

# Transfer Learning via Unsupervised Task Discovery for Visual Question Answering

Hyeonwoo Noh<sup>1,3</sup><sup>1</sup>Computer Vision Lab.  
POSTECH, KoreaTaehoon Kim<sup>2\*</sup><sup>2</sup>OpenAI  
USAJonghwan Mun<sup>1,3</sup><sup>3</sup>Computer Vision Lab., ECE & ASRI  
Seoul National University, KoreaBohyung Han<sup>3</sup><sup>1</sup>{shgusdngogo, jonghwan.mun}@postech.ac.kr <sup>2</sup>carpedm20@gmail.com <sup>3</sup>bhhan@snu.ac.kr

## Abstract

We study how to leverage off-the-shelf visual and linguistic data to cope with out-of-vocabulary answers in visual question answering task. Existing large-scale visual datasets with annotations such as image class labels, bounding boxes and region descriptions are good sources for learning rich and diverse visual concepts. However, it is not straightforward how the visual concepts can be captured and transferred to visual question answering models due to missing link between question dependent answering models and visual data without question. We tackle this problem in two steps: 1) learning a task conditional visual classifier, which is capable of solving diverse question-specific visual recognition tasks, based on unsupervised task discovery and 2) transferring the task conditional visual classifier to visual question answering models. Specifically, we employ linguistic knowledge sources such as structured lexical database (e.g. WordNet) and visual descriptions for unsupervised task discovery, and transfer a learned task conditional visual classifier as an answering unit in a visual question answering model. We empirically show that the proposed algorithm generalizes to out-of-vocabulary answers successfully using the knowledge transferred from the visual dataset.

## 1. Introduction

Human sees and understands a visual scene based on diverse visual concepts. For example, from a single image of a chair, human effortlessly recognizes diverse visual concepts such as its color, material, style, usage, and so on. Such diverse visual concepts may be associated with different questions in natural languages defining a recognition task for each of the visual concepts (e.g., what color is the chair?). Recently visual question answering (VQA) [5] is proposed as an effort to learn deep neural network models

with capability to perform diverse visual recognition tasks defined adaptively by questions.

Approaches to VQA rely on a large-scale dataset of image, question and answer triples, and train a classifier taking an image and a question as inputs and producing an answer. Despite recent remarkable progress [3, 12, 43], this direction has a critical limitation that image, question and answer triples in datasets are the only source for learning visual concepts. Such drawback may result in lack of scalability because the triplets may be collected artificially by human annotators with limited quality control and have weak diversity in visual concepts. In fact, VQA datasets [1, 13] suffer from inherent bias, which hinders learning true visual concepts from the datasets.

On the contrary, human answers a question based on visual concepts learned from diverse sources such as books, pictures, videos, and personal experience that are not necessarily associated with target questions. Even for machines, there exist more natural and scalable sources for learning visual concepts: image class labels, bounding boxes and image descriptions. Such information is already available in large-scale [9, 24, 22] and can scale further with reasonable cost [31, 32]. This observation brings a natural question; can we learn visual concepts without question annotations and transfer them for VQA?

To address this question, we introduce a VQA problem with out-of-vocabulary answers, which is illustrated in Figure 1. External visual dataset provides a set of labels  $\mathcal{A}$  and only a subset of these labels  $\mathcal{B} \subset \mathcal{A}$  appears in VQA training set as answers. The goal of this task is to handle out-of-vocabulary answers  $\mathbf{a} \in \mathcal{A} - \mathcal{B}$  successfully by exploiting visual concepts learned from external visual dataset.

This paper studies how to learn visual concepts without questions and how to transfer the learned concepts to VQA models. To learn transferable visual concepts, we train a task conditional visual classifier, whose task is defined by a task feature. The classifier is used as an answering unit, where a task feature is inferred from a question. To train the task conditional visual classifier without

\*This work is performed while at Devsisters.



Figure 1. **VQA with out-of-vocabulary answers.** Given a set of labels in visual dataset  $\mathcal{A}$  and a set of answers in VQA training set  $\mathcal{B}$ , we evaluate a model on VQA test set with answers  $\mathbf{a} \in \mathcal{A} - \mathcal{B}$ . External visual dataset provides a set of bounding box labels and visual descriptions for all answers in VQA training set and test set. See text for details.

task annotations, we propose an unsupervised task discovery technique based on linguistic knowledge sources such as structured lexical databases, *e.g.*, WordNet [10] and region descriptions [22]. We claim that the proposed transfer learning framework helps generalization in VQA with out-of-vocabulary answers. The main contribution of our paper is three-fold:

- We present a novel transfer learning algorithm for visual question answering based on a task conditional visual classifier.
- We propose an unsupervised task discovery technique for learning task conditional visual classifiers without explicit task annotations.
- We show that the proposed approach handle out-of-vocabulary answers through knowledge transfer from visual dataset without question annotations.

The rest of the paper is organized as follows. Section 2 discusses prior works related to our approach. We describe the overall transfer learning framework in Section 3. Learning visual concepts by unsupervised task discovery is described in Section 4. Section 5 analyzes experimental results and Section 6 makes our conclusion.

