptNet	39.45*
MAVEx (ours)	Wikipedia + ConceptNet + Google Images	40.28*
MAVEx (ours) (Ensemble 3)	Wikipedia + ConceptNet + Google Images	41.37*

Table 1: MAVEx outperforms current state-of-the-art approaches on OK-VQA. The middle column lists the external knowledge sources, if any, used in each system. † indicates that the system uses a pretrained model contaminated by OK-VQA test images. * indicates that the results have been reported on version 1.1 of the dataset.

For the experiments of this paper, we used version 1.1 of the dataset.

Intrinsic Evaluation

We begin with an intrinsic evaluation of MAVEx, assessing the quality of the answer candidate generation step.

Answer Candidate Accuracy. Our answer candidate generation module, based on the finetuned ViLBERT-multi-task model, outputs its top-5 answers as the candidates. We found that the best answer in this small candidate set achieves a VQA soft score of 59.7 on the test set, substantially higher than other state-of-the-art systems without data contamination. We also evaluate the score achieved by slightly larger candidate sets, consisting of the top 6, 8, and 10 candidates. These achieve VQA soft scores of 62.1, 65.1, and 67.1, respectively. Since our answer validation framework needs to retrieve and encode answer-specific knowledge, we use only top-5 answer candidates as a reasonable trade-off between efficiency, answer coverage, and overall accuracy. Note that our method cannot produce answers not in the candidate set.

Main Results

Table 1 shows that MAVEx consistently outperforms prior approaches by a clear margin. For example, MAVEx single model outperforms recent state-of-the-art models, KRISP (Marino et al. 2021), and ConceptBert (Gardères et al. 2020) by 1.4, 6.6 points, respectively. An ensemble of three MAVEx models with different initializations provides 2.47 points improvement compared to KRISP. All of our results, except for the ensemble model, are averaged across 3 different initialization seeds. The standard deviation is 0.21 for the single model computed from the three runs.

Ablation Studies

Knowledge Sources. We report the performance of using the different combinations of knowledge sources in Table 2. We see that the three sources (Wikipedia, ConceptNet, and Images) improve the performance by 3.4, 3.3, and 3.1, respectively, compared to the base ViLBERT system. This indicates the effectiveness and value of all three sources. The decontextualization technique (Choi et al. 2021) improves the performance compared to using only the Wikipedia source by 0.4%. The decontextualization partially helps address the co-reference issue since the retrieved sentences provide more information from their paragraph. Combining the three sources achieves a net performance gain of 5% over the ViLBERT baseline, supporting the intuition that the three sources together provide complementary pieces of knowledge.

We show some qualitative examples in Figure 5, where the VQA model (ViLBERT) is wrong but provides good answer candidates. Our MAVEx gathers the external knowledge from the three sources and predicts the correct answers.

Knowledge Source	System Score
ViLBERT	35.20
Wikipedia (<i>w/o</i> decontextualization)	38.21
Wikipedia	38.63
ConceptNet	38.56
Images	38.30
Wikipedia + ConceptNet	39.45
ConceptNet + Images	39.60
Wikipedia + Images	39.37
Wikipedia + ConceptNet + Images	40.28
Wikipedia + ConceptNet + Images (Oracle)	47.76

Table 2: Ablation study using different combinations of knowledge sources.

<p>What is the complimentary color to the frisbee</p> 	<p>Blue (MAVEx)</p> <p>Because orange and blue are complementary colors, life rafts and life vests are traditionally orange, to provide the highest contrast and visibility when seen from ships or aircraft over the ocean</p>	<p>Red (VQA)</p> <p>In the indian subcontinent, red is the traditional color of bridal dresses, and is frequently represented in the media as a symbolic color for married women</p>
<p>Who is the official in this sport</p> 	<p>Umpire (MAVEx)</p> <p>Umpire, related to, referee Umpire, synonym, referee Umpire, related to, baseball official</p>	<p>Pitcher (VQA)</p> 
<p>What kind of lamp is this</p> 	<p>Chandelier(MAVEx)</p> 	<p>Lava(VQA)</p> 

Figure 5: Examples where the VQA model is wrong but MAVEx with the three external knowledge sources answers correctly. The correct answer is in the green box and the incorrect answer is shown in the red box. The grey box shows the question. Sample retrieved knowledge content is shown in the boxes under the predicted answers.

Knowledge Embedding Granularity. We ablate different levels of granularity used in the MAVEx system by comparing to two baseline systems, where we replace the noun-phrase level or the question level multi-head attention module with an average pooling operation. When replacing the noun-phrase level MHAtt modules (*i.e.* three MHAtt modules corresponding to merging queries’ features for question noun phrases, question target phrase and the answer phrases), the performance reduces to 39.77. When replacing the question level MHAtt module with an average pooling operation for the question, the performance reduces to 39.60.

