

Straight to the Facts: Learning Knowledge Base Retrieval for Factual Visual Question Answering

Medhini Narasimhan, Alexander G. Schwing

University of Illinois Urbana-Champaign
{medhini2,aschwing}@illinois.edu

Abstract. Question answering is an important task for autonomous agents and virtual assistants alike and was shown to support the disabled in efficiently navigating an overwhelming environment. Many existing methods focus on observation-based questions, ignoring our ability to seamlessly combine observed content with general knowledge. To understand interactions with a knowledge base, a dataset has been introduced recently and keyword matching techniques were shown to yield compelling results despite being vulnerable to misconceptions due to synonyms and homographs. To address this issue, we develop a learning-based approach which goes straight to the facts via a learned embedding space. We demonstrate state-of-the-art results on the challenging recently introduced fact-based visual question answering dataset, outperforming competing methods by more than 5%.

Keywords: fact based visual question answering, knowledge bases

1 Introduction

When answering questions given a context, such as an image, we seamlessly combine the observed content with general knowledge. For autonomous agents and virtual assistants which naturally participate in our day to day endeavors, where answering of questions based on context and general knowledge is most natural, algorithms which leverage both observed content and general knowledge are extremely useful.

To address this challenge, in recent years, a significant amount of research has been devoted to question answering in general and Visual Question Answering (VQA) in particular. Specifically, the classical VQA tasks require an algorithm to answer a given question based on the additionally provided context, given in the form of an image. For instance, significant progress in VQA was achieved by introducing a variety of VQA datasets with strong baselines [1–8]. The images in these datasets cover a broad range of categories and the questions are designed to test perceptual abilities such as counting, inferring spatial relationships, and identifying visual cues. Some challenging questions require logical reasoning and memorization capabilities. However, the majority of the questions can be answered by solely examining the visual content of the image. Hence, numerous

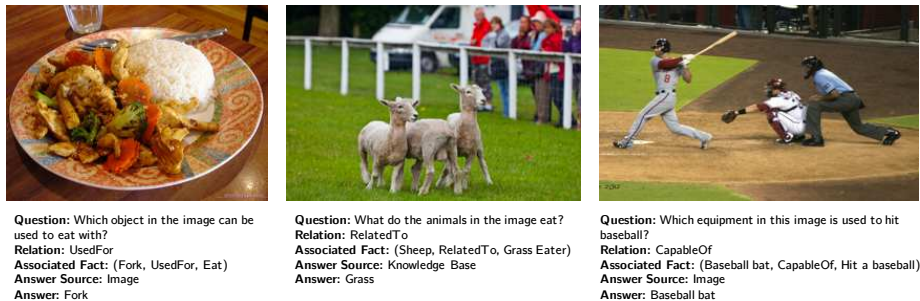


Fig. 1. The FVQA dataset expects methods to answer questions about images utilizing information from the image, as well as fact-based knowledge bases. Our method makes use of the image, and question text features, as well as high-level visual concepts extracted from the image in combination with a learned fact-ranking neural network. Our method is able to answer both visually grounded as well as fact based questions.

approaches to solve these problems [7–13] focus on extracting visual cues using deep networks.

We note that many of the aforementioned methods focus on the visual aspect of the question answering task, *i.e.*, the answer is predicted by combining representations of the question and the image. This clearly contrasts the described human-like approach, which combines observations with general knowledge. To address this discrepancy, in very recent meticulous work, Wang *et al.* [14] introduced a ‘fact-based’ VQA task (FVQA), an accompanying dataset, and a knowledge base of facts extracted from three different sources, namely WebChild [15], DBPedia [16], and ConceptNet [17]. Different from the classical VQA datasets, Wang *et al.* [14] argued that such a dataset can be used to develop algorithms which answer more complex questions that require a combination of observation and general knowledge. In addition to the dataset, Wang *et al.* [14] also developed a model which leverages the information present in the supporting facts to answer questions about an image.

To this end, Wang *et al.* [14] design an approach which extracts keywords from the question and retrieves facts that contain those keywords from the knowledge base. Clearly, synonyms and homographs pose challenges which are hard to recover from.

To address this issue, we develop a learning based retrieval method. More specifically, our approach learns a parametric mapping of facts and question-image pairs to an embedding space. To answer a question, we use the fact that is most aligned with the provided question-image pair. As illustrated in Fig. 1, our approach is able to accurately answer both more visual questions as well as more fact based questions. For instance, given the image illustrated on the left hand side along with the question, “Which object in the image can be used to eat with?”, we are able to predict the correct answer, “fork.” Similarly, the proposed approach is able to predict the correct answer for the other two exam-

ples. Quantitatively we demonstrate the efficacy of the proposed approach on the recently introduced FVQA dataset, outperforming state-of-the-art by more than 5% on the top-1 accuracy metric.

2 Related Work

We develop a framework for visual question answering that benefits from a rich knowledge base. In the following, we first review classical visual question answering tasks before discussing visual question answering methods that take advantage of knowledge bases.

Visual Question Answering. In recent years, a significant amount of research has been devoted to developing techniques which can answer a question about a provided context such as an image. Of late, visual question answering has also been used to assess reasoning capabilities of state-of-the-art predictors. Using a variety of datasets [11, 2, 8, 10, 3, 5], models based on multi-modal representation and attention [18–25], deep network architectures [26, 12, 27, 28], and dynamic memory nets [29] have been developed. Despite these efforts, assessing the reasoning capabilities of present day deep network-based approaches and differentiating them from mere memorization of training set statistics remains a hard task. Most of the methods developed for visual question answering [2, 8, 10, 18–24, 12, 27, 29–31, 6, 7, 32–34] focus exclusively on answering questions related to observed content. To this end, these methods use image features extracted from networks such as the VGG-16 [35] trained on large image datasets such as ImageNet [36]. However, it is unlikely that all the information which is required to answer a question is encoded in the features extracted from the image, or even the image itself. For example, consider an image containing a dog, and a question about this image, such as “*Is the animal in the image capable of jumping in the air?*”. In such a case, we would want our method to combine common sense and general knowledge about the world, such as the ability of a healthy dog to jump, along with features and observations from the image, such as the presence of the dog. This motivates us to develop methods that can use knowledge bases encoding general knowledge.

