

Graph Edit Distance Reward: Learning to Edit Scene Graph

Lichang Chen^{1,2†}, Guosheng Lin^{1‡}, Shijie Wang^{1,3}, and Qingyao Wu^{4,5}

¹ Nanyang Technological University

² Zhejiang University

³ Huazhong University of Science and Technology

⁴ School of Software Engineering, South China University of Technology

⁵ Key Laboratory of Big Data and Intelligent Robot, Ministry of Education
bobchen@zju.edu.cn, gslin@ntu.edu.sg

Abstract. Scene Graph, as a vital tool to bridge the gap between language domain and image domain, has been widely adopted in the cross-modality task like VQA. In this paper, we propose a new method to edit the scene graph according to the user instructions, which has never been explored. To be specific, in order to learn editing scene graphs as the semantics given by texts, we propose a Graph Edit Distance Reward, which is based on the Policy Gradient and Graph Matching algorithm, to optimize neural symbolic model. In the context of text-editing image retrieval, we validate the effectiveness of our method in CSS and CRIR dataset. Besides, CRIR is a new synthetic dataset generated by us, which we will publish soon for future use.

Keywords: Scene Graph Editing, Policy Gradient, Graph Matching

1 Introduction

Nowadays, in our daily life, more and more people have grown accustomed to typing in texts to obtain images from search engine. However, if they are not satisfied with searching results or they would like to modify the previous results, the only thing they can do is altering the language, then searching again. If users can provide their instructions about how to edit the searching results, it will be more than convenient for them to retrieve satisfied images. Figure 1 shows the simulation of such a scenario in the synthetic dataset. Vo *et al.* [27] firstly propose this text-editing image retrieval task, where the input query is composed by an input image I_{input} and an input text T_{input} , describing the desired modifications of I_{input} . Here, T_{input} can be viewed as user instructions and I_{input} can be viewed

[†] This work was done when L. Chen was an intern student in NTU.

[‡] G. Lin is the corresponding author of this paper.

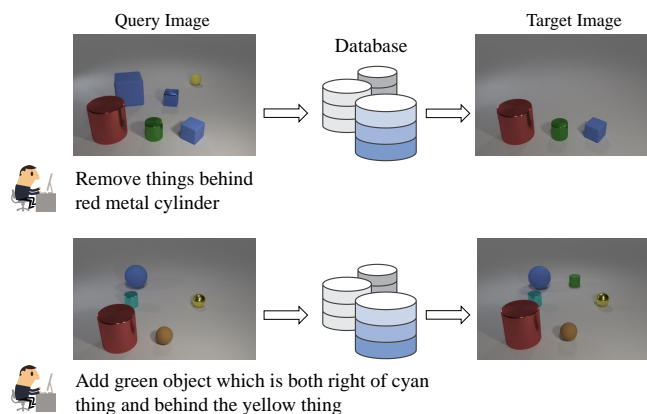


Fig. 1. Examples about cross-modal image retrieval. Given a query image and a query text describing the modification to the query image, we expect to edit the scene graph of the query image then retrieve a target image from database.

as the unsatisfactory image. We believe it is a superb platform for simulating the cross-modal task in real life.

Teaching machines to comprehend interaction between language and vision is an essential step for solving the Text+Image problem. There are some visual question answering (VQA [2, 11]) systems [12, 31] succeeding in crossing modalities by reasoning over Scene Graphs [13] following the semantics given by the texts, which proves Scene Graphs can bridge the gap between vision and language. Some image captioning models [19, 30] also adopted Scene Graphs and achieves good results, demonstrating that Scene Graphs are a fantastic representation of image scenes. We represent our scene by Scene Graph owing to the advantages of Scene Graphs. In this way, we transform the image editing problem into Scene Graph editing problem.

Due to the booming of computation resources, “old school” symbolic models lose its position from the competition with numeric model [20], which need less human intervening. However, it is undeniable that symbolism has one secret weapon, neat representation, which the high dimensional numeric does not possess. Recently, utilizing the advantage of symbolism, Yi *et al.* [31] propose applying neural symbolic model on CLEVR [11] with fully transparent reasoning. In details, firstly they use a scene parser to derender the scenes from images. Then they train a seq2seq program generator to predict the latent program from a question. Finally, their program executor will execute these predicted programs over symbolic scene representations to produce answer. Motivated by their model’s neat representations and fully explainable characteristics, we employ it to parse the user instructions. We design a Scene Graph Generation Engine to parse each image into a symbolic scene graph, whose nodes represent the objects while edges represent the relation between nodes. We also design a