## 2. Related Works

The standard VQA evaluation assumes identically distributed train and test set [5, 29, 45]. As this evaluation setting turns out to be vulnerable to models exploiting biases in training set [13], several alternatives have been proposed. One approach is to reduce observed biases either by balancing answers for individual questions [13] or by providing different biases to train and test sets intentionally [1]. Another approach is to construct compositional generalization split [2, 18] whose question and answer pairs in test set are formed by novel compositions of visual concepts and

question types appearing in the training set. This split is constructed by repurposing an existing VQA dataset [2] or by constructing a synthetic dataset [18]. The problem setting studied in this paper is similar to [19, 36] in the sense that out-of-vocabulary answers are used for testing, but unlike the prior work, we formulate the problem as a transfer learning, where out-of-vocabulary answers are learned from external visual data.

External data are often employed in VQA for better generalization. Convolutional neural networks [15, 23] pretrained on ImageNet [9] is a widely accepted standard for diverse VQA models [12, 43]. As an alternative, object detector [34] trained on the Visual Genome dataset [22] is employed to extract pretrained visual features [3]. Pretrained language models such as word embeddings [33] or sentence embeddings [21] are frequently used to initialize parameters of question encoders [12, 30, 35]. Exploiting information retrieval from knowledge base [6, 7] or external vision algorithms [40] to provide additional inputs to VQA models are investigated in [39, 40, 41]. Transfer between VQA datasets is studied in [16]. Sharing aligned image-word representations between VQA models and image classifiers has been proposed in [14] to exploit external visual data.

Transfer learning from external data to cope with out-of-vocabulary words has hardly been studied in VQA, but is actively investigated in novel object captioning [4, 28, 37, 44]. For example, [4] and [37] decompose image captioning task into visual classification and language modeling, and exploit unpaired visual and linguistic data as additional resources to train visual classifier and language model, respectively. Recent approaches incorporate pointer networks [38] and learn to point an index of word candidates [44] or an associated region [28], where the word candidates are detected by a multi-label classifier [44] or an object detector [28] trained with external visual data. However, these algorithms are not directly applicable to

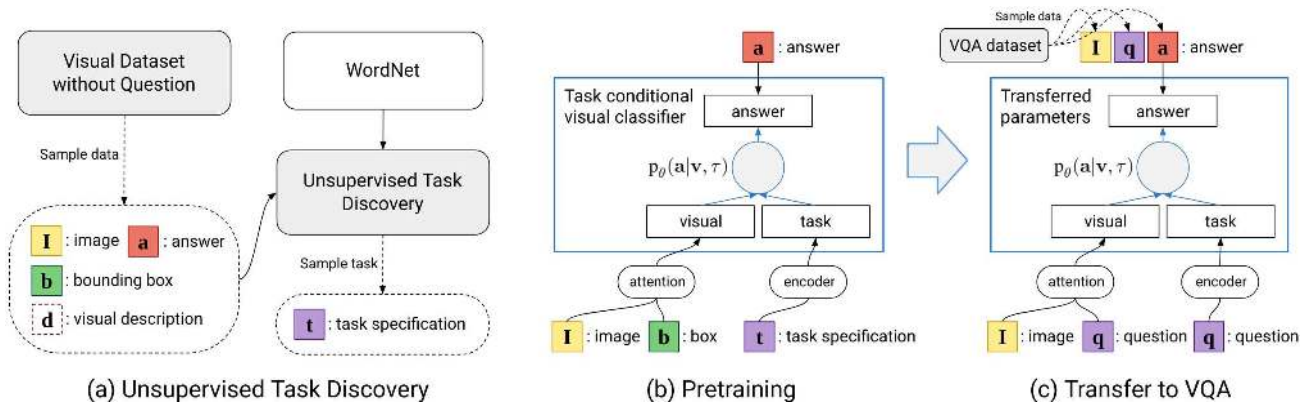


Figure 2. **Overview of the proposed algorithm.** (a) Unsupervised task discovery samples a task specification for a sampled visual data ( $\mathbf{a}, \mathbf{I}, \mathbf{b}$ ), where  $\mathbf{I}$ ,  $\mathbf{b}$  and  $\mathbf{a}$  are an image, a bounding box and an label (answer), respectively. It leverages linguistic knowledge sources such as visual description  $\mathbf{d}$  and WordNet. (b) A visual data with a task specification, denoted by  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{t})$ , is employed to pretrain a task conditional visual classifier. (c) The pretrained task conditional visual classifier is transferred to VQA and the parameters are frozen. Attention layer and question encoder are learned from scratch with VQA dataset. The terms label and answer are used interchangeably.

our problem setting because they focus on predicting object words without task specification while the task conditional visual recognition is required for VQA.

Our problem setting is closely related to zero-shot learning [11, 17, 25, 26, 42], where out of vocabulary answers are considered in classification. Zero-shot learning aims to recognize objects or classes that are unseen during training. As it aims generalization to completely unseen classes, any exposure to zero-shot classes during training is strictly prohibited [42]. On the contrary, our goal is to exploit class labels available in external dataset.

