

Visual Question Generation for Class Acquisition of Unknown Objects

Kohei Uehara¹[0000-0003-0066-319X], Antonio Tejero-De-Pablos¹,
Yoshitaka Ushiku¹, and Tatsuya Harada^{1,2}

¹ The University of Tokyo

{uehara, antonio-t, ushiku, harada}@mi.t.u-tokyo.ac.jp

² RIKEN

Abstract. Traditional image recognition methods only consider objects belonging to already learned classes. However, since training a recognition model with every object class in the world is unfeasible, a way of getting information on unknown objects (i.e., objects whose class has not been learned) is necessary. A way for an image recognition system to learn new classes could be asking a human about objects that are unknown. In this paper, we propose a method for generating questions about unknown objects in an image, as means to get information about classes that have not been learned. Our method consists of a module for proposing objects, a module for identifying unknown objects, and a module for generating questions about unknown objects. The experimental results via human evaluation show that our method can successfully get information about unknown objects in an image dataset. Our code and dataset are available at <https://github.com/mil-tokyo/vqg-unknown>

Keywords: Visual question generation, Unknown object recognition, Unknown object class acquisition, Real world recognition

1 Introduction

In recent years, in large-scale image classification tasks, image classifiers with deep convolutional neural networks (CNN) have achieved accuracies equivalent to humans [22, 7]. The recognition capabilities of these methods are limited by the object classes included in the training data. However, for an image recognition system running in the real world, for example a robot, considering all existing object classes in the world during training is unfeasible. If such a robot was able to ask for information about *objects it cannot recognize*, the robot would not have to learn all classes in advance. In this paper, we define an *unknown object* as an object belonging to a class not included in the training data. In order to acquire knowledge about the unknown object class, the most reliable way is to obtain information directly from humans. For example, the robot can present an image to a human and ask them to annotate the class of an object, as in active learning [12]. When the class is unknown, selecting the appropriate object and generating a suitable question about it is a challenging problem, and has not

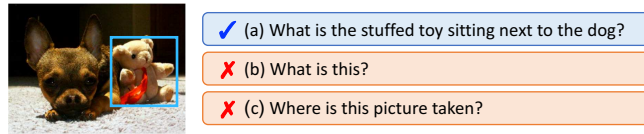


Fig. 1. Examples of suitable/unsuitable questions for unknown objects. A suitable question should specify the target object (*stuffed toy*), so the answer is the class of the unknown object (*teddy bear*). Therefore, questions such as (a) are suitable. On the other hand, simple questions such as (b) and questions about location such as (c) are unsuitable.

been tackled yet. The goal of this research is to generate questions that request information about a specific unknown object in an image. As shown in Fig. 1, compared to a simple question such as “*What is this?*”, a specific question such as “*What is the stuffed toy sitting next to the dog?*” targets better the class of the unknown object. There exist several approaches [16, 9] for general visual question generation (VQG) by using recurrent neural networks (RNN). Also, VQG with a specific target for a question has also been studied [13] by providing a *target word* (i.e., a word indicating what object the question is targeting) to an RNN as a condition. However, in these works, only the known classes are given as the target word. To the best of our knowledge, VQG targeting unknown objects has not been studied yet. Also, in order to realize a VQG method for unknown objects, first we need to detect and classify the unknown object. However, we cannot rely on object classification methods [7] nor object region proposal methods [21] if they only consider known/labeled classes (i.e., supervised learning).

In this paper, to find unknown objects in an image, we propose object regions by selective search [28], which is not based on supervised learning, and then classify whether the proposed objects are unknown or not. Since our method has to classify all the objects, in order to reduce the execution time we propose an efficient unknown object classification based on uncertainty prediction. In addition, we approach VQG for unknown objects by generating a question containing the *hypernym* of the unknown object. The hypernym of a given word is another word that is higher in the semantic hierarchy, such as “*animal*” for “*dog*.”

Contributions: (1) We propose the novel task of automatically generating questions to get information about unknown objects in images. (2) We propose a method to generate questions using the semantic hierarchy of the target word. (3) We construct the whole pipeline by combining modules of object region proposal, unknown object classification, and visual question generation, and show that it can successfully acquire information about unknown objects from humans.

The paper is organized as follows. First, we explain previous studies related to this research in Sec. 2. Next, we introduce our proposed system in Sec. 3. Then, we show experiments on our module for unknown object classification in Sec. 4, and our visual question generation module in Sec. 5. In Sec. 6, we evaluate our entire pipeline to get the class of unknown objects. Finally, in Sec. 7, we discuss the conclusions and future work.