Knowledge-based Visual Question Answering. There has been interest in the natural language processing community in answering questions based on knowledge bases (KBs) using either semantic parsing [37–47] or information retrieval [48–54] methods. However, knowledge based visual question answering is still relatively unexplored, even though this is appealing from a practical standpoint as this decouples the reasoning by the neural network from the storage of knowledge in the KB. Notable examples in this direction are work by Zhu *et al.* [55], Wu *et al.* [56], Wang *et al.* [57], Krishnamurthy and Kollar [58], and Narasimhan *et al.* [59].

The works most related to our approach include *Ask Me Anything* (AMA) by Wu *et al.* [60], Ahab by Wang *et al.* [61], and FVQA by Wang *et al.* [14]. AMA describes the content of an image in terms of a set of attributes predicted about

the image, and multiple captions generated about the image. The predicted attributes are used to query an external knowledge base, DBpedia [16], and the retrieved paragraphs are summarized to form a knowledge vector. The predicted attribute vector, the captions, and the database-based knowledge vector are passed as inputs to an LSTM that learns to predict the answer to the input question as a sequence of words. A drawback of this work is that it does not perform any explicit reasoning and ignores the possible structure in the KB. Ahab and FVQA, on the other hand, attempt to perform explicit reasoning. Ahab converts an input question into a database query, and processes the returned knowledge to form the final answer. Similarly, FVQA learns a mapping from questions to database queries through classifying questions into categories and extracting parts from the question deemed to be important. While both of these methods rely on fixed query templates, this very structure offers some insight into what information the method deems necessary to answer a question about a given image. Both these methods use databases with a particular structure: those that contain facts about visual concepts represented as tuples, for example, $(Cat, CapableOf, Climbing)$, and (Dog, IsA, Pet) . We develop our method on the dataset released as part of the FVQA work, referred to as the FVQA dataset [14], which is a subset of three structured databases – DBpedia [16], ConceptNet [17], and WebChild [15]. The method presented in FVQA [14] produces a query as an output of an LSTM which is fed the question as an input. Facts in the knowledge base are filtered on the basis of visual concepts such as objects, scenes, and actions extracted from the input image. The predicted query is then applied on the filtered database, resulting in a set of retrieved facts. A matching score is then computed between the retrieved facts and the question to determine the most relevant fact. The most correct fact forms the basis of the answer for the question.

In contrast to Ahab and FVQA, we propose to directly learn an embedding of facts and question-image pairs into a space that permits to assess their compatibility. This has two important advantages over prior work: 1) by avoiding the generation of an explicit query, we eliminate errors due to synonyms, homographs, and incorrect prediction of visual concept type and answer type; and 2) our technique is easy to extend to any knowledge base, even one with a different structure or size. We also do not require any ad-hoc filtering of knowledge, and can instead learn to transform extracted visual concepts into a vector close to a relevant fact in the learned embedding space. Our method also naturally produces a ranking of facts deemed to be useful for the given question and image.

3 Learning Knowledge Base Retrieval

In the following, we first provide an overview of the proposed approach for knowledge based visual question answering before discussing our embedding space and learning formulation.

Overview. Our developed approach is outlined in Fig. 2. The task at hand is to predict an answer y for a question Q given an image x by using an external knowl-

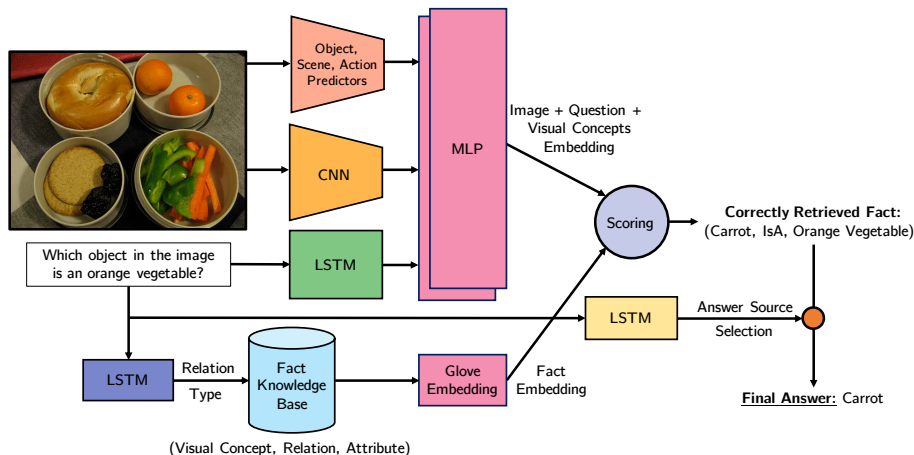


Fig. 2. Overview of the proposed approach. Given an image and a question about the image, we obtain an Image + Question Embedding through the use of a CNN on the image, an LSTM on the question, and a Multi Layer Perceptron (MLP) for combining the two modalities. In order to filter relevant facts from the Knowledge Base (KB), we use another LSTM to predict the fact relation type from the question. The retrieved structured facts are encoded using GloVe embeddings. The retrieved facts are ranked through a dot product between the embedding vectors and the top-ranked fact is returned to answer the question.

edge base KB, which consists of a set of facts f_i , *i.e.*, $\text{KB} = \{f_1, \dots, f_{|\text{KB}|}\}$. Each fact f_i in the knowledge base is represented as a Resource Description Framework (RDF) triplet of the form $f_i = (a_i, r_i, b_i)$, where a_i is a visual concept in the image, b_i is an attribute or phrase associated with the visual entity a_i , and $r_i \in \mathcal{R}$ is a relation between the two entities. The dataset contains $|\mathcal{R}| = 13$ relations $r \in \mathcal{R} = \{\text{Category}, \text{Comparative}, \text{HasA}, \text{IsA}, \text{HasProperty}, \text{CapableOf}, \text{Desires}, \text{RelatedTo}, \text{AtLocation}, \text{PartOf}, \text{ReceivesAction}, \text{UsedFor}, \text{CreatedBy}\}$. Example triples of the knowledge base in our dataset are $(\text{Umbrella}, \text{UsedFor}, \text{Shade})$, $(\text{Beach}, \text{HasProperty}, \text{Sandy})$, $(\text{Elephant}, \text{Comparative-LargerThan}, \text{Ant})$.