### 3. Transfer Learning Framework

The main goal of our work is to handle out-of-vocabulary answers in VQA by learning visual concepts from off-the-shelf visual dataset and transferring the concepts to VQA for answering the questions. Inspired by the fact that VQA can be thought of as a task conditional classification problem, where tasks are defined by questions, we introduce a task conditional visual classifier, which generates an answer from a visual input and a task specification, as a medium for learning and transferring visual concepts. Figure 2 illustrates the overall framework of the proposed approach. We pretrain the task conditional visual classifier using visual dataset without questions or task specifications via unsupervised task discovery, and adapt it to VQA models by transferring the learned parameters. We describe the task conditional visual classifier and how it is pretrained and transferred to VQA in the rest of this section.

#### 3.1. Task conditional visual classifier

Task conditional visual classifier is a function taking a visual feature  $v \in \mathbb{R}^d$  and a task feature  $\tau \in \mathbb{R}^k$  and producing a probability distribution of answers or labels

$\mathbf{a} \in [0, 1]^l$ , where the terms answer and label are used interchangeably based on context hereafter. The classifier formulated as a neural network with parameter  $\theta$  models a conditional distribution  $p_\theta(\mathbf{a}|v, \tau)$ . Note that two inputs  $v$  and  $\tau$  are typically obtained by encoders  $v_\phi(\cdot)$  and  $\tau_\eta(\cdot)$ .

In the proposed transfer learning scenario, a task conditional visual classifier is pretrained with off-the-shelf visual dataset, *e.g.*, Visual Genome [22], and transferred to VQA. In the pretraining stage, the parameters of the classifier and two feature encoders  $\theta$ ,  $\phi_{\text{pre}}$  and  $\eta_{\text{pre}}$  are jointly learned. This stage allows the task conditional visual classifier to handle diverse visual recognition tasks by learning the task feature  $\tau$ . Transfer learning to VQA is achieved by reusing parameter  $\theta$  and adapting new encoders  $v_{\phi_{\text{vqa}}}(\cdot)$  and  $\tau_{\eta_{\text{vqa}}}(\cdot)$  to the learned task conditional visual classifier.

#### 3.2. Pretraining

Learning the task conditional visual classifier is naturally formulated as the problem to maximize the following expected log likelihood:

$$\theta^*, \phi_{\text{pre}}^*, \eta_{\text{pre}}^* = \underset{\theta, \phi_{\text{pre}}, \eta_{\text{pre}}}{\operatorname{argmax}} \mathbb{E}_{p_{\mathcal{D}}} [\log p_\theta(\mathbf{a}|v_{\phi_{\text{pre}}}(\mathbf{I}, \mathbf{b}), \tau_{\eta_{\text{pre}}}(\mathbf{t}))], \quad (1)$$

where  $v_{\phi_{\text{pre}}}(\mathbf{I}, \mathbf{b})$  is a visual feature based on an image  $\mathbf{I}$  and a bounding box  $\mathbf{b}$ , and  $\tau_{\eta_{\text{pre}}}(\mathbf{t})$  is a task feature encoded from a task specification  $\mathbf{t}$ ,  $\mathbf{a}$  is an answer sampled from data distribution and it satisfies  $\mathbf{a} \in \mathcal{A}$ , and  $\{\theta, \phi_{\text{pre}}, \eta_{\text{pre}}\}$  are model parameters. We obtain  $v_{\phi_{\text{pre}}}(\mathbf{I}, \mathbf{b})$  using a learnable attention network parametrized by  $\phi_{\text{pre}}$  on top of off-the-shelf feature extractor [3], where the bounding box position  $\mathbf{b}$  is used as a key to the attention. The optimization in Eq. (1) requires a joint distribution,  $p_{\mathcal{D}}(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{t})$ , which is not accessible in the external datasets in our setting due to missing task specifications  $\mathbf{t}$ . Section 4 describes how to model the joint distribution  $p_{\mathcal{D}}(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{t})$  with the visual

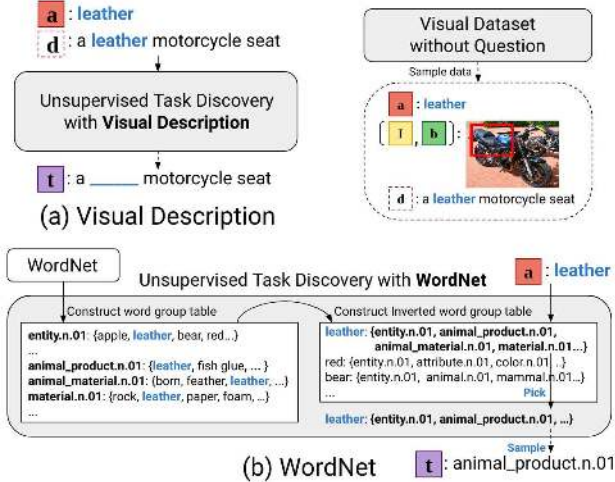


Figure 3. **Unsupervised task discovery with two different linguistic knowledge sources.** (a) For visual description, task specification is generated by replacing a visual word (label) from the description into a blank. (b) For WordNet, task specification is a synset from one of the labels hypernyms. We use *inverted wordset* to sample a synset from an input label. See Section 4 for details.

annotations and linguistic knowledge sources.

