

Interpretable Visual Reasoning via Induced Symbolic Space

Zhonghao Wang¹, Kai Wang³, Mo Yu^{2*},
Jinjun Xiong², Wen-mei Hwu¹, Mark Hasegawa-Johnson¹, Humphrey Shi^{3,1,4*}

¹UIUC, ²MIT-IBM Watson AI Lab & IBM Research, ³U of Oregon, ⁴Picsart AI Research (PAIR)

Abstract

We study the problem of concept induction in visual reasoning, *i.e.*, identifying concepts and their hierarchical relationships from question-answer pairs associated with images; and achieve an interpretable model via working on the induced symbolic concept space. To this end, we first design a new framework named *object-centric compositional attention model (OCCAM)* to perform the visual reasoning task with object-level visual features. Then, we come up with a method to induce concepts of objects and relations using clues from the attention patterns between objects' visual features and question words. Finally, we achieve a higher level of interpretability by imposing OCCAM on the objects represented in the induced symbolic concept space. Experiments on the CLEVR and GQA datasets demonstrate: 1) our OCCAM achieves a new state of the art without human-annotated functional programs; 2) our induced concepts are both accurate and sufficient as OCCAM achieves an on-par performance on objects represented either in visual features or in the induced symbolic concept space.

1. Introduction

Recent advances in Visual Question Answering (VQA) [1, 37, 19, 3, 17, 24, 16, 47, 33] usually rely on carefully designed neural attention models over images, and rely on pre-defined lists of concepts to enhance the compositional reasoning ability of the attention modules. Human prior knowledge plays an essential role in the success of the model design.

We focus on a less-studied problem in this field – given only question-answer pairs and images, induce the visual concepts that are sufficient for completing the visual reasoning tasks. By sufficiency, we hope to maintain the predictive accuracy for VQA when using the induced concepts in place of the original visual features. We consider concepts that are important for visual reasoning, including properties of objects (e.g., *red*, *cube*) and relations between objects (e.g.,

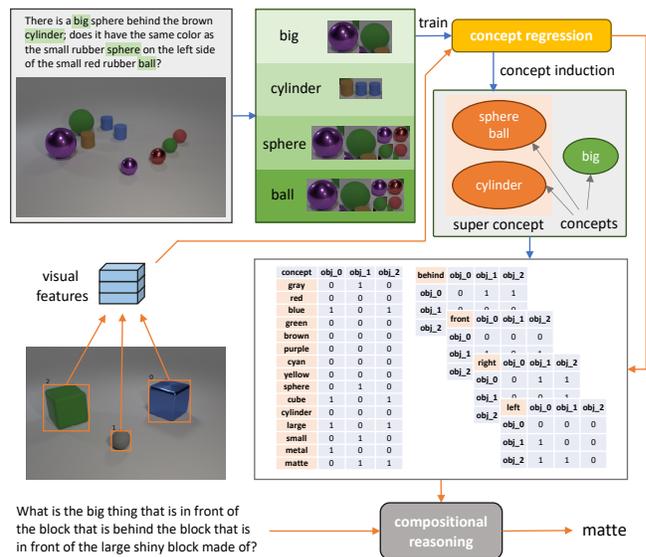


Figure 1: Illustration of our framework. Our model induces the concepts and super concepts with the attention correlation between the objects and question words in image-question pairs as the paths shown in blue arrows. Then, it answers a question about an image via compositional reasoning on the induced symbolic representations of objects and object relations, shown as the orange paths.

left, *front*). The aforementioned scope and sufficiency criterion require accurately associating the induced symbols of concepts to both visual features and words, so that each new instance of question-image pair can be transformed into the induced concept space for further computations. Additionally, it is necessary to identify super concepts, *i.e.*, hypernyms of concept subsets (e.g., *shape*). The concepts inside a super concept are exclusive, so that the system knows each object can only possess one value in each subset. This introduces structural information to the concept space (multiple one-hot vectors for each visual object) and further guarantees the accuracy of the aforementioned transformation.

The value of the study has two folds. First, our proposed problem aims to identify visual concepts, their argument patterns (properties or relations) and their hierarchy (super concepts) *without* using any concept-level su-

*Correspondence to H. Shi and M. Yu.

pervision. Solving the problem frees both the efforts of human annotations and human designs of concept schema required in previous visual reasoning works. At the same time, the problem is technically more challenging compared to the related existing problem like unsupervised or weakly-supervised visual grounding [46]. Second, by constraining the visual reasoning models to work over the induced concepts, the ability of concept induction improves the interpretability of visual reasoning models. Unlike previous interpretable visual reasoning models that rely on human-written rules to associate neural modules with given concept definitions [16, 33, 40], our method resolves the concept definitions and associations interpretability automatically in the learning process, without the need of trading off for hand-crafted model designs.

We achieve the proposed task in three steps. First, we propose a new model architecture, object-centric compositional attention model (OCCAM), that performs object-level visual reasoning instead of pixel-level by extracting object-level visual features with ResNet [14] and pooling the features according to each object’s bounding box. The object-level reasoning not only improves over the state-of-the-arts, but also provides a higher-level interpretability for concept association and identification. Second, we benefit from the trained OCCAM’s attention values over objects to create classifiers mapping visual objects to words; then derive the concepts and super concepts from the object-word cooccurrence matrices as shown in Figure 1. Finally, our concept-based visual reasoning framework predicts the concepts of objects and object relations; then performs compositional reasoning using the predicted symbolic concept embeddings instead of the original visual features.

Experiments on the CLEVR and GQA datasets confirm that our overall approach improves the interpretability of neural visual reasoning, and maintains the predictive accuracy: (1) our OCCAM improves over the previous state-of-the-art models that do not use external training data; (2) our induced concepts and concept hierarchy are accurate in human study; and (3) our induced concepts are sufficient for visual reasoning – replacing visual features with concepts leads to only $\sim 1\%$ performance drop.

