

VLC-BERT: Visual Question Answering with Contextualized Commonsense Knowledge

Sahithya Ravi^{1,2*} Aditya Chinchure^{1,2*} Leonid Sigal^{1,2} Renjie Liao¹ Vered Shwartz^{1,2}

¹ University of British Columbia ² Vector Institute for AI

{sahiravi, aditya10, lsigal, vshwartz}@cs.ubc.ca, rjliao@ece.ubc.ca

Abstract

There has been a growing interest in solving Visual Question Answering (VQA) tasks that require the model to reason beyond the content present in the image. In this work, we focus on questions that require commonsense reasoning. In contrast to previous methods which inject knowledge from static knowledge bases, we investigate the incorporation of contextualized knowledge using Commonsense Transformer (COMET), an existing knowledge model trained on human-curated knowledge bases. We propose a method to generate, select, and encode external commonsense knowledge alongside visual and textual cues in a new pre-trained Vision-Language-Commonsense transformer model, VLC-BERT. Through our evaluation on the knowledge-intensive OK-VQA and A-OKVQA datasets, we show that VLC-BERT is capable of outperforming existing models that utilize static knowledge bases. Furthermore, through a detailed analysis, we explain which questions benefit, and which don't, from contextualized commonsense knowledge from COMET. Code: <https://github.com/aditya10/VLC-BERT>

1. Introduction

Recent progress in multimodal vision-language learning has been fueled by large-scale annotated datasets for Visual Question Answering (VQA) [1, 6, 12, 37, 49], in which models are presented with questions about an image. To answer questions correctly, models are required to perform scene understanding and learn meaningful connections between the two modalities. In recent years, transformer-based vision and language (VL) models [8, 21, 44], pre-trained on large-scale multimodal corpora, have reached impressive accuracies on standard VQA datasets.

VQA often necessitates not only visual comprehension of the scene depicted by the image (e.g., “A plate with meat, potatoes and bread”) but also making inferences about plau-



Figure 1: OK-VQA [29]: Where might one buy this?

sible stories behind the image (e.g., “The plate is likely found at a restaurant”). Humans make such inferences based on prior experience and commonsense knowledge (e.g., “This is likely a lunch or dinner at a restaurant, people may be enjoying themselves...”). Most existing methods rely on world knowledge implicitly encoded by language models, which often lacks in both accuracy and coverage [32]. This is primarily due to the fact that commonsense knowledge is extremely broad, and frequently assumed. Commonsense knowledge learned from text suffers from reporting bias [11]: over-representation of exceptional facts (e.g., “people die in accidents”) in text corpora, at the expense of rarely discussed trivial facts known to everyone (e.g., “people eat”).

Several visual question answering benchmarks were proposed, in which the questions require either factual [29, 45] or commonsense knowledge [36, 49] beyond the visual scene comprehension. This prompted the development of neurosymbolic methods combining transformer-based representations with knowledge bases (KBs) [9, 28, 47]. However, retrieving relevant facts directly from a KB is challenging due to lack of coverage, and because KB facts are only appropriate in certain contexts.

In this work, we propose VLC-BERT (Vision-Language-Commonsense BERT), a model designed to incorporate contextualized commonsense knowledge into a Vision-Language transformer built on VL-BERT [41]. As an alternative to the retrieval paradigm often used in knowledge-based VQA, our model generates contextualized commonsense inferences on the question phrase combined with image object tags using COMET [2, 15], a language model

*Denotes equal contribution

trained on commonsense knowledge graphs. We augment sentence transformers [31] to rank, filter and embed the commonsense inferences. We incorporate the filtered inferences into VLC-BERT using an attention-driven fusion mechanism that learns to focus on the most important inferences for each question. Commonsense knowledge may not be necessary for answering every question, as some questions are either purely visual, factual, or straight-forward. To eliminate injecting noisy knowledge in such cases, we employ weak supervision to help us discriminate between situations when commonsense knowledge may or may not be valuable.

Our evaluations on the challenging OK-VQA [29] and A-OKVQA [36] datasets confirm that leveraging commonsense is consistently useful for knowledge-intensive visual question answering tasks. We analyze the successful predictions and show how the commonsense inferences help answering difficult questions.

2. Related Work

2.1. Vision-Language Transformer Models

Pre-trained Vision-Language models based on BERT [8] have shown impressive performances on downstream multimodal tasks such as Visual Question Answering. ViL-BERT [25] and LXMERT [42] use a two-stream architecture to first encode language and vision modalities independently, and then apply a cross-modality encoder to align textual and visual tokens. VL-BERT [41], OSCAR [22] and OSCAR+ [50] use a single-stream architecture to directly learn inter-modality interactions. Large-scale pre-training is commonly done using the Conceptual Captions [38] dataset, with objectives that are designed to encourage interaction between modalities, such as predicting masked tokens or image regions [22, 25, 41, 42], and using contrastive loss between modalities [22]. As a result, such models inherently capture some commonsense knowledge through their pre-training regime. While these models perform impressively on downstream tasks such as VQA [1], they typically perform worse on questions requiring reasoning about knowledge beyond the image content or involving multiple reasoning hops. In our work, we introduce VLC-BERT, a multimodal transformer model based on VL-BERT that explicitly incorporates external knowledge to alleviate this issue.