To answer a question Q correctly given an image x , we need to retrieve the right supporting fact and choose the correct entity, *i.e.*, either a or b . Importantly, entity a is always derived from the image and entity b is derived from the fact base. Consequently we refer to this choice as the answer source $s \in \{\text{Image}, \text{KnowledgeBase}\}$. Using this formulation, we can extract the answer y from a predicted fact $\hat{f} = (\hat{a}, \hat{r}, \hat{b})$ and a predicted answer source \hat{s} using

$$y = \begin{cases} \hat{a}, & \text{from } \hat{f} \text{ if } \hat{s} = \text{Image} \\ \hat{b}, & \text{from } \hat{f} \text{ if } \hat{s} = \text{KnowledgeBase} \end{cases}. \quad (1)$$

It remains to answer, how to predict a fact \hat{f} and how to infer the answer source \hat{s} . The latter is a binary prediction task and we describe our approach

below. For the former, we note that the knowledge base contains a large number of facts. We therefore consider it infeasible to search through all the facts $f_i \forall i \in \{1, \dots, |\text{KB}|\}$ using an expensive evaluation based on a deep net. We therefore split this task into two parts: (1) Given a question, we train a network to predict the relation \hat{r} , that the question focuses on. (2) Using the predicted relation, \hat{r} , we reduce the fact space to those containing only the predicted relation.

Subsequently, to answer the question Q given image x , we only assess the suitability of the facts which contain the predicted relation \hat{r} . To assess the suitability, we design a score function $S(g^{\text{F}}(f_i), g^{\text{NN}}(x, Q))$ which measures the compatibility of a fact representation $g^{\text{F}}(f_i)$ and an image-question representation $g^{\text{NN}}(x, Q)$. Intuitively, the higher the score, the more suitable the fact f_i for answering question Q given image x .

Formally, we hence obtain the predicted fact \hat{f} via

$$\hat{f} = \arg \max_{i \in \{j : \text{rel}(f_j) = \hat{r}\}} S(g^{\text{F}}(f_i), g^{\text{NN}}(x, Q)), \quad (2)$$

where we search for the fact \hat{f} maximizing the score S among all facts f_i which contain relation \hat{r} , *i.e.*, among all f_i with $i \in \{j : \text{rel}(f_j) = \hat{r}\}$. Hereby we use the operator $\text{rel}(f_i)$ to indicate the relation of the fact triplet f_i . Given the predicted fact using Eq. (2) we obtain the answer y from Eq. (1) after predicting the answer source \hat{s} .

This approach is outlined in Fig. 2. Pictorially, we illustrate the construction of an image-question embedding $g^{\text{NN}}(x, Q)$, via LSTM and CNN net representations that are combined via an MLP. We also illustrate the fact embedding $g^{\text{F}}(f_i)$. Both of them are combined using the score function $S(\cdot, \cdot)$, to predict a fact \hat{f} from which we extract the answer as described in Eq. (1).

In the following, we first provide details about the score function S , before discussing prediction of the relation \hat{r} and prediction of the answer source \hat{s} .

Scoring the facts. Fig. 2 illustrates our approach to score the facts in the knowledge base, *i.e.*, to compute $S(g^{\text{F}}(f_i), g^{\text{NN}}(x, Q))$. We obtain the score in three steps: (1) computing of a fact representation $g^{\text{F}}(f_i)$; (2) computing of an image-question representation $g^{\text{NN}}(x, Q)$; (3) combination of the fact and image-question representation to obtain the final score S . We discuss each of those steps in the following.

(1) *Computing a fact representation.* To obtain the fact representation $g^{\text{F}}(f_i)$, we concatenate two vectors, the averaged GloVe-100 [62] representation of the words of entity a_i and the averaged GloVe-100 representation of the words of entity b_i . Note that this fact representation is non-parametric, *i.e.*, there are no trainable parameters.

(2) *Computing an image-question representation.* We compute the image-question representation $g^{\text{NN}}(x, Q)$, by combining a visual representation $g_w^{\text{V}}(x)$, obtained from a standard deep net, *e.g.*, ResNet or VGG, with a visual concept representation $g_w^{\text{C}}(x)$, and a sentence representation $g_w^{\text{Q}}(Q)$, of the question Q , obtained using a trainable recurrent net. For notational convenience we concatenate all trainable parameters into one vector w . Making the dependence

on the parameters explicit, we obtain the image-question representation via $g_w^{\text{NN}}(x, Q) = g_w^{\text{NN}}(g_w^V(x), g_w^Q(Q), g_w^C(x))$.

More specifically, for the question embedding $g_w^Q(Q)$, we use an LSTM model [63]. For the image embedding $g_w^V(x)$, we extract image features using ResNet-152 [64] pre-trained on the ImageNet dataset [65]. In addition, we also extract a visual concept representation $g_w^C(x)$, which is a multi-hot vector of size 1176 indicating the visual concepts which are grounded in the image. The visual concepts detected in the images are objects, scenes, and actions. For *objects*, we use the detections from two Faster-RCNN [66] models that are trained on the Microsoft COCO 80-object [67] and the ImageNet 200-object [36] datasets. In total, there are 234 distinct object classes, from which we use that subset of labels that coincides with the FVQA dataset. The *scene* information (such as pasture, beach, bedroom) is extracted by the VGG-16 model [35] trained on the MIT Places 365-class dataset [68]. Again, we use a subset of Places to construct the 1176-dimensional multi-hot vector $g_w^C(x)$. For detecting *actions*, we use the CNN model proposed in [69] which is trained on the HICO [70] and MPII [71] datasets. The HICO dataset contains labels for 600 human-object interaction activities while the MPII dataset contains labels for 393 actions. We use a subset of actions, namely those which coincide with the ones in the FVQA dataset.