2. Related Works

Visual Question Answering (VQA) requires models to reason a question about an image to infer an answer. Recent VQA approaches can be partitioned into two groups: holistic models [45, 42, 1, 37, 19] and modular models [2, 3, 17, 24, 16, 47, 33], according to whether the approach has explicit sub-task structures. A typical holistic model, MAC [19], perform iterative reasoning steps with an attention mechanism on the image. A modular framework, NS-CL [33], designs multiple principle functions over the extracted features to explain the reasoning process.

Scene graph grounding requires to construct the rela-

tionship among objects in an image. [44] designs a graph R-CNN model to detect objects and classify relations among them simultaneously. [5] uses graphs to ground words and phrases to image regions. [46] proposes to link words to image concepts in an unsupervised setting. However, all these works have predefined object and relation concepts. We focus on inducing the concepts from the language compositionality to better interpret the reasoning framework.

Model interpretability aims to explain the neural model predictions. [6] proposed network dissection to quantify interpretability of CNNs. [49] explains a CNN at the semantic level with decision trees. [40] generates scene graphs from images to explicitly trace the reasoning-flow. [35, 16] focused on visual attentions to provide enhanced interpretability. Our work is closely related to *the self-explaining systems via rationalization* [31, 7, 48]. These works usually extract subsets of inputs as explanations, while our work moves one-step further by learning parts of the structural explanation definitions (i.e., our concept hierarchy) together with explanations (i.e., the concept-level reasoning flow).

Visual concept learning contributes to broad visual-linguistic applications, such as cross-modal retrieval [29], visual captioning [25], and visual-question answering [32, 4]. [47, 33] attempt to disentangle visual concept learning and reasoning. Based on the visual concepts learned from VQA, [12] learns metaconcepts, i.e., relational concepts about concepts, with augmented QA-pairs about metaconcepts. Our work differs from the previous ones in learning concepts and super concepts without external knowledge.

3. OCCAM: Object-Centric Visual Reasoning

This section introduces a new neural architecture, object-centric compositional attention model (OCCAM), that performs visual reasoning over the object-level visual features. This model not only achieves state-of-the-art performance, but also plays a key role in inducing object-wise or relational concepts as will be described in section 4.

Figure 2 shows our general framework with two training phases, each consists to the process of attaining the answers from the input images and questions. Phase 1 (black-colored paths) corresponds to the training of our OCCAM, in which we train the object-level feature extractor, the compositional reasoning module and the question embedding LSTM. Phase 2 (red-colored paths) corresponds to the induction of symbolic concepts based on the aforementioned trained neural modules, as well as the training of a concept projection module so that the induced concepts can be accommodated in the OCCAM pipeline. The figure shows the central role that the OCCAM plays in our framework.

3.1. Background on compositional reasoning

Notations. As shown in Figure 2, we name the visual vectors as vs , the output memory vector from the compositional reasoning module as m , the embedded word vectors for questions as ws , and the question embedding as q .

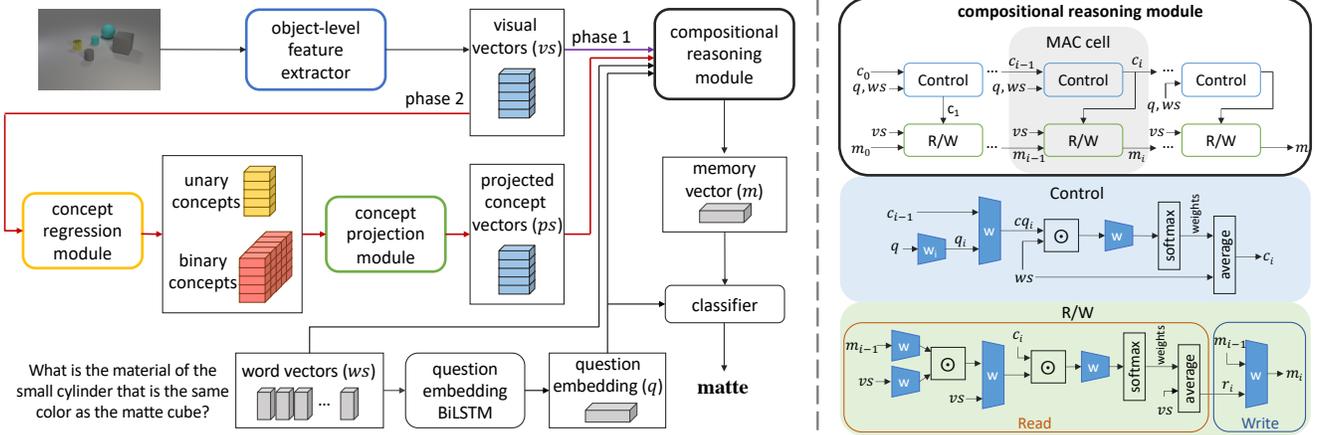


Figure 2: The framework and the compositional reasoning module. The left graph shows the general framework; The phase 1 training path is drawn in purple and the phase 2 training paths are drawn in red. The black paths are shared for both training phases. The structures of our proposed object-level feature extractor, concept regression module and concept projection module are shown in Figures 3, 4 and 7.

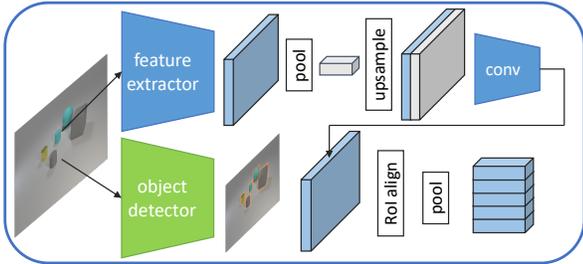


Figure 3: The architecture of our **object-level feature extractor**.

