

Reasoning over Vision and Language: Exploring the Benefits of Supplemental Knowledge

Violetta Shevchenko, Damien Teney, Anthony Dick, Anton van den Hengel
Australian Institute for Machine Learning, The University of Adelaide, Australia

{violetta.shevchenko,damien.teney,anthony.dick,anton.vandenhengel}@adelaide.edu.au

Abstract

The limits of applicability of vision-and-language models are defined by the coverage of their training data. Tasks like vision question answering (VQA) often require common-sense and factual information beyond what can be learned from task-specific datasets. This paper investigates the injection of knowledge from general-purpose knowledge bases (KBs) into vision-and-language transformers. We use an auxiliary training objective that encourages the learned representations to align with graph embeddings of matching entities in a KB. We empirically study the relevance of various KBs to multiple tasks and benchmarks. The technique brings clear benefits to knowledge-demanding question answering tasks (OK-VQA, FVQA) by capturing semantic and relational knowledge absent from existing models. More surprisingly, the technique also benefits visual reasoning tasks (NLVR2, SNLI-VE). We perform probing experiments and show that the injection of additional knowledge regularizes the space of embeddings, which improves the representation of lexical and semantic similarities. The technique is model-agnostic and can expand the applicability of any vision-and-language transformer with minimal computational overhead.

1 Introduction

The last few years have seen a surge of interest for vision and language (V&L) tasks. These include visual question answering for example (VQA), in which the machine must answer a question about an image (Fig. 1). V&L tasks require processing two modalities and reasoning over textual, visual, and abstract concepts.

The current state of the art in V&L are models based on transformers such as BERT (Devlin et al., 2019) that have been extended to handle visual inputs (e.g. Tan and Bansal (2019)). These

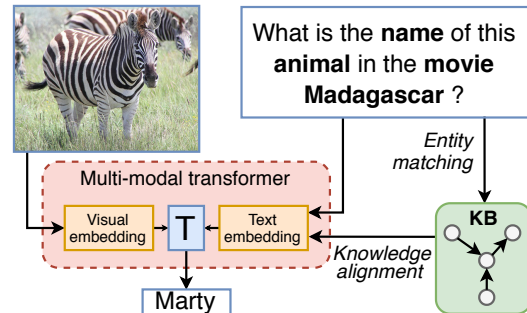


Figure 1: We inject additional information from knowledge bases (KBs) in a vision-and-language transformer. We first preprocess the KB into knowledge embeddings. When training the transformer, we use an auxiliary objective that aligns its learned representations with knowledge embeddings. In the context of visual question answering (pictured), the model is able to answer questions that require knowledge that is typically not captured in task-specific datasets.

models are usually pretrained on a collection of datasets of paired textual and visual data. Suitable datasets include image captioning data (Chen et al., 2015; Sharma et al., 2018) and VQA data (Goyal et al., 2017; Hudson and Manning, 2019). Despite their large scale, these datasets only cover limited domains. Many are based on images from COCO (Lin et al., 2014) and Visual Genome (Krishna et al., 2017) and the linguistic diversity of textual annotation is limited. The V&L tasks that we are ultimately interested in require knowledge beyond current datasets, e.g. about specific events, named entities, common sense, and abstract concepts (see Fig. 3).

This paper focuses on the expansion of the applicability of V&L models with additional knowledge. During training, we infuse the model with knowledge from an external source, distinct from datasets of paired V&L data. The challenge is that standard V&L data is not annotated or paired with such additional knowledge. Even though techniques have

been proposed to exploit additional data in NLP, including text-based question answering (Kafle et al., 2019; Lv et al., 2020; Lin et al., 2019; Rajani et al., 2019), little work has been done on the extension to V&L. Works in NLP with benchmarks of knowledge-demanding questions (Yang et al., 2015; Talmor et al., 2019; Clark et al., 2018; Miheylov et al., 2018; Sap et al., 2019) have shown that knowledge bases (KBs) contain information that can benefit large-scale pretrained models. Motivated by this line of evidence, we aim to evaluate similar mechanisms for V&L tasks.

In the context of VQA, datasets of knowledge-demanding questions (Wang et al., 2017b; Marino et al., 2019; Shah et al., 2019) have proved challenging for existing methods. For example, a question like *Is the man eating healthy food ?* requires recognizing a type of food and relating it to its nutritional quality. Learning this type of knowledge from VQA training examples would be clearly inefficient. Other examples of challenging questions involve references to named entities such as brands, locations, or movie titles. For example, the question *Levi’s is a popular brand of what item shown here ?* requires knowledge of the brand’s specialization before identifying the correct element in the image. We believe that embedding this type of knowledge in a model is a necessary step to enable progress on complex multimodal question answering.

This paper describes a technique to inject information from KBs into a transformer-based model during its training (see Fig. 1). We take inspiration from Goodwin and Demner-Fushman (2019) and adapt their regularizer from text-based models to V&L transformers. We provide an implementation of our method on top of the popular LXMERT model (Tan and Bansal, 2019). We use it to investigate the suitability of several KBs to different tasks and benchmarks.

Our contributions are summarized as follows.

- We describe a method to inject information from knowledge bases during the training of vision-and-language transformers.
- We implement the method on top of the popular LXMERT model with the ConceptNet and Wikidata knowledge bases.
- We perform an extensive empirical evaluation on four downstream tasks. We demonstrate clear improvements on knowledge-demanding VQA

(OK-VQA and FVQA datasets) and visual reasoning (NLVR2 and SNLI-VE datasets).

- We conduct an in-depth analysis, including ablations and probing experiments. They show that we improve the representation of lexical, semantic, and relational knowledge that is lacking in typical V&L models. This explains the surprising improvements on tasks that do not explicitly depend on external knowledge.

2 Related Work

Vision-and-language tasks require joint processing of visual and textual data *e.g.* for image captioning (Anderson et al., 2018; Hossain et al., 2019) and VQA (Wu et al., 2017; Teney et al., 2017). They have historically been approached with task-specific models, but recent transformer-based models (Vaswani et al., 2017) were shown to be applicable to a variety of tasks. A transformer can thus be pretrained on multiple datasets and then fine-tuned for one specific task (Tan and Bansal, 2019; Zhou et al., 2020; Alberti et al., 2019; Chen et al., 2020; Li et al., 2020; Lu et al., 2019; Su et al., 2019; Li et al., 2019; Sun et al., 2019). This paper describes a method to embed additional information in a transformer-based model (“additional” to the pretraining and fine-tuning datasets). Our implementation builds on the popular LXMERT model (Tan and Bansal, 2019).