All the three vectors $g_w^V(x), g_w^Q(Q), g_w^C(x)$ are concatenated and passed to the multi-layer perceptron $g_w^{\text{NN}}(\cdot, \cdot, \cdot)$.

(3) *Combination of fact and image-question representation.* For each fact representation $g^{\text{F}}(f_i)$, we compute a score

$$S_w(g^{\text{F}}(f_i), g_w^{\text{NN}}(x, Q)) = \cos(g^{\text{F}}(f_i), g_w^{\text{NN}}(x, Q)) = \frac{g^{\text{F}}(f_i) \cdot g_w^{\text{NN}}(x, Q)}{\|g^{\text{F}}(f_i)\| \cdot \|g_w^{\text{NN}}(x, Q)\|},$$

where $g_w^{\text{NN}}(x, Q)$ is the image question representation. Hence, the score S is the cosine similarity between the two normalized representations and represents the fit of fact f_i to the image-question pair (x, Q) .

Predicting the relation. To predict the relation $\hat{r} \in \mathcal{R} = h_{w_1}^r(Q)$, from the obtained question Q , we use an LSTM net. More specifically, we first embed and then encode the words of the question Q , one at a time, and linearly transform the final hidden representation of the LSTM to predict \hat{r} , from $|\mathcal{R}|$ possibilities using a standard multinomial classification. For the results presented in this work, we trained the relation prediction parameters w_1 independently of the score function. We leave a joint formulation to future work.

Predicting the answer source. Prediction of the answer source $\hat{s} = h_{w_2}^s(Q)$ from a given question Q is similar to relation prediction. Again, we use an LSTM net to embed and encode the words of the question Q before linearly transforming the final hidden representation to predict $\hat{s} \in \{\text{Image}, \text{KnowledgeBase}\}$. Analogous to relation prediction, we train this LSTM net’s parameters w_2 separately and leave a joint formulation to future work.

Learning. As mentioned before, we train the parameters w (score function), w_1 (relation prediction), and w_2 (answer source prediction) separately. To train w_1 , we use a dataset $\mathcal{D}_1 = \{(Q, r)\}$ containing pairs of question and the corre-

Algorithm 1 Training with hard negative mining**Input:** $(x, Q, f^*), KB$ **Output:** parameters w

-
- 1: **for** $t = 0, \dots, T$ **do**
 - 2: Create dataset $\mathcal{D}^{(t)}$ by sampling negative facts randomly (if $t = 0$) or by retrieving facts predicted wrongly with $w^{(t-1)}$ (if $t > 0$)
 - 3: Use $\mathcal{D}^{(t)}$ to obtain $w^{(t)}$ by optimizing the program given in Eq. (7)
 - 4: **end for**
 - 5: **return** $w^{(T)}$
-

sponding relation which was used to obtain the answer. To learn w_2 , we use a dataset $\mathcal{D}_2 = \{(Q, s)\}$, containing pairs of question and the corresponding answer source. For both classifiers we use stochastic gradient descent on the classical cross-entropy and binary cross-entropy loss respectively. Note that both the datasets are readily available from [14].

To train the parameters of the score function we adopt a successive approach operating in time steps $t = \{1, \dots, T\}$. In each time step, we gradually increase the difficulty of the dataset $\mathcal{D}^{(t)}$ by mining hard negatives. More specifically, for every question Q , and image x , $\mathcal{D}^{(0)}$ contains the ‘groundtruth’ fact f^* as well as 99 randomly sampled ‘non-groundtruth’ facts. After having trained the score function on this dataset we use it to predict facts for image-question pairs and create a new dataset $\mathcal{D}^{(1)}$ which now contains, along with the groundtruth fact, another 99 non-groundtruth facts that the score function assigned a high score to.

Given a dataset $\mathcal{D}^{(t)}$, we train the parameters w of the representations involved in the score function $S_w(g^F(f_i), g_w^{\text{NN}}(x, Q))$, and its image, question, and concept embeddings by encouraging that the score of the groundtruth fact f^* is larger than the score of any other fact. More formally, we aim for parameters w which ensure the classical margin, *i.e.*, an SVM-like loss for deep nets:

$$S_w(f^*, x, Q) \geq L(f^*, f) + S_w(f, x, Q) \quad \forall (f, x, Q) \in \mathcal{D}^{(t)}, \quad (3)$$

where $L(f^*, f)$ is the task loss (aka margin) comparing the groundtruth fact f^* to other facts f . In our case $L \equiv 1$. Since we may not find parameters w which ensure feasibility $\forall (f, x, Q) \in \mathcal{D}^{(t)}$, we introduce slack variables $\xi_{(f,x,Q)} \geq 0$ to obtain after reformulation:

$$\xi_{(f,x,Q)} \geq L(f^*, f) + S_w(f, x, Q) - S_w(f^*, x, Q) \quad \forall (f, x, Q) \in \mathcal{D}^{(t)}. \quad (4)$$

Instead of enforcing the constraint $\forall (f, x, Q)$ in the dataset $\mathcal{D}^{(t)}$, it is equivalent to require [72]

$$\xi_{(x,Q)} \geq \max_f \{L(f^*, f) + S_w(f, x, Q)\} - S_w(f^*, x, Q) \quad \forall (x, Q) \in \mathcal{D}^{(t)}. \quad (5)$$

Using this constraint, we find the parameters w by solving

$$\min_{w, \xi_{(x,Q)} \geq 0} \frac{C}{2} \|w\|_2^2 + \sum_{(x,Q) \in \mathcal{D}^{(t)}} \xi_{(x,Q)} \quad \text{s.t. Constraints in Eq. (5)}. \quad (6)$$

Method	Accuracy	
	@1	@3
FVQA [14]	64.94	82.42
Ours	75.4	91.97

Table 1. Accuracy of predicting relations given the question.

Method	Accuracy	
	@1	@3
Ours	97.3	100.00

Table 2. Accuracy of predicting answer source from a given question.