2 Related Works

First, we explain active learning, an information acquisition method that also considers human help for learning. Next, we introduce the research related to each of our modules, namely, object detection, unknown object classification, and visual question generation.

Active Learning. The aim of active learning is achieving efficient learning by automatically selecting data that seems to contribute the most to improve the performance of the classifier and requesting a human annotator to label them. Uncertainty Sampling [12] has been proposed to select the instances whose class is the least certain. There are three methods for Uncertainty Sampling: (1) *Least Confident* [12]: Select the instance whose classification probability is the smallest and whose class has the greatest overall classification probability. (2) *Margin Sampling* [24]: Select the instance whose difference between the most and the second most classification probabilities is the smallest. (3) *Entropy Sampling* [26, 8]: Select the instance whose distribution of classification probabilities has the largest entropy. The main difference between active learning and this research is that active learning targets only instances whose class is included in the training set, whereas we target instances whose class is not in the training set. Also, active learning only presents data to the annotator; it does not generate questions.

Object Region Proposal. Object region proposal methods detect the region surrounding objects in an image. Recent methods perform object detection that performs both object region proposal and object classification at the same time via supervised learning using CNN [21, 20]. These methods achieve accurate object detection with a huge amount of labeled data for training. However, they do not consider unknown objects. In contrast, there is some research on *objectness* that simply estimates the existence objects in a specific region of the image, without classifying the object. Alexe et al. [1] perform objectness estimation by using saliency, contrast, edge, and superpixel information. Cheng et al. [4] learn objectness from image gradients. Also, a method called selective search [28, 23] allows object region proposal by using image segmentation, and integrating similar regions with each other. Since it does not require object labels, it can propose regions without learning the object class.

Unknown Object Classification. Unknown object classification performs binary classification of objects in an image as *known* or *unknown*. Traditionally, object classification methods estimate the actual class of an object in an input image. Recent research in object classification are CNN-based methods [7, 27]. These methods assume a *closed set*, that is, they only consider the classes included in training and not unknown classes. On the other hand, there is research on the task called *open set recognition* [25] for object classification that includes unknown objects. Open set recognition is a task aimed at classifying to the correct class if the input belongs to the trained class, and if the input is unknown, classifying it as unknown. For open set recognition, methods using SVM [25] and methods extending the nearest neighbor method [2] have been proposed. Also, Bendale et al. [3] proposed open set recognition using CNN. They classify an

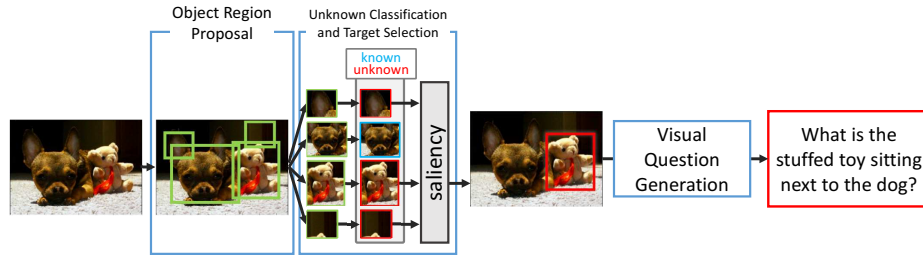


Fig. 2. Overview of the proposed method. First, regions from objects in the image (including unknown objects) are detected. Then, unknown objects are classified and the target region is selected. Finally, the target region along with the whole image is coded into a feature vector, and a question for the unknown object is generated

object as unknown if its feature distribution extracted from the CNN hidden layers is distant from known classes.

Visual Question Generation. Visual Question Generation (VQG) was recently proposed as an extension of image captioning. Whereas image captioning methods [29] generate descriptive sentences about the content of an image, VQG methods generate questions (e.g., *What color is the car?*). The common approach in VQG is encoding image features via CNN and generating a sentence by decoding those features using an RNN. Methods that use a gated recurrent unit (GRU) [16] and a long short-term memory (LSTM) [9] have been proposed. Traditional VQG methods generate questions from the whole image, without focusing in any particular image region. Only recently, methods that generate questions targeting a particular image region have been proposed. Zhang et al. [30] detect different regions to generate a variety of questions from the same image. In contrast, Li et al. [13] generate questions focusing on a specific region with the goal of distinguishing between two images. For this, they input a target word (e.g., blue) related to the region as a condition to the LSTM. In [13], target words are known classes learned in advance. To the best of our knowledge, VQG targeting an unknown object has not been approached yet.