### 3.3. Transfer learning for VQA

As illustrated in Figure 2, the proposed VQA model contains a task conditional visual classifier  $p_\theta(\mathbf{a}|v, \tau)$ . The pretrained visual concepts are transferred to VQA by sharing the learned parameters  $\theta$ . Then, learning a VQA model is now formulated as learning input representations  $v$  and  $\tau$  for  $p_\theta(\mathbf{a}|v, \tau)$ , which is given by

$$\phi_{vqa}^*, \eta_{vqa}^* = \operatorname{argmax}_{\phi_{vqa}, \eta_{vqa}} \mathbb{E}_{p_{vqa}} [\log p_\theta(\mathbf{a}|v_{\phi_{vqa}}(\mathbf{I}, \mathbf{q}), \tau_{\eta_{vqa}}(\mathbf{q}))], \quad (2)$$

where  $v_{\phi_{vqa}}(\mathbf{I}, \mathbf{q})$  is an encoded visual feature with an image  $\mathbf{I}$  and a question  $\mathbf{q}$  using an attention mechanism with parameter  $\phi_{vqa}$  and an off-the-shelf feature extractor [3]. A task feature  $\tau_{\eta_{vqa}}(\mathbf{q})$  encodes a question  $\mathbf{q}$  using parameter  $\eta_{vqa}$ . The joint distribution of a training dataset for VQA,  $p_{vqa}(\mathbf{a}, \mathbf{I}, \mathbf{q})$ , is required for optimization, where answers from the distribution satisfy  $\mathbf{a} \in \mathcal{A} - \mathcal{B}$ . We learn  $\phi_{vqa}$  and  $\eta_{vqa}$  by maximizing the likelihood of the objective in Eq. (2) while the parameter for the pretrained task conditional visual classifier  $\theta$  remains fixed.

**Weakly supervised task regression** Utilizing a pretrained task conditional visual classifier for the visual recognition specified by a question  $\mathbf{q}$  requires to infer an optimal task feature  $\tau_{\mathbf{q}}^*$ . This requirement introduces a learning problem—task regression—that optimizes an encoder  $\tau_{\eta_{vqa}}(\mathbf{q})$  to predict  $\tau_{\mathbf{q}}^*$  correctly. Because directly minimizing error  $\mathcal{E}(\tau_{\mathbf{q}}^*, \tau_{\eta_{vqa}}(\mathbf{q}))$  requires additional supervision about the tasks, we instead exploit VQA data as a

source of weak supervision. We optimize indirect loss  $-\mathbb{E}_{\tau_{\mathbf{q}}^*(\mathbf{a}|v)} [\log p_\theta(\mathbf{a}|v, \tau_{\eta_{vqa}}(\mathbf{q}))]$ , which encourages the answer distributions conditioned on  $\tau_{\mathbf{q}}^*$  and  $\tau_{\eta_{vqa}}(\mathbf{q})$  to be similar. By assuming that true task conditional answer distribution  $\tau_{\mathbf{q}}^*(\mathbf{a}|v)$  is implicitly modeled in VQA dataset, we employ Eq. (2) as the objective function for weakly supervised task regression.

**Out-of-vocabulary answering** We learn a VQA model by adapting input representations while fixing the pretrained task conditional visual classifier  $p_\theta(\mathbf{a}|v, \tau)$ . This strategy allows a model to focus on learning to infer a visual recognition task  $\tau_{\eta_{vqa}}(\mathbf{q})$  from questions, which does not require data for all possible answers. Once the task feature  $\tau$  is inferred, the learned task conditional visual classifier  $p_\theta(\mathbf{a}|v, \tau)$  can answer the pretrained visual concepts including out-of-vocabulary ones.


**Matching visual features** To reuse pretrained visual classifier in VQA without fine-tuning, semantics of visual features  $v$  should not change by learning with VQA dataset. This is fulfilled in recent approaches for VQA models that do not fine-tune pretrained visual feature extractors and focus on learning attention mechanism [20] over the extracted feature map. In our setting, we simply use the identical visual feature extractor [3] for both pretraining and VQA.

## 4. Unsupervised Task Discovery

Learning a task conditional visual classifier with the off-the-shelf visual dataset [22] is not straightforward due to missing annotation for task specifications, which is necessary to learn an encoder for a task specification vector  $\tau$ . To address this issue, we propose unsupervised task discovery which samples a task specification  $\mathbf{t}$  from a task distribution modeled by exploiting linguistic knowledge sources.