The **compositional reasoning framework** follows a VQA setting: given a question and an image as inputs, the model is required to return the correct answer choice. The target function can thus be written as:

$$\mathcal{L}(ws, vs, q) = - \sum_{k \in K} y_k \log \mathcal{F}(q_k, \mathcal{G}(ws_k, vs_k, q_k)) \quad (1)$$

$$q_k = \mathcal{Q}(ws_k), vs_k = \mathcal{I}(im_k).$$

K is the total number of image-question pairs, y is the one-hot ground truth vector, \mathcal{F} is the classifier, \mathcal{G} is the reasoning module, \mathcal{Q} is the question embedding LSTM, \mathcal{I} is the visual feature extractor and im is the image input.

The **MAC reasoning module** [19] processes visual and language inputs in a sequential way. Shown in Figure 2 (right), each MAC cell contains a control unit that uses word embedding to control what object features should be read and written into memory and an R/W (Read/Write) unit that performs reading and writing object features; the blue diagrams labeled with w stand for fully connected layers and the symbol \odot stands for Hadamard product. More details can be found in Appendix A.

3.2. Object-centric compositional attention model

Our OCCAM network is shown in Figure 2 with phase 1 path. It performs MAC-style reasoning, but over the object-level visual features generated by our proposed **object-level**

feature extractor (Figure 3). Fed with an image, the extractor produces a set of vectors vs , each encodes a single object's unary visual features and its interactions with other objects. The module works as the following steps:

(1) Following [33], we use Mask-RCNN [13] to detect all objects in an image and output the bounding boxes for them. The image is fed to a ResNet34 network [14] pre-trained on ImageNet [10] to generate the feature maps.

(2) On top of the ResNet34 feature maps, we apply a global average pooling to get a single **global feature vector** (the gray vector in the figure). We concatenate this global vector with the feature map at each location, followed by three convolution layers. This global vector is crucial since it allows the visual features to encode the interaction among objects; and the three convolution layers fuse the local and global features into a single visual vector at each position.

(3) Finally, to get object-level features from the above pixel-level fused features, we use ROI align [13] to project the objects' bounding boxes onto the fused feature vectors to generate the ROI feature maps; and average pool these ROI maps for each object to produce the object-level vs .

Our object feature extractor is jointly optimized with the reasoning module with Eqn (1) in the phase 1 training.

4. Concept Induction and Reasoning

This section describes how we achieve our goal of inducing symbolic concepts for objects and performing compositional reasoning on the induced concepts. First, we formalize the problem of concept induction (section 4.1). Second, built on the learned OCCAM network introduced in the previous section, we propose to induce concepts of both unary object properties or the binary relations between objects (section 4.2). Finally, we propose compositional reasoning over symbolic concepts by substituting the object-level features with the induced concepts (section 4.3).

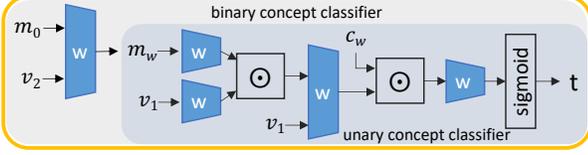


Figure 4: The structure of the **concept regression module**. v_1 and v_2 are the object-level visual vectors representing two objects respectively, and c_w is the word vector. m_0 is a fixed vector and m_w equals to m_0 for the unary concept classifier.

4.1. Problem definition

We consider identifying three types of concepts: (1) the **unary concepts** C^u that are properties of objects (e.g., *red*, *cube*, etc.); (2) the **binary concepts** C^b that are relation descriptions between any two objects (e.g., *left*, *front* etc.); and (3) the **super concepts** C^{sup} that are hypernyms of certain subsets of concepts (e.g., *color*, *shape*, etc.), subject to that each object can only possess one concept under each super concept, e.g., *cube* and *sphere*.

As questions refer to objects and describe object relations in images and, more importantly, include all the semantic information to reach an answer, it is natural to induce the concepts from question words. Therefore we assume that all the unary and binary concepts have their corresponding words; and these words are a subset of the nouns or adjectives from all the training questions. We denote the sets of words that describe unary concepts and binary concepts as M^u and M^b respectively. Therefore, the goal of concept induction consists of the following tasks:

- **Visual mapping:** for each concept $c \in C^u$ or C^b , learning a mapping from the visual feature v to c . In other words, a prediction function $f_c(v) \in \{0, 1\}$ is learned to predict the existence of concept c from the visual feature v of an object.
- **Word mapping:** for each concept $c \in C^u$ or C^b , identifying a subset of words $S_c \subset M^u$ or M^b that are synonyms representing the same concept, e.g., the concept of 'cube' corresponds to set of words $\{cube, cubes, block, blocks, \dots\}$.
- **Super concept induction:** clustering of concepts to form super concepts. Each super concept c contains a set of concepts $\{c_1, \dots, c_k\} \subset C^u$ or C^b .

4.2. Concept induction

This section describes how we achieve the aforementioned tasks of concept induction. The key idea of our approach includes: (1) benefiting from the R/W unit from the trained MAC cells to achieve the visual mapping to textual words; (2) utilizing the inclusiveness of words' visual mapping to induce each concept's multiple word descriptions; (3) clustering super concepts from the mutual exclusiveness between concepts. To achieve the above, we first train two binary classifiers that can determine if a word correctly describes an object's unique feature and if a word correctly describes a relation between two objects respectively. Then,