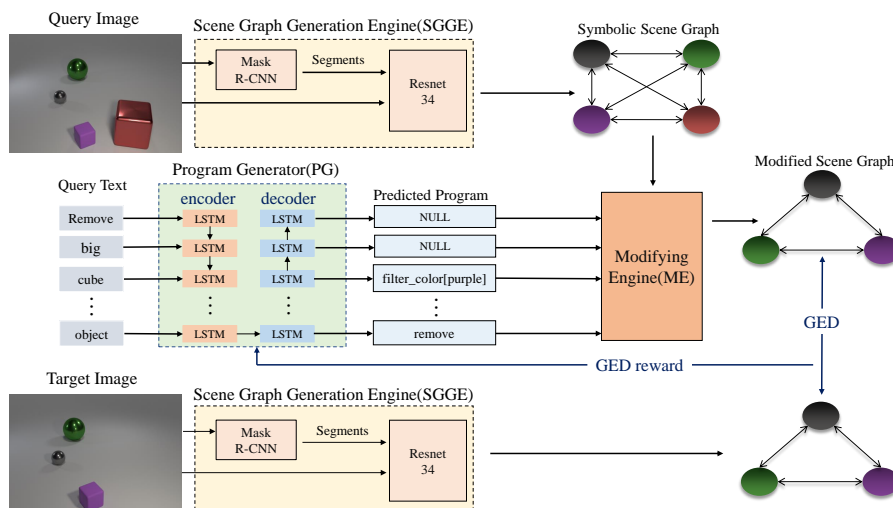


Fig. 2. Model Overview. The intact query text is “remove big cube to the right of purple object”. The intact Predicted Program is “remove”, “filter_size[large]”, “filter_shape[cube]”, “relate[right]”, “filter_color[purple]”. Because the query text is longer than the predicted program, our Program Generator will pad “NULL” program to fix the length inconsistency.

module to edit the scene graph. After that, the retrieval problem can be converted to a graph matching problem [5], which has low combinatorial space due to the symbolic representation.

We notice that REINFORCE [29] algorithm applied to finetune the Neural Symbolic Model [31] is the most vital part in achieving marvelous results. Following Johnson *et al.* [12], Yi *et al.* apply 0, 1 reward to finetune their program generator. This reward, however, is quite coarse if applying to our task. To be specific, only when exact matching [5] happens, which means all the nodes and all the edges in two generated symbolic scene graphs (one is generated from query image, another is generated from target image) matching, it is set to 1. To refine it, we propose applying Graph Edit Distance [23] based reward, which can optimize our model better.

In our experiment, we discover the CSS dataset [27] is simple because it has plenty of exact location information. To validate our model’s efficacy further, we follow CLEVR’s [11] generation process and generate our own Complex Reasoning Image Retrieval, called CRIR, which contains multi-hop queries thus the scene graphs are hard to be accurately edited. Moreover, our dataset breaks the constraints of exact location so it is more close to the real world scenes.

To summarize, our contributions are:

- We propose a GED reward to explore editing the scene graph with user instruction in the context of text-editing image retrieval. We refine the pre-




Fig. 3. (a) Image database scale – CRIR vs. CSS. (b) CSS only has zero-hop query. CRIR has 5 different types of queries. (c) Query Text Scale – CRIR vs. CSS. In CSS dataset, both training set and testing set contain 6k queries, while in CRIR dataset, both training set and testing set have 4.5k queries in each type of templates, totally 22.5k queries. (d) Query text type details. CRIR eliminates the dataset bias, each template generating the same number of queries.

vious neural symbolic model aiming at VQA and make it suitable for our cross-modal image retrieval task. To the best of our knowledge, we are the first to propose editing the scene graph.

- We propose a new dataset, CRIR, which contains abundant complex queries with full annotations of images. The dataset breaks the constraints of exact location in CSS [27], which can be used to validate the generalization ability of our model and simulate the real-world image retrieval scene better.
- Based on the policy gradient algorithm [29], we propose Graph Edit Distance(GED) reward to finetune our model as well as apply GED as retrieval metric, making it learn to edit scene graphs better. In the context of text-editing image retrieval task, we validate the efficacy of our model. Notably, we achieves new state-of-the-art performance on both CRIR and CSS [27], surpassing the previous methods by large margins.

2 Related Work

Image Retrieval. Conventional image retrieval [15] is a task to retrieve the most similar image to the query image from database, which is also viewed as image to image matching problem. In recent year, with much stronger image feature extractor [7, 26], conventional image retrieval has also made a huge progress [3, 6, 16] and has been extended to different real-world scenes like face retrieval. [18, 25].

Johnson *et al.* [13] first proposed cross-modality image retrieval. The task can be defined as given a text T_{input} , the model should retrieve an image I_{target} which is the most relevant to the T_{input} . They also give the definition of scene graph and

use it to represent the content of the scene, which builds a bridge between image domain and text domain. What’s more, cross-modality image retrieval is also extended to real-world scenes such as recipe to food image retrieval where the input query is recipe, food image to recipe retrieval, using deep metric learning method [28].

Visual Question Answering. Visual Question Answering [2] requires model to answer question given by text based on the input image. Solving this problem is a vital step to cross the modalities between language and vision.

Recently, Johnson *et al.* [11] utilize Blender to render images and they obtain a synthetic, virtual, and diagnostic dataset named CLEVR with full annotations of objects, such as 3D coordinates, size, shape, color, material, etc. The low cost of rendering virtual datasets allows for the appearance of CLEVR based datasets. For instance, Liu *et al.* [17] refine the CLEVR program generation engine and propose CLEVR-REF+ for referring expression task. They also modify the rendering engine which generates images with segments and bounding box information. There are also some real-world VQA datasets like VQA v1.0, v2.0 [2] and GQA [10], etc. Though the annotations cost is high, they provide a more realistic platform for validating the methods.