2.2. Knowledge-based Visual Question Answering

In recent years, several VQA datasets were designed specifically to require reasoning about external knowledge beyond the image, whether using factual and web information (FVQA [45], WebQA [5], a provided text passage (VLQA [34]), commonsense-driven reasoning (VCR [49]), or external commonsense knowledge (OK-VQA [29],

A-OKVQA [36]). This motivated a line of work on knowledge-enhanced VL transformer models. External knowledge is typically retrieved from a structured knowledge base like ConceptNet [40], in the form of a subgraph, and integrated into the VL transformer as an additional input [9, 20, 28, 47]. Alternative sources of knowledge include image captions [33], Google Search results [26], and textual and visual knowledge from Wikipedia, and Google Images [47]. In contrast to most of the preceding work, PICa [48] and Knowledge Augmented Transformer (KAT) [13] attempt to use GPT-3 [3] in a few-shot setting on the VQA task, by building prompts containing the caption and object tags generated using the image, followed by the question statement, asking the model to produce an answer. In our proposed model, we focus on a specific subset of the knowledge-intensive datasets that require commonsense knowledge. Our approach, that uses COMET [15], for incorporating commonsense knowledge is distinctly different, far simpler and more cost-effective.

2.3. Knowledge incorporation in NLP

Structured large-scale knowledge bases (KBs) like ConceptNet [40] and ATOMIC [35] are widely used in NLP tasks to provide additional commonsense knowledge to models. ConceptNet contains 3.4M assertions focusing on concept and entity relations (such as RelatedTo, Synonym, IsA, MadeOf). ATOMIC contains 1.33M triplets focusing on event-centric social commonsense about causes, effects, mental states of the event participants. Several approaches were proposed for incorporating symbolic knowledge from these KBs into downstream NLP tasks such as encoding subgraphs of relevant knowledge [9, 23] and pre-training on commonsense knowledge bases or tasks [51]. Despite the performance improvements, incorporating knowledge directly from KBs suffers from two limitations: lack of coverage and lack of consideration for context. Commonsense Transformer, COMET [15], attempts to alleviate these issues by fine-tuning pre-trained language models on KBs. COMET can generate inferences for the various KB relations dynamically for new inputs. It has been successfully used for generating knowledge in language tasks [4, 27, 39, 43]. Inspired by the success of these models, we chose to use COMET [15] to generate relevant contextual expansions rather than directly retrieving knowledge from KBs. To the best of our knowledge, we are the first to incorporate commonsense knowledge using COMET in VQA tasks. Newer COMET variants [30, 46] are less applicable to OK-VQA and A-OKVQA as they focus more on event commonsense than entities.

3. Method

We briefly outline the overall architecture of our model and then delve deeper into its individual components. Fig-

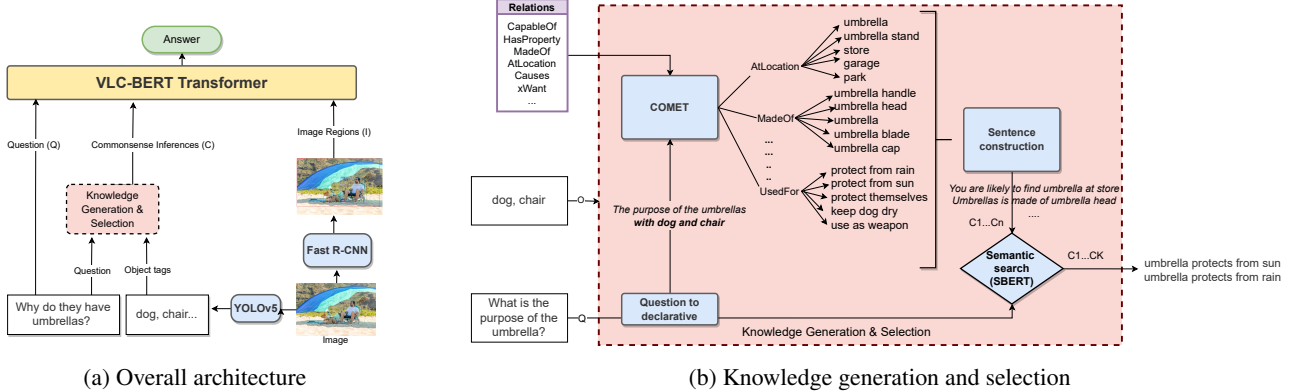


Figure 2: **Architecture of VLC-BERT**: Given an image, VLC-BERT generates commonsense inferences for the question-object phrase using COMET. These inferences are relevance ranked, and top ones are selected and fed along with image regions into a VL-Transformer in order to produce an answer. We utilize semantic similarity between Q and C to select the final K inferences that go into VLC-BERT.

Figure 2a illustrates the VLC-BERT pipeline. Given an image with corresponding image regions I precomputed using Fast RCNN [10] and a question Q related to the image, we generate commonsense inferences C on the events and entities in the question phrase and two object tags O , and select the set of commonsense inferences which is the most useful for answering the question, $C = \{C_1, C_2, \dots, C_k\}$ (§3.1). Finally, we embed Q , I and C , as input to VLC-BERT and train it to predict an answer A to Q (§3.2).