3 Proposed System

Fig. 2 shows the overview of the proposed method. First, objects in the input image are detected by the object region proposal module. Next, the unknown object classification and target selection module identifies whether each object is unknown or not, and selects an object region to be the target of the question. We refer to this region as the *target region*. Finally, the visual question generation module generates a question using features extracted from the whole image and the target region.

3.1 Object Region Proposal

Our object region proposal module detects all objects in the input image via selective search. The proposed method needs to detect unknown objects (i.e., objects never learned before), so supervised learning is not an option since it requires labels for all objects. As mentioned in Sec. 2, selective search provides candidate regions for objects without supervised learning. Thus, unknown objects can also be detected, and the number of object regions can be reduced compared with an exhaustive search. Therefore, this seems to be suitable as a method for object region proposal.

3.2 Unknown Object Classification and Target Selection

This module selects the *target object*, that is, the object to acquire information about. For that, we classify objects into *known* or *unknown*, and then select the most salient unknown object. This prevents generating questions about unimportant regions that may have been proposed by mistake by the object region proposal module. We define unknown object classification as follows: for an input object image, if its class is included in the training set, classify it to the correct class, and if not, classify it as unknown. Specifically, we perform unknown object classification on the classification results of a CNN as follows. The output of the softmax function of the CNN can be regarded as the confidence with which the input is classified into a certain class. We consider that images of unknown objects result in a low confidence value for all classes. That is, the more uniform the confidence distribution, the lower the confidence for all classes and the more possibilities the object is unknown. Therefore, we perform unknown object classification by estimating the dispersion of the probability distribution using an entropy measure, with reference to the method of Uncertainty Sampling in active learning [26, 8]. The entropy measure E is defined as:

$$E = - \sum_{j=1}^K p_j \log_2 p_j \quad (1)$$

where p_j is the output of the softmax function when a given input x is classified into class C_j ($j = 1, 2, \dots, K$). E takes the maximum value $\log_2 K$ when all p_j are all equal, that is, when $p_j = 1/K$. On the other hand, the larger the dispersion of the probability distribution is, the smaller the entropy becomes.

Also, it is necessary to select which object to generate a question about among the objects classified as unknown. For example, in some cases, the region proposed by selective search contains only the background. Background regions are likely to be classified as unknown, but they do not contain an object to ask about. In order to solve this problem, we calculate the saliency of each proposed region in the image as a criterion for selecting the target region. That is, we ask questions about objects that are unknown and particularly salient in the image. Thus, to select salient objects in the image, we propose using a saliency map. The saliency map is a plot obtained by estimating the saliency for each pixel in

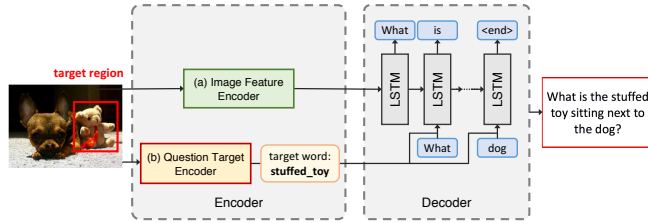


Fig. 3. Overview of the VQG module. First, we obtain the common hypernym of the prediction class of the classifier as the target word. The target word and the image features are input as conditions to the LSTM, and the question is generated

the image. We calculate the saliency map using the method of Zhu et al. [31]. This method estimates low saliency for background pixels and high saliency for foreground pixels. Therefore, it is considered to be suitable for this research. First, we preprocessed the image by applying mask based on saliency map and applied non-maximum suppression to reduce the large number of object regions. Then, the saliency of each proposed object region is expressed by:

$$I_{region} = \sum_{I(p) \geq \theta} I(p) \times \frac{S_{saliency}}{S_{region}} \quad (2)$$

where, $I(p)$ is the saliency value of each pixel, θ is the threshold value, $S_{saliency}$ is the area in the region where saliency exceeds θ , and S_{region} is the total area of the region. The threshold θ was determined using Otsu method [18]. The region with the highest saliency is selected as the target region.

3.3 Visual Question Generation

Figure 3 depicts the visual question generation module. We generate a question following the encoder-decoder methodology of Mostafazadeh et al. [16] and Li et al. [13]. The encoder extracts visual features of both the entire image and the object region (submodule (a)), and (submodule (b)) the target word into a word embedding vector representation. The decoder takes the encoded features and generates a question via LSTM.