Algorithm 1: Classifier training data generation. $ST(\cdot)$ splits a vector $\alpha \in \mathbb{R}^\beta$ to a set of β values. $GMM(\cdot)$ uses Gaussian Mixture Model to cluster a set of data points. $FB(\cdot)$ finds the decision boundary for the 2 Gaussian components. $\mathbb{1}$ is the indicator function.

```

Result:  $P^u, P^b$ 
 $P^u = \{\}, P^b = \{\}$ 
for  $x \in M^u \cup M^b$  do
   $S_x = \{\}, bd_x = 0$ 
for  $vs, ws \in DATASET$  do
  for  $c_w \in ws \cap M^u$  do
     $S_{c_w} = S_{c_w} \cup ST(\mathcal{R}(vs, c_w, m_0))$ 
  for  $c_w \in ws \cap M^b$  do
    for  $v \in vs$  do
       $S_{c_w} = S_{c_w} \cup ST(\mathcal{R}(vs, c_w, \mathcal{W}(m_0, v)))$ 
for  $x \in M^u \cup M^b$  do
   $bd_x = FB(GMM(S_x))$ 
for  $vs, ws \in DATASET$  do
  for  $v_1 \in vs$  do
    for  $c_w \in ws \cap M^u$  do
       $y = \mathbb{1}(\mathcal{R}(v_1, c_w, m_0) > bd_{c_w})$ 
       $P^u = P^u \cup \{(v_1, c_w, y)\}$ 
    for  $c_w \in ws \cap M^b$  do
      for  $v_2 \in \{vs - v_1\}$  do
         $y = \mathbb{1}(\mathcal{R}(v_1, c_w, \mathcal{W}(m_0, v_2)) > bd_{c_w})$ 
         $P^b = P^b \cup \{(v_1, v_2, c_w, y)\}$ 

```

with the help of these classifiers, we produce zero-one vectors for words that properly describe the unique features for each object and the relations between any pair of objects in single images across the dataset. Finally, we perform a clustering method on the word vectors to generalize unary and binary concepts, and the super concept sets.

Visual mapping via regression from MAC cells The concept regression module is shown in Figure 4. It is composed of a classifier for the unary concept word regression, $\mathcal{B}^u(v_1, c_w) \in [0, 1]$, and a classifier for the binary concept word regression, $\mathcal{B}^b(v_1, v_2, c_w) \in [0, 1]$. \mathcal{B}^u is expected to produce 1 if v_1 can be described by the word vector c_w . Likewise, \mathcal{B}^b is expected to produce 1 if the relation of v_1 to v_2 can be described by the word vector c_w .

We generate training data points $P^u = \{(v_{1_i}^u, c_{w_i}^u, y_i^u)\}$ and $P^b = \{(v_{1_i}^b, v_{2_i}^b, c_{w_i}^b, y_i^b)\}$ for \mathcal{B}^u and \mathcal{B}^b by utilizing the Read/Write unit (Figure 2(right)) in the reasoning module after phase 1 training. The whole generation process is described in Algorithm 1. We denote $\mathcal{R}(vs, c_i, m_{i-1}) \in \mathbb{R}^{|O|}$ for the sequence of functions before the softmax operation in the Read unit and $\mathcal{W}(m_{i-1}, r_i) \in \mathbb{R}^D$ for the function of the Write unit, where O is the set of objects in an image and D is the vector dimension.

Specifically, our algorithm first uses $\mathcal{R}(\cdot, \cdot, \cdot)$ and $\mathcal{W}(\cdot, \cdot)$ to find the attention logits on the objects corresponding to words describing the unary and binary concepts in a question as shown in Figure 5(a&b). We then use the values

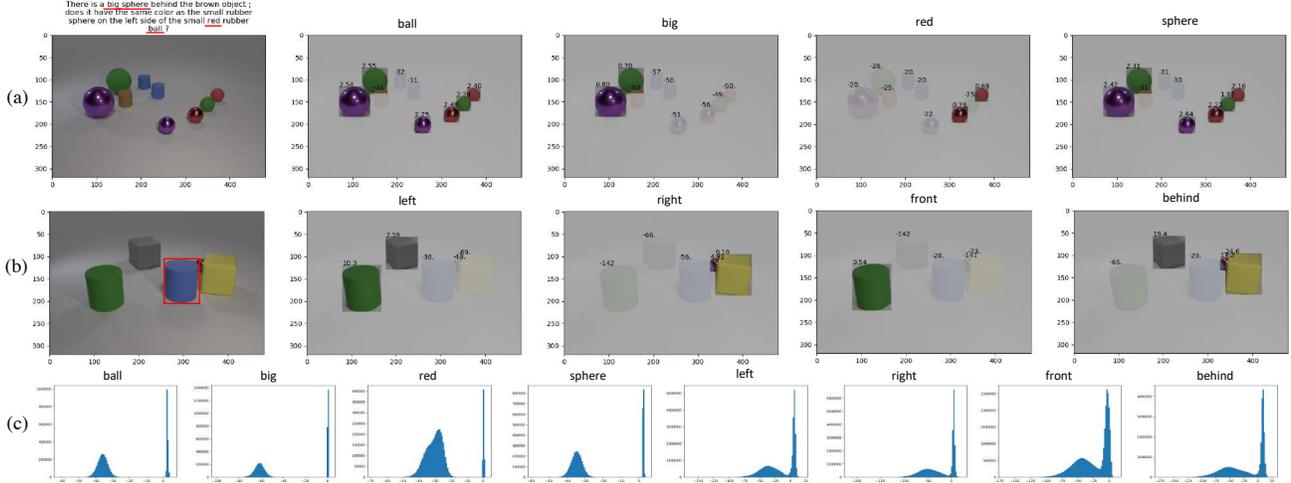


Figure 5: Attention visualization and attention logit distributions. (a) The attention visualization corresponding to the words describing the unary concepts by performing $\mathcal{R}(vs, c_w, m_0)$. Each of the words above the latter 4 images corresponds to a unique c_w and the value on each object is the attention logit (the same applies to (b)). (b) The attention visualization corresponding to the words describing the binary concepts by performing $\mathcal{R}(vs, c_w, \mathcal{W}(m_0, v_2))$. v_2 represents the object bounded by a red rectangle in the first image. (c) the attention logit distribution corresponding to each word describing a concept.

of logits to determine if the object possesses the concept of the word (positive) or not (negative). Noticing the attention logit distribution of the sampled objects for each word is a two-peak distribution (Figure 5(c)), we use a GMM [43] with two Gaussian components to model the distribution and find the decision boundary for each word’s attention logit distribution. Observe that the distribution of a binary concept word has two interfering waves, because in some cases it is hard to tell if two objects have that relation (‘front’ is inappropriate if two objects are on the same horizon). P^u and P^b are generated by classifying the data points to positives and negatives with the decision boundaries. Finally, we can train \mathcal{B}^u and \mathcal{B}^b with data P^u and P^b by minimizing the binary cross entropy loss.