Additional knowledge in NLP can help with tasks requiring commonsense or factual information (Storks et al., 2019). Various techniques have been proposed to improve transformers such as BERT (Devlin et al., 2019). Zhang et al. (2019) proposed ERNIE, which feeds graph embeddings of text entities to the model. Peters et al. (2019) proposed KnowBert, a similar technique suitable to multiple KBs. Levine et al. (2020) used WordNet to aid in the masked-word prediction objective, and improve lexical understanding in downstream tasks. Ye et al. (2019) proposed a multiple-choice question answering pretraining task and improved performance on multiple datasets requiring commonsense reasoning. Liu et al. (2020) addressed noise issues by controlling the amount of domain-specific knowledge infused into the model. Goodwin and Demner-Fushman (2019) proposed OSCAR, a regularization method to inject ontological knowledge in a pretrained language model. Wang et al. (2019) proposed to simultaneously learn knowledge rep-

representations while optimizing a masked-language objective, rather than using pretrained knowledge embeddings. All of these works were applied to NLP tasks. This paper studies the suitability of similar mechanisms to V&L tasks by applying the technique of Goodwin and Demner-Fushman (2019) to a multimodal transformer.

Knowledge-based VQA refers to benchmarks designed to require additional information (Marino et al., 2019; Shah et al., 2019; Wang et al., 2017a,b). Models have been proposed that retrieve such information from KBs based on question and image contents (Wang et al., 2017a; Wu et al., 2016a; Wang et al., 2017b; Narasimhan and Schwing, 2018). Some works use other sources of external information at training (Teney and van den Hengel, 2016) or test time (Teney and van den Hengel, 2018, 2019) but they showed limited improvements on VQA benchmarks. Recently proposed Concept-Bert (Gardères et al., 2020) model jointly learns visual, textual and knowledge embeddings to fuse commonsense information into VQA models. This paper describes a method applicable to a variety of sources of information and to tasks beyond VQA. We also show that different tasks benefit from different types of information.

3 Methodology

We describe a general, simple yet effective technique to embed additional knowledge into a transformer-based model. It is compatible with existing multimodal transformers and thus suitable to a variety of tasks. The method proceeds in three stages (see Fig. 2).

1. We preprocess the additional knowledge into a set of vector representations that we call knowledge embeddings. For example, with a relational knowledge base, we apply a graph embedding method to obtain a vector representation of every of its entities. Each is associated with a textual expression.
2. We match sentences in the V&L training data with knowledge embeddings. Matches in the training data are referred to as knowledge-rich expressions.
3. During the training of the transformer (pretraining and/or fine-tuning), we optimize an additional objective that aligns its learned representations of knowledge-rich expressions with the matching knowledge embeddings.

We now describe each stage in details.

3.1 Representations of Additional Knowledge

The versatility of the approach rests on storing the additional knowledge as a set of knowledge embeddings $V = \{v_i\}$ with $v_i \in \mathbb{R}^{d_v}$. These can be produced by preprocessing sources such as text corpora (with word embedding methods) or relational knowledge bases (with a graph embedding method, see Section 4). These knowledge embeddings capture semantic information about the relations between entities, which we will incorporate into the V&L model. Each knowledge embedding v_i is associated with an instantiation w_i in the V&L data. In this work, the w_i are purely textual (single- and multiple-word expressions) but future work could consider visual representations of concepts represented by knowledge embeddings. We denote with $W = \{w_i\}$ the vocabulary of these instantiations.

3.2 Matching V&L Training Data with Knowledge Embeddings

We match parts of the training data with entities having additional knowledge. Since our vocabulary W contains textual expressions, we use greedy longest-string matching to identify subsequences in the data that match any w_i . Possible improvements left for future work include performing named entity recognition and homonym disambiguation. In practice, entities in our KBs have unique textual representations, making homonyms a non-issue.

We represent text in the training data as a sequence of tokens $T = (t_1, \dots, t_M)$. The transformer internally maps each token t_j to a word embedding $e_j \in \mathbb{R}^{d_e}$. Each correspondence identified by the matching algorithm is of the form of a subsequence $(t_{a_k}, \dots, t_{b_k}) \subset T$ that matches $w_k \in W$. We refer to such a subsequence as a *knowledge-rich expression*. To obtain a fixed-size representation c_k of a knowledge-rich expression, we sum the word embeddings of its constituent tokens *i.e.* $c_k = e_{a_k} + \dots + e_{b_k}$.

3.3 Aligning Learned Representations with Knowledge Embeddings

The core of the method is an additional training objective that encourages the transformer to produce representations of knowledge-rich expressions (c_k) that collectively align with knowledge vectors (v_k). A vector c_k as defined above is the representation learned by the model of a knowledge-rich expression. We do not expect these vectors to correspond

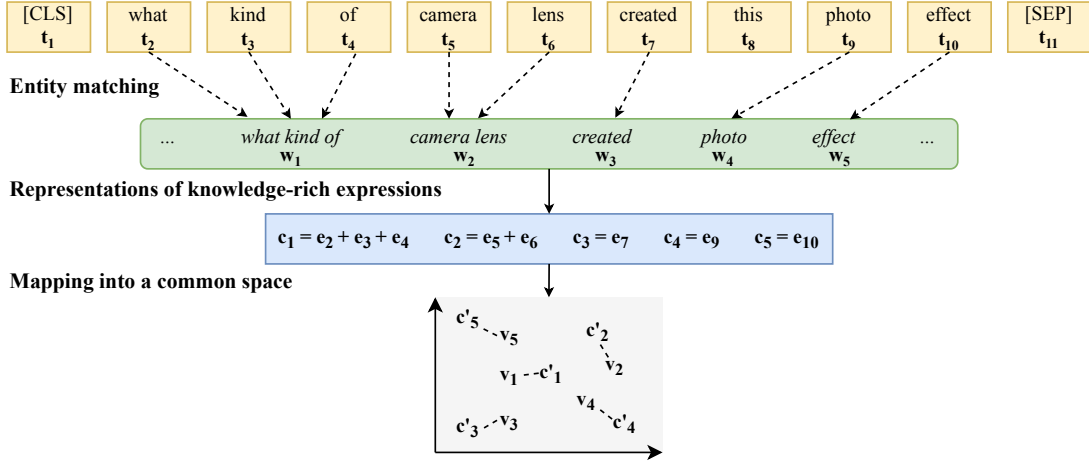


Figure 2: Summary of the approach. During training, we first match tokens in the V&L training data (yellow) with entities of the knowledge base (green). We then train the transformer with an additional loss to align the learned representations c_k (sums of word embeddings in knowledge-rich expressions from the transformer, in blue) with the knowledge embeddings v_k derived from the knowledge base.