Neural Module Network. NMN is firstly proposed by Andreas *et al.* [1]. They parsed the natural language sentences into programs and applied it to instantiate different neural modules, then executing them to obtain the answer. Following their work, Johnson *et al.* [12] and Hu *et al.* [9] extend the Neural Module Network to the CLEVR dataset. But their neural modules are numeric one with executing programs and reasoning on high dimensional features. Yi *et al.* [31] propose symbolic module network to execute programs over symbolic scenes generated by scene parser achieving new state-of-the-art result in CLEVR, 99.9%. Due to the efficacy of their model, we adapt them to text-editing image retrieval task.

3 Method

In section 3.1, we introduce our model. In section 3.2, we give the definition of Graph Edit Distance. In section 3.3, we introduce policy gradient with new method we proposed to estimate the gradient in this task.

3.1 Our Model

Figure 2 shows the overview of our model. Our model is adapted from neural symbolic reasoning model in [31]. The first step to train our model is supervised training separately on both Scene Graph Generation Engine (SGGE) and Program Generator (PG) with Images and small number of ground truth latent programs(We provide the training details in Section 5.1 and 5.2). In finetuning, the input of our model contains T_{input} and I_{input} . From SGGE, We obtain the symbolic scene graph of the I_{input} while the latent program can be parsed by

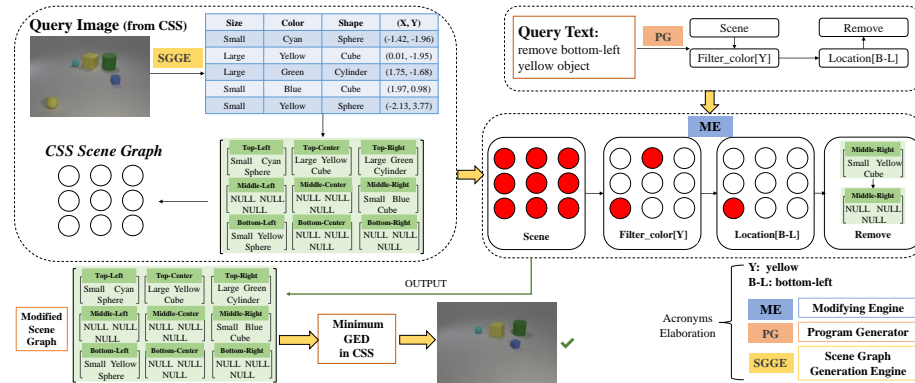


Fig. 4. A scene graph editing case in CSS dataset. CSS Scene Graph does not have any edges. So Modifying Engine just reason over nodes. In final **remove** program, we keep structure of the scene graph by just substituting the attribute values in operation. Red nodes in ME means attended nodes. Minimum GED means the minimum Graph Edit Distance between Modified Scene Graph and other Scene Graphs in CSS.

PG. The output of PG will be put into Modifying Engine (ME) with the generated scene graph. Then ME will execute programs sequentially and output the modified scene graph. Finally, the modified scene graph is compared with the target scene graph, which is generated from target image by the same SGGE, in Graph Edit Distance metric. After the computation of GED, the GED based reward can be backpropagated to finetune the PG.

The details about three main components, Scene Graph Generation Engine, Program Generator, Modifying Engine show as follows:

Scene Graph Generation Engine. We use SGGE to obtain the 3D coordinates with attribute values, including shape, size, color, material, for each objects in the image. In detail, Mask R-CNN is applied to generate segment proposals of each object in the image. Then we input the segment of each object with original image, resized to 224 by 224, into a Resnet-34 [8] to predict 3D coordinates.

Program Generator. The program generator follows the language model proposed by [31] which is improved version of [12] with attention. It is a seq2seq model. We set the dimension of input word vector to 300, 2 hidden layers in both encoder and decoder to 256 dimension. As the Figure 2 shows, we return a reward from the final part to finetune our program generator. Because the Modifying Engine (ME) execution part is symbolic, the gradient stream will be cut down there, which means the model is undifferentiable when entering ME. Therefore, we exploit small number of ground truth programs to pretrain our PG, then using GED reward to finetune it. Because training and finetuning details are different in different datasets, we will elaborate them in Experiment (Section 5).

Modifying Engine. The modifying engine will execute programs predicted by program generator over symbolic scene graphs. Thanks to the neat representation of symbolic representation, the reasoning and modifying operations over symbolic scene graphs are fully transparent and make our image retrieval process explainable. We adopt different modifying engines in two experiments. The details about 2 modifying engines show in Section 5.1 and 5.2. We implement our Modifying Engine as a set of modules in Python. Each program has one counterpart module in ME. The programs will first be converted to the corresponding modules and then be executed in Modifying Engine sequentially.