Encoding of Image Features. This submodule uses a pretrained CNN model to extract the features f_I of the entire image and the features f_R of the target region. In our method, we use the output (a 1,000-dimensional vector) of the fc layer of ResNet152 [7]. Then, in order to express the spatial information of the target region, we follow the method by Li et al. [13] to define a five-dimensional vector l_R as:

$$l_R = \left[\frac{x_{tl}}{W}, \frac{y_{tl}}{H}, \frac{x_{br}}{W}, \frac{y_{br}}{H}, \frac{S_R}{S_I} \right] \quad (3)$$

where $(x_{tl}, y_{tl}), (x_{br}, y_{br})$ is the upper left and the lower right coordinate of the target region, S_R and S_I represent the area of the target region and the entire

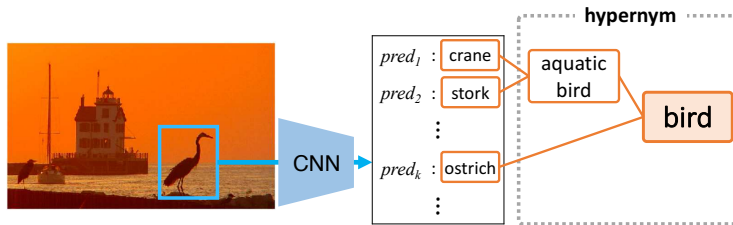


Fig. 4. Overview of the question target module for target word selection using WordNet [6]. WordNet is used to obtain a hypernym common to the predicted class labels with the highest confidence. The hypernym is then input to the visual question generation module

image respectively, and W and H denote the width and the height of the image respectively. We concatenate f_I , f_R , l_R , and let the 2,005 dimensional vector $f = [f_R, f_I, l_R]$ be the image feature encoding.

Question Target. This submodule selects a target word to represent the object in the target region and embeds it into a vector representation. Since the target object is unknown, that is, is not in the trained classes, it is not possible to use the class label as the target word as in Li et al. [13]. Therefore, we need to devise how to specify the target word. For example, if we do not know the class dog, asking a question referring to an *animal* is natural (e.g., “*What is this animal?*”). In this case, the word *animal* for *dog* is considered to be a hypernym. Such hypernym can be used as the target word. We use WordNet [6] to get the hierarchical relationship of words. Each word in WordNet is hierarchically arranged based on semantic relationships, and thus, it is possible to get the hypernym of a word by going up in the hierarchy.

As shown in Fig. 4, we use the k predicted classes ($pred_1, pred_2, \dots, pred_k$) with the highest confidence of the classification result, and we select the word with the lowest level among the common hypernyms of the k class labels. If the value of k is too large, the common hypernym becomes a very abstract word such as *whole* or *entity*, and it is not possible to designate the target appropriately. Therefore, the value of k should be chosen carefully.

Then, we use the Poincaré Embeddings [17] to embed words into feature vectors using a neural network similar to Word2Vec [14, 15]. However, unlike Word2Vec, Poincaré Embeddings are suitable for expressing a structure in which words are hierarchically represented, such as WordNet, as a vector. Let the target word embedded by Poincaré Embeddings be the vector $\sigma(v)$. Then the input to the decoder LSTM is the visual feature vector f , and the conditional input is the word embedded vector $\sigma(v)$. The decoder LSTM is trained by minimizing the negative log likelihood:

$$L = \sum -\log p(Q | f, \sigma(v); \theta) \quad (4)$$

where θ is the parameters of the LSTM, and Q denotes the generated question.

4 Evaluation of the Unknown Object Classification

Before evaluating the entire pipeline we performed experiments on independent modules to study their performance. First, we evaluated how accurately the unknown object classification module can classify whether the input image is an unknown object image or not.

4.1 Experimental Settings

We used CaffeNet [11], VGGNet [27], and ResNet152 [7], which are well-known CNN models, to study the variation in unknown classification accuracy when employing different classifiers. We pretrained our classifier with the 1,000 class dataset used in the object classification task of ILSVRC2012. We used 50,000 images of the same dataset for validation. Then, we used the dataset in the object classification task of ILSVRC2010 to create the unknown dataset. We excluded all images whose class is *known* (i.e., included in the ILSVRC2012 dataset), as well as the images whose class is a hypernym of any *known* classes. The reason for removing hypernyms is to avoid including general classes (e.g., “dog”) in the unknown dataset when the specific class is already known (e.g., “chihuahua”). Thus, we selected 50,850 images of 339 classes from ILSVRC2010 dataset, which are not included in ILSVRC2012 and its hypernyms.