Binary coding of objects With trained \mathcal{B}^u and \mathcal{B}^b , we represent an object o_1 with a binary code vector. Each dimension corresponds to a word. A dimension has value 1 if the corresponding word can describe o_1 and 0 otherwise. The binary vectors of object properties and of the relations between two objects, o_1 and o_2 can be computed with the functions $\gamma^u \in \mathbb{R}^{|M^u|}$ and $\gamma^b \in \mathbb{R}^{|M^b|}$ respectively:

$$\begin{aligned} \gamma^u &= \mathbb{1}_{i>0.5}(\mathcal{B}^u(v_1, C^u)) \\ \gamma^b &= \mathbb{1}_{i>0.5}(\mathcal{B}^b(v_1, v_2, C^b)), \end{aligned} \quad (2)$$

where v_1 and v_2 are the object-level visual vectors of o_1 and o_2 , $C^u \in \mathbb{R}^{|M^u| \times D}$ and $C^b \in \mathbb{R}^{|M^b| \times D}$ are the stacks of word embeddings in vocabulary M^u and M^b . $\mathbb{1}_\alpha(\beta)$ performs elementwise on β : return 1 if the element satisfies condition α or 0 otherwise.

By applying γ^u and γ^b to all the objects and relations in the dataset, we can attain a matrix $\Gamma^u \in \{0, 1\}^{M^u, N^u}$ and a matrix $\Gamma^b \in \{0, 1\}^{M^b, N^b}$ as shown in Figure 6, where N^u

	small	tiny	big	cube	ball	...
	1	1	0	1	0	...
	0	0	1	1	0	...
	1	1	0	0	1	...
	0	0	1	0	1	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮

(a)

	left	right	front	behind
	0	1	1	0
	1	0	0	1
	1	0	1	0
⋮	⋮	⋮	⋮	⋮

(b)

Figure 6: The zero-one matrices indicating word descriptions of objects and object relations across the dataset. (a) The matrix Γ^u indicates what words can describe objects. (b) The matrix Γ^b indicates what words can describe the relations object v_1 ’s (bounded by green rectangles) are to object v_2 ’s (bounded by red rectangles), and N^b are the total numbers of objects and co-occurred object pairs. The two matrices summarize each word’s corresponding visual objects in the whole dataset.

Concept/super-concept induction Finally, we group synonym words to unary and binary concepts and generate the super concepts. These two tasks are achieved via exploring the word inclusiveness and the concept exclusiveness captured by Γ^u and Γ^b : (1) words describing the same concept correspond to similar column vectors, e.g., Γ_{small}^u and Γ_{tiny}^u ; (2) words describing exclusive concepts have column vectors that usually do not have 1 values on same objects simultaneously, e.g., Γ_{cube}^u and Γ_{ball}^u . Based on the aforementioned ideas, we define the correlation metric between two words c_{w_1} and c_{w_2} as below:

$$\begin{aligned} \theta_{c_{w_1}, c_{w_2}} &= P(\gamma_{c_{w_1}} = 1 \mid \gamma_{c_{w_2}} = 1) + \\ &P(\gamma_{c_{w_2}} = 1 \mid \gamma_{c_{w_1}} = 1) \\ &= \frac{|\Gamma_{c_{w_1}} \odot \Gamma_{c_{w_2}}|_1}{|\Gamma_{c_{w_2}}|_1} + \frac{|\Gamma_{c_{w_1}} \odot \Gamma_{c_{w_2}}|_1}{|\Gamma_{c_{w_1}}|_1}. \end{aligned} \quad (3)$$

Algorithm 2: Concept vector generalization. $\text{MAX}(\alpha)$ and $\text{HARDMAX}(\alpha)$ return the largest value in vector α and its position as a one-hot vector, respectively.

Result: K^u, K^b
 $K^u = \mathbf{0}^{|O| \times |E^u|}, K^b = \mathbf{0}^{|O| \times |O| \times |E^b|}$

```

for  $i \in O$  do
  for  $e^u \in E^u$  do
     $K^u[i][e^u] = \text{MAX}(\mathcal{B}^u(v_i, \rho_{e^u}))$ 
  for  $l^u \in L^u$  do
     $K^u[i][l^u] = \text{HARDMAX}(K^u[i][l^u])$ 
  for  $j \in O - \{i\}$  do
    for  $e^b \in E^b$  do
       $K^b[i][j][e^b] = \text{MAX}(\mathcal{B}^b(v_i, v_j, \rho_{e^b}))$ 
    for  $l^b \in L^b$  do
       $K^b[i][j][l^b] = \text{HARDMAX}(K^b[i][j][l^b])$ 

```

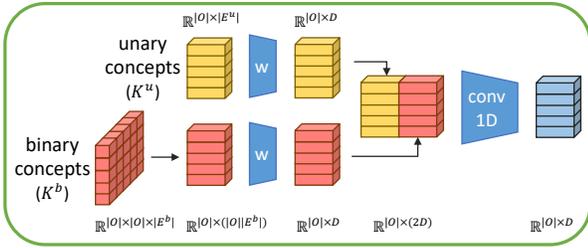


Figure 7: The structure of the **concept projection module**. We label the dimensions of matrices near them in the graph.