3.2 Graph Edit Distance.

Graph Edit Distance(GED) is a graph matching [5] approach first proposed by Sanfeliu *et al.* [23]. The concept of GED is finding the optimal set of edit operation which can transform Graph G_1 into Graph G_2 .

Definition1. (Graph G)

A graph G can be represented by a 3-tuple (V, α, β) , such that: V is a set of nodes. $\alpha : V \rightarrow L$ is the node labeling function. $\beta : V \times V \rightarrow L$ is the edge labeling function.

Definition2. (Graph Isomorphism)

A graph isomorphism between $G_1 = (V_1, \alpha_1, \beta_1)$ and $G_2 = (V_2, \alpha_2, \beta_2)$ is a bijective mapping $f : V_1 \rightarrow V_2$ such that $\alpha_1(x) = \alpha_2(f(x))$ for all $x \in V_1$ and $\beta_1((x, y)) = \beta_2((f(x), f(y)))$ for all $(x, y) \in V_1 \times V_1$.

Definition3. (Common Subgraph)

Let $G'_1 \subseteq G_1$ and $G'_2 \subseteq G_2$, if there exists a graph isomorphism between G'_1 and G'_2 , both G'_1 and G'_2 will be called a common subgraph of G_1 and G_2 . Moreover, a graph G is called a maximum common subgraph of G_1 and G_2 if G is a common subgraph of G_1 and G_2 and there exists no other common subgraph of G_1 and G_2 that has more nodes than G .

Definition4. (error-correcting graph matching)

An error-correcting graph matching from G_1 to G_2 is a bijective function $f : \hat{V}_1 \rightarrow \hat{V}_2$ where $\hat{V}_1 \subseteq V_1$ and $\hat{V}_2 \subseteq V_2$.

Definition5. (Graph Edit Distance)

Let $G_1 = (V_1, \alpha_1, \beta_1)$, $G_2 = (V_2, \alpha_2, \beta_2)$ be two graphs, the GED between these 2 graphs is defined as

$$GED(G_1, G_2) = \min_{e_1, \dots, e_k \in \gamma(f)} \sum_{i=1}^k c(e_i), \quad (1)$$

where f is an error-correcting graph matching $f : \hat{V}_1 \rightarrow \hat{V}_2$ from Graph G_1 to Graph G_2 and c denotes the cost function measuring the strength $c(e_i)$ of an edit operation e_i and $\gamma(f)$ denotes the set of edit paths transforming G_1 into G_2 . Insertions, deletions, and substitutions of both edges and nodes are 4 types of edit operations allowed. Thus, the right part of equation 1 can be denoted

as [4]

$$\min_{e_1, \dots, e_k \in \gamma(G_1, G_2)} \sum_{i=1}^k c(e_i) = \min \left(\sum_{x \in V_1 - \hat{V}_1} c_{nd}(x) + \sum_{x \in V_2 - \hat{V}_2} c_{ni}(x) + \sum_{x \in \hat{V}_1} c_{ns}(x) + \sum_{e \in \hat{E}_1} c_{es}(e) \right), \quad (2)$$

where $c_{nd}(x)$ is the cost of deleting a node $x \in V_1 - \hat{V}_1$ from G_1 . $c_{ni}(x)$ is the cost of inserting a node $x \in V_2 - \hat{V}_2$ in G_2 , $c_{ns}(x)$ is the cost of inserting a node $x \in \hat{V}_1$ by $f(x) \in \hat{V}_2$, and $c_{es}(e)$ is the cost substituting a node $x \in \hat{V}_1$ by $f(x) \in \hat{V}_2$ and $c_{es}(e)$ is the cost of substituting an egde $e = (x, y) \in \hat{V}_1 \times V_1$ by $e' = (f(x), f(y)) \in V_2 \times \hat{V}_2$. The cost of 4 operations are hyperparameter in our model. In experiment part5, we will show the best hyperparameters in different datasets.

We follow the widely used quick GED computation algorithm proposed by Riesen *et al.* [22] which is referred to as A*GED to compute GED in our paper.

3.3 Policy Gradient + GED reward

In Reinforcement Learning, unlike the image classification and segmentation, the problems are usually undifferentiable with high-dimensional discrete states and learning (state, action) pair is very hard. However, the policy, such as make a robot move left or kick a ball, is easier to learn. Thus, we can formally define a class of parametrized policies as $\Pi = \{\pi_\theta, \theta \in \mathbb{R}^m\}$. And for each policy, the value is defined as equation 3, where r_t is the reward in the time step t, γ^t is the attenuation coefficient in the time step t.

$$J(\theta) = \mathbb{E} \left[\sum_{t \geq 0} \gamma^t r_t | \pi_\theta \right] \quad (3)$$