### 4.1. Leveraging linguistic knowledge sources

A visual recognition task given by a question typically defines a mapping from visual inputs to a set of possible visual concepts, that is, a word group. For example, a question of “what is the woman holding?” defines a visual recognition task finding *holdable objects* in an image, which is a classification over a word group {ball, racket, cup, ...}. This intuition leads to a simple approach for modeling distribution of task description  $\mathbf{t}$  by treating a task as a word group (*i.e.* *holdable objects*). The main reason to use linguistic knowledge sources for the unsupervised task discovery is that the word groups are often accessible in the linguistic knowledge sources. We consider two linguistic knowledge sources: 1) visual descriptions provided with a visual data and 2) a structured lexical database called WordNet [10]. Figure 3 illustrates the overview of our approach.



task $t_w$	word groups
appliance.n.02	freezer, hair dryer, refrigerator, fridge, oven, dishwasher, washing machine, ...
opening.n.10	manhole, rear window, exit, nozzle, car window, coin slot, spout, vent, ...
food.n.02	potatoes, french fries, chicken, melted cheese, tomatoes, sausage, vegetable, ...
move.v.03	twisted, ruffled, curving, curved, wheeled, rolled, coiled, turned, rippled, ...
act.n.02	fishing, skateboarding, sailing, playing baseball, surfing, traveling, driving, ...
area.n.01	middle, in corner, center, parking space, playground, landscape, neighborhood, ...
color.n.01	yellow, pink, red, beige, royal blue, amber, aqua, dark red, olive green, teal, ...

Figure 4. **Illustration of WordNet and constructed word group table.** (Left) A subgraph of the WordNet [10]. Complex hierarchy of words reveals the diverse categorization of each words. (Right) A set of words sharing common parents in the tree is grouped as a single word group. Diverse grouping of words reveals diverse visual recognition tasks that can be defined on each word group.

## 4.2. Visual description

We use Visual Genome [22] as an off-the-shelf visual dataset, which determines a data distribution  $p_V(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d})$  based on a set of quadruples  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d})$  including visual descriptions  $\mathbf{d}$ . The description in this dataset is designed to mention answer  $\mathbf{a}$  explicitly, so that relation between the answer and the description is clear.

To this end, we define task specification  $t_d$  by replacing the answer in visual description to a special word  $\langle \text{blank} \rangle$ , which is formally denoted as  $t_d = \rho(\mathbf{d}, \mathbf{a})$ , where  $\rho(\mathbf{d}, \mathbf{a})$  is a function generating a blanked description. The subscript in  $t_d$  means that a task specification is extracted based on a visual description. Based on this definition, joint distribution,  $p_{\mathcal{D}}(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_d) = \int p(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_d, \mathbf{d}) d\mathbf{d}$ , with task specification is given by

$$p(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_d, \mathbf{d}) = p(t_d | \mathbf{d}, \mathbf{a}) p_V(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d}) \quad (3)$$

where  $p(t_d | \mathbf{d}, \mathbf{a}) = \delta(t_d, \rho(\mathbf{d}, \mathbf{a}))$  is a delta function that returns 1 if two inputs are identical and 0 otherwise. As illustrated in Figure 3a, we sample data required for pretraining  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_w)$  by first sampling a visual data  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d})$  and then sampling a task specification  $t_d$  from  $p(t_d | \mathbf{d}, \mathbf{a})$ . This procedure results in sampling description  $\mathbf{d}$  as well, but we do not care about it when we pretrain the task conditional visual classifier. For pretraining, we encode  $t_d$  into a task feature  $\tau_{\eta_{\text{pre}}}(t_d)$  based on a gated recurrent unit [8] because  $t_d$  is a sequence of words.

The main reason to use a blanked description for a task specification is that it is effective to define a set of candidate words. For example, a blanked description “*a man is holding \_\_\_\_\_*” restricts candidate words for the blank to a set of objects that can be held. Therefore, a blanked description can be used to determine a word group implicitly that represents a visual recognition task.

## 4.3. WordNet

WordNet [10] is a lexical database represented with a directed acyclic graph of disambiguated word entities, called synsets. A sample subgraph of WordNet is illustrated in Figure 4 (left). The graph represents a hierarchical structure of words, where the parents of a node correspond to

hypernyms of the word in the child. In WordNet, we define a task specification  $t_w$  as a synset for a node that is a common ancestor of multiple words because a set of words sharing a common ancestor constructs a word group and the word group may also define a visual recognition task.

A procedure for sampling a task specification  $t_w$  based on WordNet and a visual data  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d})$  is illustrated in Figure 3b. The main idea for the procedure is to model a task distribution conditioned on an answer  $p(t_w | \mathbf{a})$  as a uniform distribution over possible word groups that the answer belongs to, where a task specification  $t_w$  is a common ancestor of words in an word group. Modeling the distribution  $p(t_w | \mathbf{a})$  requires two stages: 1) constructing a word group table, which maps a task specification to a word group and 2) constructing an inverted word group table, which maps an answer word to a set of task specifications. The inverted word group table is used to retrieve a set of possible task specifications for an answer  $\mathbf{a}$  and the distribution  $p(t_w | \mathbf{a})$  is the uniform distribution over task specifications in the set. Given the distribution  $p(t_w | \mathbf{a})$ , the joint distribution,  $p_{\mathcal{D}}(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_w) = \int p(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_w, \mathbf{d}) d\mathbf{d}$ , is given by

$$p(\mathbf{a}, \mathbf{I}, \mathbf{b}, t_w, \mathbf{d}) = p(t_w | \mathbf{a}) p_V(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d}). \quad (4)$$

Therefore we sample a quadruple  $(\mathbf{a}, \mathbf{I}, \mathbf{b}, \mathbf{d})$  from the visual dataset and sample a task specification  $t_w$  subsequently. While this procedure samples a description as well, we marginalize it out. For pretraining, we encode  $t_w$  into a task specification vector  $\tau_{\eta_{\text{pre}}}(t_w)$  based on a word embedding function that is learned from scratch.