4.2 Methods

We compare unknown object classification using entropy E , which is the proposed method, with the following two methods.

Least Confident [12] We used the method of Uncertainty Sampling described in Sec. 2. We set a threshold for the softmax probability of the class label with the highest probability, that is, the most probable label. Then, if the probability is lower than the threshold, the label is considered as unknown.

Bendale et al. [3] We used the method of Bendale et al. mentioned in Sec. 2. We performed this experiment based on the code published by the authors, and change the classification models for CaffeNet, VGGNet, and ResNet.

4.3 Evaluation Metrics

We calculated the F measure as:

$$F = \frac{2TP}{2TP + FP + FN} \quad (5)$$

where TP is defined as the number of known data classified into the correct class, FN as the number of unknown data misclassified into known data, and FP as the number of misclassified known data [3]. We performed the evaluation using a five-fold cross-validation.

Also, we measured the execution time of each method. First, we measured the time required by the classifier to calculate the softmax probability distribution

Table 1. Comparison of the proposed unknown object classification method in terms of F measure results \pm standard error. We performed experiments on CaffeNet, VGGNet, and ResNet. In all three cases, the proposed method outperformed the other methods

	F measure		
	CaffeNet	VGGNet	ResNet
Ours	$0.526 \pm 1.1 \cdot 10^{-3}$	$0.602 \pm 0.2 \cdot 10^{-3}$	$0.654 \pm 0.9 \cdot 10^{-3}$
Least Confident	$0.522 \pm 1.1 \cdot 10^{-3}$	$0.590 \pm 1.5 \cdot 10^{-3}$	$0.635 \pm 1.2 \cdot 10^{-3}$
Bendale et al. [3]	$0.524 \pm 0.9 \cdot 10^{-3}$	$0.553 \pm 0.6 \cdot 10^{-3}$	$0.624 \pm 1.7 \cdot 10^{-3}$

Table 2. Comparison of the proposed unknown object classification method in terms of execution time, with CaffeNet as a classifier. We performed classification for 100 images and showed the average time per image \pm standard error

	time (sec/image)
Ours	0.0400 ± 0.0017
Least Confident	0.0365 ± 0.0019
Bendale et al. [3]	15.6 ± 0.7

of one image. Next, we experimented the calculation time per image for each method, taking the distribution of softmax probability for 100 images as input. We repeated this operation five times and calculated the average execution time.

4.4 Experimental Results

Table 1 shows the resulting F measure per classifier and method. For all three methods, the F measure increased as the classifier was changed from CaffeNet to VGGNet, and to ResNet. The higher the accuracy of the classifier, when inputting known classes, the distribution of the classification probabilities varies more largely, and thus the entropy becomes smaller. In the case of inputting unknown classes, the more accurate the classifier is, the distribution of the classification probabilities is more uniform and thus the entropy becomes larger.

Table 2 shows a comparison of the execution time for each method when CaffeNet is used as a classifier. Our method and the Least Confident method take much less time than the method of Bendale et al. This is because the method of Bendale et al. has to calculate the distance to the average distribution of 1,000 known classes for each image, so the calculation cost is large, but in the method using entropy and the method using threshold of confidence, calculation is performed only with distribution of the input image, so calculation time is shortened.

5 Evaluation of the Visual Question Generation

We studied the performance of the proposed visual question generation module given a target region and compared to other methods.

5.1 Datasets

In this experiment, we used a dataset called Visual Genome [10] with about 100,000 images, and captions and questions. There is a subset of questions that is associated with a specific region in the image. We preprocessed the data as follows. First, we removed questions not beginning with “*What.*” Furthermore, since questions about colors are not the goal of our method, we also removed questions beginning with “*What color.*” Next, for questions associated with a specific region in the image, if the word representing the object in the region was included in the answer of the question, that word was taken as the target word. For questions not associated with an image region, we searched the object included in the answer among all objects in the image. Then, if there is only one instance of the object in the image, the word and the region where the object exists are set as the target word and the target region corresponding to the question. Furthermore, in order to eliminate the imbalance in the type of questions in the data, we limited to 50 the maximum number of times the same question can be included. Through this preprocessing, we gathered 202,208 questions corresponding to specific target regions in the image (one question per region) and 528 target words.