Answer Validation Step. We consider a MAVEx baseline model that uses the retrieved knowledge (\tilde{w} , \tilde{c} , \tilde{m}) as additional inputs but without answer validation. This model achieves an overall score of 39.2, 4% higher than the ViLBERT base model and 1.1% lower than the full model, indicating that using answer-guided retrieved knowledge is helpful, and answer validation further improves performance.

Oracle Experiments

Oracle Source Selector. We report an oracle score obtained by manually choosing the best verification score $J^k(a, a)$ from the three sources $k \in \{\tilde{w}, \tilde{c}, \tilde{m}\}$ to weigh the prediction P . As a result, our answer validation framework achieves an oracle score of 47.76 as reported in Table 2. This indicates that the three knowledge sources provide complementary features, leaving further potential to improve the system.




<p>How is this form of transportation powered?</p> 	<p>What other surfaces might this sport be played on?</p> 	<p>What purpose is there to having all of these clocks on the wall?</p> 
<p>Answer candidates: electricity, diesel, gas, gasoline, engine</p> <p>Sample fact: Some locomotives use two-stroke diesel engines</p> <p>Predicted answer: diesel</p> <p>GT answers: electricity</p>	<p>Answer candidates: clay, dirt, concrete, tennis court, pavement</p> <p>Sample fact: Clay has context of tennis Clay is related to surface</p> <p>Predicted answer: tennis court</p> <p>GT answers: clay</p>	<p>Answer candidates: time zone, tell time, time, storage, to tell time</p> <p>Sample fact: The primary purpose of a clock is to display the time</p> <p>Predicted answer: tell time</p> <p>GT answers: time zone</p>

Figure 6: Some typical failure cases of our model have been shown. In these examples, the model falsely focuses on the retrieved fact (left), visual content (middle), or does not generate proper search word (right).

Failures Cases Analysis

Figure 6 shows some common types of failure examples. In the left example, the model over-relies on the retrieved fact “some locomotives use diesel engines” and ignores the key visual clue in the image (the wires above the train). In the middle example, the model relies on the visual content “tennis court” and does not use the retrieved knowledge. In the example shown on the right, the model fails to realize that the key clue is the *difference in displayed time on the clocks*.

Conclusion

We presented MAVEx, a novel approach for knowledge based visual question answering. The goal is to retrieve answer-specific textual and visual knowledge from different knowledge sources and learn what sources contain the most relevant information. Searching through the vast amount of retrieved knowledge, which is often quite noisy, is challenging. Hence, we formulate the problem as answer validation, where the goal is to learn to verify the validity of a set of candidate answers according to the retrieved knowledge. More specifically, an answer candidate validation module predicts the degree of support provided by the knowledge retrieved for each candidate, and decides which sources to trust for each candidate answer. MAVEx demonstrates the clear advantages of answer-guided knowledge retrieval, achieving the state-of-the-art performance on the OK-VQA dataset.

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-Up and Top-Down Attention for Image Captioning and VQA. In *CVPR*.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. VQA: Visual Question Answering. In *ICCV*.
- Ben-Younes, H.; Cadene, R.; Cord, M.; and Thome, N. 2017. MUTAN: Multimodal Tucker Fusion for Visual Question Answering. In *ICCV*.
- Cadene, R.; Ben-Younes, H.; Cord, M.; and Thome, N. 2019. Murel: Multimodal relational reasoning for visual question answering. In *CVPR*.