	OK-VQA	FVQA	SNLI-VE	NLVR2
Pretraining with COCO captions, Visual Genome captions, VQA v2, GQA				
Baseline (LXMERT)	37.26 ± 0.23	52.30 ± 0.09	74.05 ± 0.19	71.31 ± 0.56
With ConceptNet PT	39.04 ± 0.24	54.08 ± 0.09	75.18 ± 0.21	71.61 ± 0.24
With ConceptNet FT	36.99 ± 0.01	52.19 ± 0.15	74.80 ± 0.08	70.82 ± 0.34
With ConceptNet PT+FT	38.56 ± 0.31	54.27 ± 0.28	75.24 ± 0.22	72.59 ± 0.23
Pretraining with VQA v2, GQA				
Baseline (LXMERT)	36.71 ± 0.22	50.38 ± 0.24	73.57 ± 0.17	67.88 ± 0.66
With ConceptNet PT	38.05 ± 0.41	<u>51.94 ± 0.25</u>	74.07 ± 0.48	69.47 ± 0.20
With ConceptNet FT	36.73 ± 0.50	50.54 ± 0.08	74.24 ± 0.15	67.09 ± 0.32
With ConceptNet PT+FT	<u>38.12 ± 0.11</u>	51.53 ± 0.17	<u>74.26 ± 0.22</u>	<u>69.69 ± 0.12</u>
With Wikidata PT	37.39 ± 0.35	51.00 ± 0.03	73.48 ± 0.17	68.27 ± 0.89
With Wikidata FT	36.31 ± 0.28	50.59 ± 0.24	73.60 ± 0.26	67.64 ± 0.19
With Wikidata PT+FT	37.43 ± 0.25	50.74 ± 0.29	73.50 ± 0.04	68.25 ± 0.51

Table 1: Overall results. Our model with ConceptNet during pretraining and fine-tuning (ConceptNet PT+FT) generally proves best. We report the average accuracy (%) ± one standard deviation over three random seeds.

to their matching knowledge embeddings v , but we desire them to capture the information globally represented in the *relations* between the knowledge embeddings in V . Therefore, we define a new linear layer that maps a learned representation $c_k \in \mathbb{R}^{d_e}$ to $c'_k \in \mathbb{R}^{d_v}$:

$$c'_k = W_c c_k + b_c \quad (1)$$

where $W_c \in \mathbb{R}^{d_e \times d_v}$ and $b_c \in \mathbb{R}^{d_v}$ are learned weights and biases. We then define our alignment loss that encourages each projection c'_k to be close to its corresponding knowledge embedding v_k :

$$L_{\text{align}} = \sum_k \|c'_k - v_k\|^2. \quad (2)$$

Together, Eq. 1 and 2 encourage the global structure of the learned representations c to align with the set of knowledge embeddings V through the projection W_c . The learned representations incorporate information from the knowledge embeddings while allowing the transformer to also represent task-specific information. The model is trained for the combination of its original main loss, and the new alignment loss weighted by a hyperparameter λ *i.e.* $L = L_{\text{main}} + \lambda L_{\text{align}}$.

3.4 Training Strategies

We experimented with the application of the method during pretraining and/or fine-tuning, re-

ferred to as FT, PT, and PT+FT below. Enabling the additional objective during pretraining can benefit from the larger overlap between the training data and the additional knowledge. However, most pretraining tasks do not specifically require additional knowledge, and the model may not learn to effectively use it. During fine-tuning on knowledge-demanding tasks, the model is likely to better learn to capture and use relevant additional knowledge.

3.5 Implementation

We implemented the method on top of LXMERT (Tan and Bansal, 2019), the state-of-the-art model on multiple tasks at the onset of this project. This model is pretrained on five captioning and VQA datasets: COCO and Visual Genome captions, VQA v2 (Goyal et al., 2017), GQA (Hudson and Manning, 2019) and Visual Genome QA (Zhu et al., 2016). We also include experiments with scaled-down pretraining on two datasets (VQA v2 and GQA) which will ease the computational cost of replication.

4 Experiments

We performed a suite of experiments with multiple tasks and benchmarks. Our objective is to evaluate the suitability of two popular knowledge bases (ConceptNet and Wikidata) to these tasks.

Datasets. We used four datasets. The OK-VQA (Marino et al., 2019) and FVQA (Wang et al., 2017b) are two datasets for VQA with questions that require a wide range of commonsense and factual knowledge. The SNLI-VE (Xie et al., 2019) is a dataset for visual entailment. Each instance involves an image “premise” and text “hypothesis”. The task is to determine whether the image entails or contradicts the text, or whether they are neutral to each other. The NLVR2 dataset (Suhr et al., 2019) evaluates visual reasoning. Each instance is a pair of images and a text description. The task is to determine whether the text accurately describes the pair of images.

Knowledge bases. We experimented with two popular KBs. **ConceptNet** (Speer et al., 2017) is a knowledge graph that encodes the meaning of expressions useful for general language understanding. This decades-old project is built on a number of crowd-sourced and curated sources including dictionaries, encyclopedias, and ontologies. We use the 300-dimensional Numberbatch embeddings

distributed by the authors of ConceptNet. They are built using the technique of retrofitting (Faruqui et al., 2015) to combine relational information from the KB with distributional semantics from word2vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and OpenSubtitles (Tiedemann, 2012). We only use the subset of English expressions ($\sim 500,000$ entities). **Wikidata** (Vrandečić and Krötzsch, 2014) is a collaboratively edited database of general knowledge. While ConceptNet mostly covers commonsense knowledge, Wikidata spans a larger domain, including historical events, celebrities, locations, science facts, *etc.* We obtain 200-dimensional embeddings with the PyTorch-BigGraph graph embedding method (Lerer et al., 2019). We use a subset of entities that have links to meaningful Wikipedia pages as done in (Wang et al., 2019). We also discard entities associated with stop words (*e.g.* *the, are, there*) which are common in VQA questions but carry no important semantic information. We retain ~ 4.7 M entities associated with ~ 10 M aliases.

4.1 Overall Results

We report the overall accuracy on all datasets in Table 1. Our approach clearly outperforms the baseline, with a **higher accuracy** on the knowledge-demanding VQA datasets OK-VQA (+1.78%) and FVQA (+1.97%). We also get clear improvements on the visual reasoning datasets SNLI-VE (+1.19%) and NLVR2 (+1.28%). These do not explicitly require specific knowledge, so we hypothesized that the improvement is due to the richer linguistic representations learned by our model. We verified this hypothesis through probing experiments with the SentEval toolkit (Conneau and Kiela, 2018), which showed that our model better captures multiple semantic and syntactic properties of words (*e.g.* +9.36%, +7.72%, +3.88% for word content, tense and subject number prediction respectively). We provide example results in Fig. 3.