5.2 Methods

We split the 202,208 questions into 179,997 questions for training, 10,002 questions for validation, and 12,209 questions for testing. At training time, questions were generated by inputting images, regions, and target words. For embedding target words, we used Poincaré Embeddings trained on the tree structure of WordNet. We used the following methods as a baseline to compare our proposed method.

CNN + LSTM. As in Mostafazadeh et al. [16], we generated questions by inputting only the features of the entire image encoded by a CNN.

Retrieval. Following Mostafazadeh et al. [16], we also used retrieval method as baseline. First, we extracted features of the target regions in the training images using the *fc* layer of ResNet152. Then we retrieved the m regions with the higher cosine similarity between their features and the input target region. Then, for each question associated to the retrieved region, we calculated the similarity with the other $m - 1$ questions using the BLEU score [19], which measures textual similarity. Finally, the question with the highest BLEU score, that is, the most representative question, was taken as the final output.

5.3 Evaluation Metrics

In our experiments, we use BLEU [19] and METEOR [5] for measuring the similarity between the automatically generated questions and the ground truth. The larger the value, the more accurate the result.

Besides the automatic evaluation, we also performed human evaluation via Amazon Mechanical Turk (AMT)³. We presented an image with the target re-

³ <https://www.mturk.com/>

Table 3. Comparison between our method and the baseline in terms of automatic evaluation metrics. The proposed method outperformed baseline methods

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR
Ours	0.518	0.359	0.244	0.175	0.197
CNN + LSTM	0.456	0.296	0.175	0.110	0.163
Retrieval	0.438	0.275	0.157	0.094	0.151

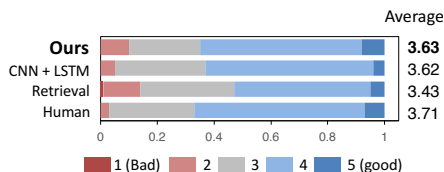


Fig. 5. (1) Human evaluation results on the naturalness of questions

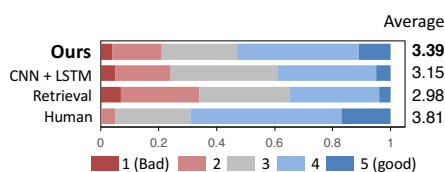


Fig. 6. (2) Human evaluation results on the relevance of questions to their region

gion and the target word to the human workers. We asked workers to blindly evaluate each method and the ground truth using a score between 5 (best) and 1 (worst). We used two criteria for evaluation: (1) whether each question is expressed naturally, and (2) whether each question is related to the target region and the target word. For the human evaluation, we used questions generated for 100 images extracted randomly from the test data.

5.4 Experimental Results

As shown in Table 3, the proposed method outperformed the baselines for all metrics. This result suggests that inputting the target object (visual features and target word condition) to the decoder LSTM allows generating more accurate questions. Figures 5 and 6 show the results of the human evaluation of our method compared to baselines. From the viewpoint of the naturalness of the question, the difference between the proposed method and CNN + LSTM was small. We believe the reason is that both methods use LSTM as the decoder. When evaluating the relevance of the question to the target region, the proposed method outperformed baselines. This is because, CNN + LSTM does not specify a target object to the decoder, so it may generate questions that are related to the image but not the target region. Also, Retrieval can generate only questions existing in the training data, so the variety of questions is limited, and thus, it may not generate questions related to the target region.

6 Evaluation of VQG for Unknown Objects

Lastly, we performed experiments using the whole pipeline, in which we generate questions to acquire knowledge about the class of unknown objects in the image.

6.1 Datasets

In order to test our VQG method for unknown objects, we used images that include unknown objects extracted from the following two datasets. First, from the test set of Visual Genome, we extracted 50 images with unknown objects, that is, not included in the 1,000 classes of ILSVRC2012. Also, from the dataset of the ILSVRC2010, 50 images of 339 unknown classes as described in Sec. 4 were extracted. In the images from the Visual Genome dataset, target regions contain an average of 8.7 objects, including small objects like “eye” and “button”. According to our method, 68.4 % of those objects were unknown. Note that we cannot indicate objectively the number of objects in the images from the ILSVRC2010 dataset since its ground truth does not include object regions.

6.2 Methods

The classifier used for unknown object classification was ResNet152, which is the method with the highest accuracy in Sec. 4. We pretrained ResNet152 with the 1,000 class data used in the object classification task of ILSVRC2012. The visual question generation module was pretrained with the dataset created in Sec. 5.1. Furthermore, as described in Sec. 3.3, when choosing a hypernym common to the top k classification results, if the value of k is too large, the target word becomes too abstract. Therefore, we performed experiments with two settings, $k = 2$ and $k = 3$. As baseline methods, we used the CNN + LSTM method and the Nearest Neighbors Retrieval method described in Sec. 5.