3.1. Structured knowledge generation and selection

3.1.1 Knowledge Generation

To generate commonsense knowledge, we employ the most recent version of COMET [15] initialized using BART [19] in a zero-shot setting. COMET is trained to complete 50 relation types from both ConceptNet [40] (such as *AtLocation*, *MadeOf*) and ATOMIC [35] (such as *xNeed*, *xWants*), thus capturing concept as well as event oriented knowledge. We generate inferences based on 30 relation types most relevant to our work and supported by COMET.¹ Consider the example shown in Figure 2b. For the given question, “*What is the purpose of the umbrella?*” we first process each question using AllenNLP’s constituency parser [17] and convert it into a declarative sentence, since COMET was mainly trained on declarative sentences. In the example shown, “*What is the purpose of the umbrella?*” is rephrased as “*The purpose of the umbrellas is*”. We then adopt a state-of-the-art object detection model, YOLOv5 [16], to translate the corresponding image into object tags that COMET can understand. We select the top two most confident object tags and combine it with the question phrase to obtain a question-object(QO) phrase, “*The purpose of the umbrella*

is, with dog and chair”. We restrict the number of the object tags used in COMET’s input to two because the addition of multiple tags make the inferences more conflated and noisy. In this manner, we can obtain inferences that can provide additional knowledge about both the visual and language inputs to VLC-BERT.

We use beam search to decode the top 5 inferences for each relation type, ranked according to the model’s confidence. Overall, we get $30 \times 5 = 150$ inferences for each input phrase. Finally, we convert each inference to a sentence in natural language using relation-specific templates as defined in [7]. In the shown example, the assertion $\langle \textit{umbrella}, \textit{Located At}, \textit{store} \rangle$ is expressed as “*You are likely to find umbrella at store*”. In order to remove redundant sentences of the same relation type, we measure the lexical overlap by measuring the percentage of common words between two given sentences. We exclude the sentences which have more than 70% overlap with previously constructed sentences of the same relation.

3.1.2 Knowledge Selection

Due to the high cost of computation, and the noise associated with feeding such a large number of text tokens, feeding up to 150 COMET inferences into the VL Transformer model is impractical. In order to rank and select the inferences, we employ semantic search based on sentence transformers (SBERT) [31], which are pre-trained on tasks that retrieve candidate answers to a search query. In this method, the question and the inferences are embedded into the same vector space using SBERT [31] and cosine similarity between the question and the inference embeddings is used to rank the inferences. We prune the set of inference sentences C by picking $K = 5$ inferences which are expected to be the most useful for answering the

¹We include the full list of relation types in the supplementary material.

question Q .

Augmented-SBERT. We augment the SBERT used for semantic search by starting with a pre-trained SBERT model and continuing to train it for 2 epochs on question-inference instances from the *training set* of our datasets. To achieve this, we label the inferences for each question with similarity scores based on the proportion of overlap with the human-annotated answers. Since SBERT is trained on corpora that are distinct from our task, the augmentation ensures that the model understands the nature of query-inference pairings in our tasks. The augmented SBERT especially helps with narrowing down the right relations to the question. For instance, the question in shown in Figure 2b benefits most from the relations that talk about what the umbrella (*UsedFor*) is used for or capable of (*CapableOf*).

3.2. VLC-BERT

We use a single-stream multimodal transformer encoder, VL-BERT [41], as the basis of VLC-BERT. VL-BERT is pre-trained on large-scale vision-language and language-only datasets with a goal of aligning the visual and linguistic features and building robust multimodal representations for downstream tasks. It is trained on the vision-language Conceptual Captions dataset [38], to predict regions-of-interest (RoIs) from language cues, and on the language-only Book-Corpus [52] and English Wikipedia corpora, with a masked language modeling objective. Figure 3 shows the VLC-BERT Transformer architecture. In the following paragraphs, we share how the input sequence is constructed and how the predicted answer is selected.

3.2.1 Inputs

Like VL-BERT, VLC-BERT accepts word token embeddings for language inputs and RoI token embeddings from the image for vision inputs. The architecture of VLC-BERT Transformer is shown in Figure 3. We use the `[CLS]` in the beginning of the sequence, `[END]` to mark the end of the sequence, and the separator token `[SEP]` between different inputs. We feed the question Q as a sequence of word tokens and the image regions I as sequences of RoIs. A `[MASK]` token is used to represent the unknown answer. In addition, we introduce a commonsense fusion token, F , to the input sequence, to incorporate our commonsense inferences.