For applicability of the standard sub-gradient descent techniques, we reformulate the program given in Eq. (6) to read as

$$\min_w \frac{C}{2} \|w\|_2^2 + \sum_{(x,Q) \in \mathcal{D}^{(t)}} \left(\max_f \{L(f^*, f) + S_w(f, x, Q)\} - S_w(f^*, x, Q) \right), \quad (7)$$

which can be optimized using standard deep net packages. The proposed approach for learning the parameters w is summarized in Alg. 1. In the following we now assess the suitability of the proposed approach.

4 Evaluation

In the following, we assess the proposed approach. We first provide details about the proposed dataset before presenting quantitative results for prediction of relations from questions, prediction of answer-source from questions, and prediction of the answer and the supporting fact. We also discuss mining of hard negatives. Finally, we show qualitative results.

Dataset and Knowledge Base. We use the publicly available FVQA dataset [14] and its knowledge base to evaluate our model. This dataset consists of 2,190 images, 5,286 questions, and 4,126 unique facts corresponding to the questions. The knowledge base, consisting of 193,449 facts, were constructed by extracting the top visual concepts for all the images in the dataset and querying for those concepts in the three knowledge bases, WebChild [15], ConceptNet [17], and DBPedia [16]. The dataset consists of 5 train-test folds, and all the scores we report are averaged across all splits.

Predicting Relations from Questions. We use an LSTM architecture as discussed in Sec. 3 to predict the relation $r \in \mathcal{R}$ given a question Q . The standard train-test split of the FVQA dataset is used to evaluate our model. Batch gradient descent with Adam optimizer was used on batches of size 100 and the model was trained over 50 epochs. LSTM embedding and word embeddings are of size 128 each. The learning rate is set to 1e-3 and a dropout of 0.7 is applied after the word embeddings as well as the LSTM embedding. Table 1 provides a comparison of our model to the FVQA baseline [14] using top-1 and top-3 prediction accuracy. We observe our results to improve the baseline by more than 10% on top-1 accuracy and by more than 9% when using the top-3 accuracy metric.

Predicting Answer Source from Questions. We assess the accuracy of predicting the answer source s given a question Q . To predict the source of the answer, we use an LSTM architecture as discussed in detail in Sec. 3. Note that for predicting the answer source, the size of the LSTM embedding and word embeddings was set to 64 each. Table 2 summarizes the accuracy of the prediction results of our model. We observe the prediction accuracy of the proposed approach to be close to perfect.

Predicting the Correct Answer. Our score function based model to retrieve the supporting fact is described in detail in Sec. 3. For the image embedding, we pass the 2048 dimensional feature vector returned by ResNet through a fully-connected layer and reduce it to a 64 dimensional vector. For the question embedding, we use an LSTM with a hidden layer of size 128. The two are then concatenated into a vector of size 192 and passed through a two layer perceptron with 256 and 128 nodes respectively. Note that the baseline doesn't use image features apart from the detected visual concepts.

The multi-hot visual concept embedding is passed through a fully-connected layer to form a 128 dimensional vector. This is then concatenated with the output of the perceptron and passed through another layer with 200 output nodes. We found a late fusion of the visual concepts to results in a better model as the facts explicitly contain these terms.

Fact embeddings are constructed using GloVe-100 vectors each, for entities a and b . If a or b contain multiple words, an average of all the embeddings is computed. We use cosine distance between the MLP and the fact embeddings to score the facts. The highest scoring fact is chosen as the answer. Ties are broken randomly.

Based on the answer source prediction which is computed using the aforementioned LSTM model, we choose either entity a or b of the fact to be the answer. See Eq. (1) for the formal description. Accuracy is computed based on exact match between the chosen entity and the groundtruth answer.

To assess the importance of particular features we investigate 5 variants of our model with varying features: two oracle approaches 'gt Question + Image + Visual Concepts' and 'gt Question + Visual Concepts' which make use of groundtruth relation type and answer type data. More specifically, 'gt Question + Image + Visual Concepts' and 'gt Question + Visual Concepts' use the groundtruth relations and answer sources respectively. We have three approaches using a variety of features as follows: 'Question + Image + Visual Concepts,' 'Question + Visual Concepts,' and 'Question + Image.' We drop either the Image embeddings from ResNet or the Visual Concept embeddings to obtain two other models, 'Question + Visual Concepts' and 'Question + Image.'

Table 3 shows the accuracy of our model in predicting an answer and compares our results to other FVQA baselines. We observe the proposed approach to outperform the state-of-the-art ensemble technique by more than 3% and the strongest baseline without ensemble by over 5% on the top-1 accuracy metric. Moreover we note the importance of visual concepts to accurately predict the answer. By including groundtruth information we assess the maximally possible

Method	Accuracy	
	@1	@3
LSTM-Question+Image+Pre-VQA [14]	24.98	40.40
Hie-Question+Image+Pre-VQA [14]	43.14	59.44
FVQA [14]	56.91	64.65
Ensemble [14]	58.76	-
Ours - Question + Image	26.68	30.27
Ours - Question + Image + Visual Concepts	60.30	73.10
Ours - Question + Visual Concepts	62.20	75.60
Ours - <i>gt</i> Question + Image + Visual Concepts	69.12	80.25
Ours - <i>gt</i> Question + Visual Concepts	70.34	82.12

Table 3. Answer accuracy over the FVQA dataset.

Method	Accuracy	
	@1	@3
FVQA-top-1 [14]	38.76	42.96
FVQA-top-3 [14]	41.12	45.49
Ours - Question + Image	28.98	32.34
Ours - Question + Image + Visual Concepts	62.30	74.90
Ours - Question + Visual Concepts	64.50	76.20

Table 4. Correct fact prediction precision over the FVQA dataset.

Iteration	# Hard Negatives	Precision	
		@1	@3
1	0	20.17	23.46
2	84,563	38.65	45.49
3	6,889	64.5	76.2

Table 5. Correct fact prediction precision with hard negative mining.

top-1 and top-3 accuracy. We observe the difference to be around 8%, suggesting that there is some room for improvement.

Question to Supporting Fact. To provide a complete assessment of the proposed approach we illustrate in Table 4 the top-1 and top-3 accuracy scores in retrieving the supporting facts of our model compared to other FVQA baselines. We observe the proposed approach to improve significantly both the top-1 and top-3 accuracy by more than 20%. We think this is a significant improvement towards efficiently including knowledge bases into visual question answering.