Our goal is to find the optimal policy $\theta^* = \arg \max_\theta J(\theta)$. And one of simple algorithms to solve this problem is REINFORCE [29] which applies gradient ascent to optimize the parameter θ . The expected future reward can be written as

$$\begin{aligned} J(\theta) &= \mathbb{E}_{\tau \sim p(\tau; \theta)} [r(\tau)] \\ &= \int_{\tau} r(\tau) p(\tau; \theta) d\tau, \end{aligned} \quad (4)$$

where the $p(\tau; \theta)$ is the probability of trajectory τ in sampling. And by differentiating the equation 4, we can obtain:

$$\nabla_\theta J(\theta) = \int_{\tau} r(\tau) \nabla_\theta p(\tau; \theta) d\tau. \quad (5)$$

$\nabla_\theta p(\tau; \theta)$ in equation 5 can be rewritten as

$$\nabla_\theta p(\tau; \theta) = p(\tau; \theta) \frac{\nabla_\theta p(\tau; \theta)}{p(\tau; \theta)} = p(\tau; \theta) \nabla_\theta \log p(\tau; \theta). \quad (6)$$

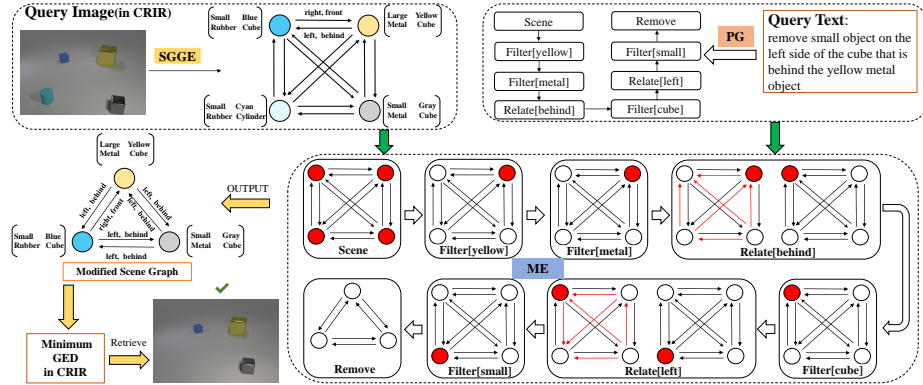


Fig. 5. A scene graph editing case in CRIR dataset. ME is Modifying Engine. PG is Program Generator. SGGE is Scene Graph Generation Engine. Symbolic Reasoning over CRIR scene graphs need utilizing edge information, as shows in “Relate” program reasoning part. Here, we display the reasoning details of ME. The reasoning process in CRIR is more complex than in CSS.

Finally, we can estimate $J(\theta)$ as equation 6:

$$\nabla_{\theta} J(\theta) \approx \sum_{t \geq 0} r(\tau) \nabla_{\theta} \log \pi_{\theta}(a_t | s_t), \tag{7}$$

where s_t is the state in time step t .

As previously state in Section 1, the simple 0, 1 reward is coarse in our task. Thanks to the graph matching algorithm, we have a better measurement about the similarity of 2 graphs. To measure the reward after a number of programs executing in Modifying Engine better, we propose a new reward (See equation 8) combining Graph Edit Distance with policy gradient algorithm named GED reward. In experiment, we will validate the efficacy of this reward when it is applied to finetune our program generator.

$$reward = 1 - GED. \tag{8}$$

4 CRIR Dataset

After exploring the Color, Shape, Size(CSS) [27] dataset carefully, we find it just adopt one simplest kind of template, zero-hop, in CLEVR’s [11] template universe, which means it does not need any reasoning process before editing the scene graphs. Thus, inheriting the CLEVR template universe, we generate more complex reasoning dataset called CRIR to validate the reasoning and generalization ability of our model.

Our dataset generation can be split into 2 parts, image generation and query text generation, which will be discussed in section 4.1 and section 4.2 respectively.

4.1 Image Generation

Firstly, we select 3k images from CSS [27] Dataset, which should have at least 4 objects in the scenes. Then we add material information to these objects, which means randomly choosing the metal or rubber for the object. According to the query answer (generation process is explained in section 4.2), we modify the scene files in CSS. Then following the image generation engine provided by Liu [17] *et al.*, we create our own image database, which contains 26k images in the training set and 26k images in the testing set. The image database is larger than the previous CSS image database, we show details in Figure 3a. Our objects contain 4-dimension information, color, shape, size, material, comparing to 3-dimension in CSS.

4.2 Query Generation

The queries’ generation can be summarized as follows:

1. We add 3 types of programs, **add**, **remove** and **make**, to the CLEVR-programs universe. **Make** programs are combined with a value chosen from 4 attributes, shape, color, material and size. For instance, in the query text “make yellow sphere in front of green cube small”, the program is “make[small]”.

2. We inherit the CLEVR templates and throw away the templates that are unsuitable for retrieval queries. Specifically, we select 5 templates from CLEVR [11]: *Zere Hop*, *One Hop*, *Two Hop*, *Three Hop*, *Single And*, which are suitable for generating our image retrieval text queries.