This guarantees that $\theta \rightarrow 0^+$ for two synonym words, $\theta \rightarrow 2^-$ for two words corresponding to exclusive concepts and $\theta \in (0, 2)$ for words corresponding to different nonexclusive concepts. We can produce the correlation sets for the words describing the unary concepts and the binary concepts respectively with Eqn (4).

$$\Theta^x = \{\theta_{c_{w_1}, c_{w_2}}\}; c_{w_1}, c_{w_2} \in M^x; x \in \{u, b\} \quad (4)$$

Our final step fits two GMM on Θ^u and Θ^b respectively. Each GMM has three components $\mathcal{N}_0, \mathcal{N}_1$ and \mathcal{N}_2 , with their mean values initialized with 0, 1 and 2. We then induce the unary and binary concepts, where each concept consists of synonym words whose mutual correlation is clustered to the Gaussian component \mathcal{N}_0 . Similarly, we induce the super concepts, where each super concept contains multiple concepts and any two words from different concepts have correlation clustered to the Gaussian component of \mathcal{N}_2 .

We denote the set of words corresponding to a concept e as ρ_e , the set of the super concept sets as L , the set of all concepts as E . Then, we can represent all the objects in an image with a unary concept matrix K^u and represent all the relations between any two objects in an image with a binary concept matrix K^b with Algorithm 2.

4.3. Concept compositional reasoning

Our ultimate goal is to perform compositional reasoning to answer a question with the generated concept representa-

tions K^u and K^b for an image; so as to confirm that our induced concepts are accurate and sufficient. We achieve this with the phase 2 training process in Figure 2. The key idea is to transplant the learned compositional reasoning module from manipulating the visual features to manipulating K^u and K^b , for attaining the answer to a question.

To this end, first, we project K^u and K^b to the same vector space with vs with the concept projection module shown in Figure 7, so that the compositional module can perform the reasoning steps on the projected concept vectors. Specifically, we first reduce the dimension of K^b from $\mathbb{R}^{|O| \times |O| \times |E^b|}$ to $\mathbb{R}^{|O| \times |O| \times |E^b|}$, resulted in \hat{K}^b , because K^b can be understood as the relations to other objects for each object in an image. Then, we use two separate fully connected networks to project K^u and \hat{K}^b respectively, concatenate and use a sequence of 1D convolution layers to project the results to the same dimension of vs 's.

Second, to minimize the discrepancy between the distribution of our projected vectors and that of the original visual vectors vs , we fix the weights of other modules in the framework and only train the concept project module by optimizing the target function Eqn. (1). Then, we train the concept projection module and the compositional reasoning module with other modules' weights fixed to better optimize Eqn. (1). The result is a compositional reasoning model that works on the induced concepts only.

5. Experiment

5.1. Settings

Datasets (1) We first evaluate our model on the **CLEVR** [23] dataset. The dataset comprises images of synthetic objects of various shapes, colors, sizes and materials and question/answer pairs about these images. The questions require multi-hop reasoning, such as finding the transitive relations, counting numbers, comparing properties, to attain correct answers. Each question corresponds to a ground truth human-written programs. Because the programs rely on pre-defined concepts thus do not fit our problem, we let our framework learn from scratch *without* using the program annotations. There are 70k/15k images and $\sim 700k/\sim 150k$ questions in the training/validation sets. We follow the previous works [47, 19, 33] to train our model on the whole training set and test on the validation set.

(2) To demonstrate the generalizability of our approach, we further evaluate on the **GQA** dataset. GQA is a real-world visual reasoning benchmark. It consists of 113K images collected from the Visual Genome dataset [30] and 22M questions. It has a train split for model training and three test splits (val, test, test-dev) [20]. The dataset provides the detected object features extracted from a Faster RCNN detector [38], so each object is represented as a 2048-dimensional vector.

Implementation details We include the checklist of our implementation details in Appendix B.

Table 1: The comparison of our OCCAM framework to the state-of-the-art methods on CLEVR (left) and GQA (right) datasets. † means training with additional program supervision. ‡ means pretraining on larger visual+language corpora. § means pretraining a scene-graph extraction model with additional rich annotated data.

(a) CLEVR							(b) GQA			
method	overall	count	exist	comp num	query attr	comp attr	method	val	test-dev	test
RN [39]	95.5	90.1	93.6	97.8	97.1	97.9	MAC [3]	57.5	-	54.1
FiLM [37]	97.6	94.5	93.8	99.2	99.2	99.0	LXMERT [41]	-	50.0	-
MAC [19]	98.9	97.2	99.4	99.5	99.3	99.5	LCGN [18]	63.9	55.8	56.1
NS-CL [33]	98.9	98.2	99.0	98.8	99.3	99.1	OCCAM (ours)	64.5	56.2	56.2
OCCAM (ours)	99.4	98.1	99.8	99.0	99.9	99.9	MMN [8] [†]	-	60.4	60.8
NS-VQA [†] [47]	99.8	99.7	99.9	99.9	99.8	99.8	NSM [21] [§]	-	63.0	63.2
Human [24]	92.6	86.7	96.6	86.5	95.0	96.0	LXMERT [41] [‡]	-	60.0	60.3
							ViLT [27] ^{‡§}	-	65.1	64.7

Table 2: Effect of the choice of reasoning steps for our model.

steps	4	8	12	16
accuracy (CLEVR)	94.3	98.6	99.4	99.1
accuracy (GQA test-dev)	55.1	55.6	55.2	56.2

5.2. Object-level reasoning

We first perform the end-to-end phase 1 training shown in Figure 2, i.e., our OCCAM model. The performance comparison of our model to the state-of-the-art models is shown in Table 1. Under the setting that no external human-labeled programs and no pretraining are used, our model achieves state-of-the-arts compared with published results on both CLEVR and GQA datasets. For comparison with models on GQA leaderboard, we also train our OCCAM model on train-all split and achieves an accuracy of 58.5% on the test-standard split of GQA dataset, which outperforms other popular models (e.g. MCAN, BAN and LCGN) trained with no additional data (accuracies are 57%~58%). While transformer-based methods with pre-training phase boost the performance, however, they undermine the model’s explainability and make it difficult to induce concepts. On CLEVR, our model also has an on-par performance with the best model [47] that uses external human-labeled programs.