Mining Hard Negatives. We trained our model over three iterations of hard negative mining, *i.e.*, $T = 2$. In iteration 1 ($t = 0$), all the 193,449 facts were used to sample the 99 negative facts during train. At every 10th epoch of training, negative facts which received high scores were saved. In the next iteration, the

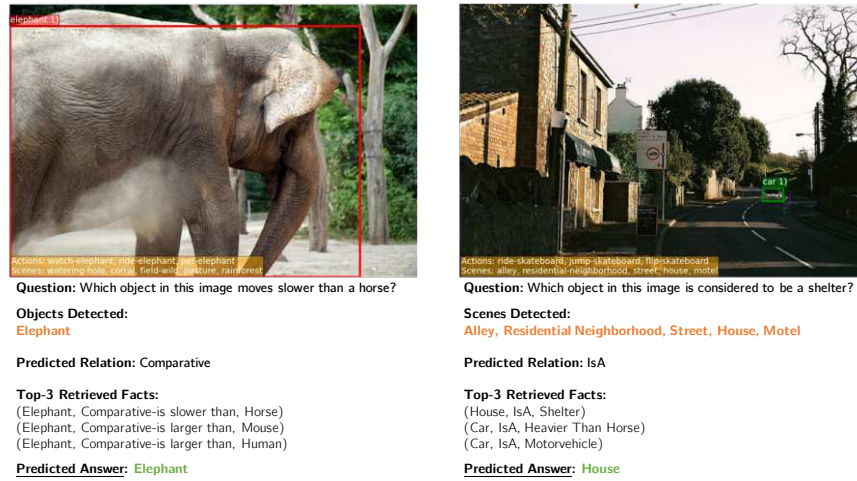


Fig. 3. Examples of Visual Concepts (VCs) detected by our framework. Here, we show examples of detected objects, scenes, and actions predicted by the various networks used in our pipeline. There is a clear alignment between useful facts, and the predicted VCs. As a result, including VCs in our scoring method helps improve performance.

trained model along with the negative facts is loaded and we ensure that the 99 negative facts are now sampled from the hard negatives. Table 5 shows the Top-1 and Top-3 accuracy for predicting the supporting facts over each of the three iterations. We observe significant improvements due to the proposed hard negative mining strategy. While naïve training of the proposed approach yields only 20.17% top-1 accuracy, two iterations improve the performance to 64.5%.

Synonyms and Homographs. Here we show the improvements of our model compared to the baseline with respect to synonyms and homographs. To this end, we run additional tests using Wordnet to determine the number of question-fact pairs which contain synonyms. The test data contains 1105 such pairs out of which our model predicts 91.6% (1012) correctly, whereas the FVQA model predicts 78.0% (862) correctly. In addition, we manually generated 100 synonymous questions by replacing words in the questions with synonyms (*e.g.*, “What in the bowl can you eat?” is rephrased to “What in the bowl is edible?”). Tests on these 100 new samples find that our model predicts 89 of these correctly, whereas the key-word matching FVQA technique [14] gets 61 of these right. With regards to homographs, the test set has 998 questions which contain words that have multiple meanings across facts. Our model predicts correct answers for 79.4% (792), whereas the FVQA model gets 66.3% (662) correct.

Qualitative Results. Fig. 3 shows the Visual Concepts (VCs) detected for a few samples along with the top 3 facts retrieved by our model. Providing these predicted VCs as input to our fact-scoring MLP helps improve supporting fact retrieval as well as answer accuracy by a large margin of over 30% as seen in Tables 3 and 4. As can be seen in Fig. 3, there is a close alignment between



Question: What is a bookshelf used for?

Predicted Relation: UsedFor
Predicted Supporting Fact:
 (Bookshelf, UsedFor, Carrying Books)
Predicted Answer Source: KB

Predicted Answer: Carrying books
GT Answer: Carrying books



Question: What object in this image is capable of flying?

Predicted Relation: CapableOf
Predicted Supporting Fact:
 (Frisbee, CapableOf, Flying)
Predicted Answer Source: Image

Predicted Answer: Frisbee
GT Answer: Frisbee



Question: Which property does the place in the image have?

Predicted Relation: HasProperty
Predicted Supporting Fact:
 (Beach, HasProperty, Sandy)
Predicted Answer Source: KB

Predicted Answer: Sandy
GT Answer: Sandy



Question: What object in this image is cheaper than a taxi?

Predicted Relation: Comparative
Predicted Supporting Fact:
 (Bus, Comparative-cheaper, Taxi)
Predicted Answer Source: Image

Predicted Answer: Bus
GT Answer: Bus



Question: Which kind of food is sweet in this image?

Predicted Relation: HasProperty
Predicted Supporting Fact:
 (Cake, HasProperty, Sweet)
Predicted Answer Source: Image

Predicted Answer: Cake
GT Answer: Cake



Question: What is an item of office equipment in this image?

Predicted Relation: Category
Predicted Supporting Fact:
 (Monitor, Category, Office Equipment)
Predicted Answer Source: Image

Predicted Answer: Monitor
GT Answer: Monitor



Question: What object in this image is round?

Predicted Relation: HasProperty
Predicted Supporting Fact:
 (Person, HasProperty, Alive)
GT Supporting Fact:
 (TennisBall, HasProperty, Round)

Predicted Answer Source: Image
GT Answer Source: Image

Predicted Answer: Person
GT Answer: TennisBall



Question: Which action is less strenuous than the action in the image?

Predicted Relation: Comparative
Predicted Supporting Fact:
 (Jumping, Comparative-more strenuous, Dressage)

Predicted Answer Source: Image
GT Answer Source: KB

Predicted Answer: Jumping
GT Answer: Dressage



Question: What sort of food can you see in the image?