The best training strategy is generally to use the additional objective **during both pretraining and fine-tuning** (PT+FT). Only on OK-VQA did the PT strategy perform slightly better. Comparing PT alone with FT alone (the former being superior on all datasets) shows that fine-tuning the representations is not sufficient. The additional objective induces a globally different organization of the embedding space, which cannot be reorganized during fine-tuning alone. It is also interesting to note that

	OK-VQA										
	VT	BCP	OMC	SR	CF	GHLC	PEL	PA	ST	WC	Other
Baseline (LXMERT)	34.91	29.34	35.22	46.41	40.16	32.58	32.68	35.75	35.79	<u>48.68</u>	35.96
With ConceptNet PT	37.40	31.94	<u>37.66</u>	47.52	<u>40.82</u>	38.39	35.06	37.20	31.98	48.63	38.42
With ConceptNet FT	34.46	29.03	35.89	46.02	39.73	33.66	32.57	35.51	<u>35.16</u>	45.84	35.75
With ConceptNet PT+FT	<u>37.11</u>	<u>31.47</u>	39.06	<u>46.52</u>	40.89	<u>36.50</u>	<u>34.08</u>	<u>36.16</u>	33.73	49.56	<u>37.05</u>

Table 2: Accuracy (%) on OK-VQA per question category. The method obtains clear improvements on categories related to commonsense knowledge such as **BCP**: brands, companies and products, **OMC**: objects, material and clothing, **PEL**: people and everyday life, **VT**: vehicles and transportation.

the knowledge injection during pretraining is effective despite the pretraining tasks not specifically requiring additional knowledge.

4.2 Nearest Neighbors in Embedding Space

To better understand how the regularization with additional knowledge changes the structure of the learned embedding space, we examine nearest neighbors of embeddings of individual words (see Table 3). Our model clearly captures more lexical and semantic information. For example, the nearest neighbors for *vitamin* are *calcium* and *supplements* with our model, but they are *margarita* and *amphibious* with the baseline. This stark qualitative difference was observed across the board and is not confined to cherry-picked examples. The lack of linguistic information encoded by baseline V&L transformers is unsurprising since their training data has limited linguistic diversity. Initializing a V&L model with a pretrained BERT has been proposed to address this deficiency, but it was reported to give lower downstream performance by the authors of LXMERT. In comparison, our method brings linguistic information while improving downstream performance.

4.3 Results on OK-VQA

We examine the accuracy on question categories of OK-VQA in Table 2. Each category corresponds to a type of knowledge required. We get **high gains on categories that correspond to a type of knowledge covered in ConceptNet** (objects properties and features, behavior of people, *etc.*). These include OMC (objects, material and clothing), PEL (people and everyday life), BCP (brands, companies and products) and VT (vehicles and transportation). The only category with a drop in accuracy is ST (science and technology), which is also the smallest (84 questions). The category with the largest gain is GHLC (geography, history, language,

and culture), but it contains only 141 questions, and some are distant from these topics (e.g. *What fruit come from these trees ?*). These results should not be over interpreted because of the small size of these categories. Some questions also have imprecise labels. For example, the question *What activity are they doing ?* is labeled with the correct answer *video game*, and our model’s answer *play video game* is considered incorrect.

The **hardest questions for our model** are those referring to exact facts and entities such as place names, famous people, or historical dates. Such precise facts are more difficult to represent and recall than “soft” commonsense knowledge. For example, the question *What year was this picture taken ?* requires to recognize a specific event and fetch a precise related fact. Other difficult questions refer to precise visual cues, while the text of the question is generic, like *What language is on the sign ?*, *Can you guess the place shown in this picture ?* or *Which season is it ?*. Recent analysis showed that V&L transformers rely primarily on the textual modality (Cao et al., 2020; Singh et al., 2020). Additional mechanisms would be needed to allow the recall of facts solely from visual cues.

4.4 ConceptNet vs Wikidata

We now examine the suitability of ConceptNet and Wikidata to the four datasets considered. **ConceptNet provides larger improvements than Wikidata on every dataset**. Wikidata shows improvements on knowledge-driven tasks (OK-VQA and FVQA) but fails to improve over the baseline on the visual reasoning ones (SNLI-VE, NLVR2). Wikidata is almost ten times larger than ConceptNet, but it contains more redundant and noisy information due to its open-source nature. ConceptNet, in contrast, is based on a collection of mostly curated sources. Finally, the vector representations of ConceptNet used were obtained through an advanced

Word	Nearest neighbors with the baseline	Nearest neighbors with our model
argentina	questioned, [PAD], neutron, dorset	argentine, uruguay, paraguay, mendoza
behaviour	[PAD], absurd, authoritative, mba	behavior, behaviors, demeanor, behavioral
bowling	boxing, smashed, dancing, 75	bowler, bowled, cricket, tennis
cottage	farmhouse, scan, condo, ###tta	cottages, bungalow, farmhouse, ##ode
facebook	jade, brady, institution, utrecht	myspace, twitter, youtube, dit
genes	[PAD], oro, subsistence, ##vah	gene, genetic, genetics, genome
lecturer	greenberg, [unused983], avoidance, ##mour	professor, prof, professors, lectures
playstation	splendid, indo, financial, tapping	xbox, wii, sega, consoles

Table 3: Examples of nearest neighbors in the space of word embeddings learned by the baseline and by our model. Our model better captures lexical and semantic similarity.