The word group table is constructed by selecting a synset of a node in WordNet as a task specification  $t_w$  and mapping it to a set of words (a word group) corresponding to all its descendants. Any word group can be defined regardless of its level in WordNet hierarchy and the part-of-speech of its members; the biggest word group contains all words in WordNet and its task specification corresponds to the root of WordNet. We illustrate the constructed word group table in Figure 4. The inverted word group table is constructed in a similar way to an inverted index of the word group table, but the range of mapping is not a set of indices but a set of task specifications.




Figure 5. **Model comparisons.** Exploiting external data with the unsupervised task discovery boosts performance of the proposed model and separable classifier significantly while the separable classifier shows limited gains on attribute answers with large variations.




Figure 6. **Data comparisons.** Using visual description and WordNet shows different generalization characteristics and combining them brings additional improvement.




Figure 7. **Complementary characteristics of data.** Visual description and WordNet show complementary characteristics in terms of VQA score for different answers.

## 5. Experiments

We evaluate how effectively the proposed framework leverages the external data without questions to answer out-of-vocabulary words in visual question answering. We compare the proposed method with the baselines equipped with idea for zero-shot image classification [11] and novel object captioning [4, 37], which are related to the proposed problem. We also analyze the impact of the external data used for pretraining, and visualize the mapping between questions and task specifications learned by weakly supervised

task regression. Notably, the evaluation setting is not comparable to that from zero-shot learning literature [42], as our objective is not generalization to completely unseen classes, but exploiting class labels exposed in the external data. We open-sourced all the codes and datasets used in the experiments to facilitate reproducing the result in this paper<sup>1</sup>.

### 5.1. Datasets

**Pretraining** We learn visual concepts of most frequently observed 3,000 objects and 1,000 attributes in the Visual Genome dataset [22]. To pretrain the task conditional visual classifier, we construct external visual data with bounding box annotations, which are provided with region descriptions. Then, visual words (answers) are extracted from region descriptions to construct visual data quadruples (a, I, b, d). We use 1,169,708 regions from 80,602 images to construct the training data. To use WordNet [10], we map visual words to synsets using synset annotations from Visual Genome dataset, and the words that are not covered by the annotations are mapped to synsets using Textblob [27].

**Dataset construction** We repurpose VQA v2 dataset to construct a training/test split as illustrated in Figure 1. We use training and validation set of VQA v2. To ensure that

<sup>1</sup>[https://github.com/HyeonwooNoh/vqa\\_task\\_discovery](https://github.com/HyeonwooNoh/vqa_task_discovery)



Figure 8. **Out-of-vocabulary answers with diverse types of concepts.** Green color denotes correct answers. All predicted answers are from out-of-vocabulary answers. The proposed model successfully predicts diverse out-of-vocabulary answers depending on questions.

Table 1. **Result of weakly supervised task discovery.** We retrieve questions for each task specification based on the similarity score with task features. Results show that appropriate task specifications are regressed from each question. Note that no explicit supervision is used for learning the mapping between questions and task specifications.

Task specification $t_w$	Questions
organic_process.n.01	what are the giraffes doing? / what are the animals doing? / what are the giraffes doing in this picture?
athletic_game.n.01	what type of sport ball is shown? / what type of sport are the men participating in?
furniture.n.01	what piece of furniture are the cats sitting on? / what furniture is the cat sitting on?
fruit.n.01	what type of fruit is the animal eating? / what type of fruit juice is on the counter?
time_period.n.01	what kind of season is it? / what type of season is it? / what holiday season is it? / which season is it?
tool.n.01	what utensil is in the person 's hand? / what utensil is laying next to the bread?
hair.n.01	what hairstyle does the surfer have ? / what type of hairstyle does this man have ?

every out-of-vocabulary answer appears during pretraining, we randomly select out-of-vocabulary answers from all the visual words used for pretraining (954 out of 3,813). Since we focus on transferability of visual words, answers about yes/no and numbers are not considered in our evaluation. Based on the selected out-of-vocabulary answers, we generate 3 questions splits—462,788 for training, 51,421 for validation, and 20,802 for testing. The training and validation splits do not contain out-of-vocabulary answers while the test split consist of out-of-vocabulary answers only. For evaluation of models, we follow the standard VQA protocol with 10 ground-truth answers for each question [5].

## 5.2. Baselines

Since leveraging external visual data for visual question answering with out-of-vocabulary answers has hardly been explored, there is no proper evaluation benchmark and we employ the following baselines to compare with our algorithm: 1) *answer embedding*, which employs idea from zero-shot image classification [11] that learns mapping from visual features to pretrained answer embedding, where we use GloVe [33] to embed each answer, and 2) *separable classifier*, which adopts idea from novel object captioning [4, 37] that learns visual and language classifier separately and combines them by element-wise sum of logits for joint inference. Note that the separable classifier and our proposed model are trained with the same data.