A straightforward way to leverage the commonsense inferences $C = \{C_1, C_2, \dots, C_k\}$ is to embed each word token in every inference sentence as an input token. However, this would lead to a very long input sequence, where the majority of inputs consist of inferences, thus potentially drawing the model’s attention away from the other inputs. To overcome the challenge, we summarize the information

contained in each inference sentence C_i into a single token representation \vec{C}_i , by embedding the inference using SBERT [31]:

$$\vec{C}_i = \text{SBERT}(C_i) \quad (1)$$

Next, in order to obtain a fused representation of the k commonsense inferences, we attend to the corresponding SBERT embeddings, $[\vec{C}_1 \dots \vec{C}_k]$ against the SBERT embedding of the question, $\vec{Q} = \text{SBERT}(Q)$. The intuition behind this approach is that the model learns to assign a higher score to the most important inference to the question. The key (K_A), query (Q_A) and value (V_A) are assigned as shown below,

$$K_A = \vec{Q} \quad (2)$$

$$Q_A, V_A = \text{append}([\vec{C}_1 \dots \vec{C}_k], \vec{Q}) \quad (3)$$

$$\vec{F} = \text{MHA}(K_A, Q_A, V_A) \quad (4)$$

where MHA is the standard multi-head attention [44], that delivers a *single vector* incorporating all relevant commonsense knowledge required to answer the question. Note that we append the question embedding \vec{Q} to list of commonsense inference embeddings for Q and V because there may be cases where none of the inferences are useful to answer the question. In such a case, the model may choose to ignore the inferences by attending to the question embedding \vec{Q} instead.

Weak Supervision. In order to train the MHA block effectively, we employ weak supervision on the attention weights. For a small subset of the questions in the training set, we obtain label attention weights by following these steps: (1) we initialize a vector \hat{A} of length $k + 1$ where all values are 0.05, (2) for each C_i , if C_i contains a word in the ground-truth answer list, then we set the \hat{A}_i to 0.8, (3) if none of the C inferences contain answer words, we assign a weight of 0.8 to \hat{A}_{k+1} so that the question has the largest weight, and (4) we normalize \hat{A} so that its values sum up to 1. We then apply cross-entropy loss between the predicted attention weights from MHA and our label attention weights \hat{A} , and sum this with the answer prediction loss.

Finally, a positional encoding is added to all input tokens following the method described in VL-BERT. In addition, a different segment type encoding is applied to the four segments in the input sequence: the question segment, the commonsense segment, the masked answer segment, and the image region segment.

3.2.2 Answer Selection

We use the encoded `[MASK]` token to represent the answer, thereby making VQA a masked language modelling task

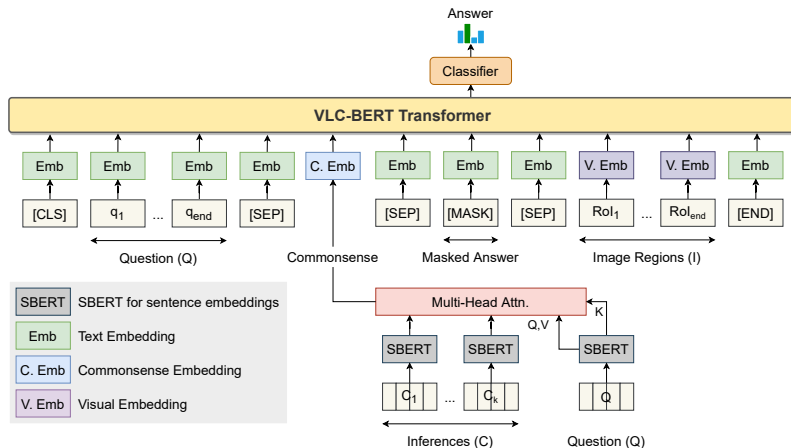


Figure 3: **VLC-BERT Transformer** is a single-stream Transformer that can attend across language, vision, and commonsense representations. We use the MHA block to fuse commonsense inferences into a useful commonsense representation.

with visual cues. To predict the final answer, we apply a classifier over the entire answer vocabulary, as done in VL-BERT. During training, we follow VL-BERT and use a cross-entropy loss over picking the correct answer from an answer vocabulary.

4. Datasets

We perform experiments on the OK-VQA [29] and A-OKVQA [36] datasets. In order to utilize the existing VL-BERT model effectively, we pre-train VLC-BERT on the larger VQA 2.0 [12].

OK-VQA. In the Outside-Knowledge VQA dataset, questions require external knowledge in addition to the information in the images. The dataset is composed of 14,031 images and 14,055 questions, and the crowdsourced questions are divided into ten knowledge categories: Vehicles and Transportation; Brands, Companies and Products; Objects, Materials and Clothing; Sports and Recreation; Cooking and Food; Geography, History, Language and Culture; People and Everyday Life, Plants and Animals; Science and Technology; and Weather and Climate. OK-VQA only contains open-ended questions with five human-provided answers. Since OK-VQA does not have a validation set, we dedicate 1,000 of the 9,009 training questions for validation.

A-OKVQA. A-OKVQA [36] is the augmented successor to OK-VQA and consists of 25K questions that require a combination of commonsense, visual, and physical knowledge. In contrast to other knowledge-based visual question answering datasets, the questions in A-OKVQA are conceptually diverse, involving knowledge that is not contained in the image, and cannot be resolved by a

simple knowledge base query. A-OKVQA is split into training, validation, and test sets based on images used from the COCO 2017 [24] dataset. Moreover, all questions in the dataset have human annotated direct answers as well as multiple-choice options, but we focus on the direct answers. The A-OKVQA test set is blind, requiring us to submit to the leaderboard to obtain a test accuracy.