3. We modify the CLEVR question generation engine, whose original output is the answer in CLEVR answer set. To be specific, we force our engine to output the objects’ indices which need to be modified with the text query suitable for this task, unlike the questions in CLEVR. Firstly, we generate **remove**, **make** queries and maintain the number of **remove** queries 2 times more than **make** queries. Then we use these queries to change the image scenes according to the **remove** or **make** type and objects’ indices. After that, we apply these scenes to render new images. For **add** query generation, we compare the original scene with the generated scene one by one, if it reduces one object, we put it to an addition candidate list. After that, we randomly choose pairs from the list then change query type from **add** to **remove** and convert the I_{target} to I_{input} until the number of **add** query text is equal to the **remove** query text. Finally, we obtain 22.5k queries in training set and 22.5k queries in testing set.

Results. We compare CRIR dataset with CSS dataset which shows in Figure 3. To conclude, CRIR’s image database is larger than CSS, Figure 3a. CRIR’s query text is more versatile than CSS’s, Figure 3b. Additionally, CRIR’s scale in query texts is larger than CSS’s, Figure 3c. In Figure 3d, we show the details about the types of query text in CRIR. All these perspectives above demonstrate

Method	ImageO	TextO	Concate	MRN [14]	Relat [24]	FiLM [21]	TIRG [27]	Ours
Add	0.1	0.1	73.6	72.9	74.7	75.9	83.6	99.7
Remove	0.2	0.1	45.3	42.8	49.8	52.3	64.3	99.8
Make	9.1	0.1	63.2	64.6	61.8	68.6	73.3	99.9
Overall	6.3	0.1	60.6	60.1	62.1	65.6	73.7	99.8

Table 1. Quantitative results in CSS. The retrieval metric is Recall at rank 1. ImageO and TextO means Image Only and Text Only, respectively. Concate is a method just concatenating language feature and image feature without other processing. As show above, our model with GED reward surpasses all other models by large margins in all three types of queries.

we generate a more complex reasoning and hard-to-edit dataset for text-editing image retrieval task.

5 Experiment

To display the effectiveness of our model on editing the scene graphs, we carry out experiments on CSS and CRIR. Our measurement is set to recall at rank 1 in order to better compare with other methods.

5.1 Experiment on CSS Dataset

Because nearly 85% percent query texts in CSS dataset has exact location word, we create exact location program for this dataset, such as “Location[B-L]” program in Figure 4, which can attend the object at Bottom-Left grid. The location boundaries used to split the 3 by 3 grids are (-0.99, 0.86) horizontally and (-0.47, 2.35) vertically. In CSS training, we randomly select 30 text queries (add, remove, make, each 10) and annotate them for pretraining our program generator, with learning rate 6×10^{-4} , 12000 iterations, and batch size 64. Especially, in CSS, we keep the structure of Scene Graph when ME executing programs to modify it. We show an example of ME’s reasoning details in Figure 4.

In Modifying Engine, we inherit programs which have prefix `filter` from CLEVR and create corresponding modules for them. Here we list some special programs outside CLEVR-programs universe and its corresponding modules’ operation details, which shows as follows:

Location Modules. Location modules are instantiated from location programs and they will select the exact location like “Top-Left” for operating modification later on. There are totally 9 different location modules corresponding to 3 by 3 grids.

Remove, Add, Make Modules. According to the start word of the query text we can obtain `Remove`, `Add`, or `Make` programs to instantiate modules and execute operations over structural scene graph. All these three types of modules will be executed last. We also create 2 reasoning modes, “normal” and “add”.

Mode “normal” is created for **Remove** and **Make** modules while mode “add” for **Add** module. **Remove** module is applied to change every attribute value of the attended nodes into **NULL** while **Make** modules change the specific attribute that attended by the previous modules. But in **Add** type queries, some attributes will not be assigned value, for example, in query text “Add small cube to the top-left”, the color information is not assigned. Thus we create value 1 for these unassigned attributes when modifying the node in **Add**. In Graph Edit Distance [23] computation, value 1 can match any value except **NULL**.

Graph Edit Distance in CSS. In the CSS settings, scene graphs are structural and each graph has 9 nodes with rigid location. Thus in equation 2, only substituting cost exists. The Graph Edit Distance [23] computation can be simplified as:

$$GED = \min \left(\sum_{x \in \hat{V}_1} c_{ns}(x) \right). \quad (9)$$

For every node having 3 attributes, we set the cost of substituting every attribute to 1/3 and force the reward to be 0 if GED between 2 graphs is larger than 1. In the testing, we apply graph edit distance as our retrieval metric. If GED of two pairs’ graphs happen to be the same, we will randomly choose one as our target image. We return this reward to finetune the program generator with 50,000 iterations and early-stopping strategy. Batch size is fixed to 64 in finetuning.

Results and Analysis. As table 1 shows, our model achieves state-of-the-art result and surpass the previous baseline model TIRG [27] drastically no matter in which query text type. Additionally, We give an example in CSS and display our symbolic model’s transparent and powerful reasoning ability in Figure 4. The results and the visualization prove our method’s effectiveness on editing symbolic scene graph according to the instructions.