Compared to the original MAC [19] framework which uses image-level attentions, our model proves that the constraint of attentions on the objects are useful for improving the performance on both datasets, with significant improvement on the validation sets. We do not use the position embedding to explicitly encode the positions of objects for relational reasoning; however, we use the global features to enhance the model’s understanding of inter-object relations. This shows that the relations among objects are learnable concepts without external knowledge for the deep network.

Table 2 further gives an ablation study on the numbers of reasoning steps, i.e., the number of MAC modules, for

our model. The reasoning model with 4 steps has a performance gap to the models with 8, 12 or 16 steps, while the latter three models have on-par performances. We conjecture that the model with low reasoning steps may not be able to capture multiple hops of a question and the model performance converges with an increasing number of reasoning steps. We also did ablation study on the contribution of object-level feature extractor on CLEVR dataset. With pretrained ResNet101 features, learnable ResNet34 features, learnable ResNet34 features plus global features respectively, the model achieves 97.9%, 99.0% and 99.4% on the validation set. It shows the importance of enhancing global context understanding at object level.

5.3. Concept induction and reasoning

Next we evaluate the performance of our concept induction method, i.e., the phase 2 training in Figure 2. To qualitatively show that our induced concepts capture sufficient and accurate information for visual reasoning, we replace the visual inputs to the objects’ induced concepts according to Section 4.3. The resulted model, denoted as OCCAM_{concept}, is expected to perform closely to the original OCCAM with high-quality induced concepts.

Table 3 gives the results. To achieve the balance of the performance and the interpretability, we make the OCCAM model run 8 reasoning steps for both concept induction and reasoning. It is observed that our concept-based OCCAM (with induced concept features only) achieves on-par performance with the original OCCAM model (with full input visual features). We also visualize the reasoning steps for the OCCAM_{concept} model in Appendix F.

Human study of concept induction We present the unary concept correlations Θ^u in Figure 8 and 9 for both CLEVR and GQA. Since GQA consists of a huge vocabulary with highly-correlated concepts, we demonstrated a sub-set of concepts associated to general words/phrases.

On CLEVR, the concept definition from the data generator can be perfectly recovered by our approach: from Fig-

Table 3: Comparison of our visual feature-based OCCAM and our concept-only OCCAM. The number of reasoning steps is 8.

(a) CLEVR							(b) GQA		
method	overall	count	exist	comp num	query attr	comp attr	method	val	test-dev
OCCAM _{visual}	98.6	95.9	99.8	96.2	99.8	99.7	OCCAM _{visual}	63.8	55.6
OCCAM _{concept}	97.9	95.6	98.7	97.3	98.4	99.3	OCCAM _{concept}	63.1	54.2

ure 8, the correlation between any pair of synonyms is close to 2, the correlation between words belonging to the same super concept set is close to 0, and the correlation between words belonging to two different super concept sets is in the middle of the range[0, 2]. Appendix C provides the full generated concept hierarchy, which perfectly matches the definition in CLEVR generator and human prior knowledge, i.e., 100% accuracy according to our human investigation.

On GQA, the correlations between words are complicated. We present a subset of word correlations in Figure 9. Instead of using Eqn. (3), here we show the conditional probability that a column attribute exists given that the row attribute exists. We observe that words describing a similar property have a high positive correlation, such as ‘yellow’ and ‘orange’, ‘concrete’ and ‘stone’. They can be grouped into a single concept. Words of exclusive meanings negatively correlates with each other, such as ‘flying’ and ‘standing’, ‘pointy’ and ‘sandy’. They can be grouped into a super concept. However, the real-world data makes it difficult to induce some commonsense super concepts. For example, the same object can have multiple colors (e.g., sky can be both gray and blue). Also, object concept can have degrees (light or dark blue), so we have to use soft values to represent the concepts. We additionally conduct human studies on the pairwise accuracy of detected concept and super concept clusters, which can be found in Appendix G.

We provide more visualization results in Appendix D and E, including the extension of word analogy [34] (e.g., “Madrid” - “Spain” + “France” → “Paris”) to multimodality and the quantification of distance between two visual objects with the super concept space.

6. Conclusions

Our proposed OCCAM framework performs pure object-level reasoning and achieves a new state-of-the-art without human-annotated functional programs on the CLEVR dataset. Our framework makes the object-word cooccurrence information available, which enables induction of the concepts and super concepts based on the inclusiveness and the mutual exclusiveness of words’ visual mappings. When working on concepts instead of visual features, OCCAM achieves comparable performance, proving the accuracy and sufficiency of the induced concepts. For future works, our method can be extended to more sophisticated induction tasks, such as inducing concepts from phrases, with more complicated hierarchy, with degrees of features (e.g., *dark blue*, *light blue*) and inducing compli-

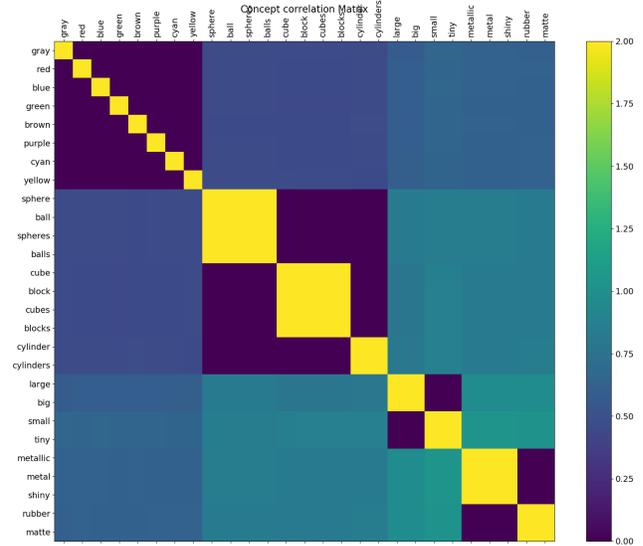


Figure 8: The CLEVR unary concept correlations Θ^u .