- Chen, Y.-C.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020. Uniter: Universal image-text representation learning. In *ECCV*.
- Choi, E.; Palomaki, J.; Lamm, M.; Kwiatkowski, T.; Das, D.; and Collins, M. 2021. Decontextualization: Making Sentences Stand-Alone. *TACL*.
- Demszky, D.; Guu, K.; and Liang, P. 2018. Transforming question answering datasets into natural language inference datasets. *arXiv preprint arXiv:1809.02922*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Gardères, F.; Ziaeeafard, M.; Abeloos, B.; and Lecue, F. 2020. ConceptBert: Concept-Aware Representation for Visual Question Answering. In *EMNLP*.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: a new dataset for compositional question answering over real-world images. In *CVPR*.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. ReferItGame: Referring to Objects in Photographs of Natural Scenes. In *EMNLP*.
- Khot, T.; Sabharwal, A.; and Clark, P. 2017. Answering Complex Questions Using Open Information Extraction. In *ACL*.
- Kim, J.; Ma, M.; Kim, K.; Kim, S.; and Yoo, C. D. 2019. Progressive Attention Memory Network for Movie Story Question Answering. In *CVPR*.
- Kim, J.-H.; Jun, J.; and Zhang, B.-T. 2018. Bilinear Attention Networks. In *NeurIPS*.
- Lee, K.; He, L.; Lewis, M.; and Zettlemoyer, L. 2017. End-to-end Neural Coreference Resolution. In *EMNLP*.
- Li, G.; Duan, N.; Fang, Y.; Gong, M.; Jiang, D.; and Zhou, M. 2020. Unicoder-VL: A Universal Encoder for Vision and Language by Cross-Modal Pre-Training. In *AAAI*.
- Li, G.; Wang, X.; and Zhu, W. 2020. Boosting Visual Question Answering with Context-aware Knowledge Aggregation. In *ACM Conference on Multimedia*.
- Li, L. H.; Yatskar, M.; Yin, D.; Hsieh, C.-J.; and Chang, K.-W. 2019. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Liu, B.; Huang, Z.; Zeng, Z.; Chen, Z.; and Fu, J. 2019. Learning Rich Image Region Representation for Visual Question Answering. *arXiv preprint arXiv:1910.13077*.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*.
- Lu, J.; Goswami, V.; Rohrbach, M.; Parikh, D.; and Lee, S. 2020. 12-in-1: Multi-Task Vision and Language Representation Learning. In *CVPR*.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016. Hierarchical Question-Image Co-attention for Visual Question Answering. In *NeurIPS*.
- Marino, K.; Chen, X.; Parikh, D.; Gupta, A.; and Rohrbach, M. 2021. KRISP: Integrating Implicit and Symbolic Knowledge for Open-Domain Knowledge-Based VQA. In *CVPR*.
- Marino, K.; Rastegari, M.; Farhadi, A.; and Mottaghi, R. 2019. OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge. In *CVPR*.
- Narasimhan, M.; Lazebnik, S.; and Schwing, A. 2018. Out-of-The-Box: Reasoning with Graph Convolution Nets for Factual Visual Question Answering. In *NeurIPS*.
- Narasimhan, M.; and Schwing, A. G. 2018. Straight to the facts: Learning knowledge base retrieval for factual visual question answering. In *ECCV*.
- Qu, C.; Zamani, H.; Yang, L.; Croft, W. B.; and Learned-Miller, E. 2021. Passage Retrieval for Outside-Knowledge Visual Question Answering. In *ACM SIGIR*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *JMLR*.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*.
- Shevchenko, V.; Teney, D.; Dick, A.; and van den Hengel, A. 2021. Reasoning over Vision and Language: Exploring the Benefits of Supplemental Knowledge. In *Proceedings of the Third Workshop on Beyond Vision and LAnguage: Integrating Real-world Knowledge (LANTERN)*.
- Tan, H.; and Bansal, M. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *EMNLP*.
- Tompson, J. J.; Jain, A.; LeCun, Y.; and Bregler, C. 2014. Joint training of a convolutional network and a graphical model for human pose estimation. In *NeurIPS*.
- Turc, I.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*.
- Wang, P.; Wu, Q.; Shen, C.; Dick, A.; and van den Hengel, A. 2018. Fvqa: Fact-based visual question answering. *TPAMI*.
- Wang, P.; Wu, Q.; Shen, C.; Hengel, A. v. d.; and Dick, A. 2017. Explicit knowledge-based reasoning for visual question answering. In *IJCAI*.
- Wu, J.; Chen, L.; and Mooney, R. J. 2020. Improving VQA and its Explanations by Comparing Competing Explanations. *arXiv preprint arXiv:2006.15631*.
- Wu, Q.; Wang, P.; Shen, C.; Dick, A.; and van den Hengel, A. 2016. Ask Me Anything: Free-Form Visual Question Answering Based on Knowledge From External Sources. In *CVPR*.
- Yu, Z.; Yu, J.; Cui, Y.; Tao, D.; and Tian, Q. 2019. Deep modular co-attention networks for visual question answering. In *CVPR*.
- Zhang*, T.; Kishore*, V.; Wu*, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *ICLR*.
- Zhang, W.; Deng, Y.; Ma, J.; and Lam, W. 2020. AnswerFact: Fact Checking in Product Question Answering. In *EMNLP*.

Zhang, Z.; Vu, T.; and Moschitti, A. 2021. Joint Models for Answer Verification in Question Answering Systems. In *ACL-IJCNLP*.

Zhou, L.; Palangi, H.; Zhang, L.; Hu, H.; Corso, J. J.; and Gao, J. 2020. Unified Vision-Language Pre-Training for Image Captioning and VQA. In *AAAI*.

Zhu, Y.; Groth, O.; Bernstein, M.; and Fei-Fei, L. 2016. Visual7w: Grounded Question Answering in Images. In *CVPR*.

Zhu, Z.; Yu, J.; Wang, Y.; Sun, Y.; Hu, Y.; and Wu, Q. 2020. Mucko: Multi-Layer Cross-Modal Knowledge Reasoning for Fact-based Visual Question Answering. In *IJCAI*.