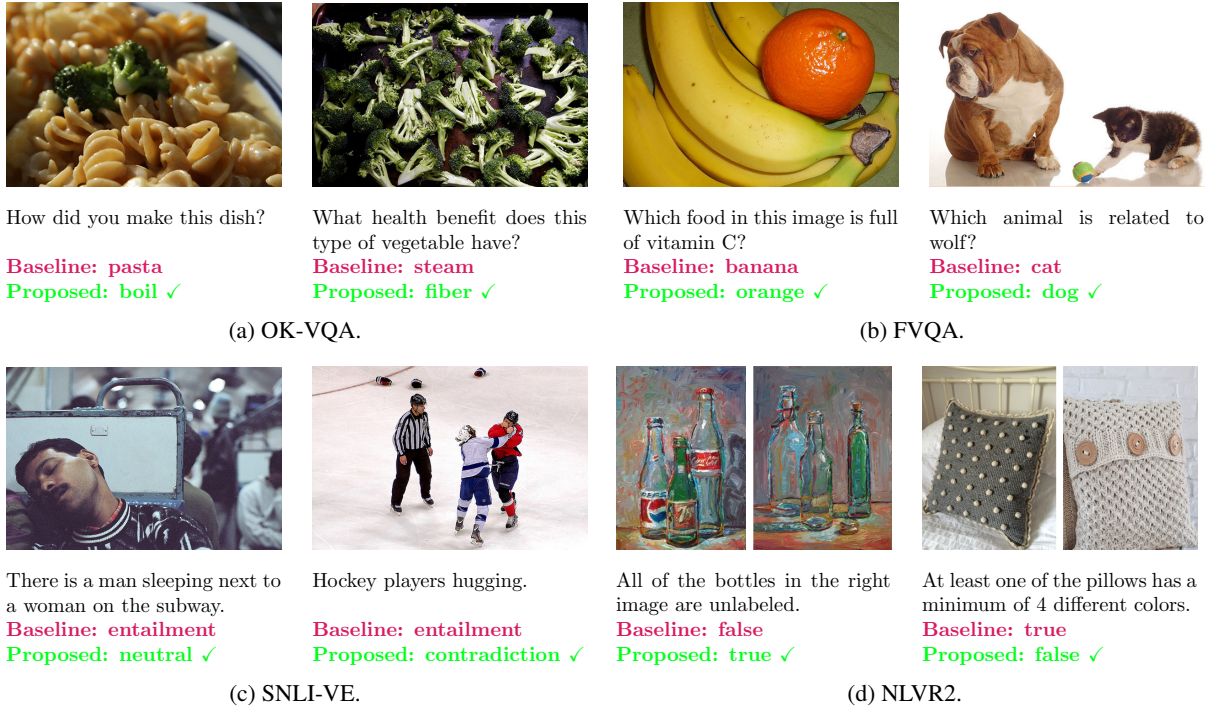


Figure 3: Test cases on which our model (Pretraining with VQA v2, GQA w/ ConceptNet PT+FT) produces better predictions than the baseline. On OK-VQA and FVQA, the questions require factual knowledge of the type found in ConceptNet. On SNLI-VE and NLVR2, the baseline frequently produces common, but incorrect answers.

and proven procedure that involves ConceptNet as well as other pretrained word representations. The representations of Wikidata used are a more direct representation of the knowledge graph.

4.5 Comparison with Existing Methods

In Table 4 we compare our results with the top entries from the leaderboards of OK-VQA¹ and NLVR2². On OK-VQA, the best accuracy is shown by ConceptBert model pretrained on a captioning dataset. The other reported results are from the traditional VQA models BAN (Kim et al., 2018) and MUTAN (Ben-Younes et al., 2017). The

¹<https://okvqa.allenai.org>

²<http://lil.nlp.cornell.edu/nlvr/>

BAN and MUTAN models supplemented with ArticleNet (Marino et al., 2019) obtain each a small improvement (+.44 and +1.43%). This component retrieves Wikipedia articles from which it extracts an answer for each question. These models perform much worse than the LXMERT baseline, which is trained on multiple datasets. We include an LXMERT model provided by its authors (LXMERT-Paper) and one retrained with code they provide (LXMERT-GitHub). The latter uses a simplified training strategy, hence a slight discrepancy (e.g. 69.50% on VQA v2 with their model and 68.52% with the retrained one). Our model brings a clear improvement over LXMERT-github, but it does not surpass LXMERT-paper that

	OK-VQA	NLVR2
MLP	20.67	–
BAN	25.17	–
BAN + ArticleNet	25.61	–
MUTAN	26.41	–
MUTAN + ArticleNet	27.84	–
ConceptBert (Conceptual Cap.)	33.66	–
FiLM	–	52.1
VisualBERT (COCO captions)	–	67.00
UNITER (COCO, VG, Conceptual Cap., SBU)	–	79.50
LXMERT-Paper (cannot be reproduced, see text)	42.94	<u>74.50</u>
This paper: (COCO and VG captions, VQA v2, GQA, VG QA)		
LXMERT-GitHub	37.26	71.31
LXMERT-GitHub with ConceptNet	<u>39.04</u>	72.59

Table 4: Comparison with existing methods. Datasets used for pretraining are given in parentheses.

we could not reproduce.

On NLVR2, the classical method FiLM (Perez et al., 2018) expectedly performs worse than transformers pretrained on multiple datasets. The LXMERT baseline surpasses VisualBERT (Li et al., 2019), and our knowledge injection brings a small improvement. The state of the art on NLVR2 is obtained by UNITER (Chen et al., 2020) which is pretrained on a much greater amount of captioning data and uses a significantly larger architecture.

4.6 Knowledge Ablation

The main source of improvements on knowledge-demanding tasks like OK-VQA is the representation of knowledge relevant to test questions. To illustrate this claim, we create a small knowledge test with OK-VQA to examine how removing certain pieces of knowledge affects model performance. We select, by keyword search, a small subset of 19 test questions that focus on nutrition, on which our model obtains an accuracy of 91.11% vs 71.11% for the baseline. We then identify all entities related to nutrition (*e.g. health benefit, fiber, protein, vitamin, etc.*) and remove them from the knowledge base. After retraining our model with the pruned KB, the performance on nutrition questions drops markedly to 68.89%. The overall accuracy (on mostly non-nutrition-related questions) is maintained. This confirms that the withheld knowledge was indeed responsible for the high performance on related questions. The automated construction of diagnostic tests of this type could allow a quantitative evaluation and would be an interesting direction for future work.

4.7 Discussion and Limitations

Our experiments showed that KBs can expand the domain of applicability of V&L models. The tested approach proved robust in a range of settings with ConceptNet but the larger Wikidata did not fulfill our expectations. Our results suggest that the noise and ambiguities in Wikidata prevent realizing its full potential. Methods for better text-to-knowledge matching such as named entity recognition and homonym disambiguation are promising solutions to investigate.

We also identified the grounding of knowledge with visual evidence to be a limitation currently for certain tasks. Questions in OK-VQA about specific places or famous people for example require the model to recall specific facts on the basis of precise visual cues. Although such facts are stored in Wikidata, the tested model did not prove effective at recalling them. Improvements of the visual grounding could also help with visual reasoning tasks like SNLI-VE where the correct interpretation of a text input is heavily dependent on the visual input.

5 Conclusions

In this paper we described a general-purpose technique to inject external information from knowledge bases into multimodal transformers for vision-and-language tasks. The current prevailing paradigm is to pretrain large models on collections of datasets. Our experiments demonstrate that some types of commonsense and factual knowledge are not captured within these models. Knowledge bases like ConceptNet and Wikidata can fill in these deficiencies. We showed clear improvements in performance on a variety of tasks and benchmarks requiring visual and multimodal reasoning, demonstrating the versatility of the procedure.

The value of these results for future research are twofold. On the one hand, they indicate that the combination of heterogeneous sources of information is a promising way to expand the applicability of current machine learning models. On the other hand, by improving the availability of supporting knowledge, the approach opens the door to future advances on reasoning procedures that process this information. Advances on this front would lead to improved capabilities on tasks that require high-level or multi-hop reasoning, for example.