VQA 2.0. The Visual Question Answering (v2.0) dataset contains 1.1 million crowdsourced questions about 204,721 images from the COCO dataset [24]. Each question is annotated with 10 ground truth answers obtained using Amazon Mechanical Turk. A majority of the questions in this dataset do not require external commonsense knowledge.

4.1. Evaluation Metric

Both datasets use the same accuracy-based evaluation metric. Each question has a set of 10 ground truth answers provided by different annotators. Accuracy is calculated as the percentage of predicted answers that were proposed by at least 3 human annotators: $\text{acc} = \min(\frac{\# \text{ humans gave the answer}}{3}, 1)$.²

5. Implementation Details

The implementation of our model builds on VL-BERT [41]. To that end, we follow the fine-tuning steps provided in the official codebase of the VL-BERT model for VQA 2.0, and modify it to support the OK-VQA and A-OKVQA datasets. We maintain the recommended hyperparameter values, and train the $BERT_{BASE}$ size of the model, with a hidden feature dimension of 768. The model is trained for

²Following the same evaluation, each of the 5 answers in OK-VQA is used twice

20 epochs on the OK-VQA and A-OKVQA datasets. For all models, we use a batch size of 16 and gradient accumulation step size of 4. We train the models presented in the main result thrice and report the average test accuracy on the OK-VQA dataset, and the best (leaderboard) test accuracy on the A-OKVQA dataset.

Answer Vocabulary. Due to the large number of unique answers to questions in visual question answering datasets, it is infeasible to use all answers in the answer vocabulary. For the OK-VQA dataset, following KRISP [28], we build an answer vocabulary of 2,249 answers by selecting all answers in the training set that appear at least 10 times. This answer vocabulary ignores the empty space answer, and includes an `<UNK>` answer token. During training, if a ground truth answer is not present in the answer vocabulary, we assign it to the (`<UNK>`) token. For the A-OKVQA dataset, we use the answer dictionary that is already provided in the dataset [36].

VQA Pre-Training (VQA P.T). Following the idea that pre-training is beneficial for Transformer models, we initialize VLC-BERT with weights obtained after fine-tuning VL-BERT on the VQA 2.0 dataset for 5 epochs. Note that KRISP [28] benefits from pre-training on the VQA 2.0 dataset, and PICa [48] and KAT [14] utilize GPT-3, a large-scale pre-trained model, for external commonsense. Furthermore, because OK-VQA and A-OKVQA are significantly smaller than VQA 2.0, this initialization favourably benefits the training process and gives us a stronger baseline to work with.

6. Evaluation

In this section, we focus on evaluating VLC-BERT on the OK-VQA and A-OKVQA datasets and comparing against existing state-of-the-art models for VQA with external commonsense knowledge. Table 1 highlights our performance improvements on the test set for OK-VQA and A-OKVQA against other models. Later in this section, we ablate on the components of our model.

6.1. Main Results

Table 1 specifies which knowledge sources each model leverages. In the top section, we consider models that utilize knowledge bases such as ConceptNet and Wikipedia, as well as models that utilize web search APIs to obtain external knowledge. VLC-BERT incorporates COMET, which is trained on ConceptNet and ATOMIC, and we compare favourably against these models. Notably, VLC-BERT achieves an accuracy of 43.14 on OK-VQA, outperforming KRISP (Wikipedia + ConceptNet + VQA P.T.) by over 4 points, and MAVEx (Wikipedia + ConceptNet + Google Images) by about 2 points. While our model

clearly outperforms previous methods that use knowledge bases, it does not outperform models with large-scale pre-training and large number of parameters such as GPT-3 [3] and GPV2 [18], which incorporate implicit commonsense knowledge and require extensive resources to train. However, on OK-VQA, we achieve very similar results to PICa-Base [48], despite not having access to GPT-3. We expect that the use of a large pre-trained model like GPT-3 can further boost the performance of VLC-BERT.

6.2. Ablation Tests

We perform comprehensive ablations on the validation set of the A-OKVQA dataset, as represented in Table 2.³

VQA P.T. We begin by training A-OKVQA on the baseline VL-BERT model without VQA pre-training. This gives us a score of 36.24. Next, obtain a new baseline for our model with VQA pre-training, where we then initialize VLC-BERT with pre-trained weights on the VQA 2.0 dataset, and further train it on the A-OKVQA dataset. This results in a score of 43.46, over 7 points better, highlighting the impact of pre-training with a large-scale dataset. This model is a strong baseline for our VQA tasks.

Comm. Inference Representation. In the full model, we use SBERT to summarize each commonsense inference into a single vector, and use the multi-head attention block to capture useful information from the list of inference vectors. To test the effectiveness of our commonsense inference representation method, we first ablate SBERT, i.e., we incorporate all inferences as an additional text input for VLC-BERT, feeding them token-by-token. This results in an accuracy score of 43.44, which is slightly lower than our baseline with VQA pre-training. Next, we use SBERT to summarize inferences, and feed the SBERT embeddings directly into VLC-BERT with only a linear projection layer rather than the MHA block. This variant performs worse than the model with the MHA block by 1.25 points.