6.3 Evaluation Metrics

Since there is no ground truth in this experiment, it is not possible to perform automatic evaluation by comparison with the ground truth. Therefore, we performed only human evaluation via AMT, which consists of the following two tasks.

(1) We presented to three workers images and the questions generated automatically by our method and the baselines, and asked them to answer the generated questions. When they cannot understand the meaning of the question, we instructed them to answer “*Do not understand.*” Note that this task did not present a target region.

(2) Also, we evaluated the question and the answer obtained in task (1). Specifically, we presented to three workers with the question, the answer of each worker in task (1), and the image with the target region, and asked them whether the question and the answer are related to the target region in a 5-point scale. We evaluated only answers different from “*Do not understand.*”

As the evaluation value for task (2), we used the median of the evaluation values of the three workers.

Lastly, we evaluated to what extent the generated questions are able to successfully acquire information on unknown objects. We counted only the questions whose answers (task (1)) are not included in the known classes of the classifier,

		target word	generated question			target word	generated question			target word	generated question			
		k=2	camera	What is the woman holding in her right hand?			k=2	garment	What type of shirt is the man wearing?			k=2	instrumentality	What is on the man's lap?
		k=3	equipment	What is the woman looking at?			k=3	artifact	What is the man holding in the right hand?			k=3	instrumentality	What is on the man's lap?
		CNN + LSTM		What is the man wearing on his face?			CNN + LSTM		What is the man wearing on his hand?			CNN + LSTM		What is the man wearing on his face?
		retrieval		What is in between the people?			retrieval		What is the man closest to the camera sitting on?			retrieval		What is attached to the skateboard?

Fig. 7. Examples of input images (upper), the target words and generated questions by our proposed VQG method for unknown objects (middle), and the generated questions by the *CNN + LSTM* and *retrieval* baselines (lower).

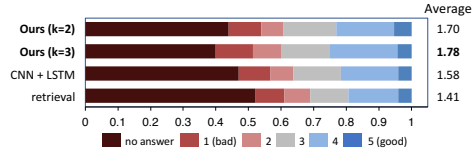


Fig. 8. Comparison of our method with the baseline in terms of the human evaluation in task (2). Task (2) evaluates whether or not the generated question, the image region, and the obtained answer are related. The greater the score, the higher the relevance.

Table 4. The number of generated questions that successfully allowed acquiring information on unknown objects (out of 300). We counted only the questions whose answers (task (1)) are not included in the known classes of the classifier, and the relevance of the question and target region in the image (task (2)) is four or more.

Ours($k = 2$)	61
Ours($k = 3$)	49
CNN + LSTM	46
Retrieval	45

and the relevance of the question and target region in the image (task (2)) is four or more.

6.4 Experimental Results

Fig. 7 shows our qualitative results. In Fig. 7 (a) and (b), when $k = 2$, the target word is a concrete word (i.e., “camera” and “garment”), and the generated question refers to an object in the target region. In the case of $k = 3$, the target word is an abstract word such as “equipment” and “artifact”, and the generated question is not related to the region. Fig. 7 (c) shows an example where the object region proposal is not performed properly, and thus, it is not possible to generate the question accurately. The lower part of the image shows examples of questions generated by the baselines.

Figure 8 shows the results of the human evaluation. The answer “Do not understand” in task (1) is shown as “no answer.” The average of evaluation

values is calculated by assigning “no answer” a score of 0. The proposed method outperformed the baseline in terms of relevance to the region and relevance to the answer. The reason is that the proposed method specifies a target object to the LSTM to generate the question, whereas the baselines do not consider any target. Regarding the number k of class labels used to select the target word, the average score is higher when $k = 3$, but the ratio of the highest score 5 is higher when $k = 2$. Also, the proportion of “no answer” is higher when $k = 2$. This means that, when $k = 3$, the target word becomes more generic and the relevance with the target region is less clear than when using $k = 2$. On the other hand, a value of $k = 2$ is more likely to specify a wrong target word for the visual question generation.