## 5.3. Results

**Model comparisons** Figure 5 illustrates model comparison results with the baselines. For this experiment, we perform VQA adaptation with 6 different random seeds and plot their mean and standard deviation. The standard VQA model fails to predict any out-of-vocabulary answers (*i.e.*, 0 VQA score) because there is no clue for inferring out-of-vocabulary answers. Answer embedding baseline generalizes slightly better by exploiting the similarity of answer words in the embedding space, but the improvement is marginal. Using off-the-shelf visual data and task specifications from linguistic knowledge sources dramatically improves performance both on the separable classifier and the proposed model. However, independent consideration of visual data and task specifications in separable classifier has a critical limitation of modeling joint interaction between task specifications and visual features. Especially, this baseline shows substantially lower performance on attribute answers, which have significant variations depending on tasks. Note that the bias in the VQA training set cannot be exploited in the proposed evaluation setting, as the evaluation is performed with out-of-distribution answers only.

**Data comparisons** Figure 6 illustrates the effect of different linguistic sources in our algorithm. Two models learned using visual description and WordNet, respectively, have complementary characteristics and additional improvement is achieved by exploiting both data. More detailed comple-




Figure 9. **Combining knowledge from VQA and external visual data.** Evaluation results on a test set containing both out-of-vocabulary answers and trained answers. The proposed model shows relatively lower performance on trained answers but significantly better performance on out-of-vocabulary answers. In total, the proposed model shows the best performance.

mentary characteristics of models are illustrated in Figure 7, where we visualize average VQA scores for 20 answers.

**Qualitative results** Figure 8 shows examples of predicted answers from the proposed model. The proposed model correctly predicts out-of-vocabulary answers for questions asking diverse visual concepts such as type of flooring, material, type of sport and brand.

**Weakly supervised task regression** Given that task specifications extracted from WordNet models diverse visual recognition tasks, matching them to relevant questions is useful for categorization of VQA data and model interpretation. As we learn VQA models by task regression, this matching can be performed by comparing the encoded task feature from a question  $\tau_{\eta_{\text{vqa}}}(\mathbf{q})$  and the vector from a task specification  $\tau_{\eta_{\text{pre}}}(\mathbf{t}_w)$ . For each  $\tau_{\eta_{\text{pre}}}(\mathbf{t}_w)$ , we sorted questions in a descending order of dot product similarity between  $\tau_{\eta_{\text{pre}}}(\mathbf{t}_w)$  and  $\tau_{\eta_{\text{vqa}}}(\mathbf{q})$ . In the sorted question list, the most similar questions are visualized in Table 1. The visualization shows that the weakly supervised task regression successfully trains a question encoder that match a question to a relevant task feature.

#### 5.4. Combining knowledge learned by VQA

While we focus on learning visual concepts from external visual data, VQA dataset is still a valuable source of learning diverse knowledges. Especially, some answers are not visual words and require visual reasoning. For example, yes and no are one of the most frequent answers in the VQA dataset [5] but it is not straightforward to learn these answers only with the external visual data. Therefore, we consider combining knowledge learned from VQA dataset and from external visual data.

We construct a split of the VQA dataset consisting of 405,228 training, 37,031 validation, and 172,681 test questions. The training and validation set do not contain any out-of-vocabulary answers and test set contains out-of-vocabulary answers. However, contrary to the main ex-

periment, the test set also contains training answers which include logical answers, numbers and visual words. The list of out-of-vocabulary answers are identical to that of the main experiment. Among 172,681 test questions, 103,013 questions can be answered only with the training answers.

To combine knowledges from VQA dataset and external visual data, we learn a VQA model with two task conditional visual classifiers; we fine-tune one classifier for adapting answers requiring visual reasoning (*i.e.*, numbers and yes/no) and fix the other classifier for visual answers including out-of-vocabulary answers. After training the VQA model, we combine two logits by element-wise sum and pick the answer with the highest score in the inference.

The results are presented in Figure 9. Models in each method are trained with 6 different random seeds and their mean and standard deviation are plotted. Overall, the proposed model performs the best. While the standard VQA model achieves the best performance for training answers, it fails to predict any out-of-vocabulary answers. The answer embedding baseline somewhat generalizes to out-of-vocabulary answers, but constraints in the answer embedding degrade its performance on answers in training set.

#### 6. Conclusion

We present a transfer learning approach for visual question answering with out-of-vocabulary answers. We pre-train a task conditional visual classifier with off-the-shelf visual and linguistic data based on unsupervised task discovery. The pretrained task conditional visual classifier is transferred to VQA adaptively. The experimental results show that exploiting external visual and linguistic data boosts performance in the proposed setting and training with unsupervised task discovery is important to model interaction between visual features and task specifications.

**Acknowledgments** This research was partly supported by Kakao and Kakao Brain and Korean ICT R&D program of the MSIP/IITP grant [2016-0-00563, 2017-0-01778].