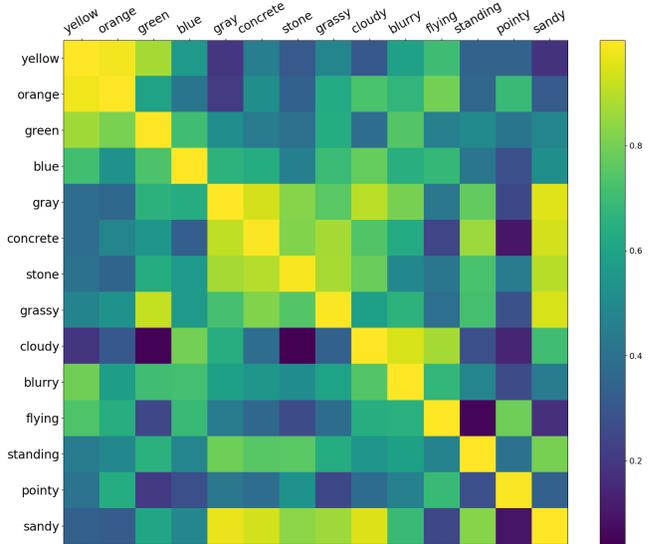


Figure 9: The subset of GQA concept correlations.

cated relations between objects (e.g. *a little bigger*).

Acknowledgments This work is in part supported by IBM-Illinois Center for Cognitive Computing Systems Research (C3SR) - a research collaboration as part of the IBM AI Horizons Network.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 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 (CVPR)*, pages 6077–6086, 2018.
- [2] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Learning to compose neural networks for question answering. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2016.
- [3] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 39–48, 2016.
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2425–2433, 2015.
- [5] Mohit Bajaj, Lanjun Wang, and Leonid Sigal. G3raphground: Graph-based language grounding. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 4281–4290, 2019.
- [6] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6541–6549, 2017.
- [7] Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *International Conference on Machine Learning (ICML)*, pages 883–892, 2018.
- [8] Wenhu Chen, Zhe Gan, Linjie Li, Yu Cheng, William Wang, and Jingjing Liu. Meta module network for compositional visual reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 655–664, 2021.
- [9] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *International Conference on Learning Representations (ICLR)*, 2016.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [11] Alex Graves and Jürgen Schmidhuber. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5-6):602–610, 2005.
- [12] Chi Han, Jiayuan Mao, Chuang Gan, Josh Tenenbaum, and Jiajun Wu. Visual concept-metaconcept learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5001–5012, 2019.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2961–2969, 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [16] Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. Explainable neural computation via stack neural module networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 53–69, 2018.
- [17] Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 804–813, 2017.
- [18] Ronghang Hu, Anna Rohrbach, Trevor Darrell, and Kate Saenko. Language-conditioned graph networks for relational reasoning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 10294–10303, 2019.
- [19] Drew A Hudson and Christopher D Manning. Compositional attention networks for machine reasoning. In *International Conference on Learning Representations (ICLR)*, 2018.
- [20] Drew A Hudson and Christopher D Manning. 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 (CVPR)*, pages 6700–6709, 2019.
- [21] Drew A Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. *arXiv preprint arXiv:1907.03950*, 2019.
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015.
- [23] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2901–2910, 2017.
- [24] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Inferring and executing programs for visual reasoning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2989–2998, 2017.
- [25] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3128–3137, 2015.
- [26] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [27] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. *arXiv preprint arXiv:2102.03334*, 2021.

- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [29] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models, 2014.
- [30] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)*, volume=123, number=1, pages=32–73, year=2017, publisher=Springer.
- [31] Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 107–117, 2016.
- [32] Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1–9, 2015.
- [33] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations (ICLR)*, 2018.
- [34] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3111–3119, 2013.
- [35] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Attentive explanations: Justifying decisions and pointing to the evidence. *arXiv preprint arXiv:1612.04757*, 2016.
- [36] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [37] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C Courville. Film: Visual reasoning with a general conditioning layer. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2018.
- [38] Shaoqing Ren, Kaiming He, Ross B Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [39] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4967–4976, 2017.
- [40] Jiaxin Shi, Hanwang Zhang, and Juanzi Li. Explainable and explicit visual reasoning over scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8376–8384, 2019.
- [41] Hao Tan and Mohit Bansal. 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 (EMNLP-IJCNLP)*, pages 5103–5114, 2019.
- [42] Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 451–466. Springer, 2016.
- [43] Guorong Xuan, Wei Zhang, and Peiqi Chai. Em algorithms of gaussian mixture model and hidden markov model. In *International Conference on Image Processing (ICIP)*, volume 1, pages 145–148. IEEE, 2001.
- [44] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 670–685, 2018.
- [45] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21–29, 2016.
- [46] Raymond A Yeh, Minh N Do, and Alexander G Schwing. Unsupervised textual grounding: Linking words to image concepts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6125–6134, 2018.
- [47] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1031–1042, 2018.
- [48] Mo Yu, Shiyu Chang, Yang Zhang, and Tommi Jaakkola. Re-thinking cooperative rationalization: Introspective extraction and complement control. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4085–4094, 2019.
- [49] Quanshi Zhang, Yu Yang, Haotian Ma, and Ying Nian Wu. Interpreting cnns via decision trees. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6261–6270, 2019.