One biggest problem in CSS is that if our model predicts the correct location program, it will attend to the correct object because of the exact location information text provided. Besides, all queries in CSS are *Zero Hop*, which does not need any edge information for reasoning. Only through attribute **filter** or location programs can the model attend to correct nodes. Therefore, we propose a more complex dataset to validate our model’s reasoning ability: **1.** To simulate the real-world scene, the dataset should just provide the relational location information, not exact locations. **2.** The dataset should also contain more complex queries, not just *Zero Hop*. Aiming at this, we generate CRIR. Details can be referred to Section 4.2.

5.2 Experiment on CRIR Dataset

In CRIR dataset, we remove the location programs to simulate a more realistic situation. Though we can still generate the structural scene graph shown as Figure 4, we consider that 3 by 3 grid structural scene graph is a special case in the real world. Therefore, we cancel this rigid location and generate relational scene graphs for the images in the CRIR dataset then editing relational scene graphs. In training procedure, we randomly select 300 programs (add, remove,

Model	Zero_Hop	One_Hop	Two_Hop	Three_Hop	Single_And	Overall
Concatenate	45.1	25.2	16.7	3.7	15.8	21.3
MRN [14]	44.8	28.3	19.4	3.5	19.6	23.1
Relat [24]	51.9	28.2	21.3	4.0	18.1	24.7
FiLM [21]	50.2	28.1	21.1	3.6	16.3	23.9
TIRG [27]	53.8	29.6	22.9	4.8	19.4	26.1
0,1 Reward	95.9	95.4	95.8	95.3	94.5	95.4
GED Reward	99.0	98.4	98.2	97.9	98.4	98.4

Table 2. Quantitative results in CRIR. The retrieval metric is Recall at rank 1. We adapt Vo’s [27] code to train TIRG and Concat models in CRIR. 0, 1 means applying 0, 1 reward to finetune our model. GED means applying GED reward to optimize our model. In particular, our model with GED reward optimizing surpass all the other models in every type of query.

make type each 100 programs) to pretrain our program generator with 16,000 iterations and 7×10^{-4} learning rate. **Add**, **Remove** modules in CRIR settings are different from CSS. To illustrate, **Add** module in CRIR Modifying Engine will add a pseudo node to the scene graph, unlike **Add** module in CSS just changing the attended node. We give an instance here. Assuming the query text is “Add red large sphere to the left of large yellow cylinder”, our Modifying Engine will execute **Add** module to create one pseudo node. The pseudo node will create edges linking to all other nodes in scene graphs. Then if PG outputs “Relate[left]” program, it will assign edge linking to the “large yellow cylinder” **Left** value. For **Remove** module, unlike the example shows in Figure 4 just changing the attributes to NULL, in CRIR, **Remove** module will remove the attended objects and all the edges of the nodes as example shows in Figure 5.

Graph Edit Distance in CRIR. We apply graph edit distance to finetune our program generator. Unlike the graph edit distance in CSS dataset without edge information, in this part, our graph edit distance has all costs (four parts) in equation 2. And we set the cost of inserting a node to 1, the cost of deleting to 1, substituting an attribute value to 1/4, and substituting an edge to 1/16. We only return the positive reward and force negative reward to 0. We return this reward to finetune PG with 60, 000 iterations, early-stopping training strategy and batch size is fixed to 64.

Results and Analysis. We provide the quantitative results in table 2. As it shows, our model can still maintain high-level performance in all five types of queries while TIRG [27] and Concat models drop its performance on CRIR drastically. *Zero Hop* is the easiest type query text so all the model achieve its best performance on this type of query text. From table 2, we also discover that *Three Hop* is so complicated that TIRG and Concat model can not learn this type of query text well, but our models are able to comprehend this type of queries. Notably, Our GED reward is more efficient than simple 0, 1 reward according to the last 2 columns of table 2. The performance of our model with GED reward surpasses all other models in all types of queries, which proves

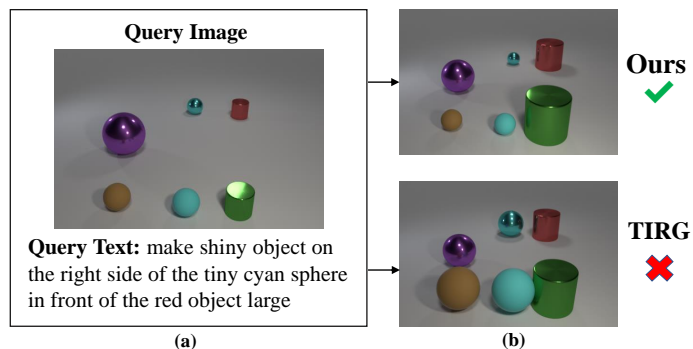


Fig. 6. The qualitative results in CRIR. (a) represents the input query while (b) is the retrieval results by applying our model (with GED reward optimizing) and TIRG [27]

the GED reward is a precise measurement of the final modified graph to the target graph. We also show one case of our model’s (optimized by GED reward) symbolic reasoning insight in Figure 5.

We compare the performance of feature combination model and ours qualitatively and show a case in Figure 6. As case shows, the TIRG [27] just retrieve the image with similar objects while our scene-graph-editing based model can understand the true meaning of this text and retrieve the correct image in CRIR dataset. In conclusion, our model has more powerful generalization ability than TIRG [27].