Augmented SBERT. In order to familiarize SBERT with our question-inference pairs, we fine-tune SBERT on the training set of A-OKVQA and OK-VQA (Sec 3.1.2). We perform an ablation by evaluating our model on SBERT that has never been exposed to the question-inference-pairs. This results in a drop of 0.85 points in accuracy, which shows that our augmentation of SBERT is effective.

7. Analysis

7.1. Commonsense subsets

Questions in OK-VQA and A-OKVQA datasets are diverse and require commonsense reasoning, visual under-

³We present additional ablations in supplementary material Sec 2.3

Table 1: Accuracy of our model against other models for OK-VQA and A-OKVQA datasets. Our model improves upon existing knowledge base based models due to the contextualized commonsense inferences from COMET, which is trained on ConceptNet and ATOMIC. We compare favourably against the highlighted models that utilize external knowledge bases. Note: P.T. stands for Pre-Training.

Method	Knowledge Sources	OK-VQA	A-OKVQA	Approx. Params
ViLBERT [36]	-	-	25.85	116M
LXMERT [36]	-	-	25.89	-
BAN + AN [29]	Wikipedia	25.61	-	-
BAN + KG-AUG [20]	Wikipedia + ConceptNet	26.71	-	-
MUTAN + AN [29]	Wikipedia	27.84	-	-
ConceptBert [9]	ConceptNet	33.66	-	118M
KRISP [28]	Wikipedia + ConceptNet	32.31	27.1	116M
KRISP [28]	Wikipedia + ConceptNet + VQA P.T.	38.9	-	116M
Visual Retriever-Reader [26]	Google Search	39.2	-	-
MAVEx [47]	Wikipedia + ConceptNet + Google Images	41.37	-	-
GPV2 [18, 36]	Web Search (Web10k) + COCO P.T.	-	40.7	220M
PICa-Base [48]	GPT-3	43.3	-	175B
PICa-Full [48]	GPT-3	48.0	-	175B
KAT [14]	Wikidata + GPT-3	54.41	-	175B
VLC-BERT (Ours)	VQA P.T. + COMET	43.14	38.05	118M

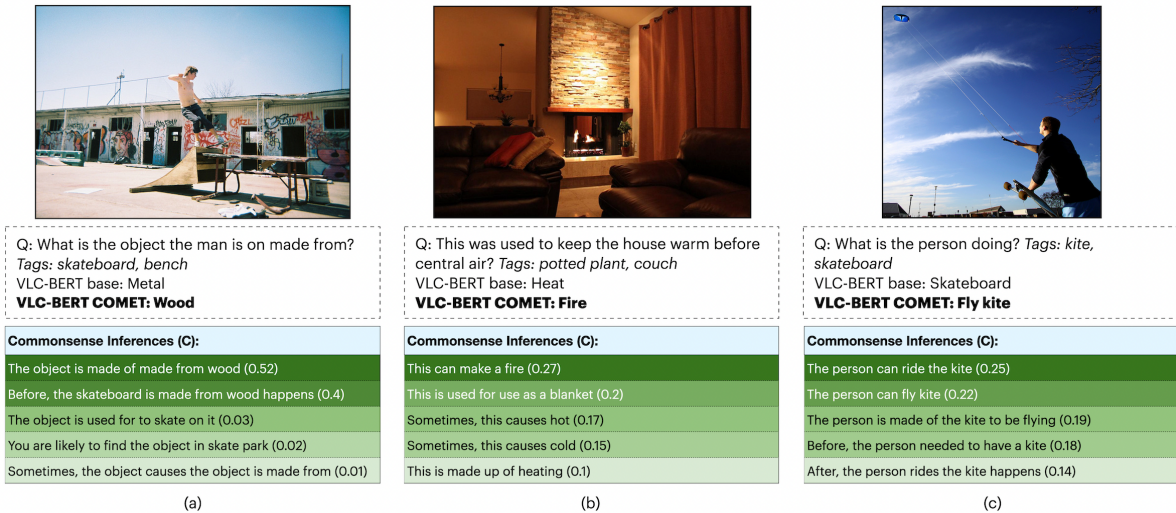


Figure 4: **Attention analysis:** (a) is from A-OKVQA, and (b) and (c) are from OK-VQA. We observe that the weakly supervised attention layer in VLC-BERT accurately picks useful commonsense inferences. In (c), we observe how object tags are useful to guide COMET to produce contextualized knowledge.

standing, as well as factual knowledge. While COMET can generate contextualized commonsense knowledge, it does not help with questions that require scene understanding (e.g., “What is to the left of the computer?”), factual knowledge (e.g., “Where was this food invented?”), or text/symbol recognition (e.g., “What does this sign say?”). Moreover, averaging results on the entirety of OK-VQA and A-OKVQA obfuscates the improvements brought about to a

subset of questions that truly require commonsense knowledge. We propose subsets to assess the performance of our model on questions that are more likely to require external commonsense knowledge. We obtain the subsets by eliminating questions that are mostly factual or visual, and hence do not require commonsense, following these conditions: (1) *factual*: The question or answer contains named entities (e.g., “USA”); (2) *numerical*: The answers contain numbers

Table 2: Ablation of various components in VLC-BERT, evaluated on the A-OKVQA validation set. We observe that all the components of our model play a critical role in empirical performance.