References

- Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. 2019. Fusion of detected objects in text for visual question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2131–2140.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 3, page 6.
- Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. 2017. Mutan: Multimodal tucker fusion for visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2612–2620.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. 2020. Behind the scene: Revealing the secrets of pre-trained vision-and-language models. In *Proceedings of the European Conference on Computer Vision*, pages 565–580. Springer.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *Proceedings of the European Conference on Computer Vision*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? Try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Alexis Conneau and Douwe Kiela. 2018. Senteval: An evaluation toolkit for universal sentence representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 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.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah A Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1606–1615.
- François Gardères, Maryam Ziaeeffard, Baptiste Abe-loos, and Freddy Lecue. 2020. Conceptbert: Concept-aware representation for visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 489–498.
- Travis R. Goodwin and Dina Demner-Fushman. 2019. Bridging the knowledge gap: Enhancing question answering with world and domain knowledge. *arXiv preprint arXiv:1910.07429*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913.
- MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. 2019. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6):1–36.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709.
- Sabin Kafle, Nisansa de Silva, and Dejing Dou. 2019. An overview of utilizing knowledge bases in neural networks for question answering. In *2019 IEEE 20th International Conference on Information Reuse and Integration for Data Science (IRI)*, pages 326–333.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, pages 1564–1574.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, and Others. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73.

- Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. PyTorch-BigGraph: A large-scale graph embedding system. In *Proceedings of the 2nd SysML Conference*, Palo Alto, CA, USA.
- Yoav Levine, Barak Lenz, Or Dagan, Ori Ram, Dan Padnos, Or Sharir, Shai Shalev-Shwartz, Amnon Shashua, and Yoav Shoham. 2020. Sensebert: Driving some sense into bert. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4656–4667.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11336–11344.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2822–2832.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, pages 740–755. Springer.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-bert: Enabling language representation with knowledge graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2901–2908.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*.
- Shangwen Lv, Daya Guo, Jingjing Xu, Duyu Tang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, and Songlin Hu. 2020. Graph-based reasoning over heterogeneous external knowledge for commonsense question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8449–8456.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. OK-VQA: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3195–3204.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Medhini Narasimhan and Alexander G. Schwing. 2018. Straight to the facts: Learning knowledge base retrieval for factual visual question answering. In *Proceedings of the European Conference on Computer Vision*, pages 451–468.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Matthew E Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. 2019. Knowledge enhanced contextual word representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 43–54.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 4453–4463.
- Sanket Shah, Anand Mishra, Naganand Yadati, and Partha Pratim Talukdar. 2019. KVQA: Knowledge-aware visual question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.

- Amanpreet Singh, Vedanuj Goswami, and Devi Parikh. 2020. Are we pretraining it right? Digging deeper into visio-linguistic pretraining. *arXiv preprint arXiv:2004.08744*.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Shane Storks, Qiaozi Gao, and Joyce Y Chai. 2019. Commonsense reasoning for natural language understanding: A survey of benchmarks, resources, and approaches. *arXiv preprint arXiv:1904.01172*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7464–7473.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 5103–5114.
- Damien Teney and Anton van den Hengel. 2016. Zero-shot visual question answering. *arXiv preprint arXiv:1611.05546*.
- Damien Teney and Anton van den Hengel. 2018. Visual question answering as a meta learning task. In *Proceedings of the European Conference on Computer Vision*, pages 219–235.
- Damien Teney and Anton van den Hengel. 2019. Actively seeking and learning from live data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Damien Teney, Qi Wu, and Anton van den Hengel. 2017. Visual question answering: A tutorial. *IEEE Signal Processing Magazine*, 34:63–75.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Henge. 2017a. Explicit knowledge-based reasoning for visual question answering. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1290–1296.
- Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2017b. Fvqa: Fact-based visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2019. KEPLER: A unified model for knowledge embedding and pre-trained language representation. *arXiv preprint arXiv:1911.06136*.
- Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2017. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*.
- Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2016a. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4622–4630.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016b. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018.

- Zhi-Xiu Ye, Qian Chen, Wen Wang, and Zhen-Hua Ling. 2019. Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models. *arXiv preprint arXiv:1908.06725*.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. Ernie: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4995–5004.

Reasoning over Vision and Language: Exploring the Benefits of Supplemental Knowledge

Appendices

A Implementation Details

Most transformer-based V&L models map the input text tokens to word, segment, and position embeddings, that are ultimately combined. Our approach is applied on the word embeddings.

The implementation of our method builds on top of the official implementation of LXMERT³ (Tan and Bansal, 2019). We use values of hyperparameters recommended in the code, including a reduced number of 12 training epochs (compared to 20 mentioned in the paper) and single-stage training strategy.

The loss weight λ was selected by cross-validation for maximum accuracy on the OK-VQA validation set. From the set $\{0.1, 0.5, 1, 10, 100, 200\}$, the optimal values were found to be 0.5 for Wikidata and 100 for ConceptNet. The optimal λ seems to vary in inverse proportion with the size of the KB (based on Wikipedia, ConceptNet, and three other KBs used in preliminary experiments and not included in the paper).

The fine-tuning experiments used the same batch size of 32 and learning rate of 0.00005 as the original LXMERT implementation. The number of fine-tuning epochs was adjusted according to the dataset size: 25 for OK-VQA and FVQA, 8 for NLVR2 and SNLI-VE.

The pretraining experiment used two NVIDIA Tesla M40 GPUs and took about four days. To pretrain the model on the full five datasets we used four GPUs and each experiment took up to eight days. The fine-tuning experiments used a single GPU and took about 1h for FVQA, 3h for OK-VQA, 10h for NLVR2, and 22h for SNLI-VE.

Since the matching of KB entities with the V&L training data proceeds by exact string matching, we processed the KB entities with the same WordPiece tokenizer (Wu et al., 2016b) as the LXMERT model does for the V&L training data. To enable fast indexing of the KB, we store the knowledge embeddings V as a hash table indexed by textual expressions W .

³<https://github.com/airsplay/lxmert>

B Negative Results

We explored a variety of architectural choices, but some of them did not gain any improvement over the baseline. To measure the distance between learned projections and their embeddings in Eq. 2 we tried two other options: smooth L1 loss and cosine distance, but they showed worse results than the used mean squared error. To obtain representations for knowledge-rich expressions we first included additional knowledge embedding layer as a fourth component of textual representations along with word, segment and position embeddings. This technique proved to be inefficient, so eventually we used word embeddings as a target for knowledge alignment.

Further, we tried to leverage information about objects detected in the image. We used predicted labels for each image object as a knowledge-rich expression and added another learning objective to align visual embeddings with corresponding knowledge. With this objective the model converged to a smaller loss during pretraining but the accuracy on all fine-tuning tasks was lower than the baseline results thereby indicating possible overfitting. We think that a better way to exploit visual information would be to use object labels as additional supervision to enhance entity matching.

C Linguistic Probing Analysis

We analyzed the representations learned by our model with probing tasks using the SentEval toolkit (Conneau and Kiela, 2018). This complements the evaluation on downstream tasks presented in the main paper.