Predicted Relation: IsA
GT Relation: Category

Predicted Supporting Fact: (Lemon, isA, Fruit)
GT Supporting Fact: (Fruits, Category, Food)

Predicted Answer Source: Image

Predicted Answer: Lemon
GT Answer: Fruits

Fig. 4. Success and failure cases of our method. In the top two rows, our method correctly predicts the relation, the supporting fact, and the answer source to produce the correct answer for the given question. The bottom row of examples shows the failure modes of our method.

relevant facts and predicted VCs, as VCs provide a high-level overview of the salient content in the images.

In Fig. 4, we show success and failure cases of our method. There are 3 steps to producing the correct answer using our method: (1) correctly predicting the relation, (2) retrieving supporting facts containing the predicted relation, and relevant to the image, and (3) choosing the answer from the predicted answer source (Image/Knowledge Base). The top two rows of images show cases where all the 3 steps were correctly executed by our proposed method. Note that our method works for a variety of relations, objects, answer sources, and varying difficulty. It is correctly able to identify the object of interest, even when it is not the most prominent object in the image. For example, in the middle image of the first row, the frisbee is smaller than the dog in the image. However, we were correctly able to retrieve the supporting fact about the frisbee using information from the question, such as ‘*capable of*’ and ‘*flying*.’

A mistake in any of the 3 steps can cause our method to produce an incorrect answer. The bottom row of images in Fig. 4 displays prototypical failure modes. In the leftmost image, we miss cues from the question such as ‘*round*,’ and instead retrieve a fact about the person. In the middle image, our method makes a mistake at the final step and uses information from the wrong answer source. This is a very rare source of errors overall, as we are over 97% accurate in predicting the answer source, as shown in Table 2. In the rightmost image, our method makes a mistake at the first step of predicting the relation, making the remaining steps incorrect. Our relation prediction is around 75%, and 92% accurate by the top-1 and top-3 metrics, as shown in Table 1, and has some scope for improvement. For qualitative results regarding synonyms and homographs we refer the interested reader to the supplementary material.

5 Conclusion

In this work, we addressed knowledge-based visual question answering and developed a method that learns to embed facts as well as question-image pairs into a space that admits efficient search for answers to a given question. In contrast to existing retrieval based techniques, our approach learns to embed questions and facts for retrieval. We have demonstrated the efficacy of the proposed method on the recently introduced and challenging FVQA dataset, producing state-of-the-art results. In the future, we hope to address extensions of our work to larger structured knowledge bases, as well as unstructured knowledge sources, such as online text corpora.

Acknowledgments: This material is based upon work supported in part by the National Science Foundation under Grant No. 1718221, Samsung, and 3M. We thank NVIDIA for providing the GPUs used for this research. We also thank Arun Mallya and Aditya Deshpande for their help.

References

1. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV* (2017)
2. Ren, M., Kiros, R., Zemel, R.: Exploring models and data for image question answering. In: *NIPS*. (2015)
3. Zhu, Y., Groth, O., Bernstein, M., Fei-Fei, L.: Visual7W: Grounded Question Answering in Images. In: *CVPR*. (2016)
4. Malinowski, M., Fritz, M.: Towards a visual turing challenge. In: *NIPS*. (2014)
5. Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C.L., Girshick, R.: Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In: *CVPR*. (2017)
6. Jabri, A., Joulin, A., van der Maaten, L.: Revisiting Visual Question Answering Baselines. In: *ECCV*. (2016)
7. Yu, L., Park, E., Berg, A., Berg, T.: Visual Madlibs: Fill in the blank image generation and question answering. In: *ICCV*. (2015)
8. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: VQA: Visual Question Answering. In: *ICCV*. (2015)
9. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In: *CVPR*. (2017)
10. Gao, H., Mao, J., Zhou, J., Huang, Z., Wang, L., Xu, W.: Are you talking to a machine? Dataset and Methods for Multilingual Image Question Answering. In: *NIPS*. (2015)
11. Malinowski, M., Fritz, M.: A Multi-World Approach to Question Answering about Real-World Scenes based on Uncertain Input. In: *NIPS*. (2014)
12. Malinowski, M., Rohrbach, M., Fritz, M.: Ask your neurons: A neural-based approach to answering questions about images. In: *ICCV*. (2015)
13. Hu, R., Andreas, J., Rohrbach, M., Darrell, T., Saenko, K.: Learning to reason: End-to-end module networks for visual question answering. *CoRR*, abs/1704.05526 **3** (2017)
14. Wang, P., Wu, Q., Shen, C., Dick, A., v. d. Hengel, A.: Fvqa: Fact-based visual question answering. *TPAMI* (2018)
15. Tandon, N., de Melo, G., Suchanek, F., Weikum, G.: Webchild: Harvesting and organizing commonsense knowledge from the web. In: *WSDM*. (2014)
16. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: Dbpedia: A nucleus for a web of open data. In: *ISWC/ASWC*. (2007)
17. Speer, R., Chin, J., Havasi, C.: Conceptnet 5.5: An open multilingual graph of general knowledge. In: *AAAI*. (2017)
18. Lu, J., Yang, J., Batra, D., Parikh, D.: Hierarchical question-image co-attention for visual question answering. In: *NIPS*. (2016)
19. Yang, Z., He, X., Gao, J., Deng, L., Smola, A.: Stacked attention networks for image question answering. In: *CVPR*. (2016)
20. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Deep compositional question answering with neural module networks. In: *CVPR*. (2016)
21. Das, A., Agrawal, H., Zitnick, C.L., Parikh, D., Batra, D.: Human attention in visual question answering: Do humans and deep networks look at the same regions? In: *EMNLP*. (2016)