VQA P.T.	Aug. SBERT	SBERT	Attn.	Val
VQA Pre-training				
-	-	-	-	36.24
✓	-	-	-	43.46
Comm. Inference Representation				
✓	✓	-	-	43.44
✓	✓	✓	-	43.64
✓	✓	✓	✓	44.95
Augmentation of SBERT				
✓	-	✓	✓	44.10
✓	✓	✓	✓	44.95

Table 3: Evaluation on the subsets of OK-VQA test (OK_s) and A-OKVQA validation (A-OK_s) sets, where factual, numerical and visual questions are pruned. The performance gain observed on the subsets shows a better picture of where external commonsense is effective.

Method	OK	OK _s	A-OK	A-OK _s
Base	42.29	47.4	43.46	46.52
w/ COMET	43.14	48.21	44.95	49.53

or number words (e.g., “twenty”) or the question has date or time words (e.g., “century”); (3) *visual*: The question contains directional words (e.g., “left of”) and words referring to symbols (e.g., “mascot”).

In Table 3, we show that VLC-BERT with COMET performs 3 points better on the A-OKVQA subset, and maintains an 0.8 point improvement on the OK-VQA subset. This substantiates our claim that utilizing our COMET pipeline substantially increases VLC-BERT’s ability to answer questions that require external knowledge.

7.2. Attention Analysis

In this section, we show qualitative examples to demonstrate questions where VLC-BERT benefits from contextualized commonsense knowledge from COMET. We also show the corresponding attention weights, to show the effectiveness of the proposed weakly-supervised attention mechanism. Fig 4a shows an example from A-OKVQA, where COMET’s inferences on the question and the object tags, weighted by the attention score, results in the correct answer. Fig 4b shows an example from OK-VQA where VLC-BERT COMET exhibits higher attention towards the fire despite the object tags missing the fireplace. This is

an example where deriving inferences from the question phrase is equally important as doing so with the object tags. Fig 4c shows that inferences on the object tag *kite* drove the model to answer correctly. The supplementary material includes additional examples of improvements and failures.

8. Conclusions

We presented Vision-Language-Commonsense BERT (VLC-BERT) for external knowledge-driven VQA tasks. VLC-BERT outperforms previous models based on knowledge bases on the OK-VQA and A-OKVQA datasets by incorporating contextualized commonsense knowledge from COMET and combining it with visual and linguistic inputs. Through our evaluation, we show the effectiveness of our knowledge generation, selection, and incorporation strategies, and the positive impact of VQA pre-training.

Our analysis of VLC-BERT highlighted a few limitations of our model and the datasets we evaluate on. First, some questions require a deeper understanding and linking of multiple entities and events in the image, that object tags lack, for deriving relevant commonsense inferences. Second, condensing the commonsense inferences using SBERT and MHA leads to a compressed representation that may cause the model to lose some information. Finally, our model is limited by COMET, and the knowledge bases it is trained on, as we observe that large-scale models like GPT-3 outperform it.

We view our work as a first step in analyzing the potential of *generative commonsense incorporation*, and exploring approaches to *decide when commonsense is needed*. In the future, our goal is to work towards creating a version of COMET that can utilize image context concerning multiple entities and events. We also plan to investigate the potential of multi-hop reasoning with COMET to bridge the question and image-based expansions closer.

9. Acknowledgments

This work was funded, in part, by the Vector Institute for AI, Canada CIFAR AI Chair, NSERC CRC, NSERC DG and Accelerator Grants, and a research gift from AI2. Hardware resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute⁴. Additional hardware support was provided by John R. Evans Leaders Fund CFI grant and Compute Canada under the Resource Allocation Competition award. Finally, we sincerely thank Prof. Giuseppe Carenini for valuable feedback and discussions.