Table 4 shows the number of generated questions that successfully allowed acquiring information on unknown objects. We consider successful questions whose answers were not included in the known class of the unknown classifier neither in their hypernym, and whose relevance score in task (2) was 4 or more. We obtained the highest number of successful questions using our method with $k = 2$, since the selected target word is more concrete. On the other hand, our method with $k = 3$ generates questions for a more generic target word, and thus, it does not necessarily get the expected answer, but is still partly related to the target region.

We can conclude that the proposed method can successfully generate questions that allow acquiring information about unknown objects in an image.

7 Conclusions

In this paper, we presented a novel visual question generation (VQG) task to acquire class information of unknown (i.e., not learned previously) objects in the image, and proposed a method that automatically generates questions that target unknown objects in the image. To the best of our knowledge, this is the first research that approaches acquiring unknown information via VQG. The evaluation of our method shows that it can successfully acquire class information of unknown objects from humans. We believe this research will help other researchers in tackling this novel task.

Our future work includes feeding back the acquired information about the unknown object to the system, and learning it as new knowledge. For example, our method could be combined with recent works in few-shot learning and incremental learning to re-train the classifier with the new class. In addition, it is considered that the answers from humans will be noisy, so a system that can use noisy answers for re-training is necessary. If the answer obtained from humans is not the expected, it can be useful to generate multiple questions as necessary.

8 Acknowledgement

This work was supported by JST CREST Grant Number JPMJCR1403, Japan.

References

1. Alexe, B., Deselaers, T., Ferrari, V.: Measuring the Objectness of Image Windows. *IEEE PAMI* (2012)
2. Bendale, A., Boulton, T.E.: Towards Open World Recognition. In: *CVPR* (2015)
3. Bendale, A., Boulton, T.E.: Towards Open Set Deep Networks. In: *CVPR* (2016)
4. Cheng, M.M., Zhang, Z., Lin, W.Y., Torr, P.: BING: Binarized Normed Gradients for Objectness Estimation at 300fps. In: *CVPR* (2014)
5. Denkowski, M., Lavie, A.: Meteor universal: Language specific translation evaluation for any target language. In: *EACL Workshop on Statistical Machine Translation* (2014)
6. Fellbaum, C.: *WordNet : an electronic lexical database*. MIT Press (1998)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: *CVPR* (2016)
8. Hwa, R.: Sample selection for statistical parsing. *Computational Linguistics* (2004)
9. Jain, U., Zhang, Z., Schwing, A.G.: Creativity: Generating diverse questions using variational autoencoders. In: *CVPR* (2017)
10. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalanidis, Y., Li, L.J., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *IJCV* (2017)
11. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *NIPS* (2012)
12. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: *ACM SIGIR* (1994)
13. Li, Y., Huang, C., Tang, X., Loy, C.C.: Learning to disambiguate by asking discriminative questions. In: *ICCV* (2017)
14. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: *ICLR* (2013)
15. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *NIPS* (2013)
16. Mostafazadeh, N., Misra, I., Devlin, J., Mitchell, M., He, X., Vanderwende, L.: Generating natural questions about an image. In: *ACL* (2016)
17. Nickel, M., Kiela, D.: Poincaré embeddings for learning hierarchical representations. In: *NIPS* (2017)
18. Otsu, N.: A Threshold Selection Method from Gray-Level Histograms. *IEEE SMC* (1979)
19. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: *ACL* (2001)
20. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You Only Look Once: Unified, Real-Time Object Detection. In: *CVPR* (2016)
21. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *NIPS* (2015)
22. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *IJCV* (2015)
23. van de Sande, K.E.A., Uijlings, J.R.R., Gevers, T., Smeulders, A.W.M.: Segmentation as selective search for object recognition. In: *ICCV* (2011)
24. Scheffer, T., Decomain, C., Wrobel, S.: Active hidden markov models for information extraction. In: *CAIDA* (2001)

25. Scheirer, W.J., de Rezende Rocha, A., Sapkota, A., Boult, T.E.: Toward Open Set Recognition. *IEEE PAMI* (2013)
26. Settles, B., Craven, M.: An analysis of active learning strategies for sequence labeling tasks. In: *EMNLP* (2008)
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *ICLR* (2015)
28. Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., Smeulders, A.W.M.: Selective Search for Object Recognition. *IJCV* (2013)
29. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. *CVPR* (2015)
30. Zhang, S., Qu, L., You, S., Yang, Z., Zhang, J.: Automatic generation of grounded visual questions. In: *IJCAI* (2017)
31. Zhu, W., Liang, S., Wei, Y., Sun, J.: Saliency Optimization from Robust Background Detection. In: *CVPR* (2014)