The probing tasks aim to identify the linguistic information encoded in the learned representations. Following (Cao et al., 2020), we tested our model on the following ten linguistic probing tasks (Conneau et al., 2018):

- **SentLen:** predict the length of sentences;
- **WC:** word content task requires to predict which words appear in the sentence;
- **TreeDepth:** categorize the sentence according to the depth of its syntax tree;
- **TopConst:** predict the sequence of top constituents;
- **BShift:** check whether two random adjusted words have been inverted in the sentence;

- **Tense**: predict the tense of the main-clause verb;
- **SubjNum**: predict if the main-clause subject is in singular or plural form;
- **ObjNum**: same as SubjNum but for the main-clause object;
- **SOMO**: semantic odd man count task is to determine if a random noun or verb has been replaced.
- **CoordInv**: check whether two coordinate clauses have been inverted.

These tasks are designed to evaluate the quality of sentence embeddings, but LXMERT model learns separate embeddings for each token. To obtain sentence representations, we thus take outputs of 9 intermediate layers in the language encoder and average each of them across all tokens. Every layer output is used as a separate embedding and we report the best result across all layers (Table 5). Since the first 9 layers perform attention over textual modality only, no image input is required.

The results show that our model surpasses the baseline in most of the tasks. The highest gain (+9.36%) is seen in WC category meaning that our model better captures the content of words in a sentence. Importantly, best results for this task are obtained with the first layer outputs. WC accuracy drops for every subsequent layer output til it reach 16.11% and 18.65% for the baseline and our model respectively. It implies that information about individual words is well encoded in early layers but gets washed off with further self-attention. Surprisingly, we observe a noticeable gain in accuracy for predicting tense and plurality of words (Tense, SubjNum and ObjNum). Slight improvements in BShift and CoordInv tasks may be caused by additional constrains imposed on the word order by our model where the order of tokens that constitute a knowledge-rich expression is important. Finally, SentLen and TreeDepth probes highly rely on length of sentences and since both tested models limit textual input to 20 tokens, the results may not be reliable.

D Nearest Neighbors in Embedding Space

We provide in Table 6 a random selection of words and their nearest neighbors in the space of learned embeddings. We computed these neighbors using the L2 distance. Using the cosine distance gives qualitatively similar results.

E Results on NLVR2

We examine the extended metrics on NLVR2 in Table 7. The “consistency” reflects whether the model answers a given question correctly for all related pairs of images. Our model shows a **clear improvement in consistency** over the baseline (34.65 \rightarrow 37.56). This suggests that the model can better relate a given textual input to different image contexts. We also report the accuracy on the balanced and unbalanced test sets, designed to evaluate a model’s reliance on visual biases. In the balanced set, every image pair appears twice, one with each label (true/false). A drop in performance from the standard test set (“All”) to the balanced set (“Bal.”) would indicate that the method exploits biases. Neither the baseline nor any of the described models show an undesirable reliance on image biases.

F Description of Datasets

OK-VQA (Marino et al., 2019) is an open-ended VQA dataset where all questions require some sort of outside knowledge. It comprises about 14,000 questions about images from the MS COCO (Lin et al., 2014) dataset. All questions are produced by human annotators based either on information found in Wikipedia, or commonsense knowledge and visual evidence from the images. The questions are divided into ten categories (see Table 8) according to the type of knowledge needed to answer them. OK-VQA is one of the most diverse VQA datasets currently available that requires general knowledge. We use it accordingly as a primary benchmark in this study.

FVQA (Wang et al., 2017b) is a VQA dataset that contains about 5,000 questions that probe for commonsense knowledge. The questions are produced by annotators in procedure that forces the question to involve facts found in a reference KB. Each question in the dataset is therefore associated with one specific “supporting” fact that describes a relation between concepts in the question and/or image. We do not use the annotations of these supporting-facts. FVQA provides five different training/test splits. We report results that correspond to the average across the five splits. The quality of the questions in FVQA is mediocre in comparison to OK-VQA, and it is also much smaller. We include it in this study because it previously served to evaluate other models designed to use KBs for VQA.

SNLI-VE (Xie et al., 2019) is a dataset for a vi-

	SentLen	WC	TreeDepth	TopConst	BShift	Tense	SubjNum	ObjNum	SOMO	CoordInv
Baseline	69.10 (6)	56.83 (1)	32.71 (7)	67.04 (4)	66.19 (9)	74.38 (6)	76.38 (8)	75.85 (7)	51.01 (7)	59.52 (6)
Ours	70.36 (5)	66.19 (1)	33.35 (5)	71.49 (6)	68.23 (8)	82.10 (1)	80.26 (7)	79.02 (7)	50.70 (5)	60.24 (7)

Table 5: Results on linguistic probing tasks. For each, we report the maximum score obtained across layers, with the corresponding layer number in parentheses. Our model (ConceptNet PT) outperforms the baseline (LXMERT) in most probing tasks.

Word	Nearest neighbors with the baseline	Nearest neighbors with our model
birth	##gm, dat, sensitive, incorporated	births, childbirth, born, newborn
boxers	dare, indo, briefs, aiding	boxer, briefs, underwear, panties
creeping	goalscorer, fertility, ineffective, [PAD]	crept, creep, crawling, sneaking
dependent	bacterial, [PAD], ##idad, outlaws	depended, depend, depends, dependency
displaced	neptune, roche, peterborough, norway	displacement, refugees, relocated, atletico
down	up, MASK, ##combe, ##ending	up, on, out, ##s
equity	[PAD], implementations, eurasian, newfound	investors, investments, investor, investment
ghosts	[PAD], germain, combustion, ##ignment	ghost, ghostly, haunted, phantom
indication	neptune, musee, converting, legion	indications, indicating, signaled, indicative
limb	stump, branch, limbs, thorn	limbs, branch, leg, ##wara
policemen	cowboys, firefighters, youths, 37	policeman, police, cops, constabulary
quebec	sutton, [PAD], monasteries, frederick	montreal, laval, ontario, sudbury
smells	[PAD], aiding, preston, quentin	smelled, smell, odor, scent
successes	[PAD], kilometres, tina, marne	success, achievements, accomplishments, successful
sugar	yeast, powder, memo, coating	chocolate, butter, candy, celaena
taste	feel, smell, fade, become	tastes, flavor, tasted, flavors
unmarried	[PAD], [unused285], [unused685], ##dium	divorced, childless, widowed, marrying
vintage	retro, antique, victorian, rustic	antique, retro, old, ##60

Table 6: Selection of words and four nearest neighbors in the space of word embeddings learned by the LXMERT baseline and our model. The embedding space of our model better reflects lexical and semantic similarity.