⁴www.vectorinstitute.ai/partners

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiaseen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, 2015.
- [2] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. COMET: Commonsense transformers for automatic knowledge graph construction. In *57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 1877–1901, 2020.
- [4] Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. It’s not rocket science : Interpreting figurative language in narratives. *Transactions of the Association for Computational Linguistics (TACL)*, 2022.
- [5] Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [6] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M.F. Moura, Devi Parikh, and Dhruv Batra. Visual Dialog. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [7] Joe Davison, Joshua Feldman, and Alexander Rush. Commonsense knowledge mining from pretrained models. In *2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1173–1178, 2019.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [9] François Gardères, Maryam Ziaefard, Baptiste Abelois, and Freddy Lécué. Conceptbert: Concept-aware representation for visual question answering. In *FINDINGS*, 2020.
- [10] Ross Girshick. Fast R-CNN. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [11] Jonathan Gordon and Benjamin Van Durme. Reporting bias and knowledge acquisition. In *Workshop on Automated Knowledge Base Construction*, page 25–30, 2013.
- [12] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [13] Liangke Gui, Borui Wang, Qiuyuan Huang, Alex Hauptmann, Yonatan Bisk, and Jianfeng Gao. Kat: A knowledge augmented transformer for vision-and-language, 2021.
- [14] Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander Hauptmann, Yonatan Bisk, and Jianfeng Gao. KAT: A knowledge augmented transformer for vision-and-language. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, July 2022.
- [15] Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. Comet-atomic 2020: On symbolic and neural commonsense knowledge graphs. In *AAAI*, 2021.
- [16] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, TaoXie, Jiacong Fang, imyhxy, Kalen Michael, Lorna, Abhiram V, Diego Montes, Je-bastin Nadar, Laughing, tkianai, yxNONG, Piotr Skalski, Zhiqiang Wang, Adam Hogan, Cristi Fati, Lorenzo Mamma, AlexWang1900, Deep Patel, Ding Yiwei, Felix You, Jan Hajek, Laurentiu Diaconu, and Mai Thanh Minh. ultralytics/yolov5: v6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference, Feb. 2022.
- [17] V. Joshi, Matthew E. Peters, and Mark Hopkins. Extending a parser to distant domains using a few dozen partially annotated examples. In *ACL*, 2018.
- [18] Amita Kamath, Christopher Clark, Tanmay Gupta, Eric Kolve, Derek Hoiem, and Aniruddha Kembhavi. Webly supervised concept expansion for general purpose vision models, 2022.
- [19] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics.
- [20] Guohao Li, Xin Wang, and Wenwu Zhu. Boosting visual question answering with context-aware knowledge aggregation. *28th ACM International Conference on Multimedia*, 2020.
- [21] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- [22] Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantic aligned pre-training for vision-language tasks. *European Conference on Computer Vision (ECCV)*, 2020.

- [23] Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *EMNLP-IJCNLP*, 2019.
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755. Springer International Publishing, 2014.
- [25] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [26] Man Luo, Yankai Zeng, Pratyay Banerjee, and Chitta Baral. Weakly-supervised visual-retriever-reader for knowledge-based question answering. In *EMNLP*, 2021.
- [27] Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. Like hiking? you probably enjoy nature: Persona-grounded dialog with commonsense expansions. In *2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9194–9206. Association for Computational Linguistics, Nov. 2020.
- [28] Kenneth Marino, Xinlei Chen, Devi Parikh, Abhinav Gupta, and Marcus Rohrbach. KRISP: Integrating implicit and symbolic knowledge for open-domain knowledge-based vqa. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14111–14121, 2021.
- [29] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. OK-VQA: A visual question answering benchmark requiring external knowledge. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [30] Jae Sung Park, Chandra Bhagavatula, Roozbeh Mottaghi, Ali Farhadi, and Yejin Choi. Visualcomet: Reasoning about the dynamic context of a still image. In *ECCV*, 2020.
- [31] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019.
- [32] Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2020.
- [33] Ander Salaberria, Gorka Azkune, Oier Lopez de Lacalle, Aitor Soroa, and Eneko Agirre. Image captioning for effective use of language models in knowledge-based visual question answering. *CoRR*, 2021.
- [34] Shailaja Keyur Sampat, Yezhou Yang, and Chitta Baral. Visuo-linguistic question answering (VLQA) challenge. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Nov. 2020.
- [35] Maarten Sap, Ronan LeBras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. Atomic: An atlas of machine commonsense for if-then reasoning. In *AAAI*, 2019.
- [36] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. *arXiv*, 2022.
- [37] Meet Shah, Xinlei Chen, Marcus Rohrbach, and Devi Parikh. Cycle-consistency for robust visual question answering. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6642–6651, 2019.
- [38] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of ACL*, 2018.
- [39] Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Unsupervised commonsense question answering with self-talk. In *2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4615–4629. Association for Computational Linguistics, Nov. 2020.
- [40] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. In *AAAI*, 2017.
- [41] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020.
- [42] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Conference on Empirical Methods in Natural Language Processing*, 2019.
- [43] Yufei Tian, Arvind krishna Sridhar, and Nanyun Peng. HypoGen: Hyperbole generation with commonsense and counterfactual knowledge. In *Association for Computational Linguistics: EMNLP 2021*, pages 1583–1593, 2021.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [45] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Fvqa: Fact-based visual question answering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(10):2413–2427, oct 2018.
- [46] Peter West, Chandrasekhar Bhagavatula, Jack Hessel, Jena D. Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. Symbolic knowledge distillation: from general language models to commonsense models. In *NAACL*, 2022.
- [47] Jialin Wu, Jiasen Lu, Ashish Sabharwal, and Roozbeh Mottaghi. Multi-Modal Answer Validation for Knowledge-based VQA. In *AAAI*, 2022.
- [48] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. An empirical study of gpt-3 for few-shot knowledge-based vqa, 2021.
- [49] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reason-

ing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

- [50] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [51] Wanjun Zhong, Duyu Tang, Nan Duan, Ming Zhou, Jiahai Wang, and Jian Yin. Improving question answering by commonsense-based pre-training. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 16–28. Springer, 2019.
- [52] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *IEEE International Conference on Computer Vision (ICCV)*, pages 19–27, 2015.