22. Fukui, A., Park, D.H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multi-modal compact bilinear pooling for visual question answering and visual grounding. In: EMNLP. (2016)
23. Shih, K.J., Singh, S., Hoiem, D.: Where to look: Focus regions for visual question answering. In: CVPR. (2016)
24. Xu, H., Saenko, K.: Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In: ECCV. (2016)
25. Schwartz, I., Schwing, A.G., Hazan, T.: High-Order Attention Models for Visual Question Answering. In: NIPS. (2017)
26. Ben-younes, H., Cadene, R., Cord, M., Thome, N.: Mutan: Multimodal tucker fusion for visual question answering. In: ICCV. (2017)
27. Ma, L., Lu, Z., Li, H.: Learning to answer questions from image using convolutional neural network. In: AAAI. (2016)
28. Jain, U., Zhang, Z., Schwing, A.G.: Creativity: Generating Diverse Questions using Variational Autoencoders. In: CVPR. (2017)
29. Xiong, C., Merity, S., Socher, R.: Dynamic memory networks for visual and textual question answering. In: ICML. (2016)
30. Kim, J.H., Kwak, S.W.L.D.H., Heo, M.O., Kim, J., Ha, J.W., Zhang, B.T.: Multimodal residual learning for visual qa. In: NIPS. (2016)
31. Zitnick, C.L., Agrawal, A., Antol, S., Mitchell, M., Batra, D., Parikh, D.: Measuring machine intelligence through visual question answering. AI Magazine (2016)
32. Zhou, B., Tian, Y., Sukhbaatar, S., Szlam, A., Fergus, R.: Simple baseline for visual question answering. In: arXiv:1512.02167. (2015)
33. Wu, Q., Shen, C., van den Hengel, A., Wang, P., Dick, A.: Image captioning and visual question answering based on attributes and their related external knowledge. In: arXiv:1603.02814. (2016)
34. Jain, U., Lazebnik, S., Schwing, A.G.: Two can play this Game: Visual Dialog with Discriminative Question Generation and Answering. In: CVPR. (2018)
35. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
36. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. IJCV (2015)
37. Zettlemoyer, L.S., M.Collins: Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In: UAI. (2005)
38. Zettlemoyer, L.S., M.Collins: Learning context-dependent mappings from sentences to logical form. In: ACL. (2005)
39. Berant, J., Chou, A., Frostig, R., Liang, P.: Semantic Parsing on Freebase from Question-Answer Pairs. In: EMNLP. (2013)
40. Cai, Q., Yates, A.: Large-scale Semantic Parsing via Schema Matching and Lexicon Extension. In: ACL. (2013)
41. Liang, P., Jordan, M.I., Klein, D.: Learning dependency-based compositional semantics. In: Computational Linguistics. (2013)
42. Kwiatkowski, T., Choi, E., Artzi, Y., Zettlemoyer, L.: Scaling semantic parsers with on-the-fly ontology matching. In: EMNLP. (2013)
43. Berant, J., Liang, P.: Semantic parsing via paraphrasing. In: ACL. (2014)
44. Fader, A., Zettlemoyer, L., Etzioni, O.: Open question answering over curated and extracted knowledge bases. In: KDD. (2014)
45. Yih, W., Chang, M.W., He, X., Gao, J.: Semantic parsing via staged query graph generation: Question answering with knowledge base. In: ACL-IJCNLP. (2015)

46. Reddy, S., Täckström, O., Collins, M., Kwiatkowski, T., Das, D., Steedman, M., Lapata, M.: Transforming dependency structures to logical forms for semantic parsing. In: ACL. (2016)
47. Xiao, C., Dymetman, M., Gardent, C.: Sequence-based structured prediction for semantic parsing. In: ACL. (2016)
48. Unger, C., Bühmann, L., Lehmann, J., Ngomo, A.C.N., Gerber, D., Cimiano, P.: Template-based question answering over RDF data. In: WWW. (2012)
49. Kolomyiets, O., Moens, M.F.: A survey on question answering technology from an information retrieval perspective. In: Information Sciences. (2011)
50. Yao, X., Durme, B.V.: Information extraction over structured data: Question answering with Freebase. In: ACL. (2014)
51. Bordes, A., Chopra, S., Weston, J.: Question answering with sub-graph embeddings. In: EMNLP. (2014)
52. Bordes, A., Weston, J., Usunier, N.: Open question answering with weakly supervised embedding models. In: ECML. (2014)
53. Dong, L., Wei, F., Zhou, M., Xu, K.: Question answering over freebase with multi-column convolutional neural networks. In: ACL. (2015)
54. Bordes, A., Usunier, N., Chopra, S., Weston, J.: Large-scale simple question answering with memory networks. In: ICLR. (2015)
55. Zhu, Y., Zhang, C., Ré, C., Fei-Fei, L.: Building a large-scale multimodal Knowledge Base for Visual Question Answering. In: CoRR. (2015)
56. Wu, Q., Wang, P., Shen, C., van den Hengel, A., Dick, A.: Ask Me Anything: Free-form Visual Question Answering Based on Knowledge from External Sources. In: CVPR. (2016)
57. Wang, P., Wu, Q., Shen, C., van den Hengel, A., Dick, A.: Explicit Knowledge-based Reasoning for Visual Question Answering. In: IJCAI. (2017)
58. Krishnamurthy, J., Kollar, T.: Jointly learning to parse and perceive: Connecting natural language to the physical world. In: ACL. (2013)
59. Narasimhan, K., Yala, A., Barzilay, R.: Improving information extraction by acquiring external evidence with reinforcement learning. In: EMNLP. (2016)
60. Wu, Q., Wang, P., Shen, C., Dick, A., van den Hengel, A.: Ask me anything: Free-form visual question answering based on knowledge from external sources. In: CVPR. (2016)
61. Wang, P., Wu, Q., Shen, C., Dick, A., Van Den Henge, A.: Explicit knowledge-based reasoning for visual question answering. In: IJCAI. (2017)
62. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP. (2014)
63. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* (1997)
64. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. (2016)
65. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. (2009)
66. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS. (2015)
67. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C. Lawrence, e.D., Pajdla, T., Schiele, B., Tuytelaars, T.: Microsoft coco: Common objects in context. In: ECCV. (2014)
68. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *TPAMI* (2017)

69. Mallya, A., Lazebnik, S.: Learning models for actions and person-object interactions with transfer to question answering. In: ECCV. (2016)
70. Chao, Y.W., Wang, Z., He, Y., Wang, J., Deng, J.: Hico: A benchmark for recognizing human-object interactions in images. In: ICCV. (2015)
71. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: CVPR. (2014)
72. Tsochantaridis, I., Joachims, T., Hofmann, T., Altun, Y.: Large Margin Methods for Structured and Interdependent Output Variables. JMLR (2005)