sual entailment task. The task is an extension of the classical task of natural language inference (NLI). The visual version involves an image “premise” and a text “hypothesis”. The model must determine whether the hypothesis contradicts the information shown in the image, entails it, or whether there are not enough clues to draw any conclusion. The SNLI-VE dataset contains about 560,000 instances, which were constructed from captions from SNLI (Bowman et al., 2015) paired with images from Flickr30k (Young et al., 2014). Despite

	NLVR2 Test-P			
	All	Consist.	Bal.	Unbal.
Baseline (LXMERT)	71.31	34.65	70.34	72.45
With ConceptNet PT	71.61	34.84	71.11	72.60
With ConceptNet FT	70.82	34.42	69.86	72.35
With ConceptNet PT+FT	72.59	37.56	72.12	73.71

Table 7: Detailed metrics on NLVR2: consistency (%) and accuracy (%) on balanced and unbalanced subsets. The method brings a clear improvement in consistency.

similarities in the skills required for this task and for VQA, existing VQA models show relatively poor performance on SNLI-VE. The authors of the dataset attribute it to the need for more fine-grained visual understanding and reasoning, and they suggested the use of external knowledge to improve performance, hence its inclusion in this study.

NLVR2 (Suhr et al., 2019) is a dataset that evaluates visual reasoning over pairs of images. Each of the $\sim 107,000$ instances in the dataset consists of two images with a statement in natural language. The model must predict whether the statement accurately describes the pair of images. The creation of the dataset emphasized the linguistic diversity of the sentences, with the objective for the task to require some compositional reasoning. We use this task in our study to evaluate the suitability of our model to perform compositional reasoning on the knowledge injected in them from KBs.

Alias	Category	Example
VT	Vehicles and Transportation	What is the title of the person driving this vehicle ?
BCP	Brands, Companies and Products	Name the laptop model shown in this picture ?
OMC	Objects, Material and Clothing	What sort of room is this woman sitting in ?
SR	Sports and Recreation	What is this baseball player doing ?
CF	Cooking and Food	Which of the foods here have the highest saturated fats ?
GHLC	Geography, History, Language and Culture	What city is this meeting taking place ?
PEL	People and Everyday Life	What kind of hairstyle does the woman in the black shirt have ?
PA	Plants and Animals	What sound does this animal make ?
ST	Science and Technology	Which rodent has a similar name to the technical device seen ?
WC	Weather and Climate	What kind of clouds are shown ?

Table 8: Categories of questions in the OK-VQA dataset corresponding to the type of knowledge required.

	Images	Questions	Answers	Type
OK-VQA	14,031	14,055	14,456	knowledge-based VQA
FVQA	2,190	5,826	954	knowledge-based VQA
SNLI-VE	31,783	565,286	3	entailment
NLVR2	141,480	107,292	2	reasoning

Table 9: Summary of datasets used in the experiments.

G Additional Qualitative Results



What are they riding on?

Baseline: ski ✓
Proposed: ski ✓
GT: ski, snowboard



What is outside of the window?

Baseline: pane
Proposed: tree ✓
GT: bench, tree, table



What is the blue eagle on the buffet made out of?

Baseline: metal
Proposed: plastic
GT: ice



What type of business is this picture taken in?

Baseline: hotel ✓
Proposed: hotel ✓
GT: hotel



What season are these gourds typically harvested in?

Baseline: fall ✓
Proposed: summer
GT: fall




What is the man in that ad riding?

Baseline: ski
Proposed: snowboard ✓
GT: snowboard, microphone, board



How many chromosomes do these creatures have?

Baseline: 3
Proposed: 46 ✓
GT: 46, 23 pair, 23



What website is the left computer currently on?

Baseline: flickr
Proposed: flickr
GT: googl, weathervgov, yahooom, amazon

(a) OK-VQA.



What is the large object in the right of this image used for?

Baseline: tennis
Proposed: playing tennis
GT: play tennis



Where can people find fish?

Baseline: fish
Proposed: fish
GT: lakes rivers and ocean




What is the brass object in this image?

Baseline: cello
Proposed: trombone ✓
GT: trombone



What the woman is using to insert screw?

Baseline: toothbrush
Proposed: screwdriver ✓
GT: screwdriver



Which kind of outdoor recreation are shown in this image?

Baseline: skis ✓
Proposed: skis ✓
GT: skis



Which object in this image is used for sitting?

Baseline: couch
Proposed: couch
GT: sofa



Which objects in this image may be known as avians?

Baseline: boats
Proposed: bird
GT: birds



What object in this image is commonly eaten for lunch?

Baseline: sandwich ✓
Proposed: sandwich ✓
GT: sandwich

(b) FVQA.

Figure 4: Random selection of test instances from OK-VQA and FVQA datasets, with predictions of the baseline (LXMERT) and of our model (Pretraining with VQA v2, GQA w/ ConceptNet PT+FT).



A man plays violin.

Baseline: contradiction
Proposed: entailment ✓
GT: entailment



There are many children.

Baseline: entailment ✓
Proposed: entailment ✓
GT: entailment



A group of casually dressed people stand in the room.

Baseline: contradiction ✓
Proposed: contradiction ✓
GT: contradiction



The boy is wearing overalls.

Baseline: contradiction
Proposed: neutral ✓
GT: neutral



A man plays with a dog.

Baseline: entailment ✓
Proposed: contradiction ✓
GT: entailment



A cat sleeps indoors.

Baseline: contradiction ✓
Proposed: contradiction ✓
GT: contradiction



The man is cutting steak on the counter.


Baseline: contradiction ✓
Proposed: contradiction ✓
GT: contradiction



The Broadway Rite Aid captures the attention of everyone who walks by due to the good sales.

Baseline: neutral ✓
Proposed: neutral ✓
GT: neutral

(a) SNLI-VE.



An image shows a forward-facing non-standing hound with a paw on some type of toy.

Baseline: true ✓
Proposed: false
GT: true



At least one image is of a multi-serving trifle bowl.

Baseline: false ✓
Proposed: false ✓
GT: false




One of the images features exactly three musicians.

Baseline: false
Proposed: false
GT: true



An image shows exactly one fragrance bottle displayed on the right of an upright black box.

Baseline: false
Proposed: true ✓
GT: true



All the instruments are standing on their ends.

Baseline: true
Proposed: false ✓
GT: false




A beetle on top of a dungball is facing left.

Baseline: true
Proposed: false ✓
GT: false



There is a plant near the cabinet in the image on the left.

Baseline: true ✓
Proposed: true ✓
GT: true



Each sled driver is driving a group of at least five dogs.

Baseline: false ✓
Proposed: false ✓
GT: false

(b) NLVR2.

Figure 5: Random selection of test instances from SNLI-VE and NLVR2 datasets, with predictions of the baseline (LXMERT) and of our model (Pretraining with VQA v2, GQA w/ ConceptNet PT+FT).