6 Conclusion

In this paper, we have proposed a GED reward for learning to edit scene graph in the context of text-editing image retrieval. By exploiting CLEVR toolkit, we generate a Complex Reasoning Image Retrieval dataset, simulating a harder case of editing. To the best of our knowledge, we are the first to apply Graph Edit Distance as retrieval metric. Furthermore, we validate our model on both CSS and CRIR dataset and achieve nearly perfect results, surpassing other models by large margins, proving the effectiveness of editing scene graphs in the context of text-editing image retrieval. However, our model still need prior information like the annotations of some latent programs to pretrain the PG. We expect our machine can own the reasoning ability without using any prior information in the future.

Acknowledgments. This research was supported by the National Research Foundation Singapore under its AI Singapore Programme (Award Number: AISG-RP-2018-003) and the MOE Tier-1 research grants: RG28/18 (S) and RG22/19 (S). Q. Wu’s participation was supported by NSFC 61876208, Key-Area Research and Development Program of Guangdong 2018B010108002.

References

1. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Neural module networks. In: CVPR (2016)
2. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., Parikh, D.: Vqa: Visual question answering. In: CVPR (2015)
3. Babenko, A., Lempitsky, V.: Aggregating local deep features for image retrieval. In: CVPR (2015)
4. Bunke, H.: Error correcting graph matching: On the influence of the underlying cost function. TPAMI **21**(9), 917–922 (1999)
5. Conte, D., Foggia, P., Sansone, C., Vento, M.: Thirty years of graph matching in pattern recognition. International journal of pattern recognition and artificial intelligence **18**(03), 265–298 (2004)
6. Gordo, A., Almazán, J., Revaud, J., Larlus, D.: Deep image retrieval: Learning global representations for image search. In: ECCV (2016)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
9. Hu, R., Andreas, J., Rohrbach, M., Darrell, T., Saenko, K.: Learning to reason: End-to-end module networks for visual question answering. In: ICCV (2017)
10. Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. In: CVPR (2019)
11. Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In: CVPR (2017)
12. Johnson, J., Hariharan, B., van der Maaten, L., Hoffman, J., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: Inferring and executing programs for visual reasoning. In: CVPR (2017)
13. Johnson, J., Krishna, R., Stark, M., Li, L.J., Shamma, D., Bernstein, M., Fei-Fei, L.: Image retrieval using scene graphs. In: CVPR (2015)
14. Kim, J.H., Lee, S.W., Kwak, D., Heo, M.O., Kim, J., Ha, J.W., Zhang, B.T.: Multimodal residual learning for visual qa. In: Advances in neural information processing systems. pp. 361–369 (2016)
15. Lin, K., Yang, H.F., Hsiao, J.H., Chen, C.S.: Deep learning of binary hash codes for fast image retrieval. In: CVPR Workshops (2015)
16. Liu, H., Wang, R., Shan, S., Chen, X.: Deep supervised hashing for fast image retrieval. In: CVPR (2016)
17. Liu, R., Liu, C., Bai, Y., Yuille, A.L.: Clevr-ref+: Diagnosing visual reasoning with referring expressions. In: CVPR (2019)
18. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: CVPR (2016)
19. Luo, R., Price, B., Cohen, S., Shakhnarovich, G.: Discriminability objective for training descriptive captions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6964–6974 (2018)
20. Minsky, M.L.: Logical versus analogical or symbolic versus connectionist or neat versus scruffy. AI magazine **12**(2), 34–34 (1991)
21. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)

22. Riesen, K., Fankhauser, S., Bunke, H.: Speeding up graph edit distance computation with a bipartite heuristic. In: *MLG (2007)*
23. Sanfeliu, A., Fu, K.S.: A distance measure between attributed relational graphs for pattern recognition. *IEEE transactions on systems, man, and cybernetics* **SMC-13**(3), 353–362 (1983)
24. Santoro, A., Raposo, D., Barrett, D.G., Malinowski, M., Pascanu, R., Battaglia, P., Lillicrap, T.: A simple neural network module for relational reasoning. In: *Advances in neural information processing systems*. pp. 4967–4976 (2017)
25. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: *CVPR (2015)*
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
27. Vo, N., Jiang, L., Sun, C., Murphy, K., Li, L.J., Fei-Fei, L., Hays, J.: Composing text and image for image retrieval-an empirical odyssey. In: *CVPR (2019)*
28. Wang, H., Sahoo, D., Liu, C., Lim, E.p., Hoi, S.C.H.: Learning cross-modal embeddings with adversarial networks for cooking recipes and food images. In: *CVPR (2019)*
29. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8**(3-4), 229–256 (1992)
30. Yang, X., Tang, K., Zhang, H., Cai, J.: Auto-encoding scene graphs for image captioning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 10685–10694 (2019)
31. Yi, K., Wu, J., Gan, C., Torralla, A., Kohli, P., Tenenbaum, J.B.: Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In: *NIPS (2018)*