## References

- [1] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don't Just Assume; Look and Answer: Overcoming Priors for Visual Question Answering. In *CVPR*, 2018. 1, 2
- [2] Aishwarya Agrawal, Aniruddha Kembhavi, Dhruv Batra, and Devi Parikh. C-VQA: A Compositional Split of the Visual Question Answering (VQA) v1.0 Dataset. *arXiv preprint arXiv:1704.08243*, 2017. 2
- [3] 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 *CVPR*, 2018. 1, 2, 3, 4
- [4] Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond Mooney, Kate Saenko, Trevor Darrell, Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, et al. Deep Compositional Captioning: Describing Novel Object Categories without Paired Training Data. In *CVPR*, 2016. 2, 6, 7
- [5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *ICCV*, 2015. 1, 2, 7, 8
- [6] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. DBpedia: A Nucleus for A Web of Open Data. In *The Semantic Web*, pages 722–735. Springer, 2007. 2
- [7] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *ACM SIGMOD International Conference on Management of Data*, 2008. 2
- [8] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv preprint arXiv:1412.3555*, 2014. 5
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 1, 2
- [10] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. Bradford Books, 1998. 2, 4, 5, 6
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. DeViSE: A Deep Visual-Semantic Embedding Model. In *NIPS*, 2013. 3, 6, 7
- [12] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. In *EMNLP*, 2016. 1, 2
- [13] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *CVPR*, 2017. 1, 2
- [14] Tanmay Gupta, Kevin Shih, Saurabh Singh, and Derek Hoiem. Aligned Image-Word Representations Improve Inductive Transfer across Vision-Language Tasks. In *CVPR*, 2017. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 2
- [16] Hexiang Hu, Wei-Lun Chao, and Fei Sha. Learning Answer Embeddings for Visual Question Answering. In *CVPR*, 2018. 2
- [17] Sheng Huang, Mohamed Elhoseiny, Ahmed Elgammal, and Dan Yang. Learning Hypergraph-Regularized Attribute Predictors. In *CVPR*, 2015. 3
- [18] 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 *CVPR*, 2017. 2
- [19] Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. DVQA: Understanding Data Visualizations via Question Answering. In *CVPR*, 2018. 2
- [20] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard Product for Low-Rank Bilinear Pooling. In *ICLR*, 2016. 4
- [21] Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Skip-Thought Vectors. In *NIPS*, 2015. 2
- [22] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yanis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual Genome: Connecting Language and Vision using Crowdsourced Dense Image Annotations. *IJCV*, 123(1):32–73, 2017. 1, 2, 3, 4, 5, 6
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012. 2
- [24] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The Open Images Dataset V4: Unified Image Classification, Object Detection, and Visual Relationship Detection at Scale. *arXiv preprint arXiv:1811.00982*, 2018. 1

- [25] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-Based Classification for Zero-Shot Visual Object Categorization. *TPAMI*, 2014. 3
- [26] Hugo Larochelle, Dumitru Erhan, and Yoshua Bengio. Zero-Data Learning of New Tasks. In *AAAI*, 2008. 3
- [27] Steven Loria. TextBlob: Simplified Text Processing. <http://textblob.readthedocs.io/en/dev/>, 2018. 6
- [28] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural Baby Talk. In *CVPR*, 2018. 2
- [29] Mateusz Malinowski and Mario Fritz. A Multi-World Approach to Question Answering about Real-World Scenes Based on Uncertain Input. In *NIPS*, 2014. 2
- [30] Hyeonwoo Noh, Paul Hongsuck Seo, and Bohyung Han. Image Question Answering using Convolutional Neural Network with Dynamic Parameter Prediction. In *CVPR*, 2016. 2
- [31] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. We Dont Need No Bounding-Boxes: Training Object Class Detectors Using Only Human Verification. In *CVPR*, 2016. 1
- [32] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. Extreme Clicking for Efficient Object Annotation. In *ICCV*, 2017. 1
- [33] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global Vectors for Word Representation. In *EMNLP*, 2014. 2, 7
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NIPS*, 2015. 2
- [35] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. Tips and Tricks for Visual Question Answering: Learnings from the 2017 Challenge. In *CVPR*, 2018. 2
- [36] Damien Teney and Anton van den Hengel. Zero-Shot Visual Question Answering. *arXiv preprint arXiv:1611.05546*, 2016. 2
- [37] Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond Mooney, Trevor Darrell, and Kate Saenko. Captioning Images With Diverse Objects. In *CVPR*, 2017. 2, 6, 7
- [38] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer Networks. In *NIPS*, 2015. 2
- [39] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. FVQA: Fact-Based Visual Question Answering. *TPAMI*, 2017. 2
- [40] Peng Wang, Qi Wu, Chunhua Shen, and Anton van den Hengel. The VQA-Machine: Learning How to Use Existing Vision Algorithms to Answer New Questions. In *CVPR*, 2017. 2
- [41] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Ask Me Anything: Free-Form Visual Question Answering based on Knowledge from External Sources. In *CVPR*, 2016. 2
- [42] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Shot Learning-The Good, the Bad and the Ugly. In *CVPR*, 2017. 3, 6
- [43] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked Attention Networks for Image Question Answering. In *CVPR*, 2016. 1, 2
- [44] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Incorporating Copying Mechanism in Image Captioning for Learning Novel Objects. In *CVPR*, 2017. 2
- [45] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7W: Grounded Question Answering in Images. In *CVPR*, 2016. 2