

# Image-to-Image Retrieval by Learning Similarity between Scene Graphs

Sangwoong Yoon<sup>1\*</sup>, Woo Young Kang<sup>2</sup>, Sungwook Jeon<sup>3</sup>, SeongEun Lee<sup>3</sup>, Changjin Han<sup>3</sup>,  
Jonghun Park<sup>3</sup>, and Eun-Sol Kim<sup>2</sup>

<sup>1</sup>Seoul National University Robotics Lab, <sup>2</sup>Kakao Brain, <sup>3</sup>Seoul National University Information Management Lab  
swyoon@robotics.snu.ac.kr, edwin.kang@kakaobrain.com, {wookee3,ryuha96,changjin9653,jonghun}@snu.ac.kr,  
eunsol.kim@kakaobrain.com

## Abstract

As a scene graph compactly summarizes the high-level content of an image in a structured and symbolic manner, the similarity between scene graphs of two images reflects the relevance of their contents. Based on this idea, we propose a novel approach for image-to-image retrieval using scene graph similarity measured by graph neural networks. In our approach, graph neural networks are trained to predict the proxy image relevance measure, computed from human-annotated captions using a pre-trained sentence similarity model. We collect and publish the dataset for image relevance measured by human annotators to evaluate retrieval algorithms. The collected dataset shows that our method agrees well with the human perception of image similarity than other competitive baselines.

## Introduction

Image-to-image retrieval, the task of finding similar images to a query image from a database, is one of the fundamental problems in computer vision and is the core technology in visual search engines. The application of image retrieval systems has been most successful in problems where each image has a clear representative object, such as landmark detection and instance-based retrieval (Gordo et al. 2016; Mohedano et al. 2016; Radenović, Tolias, and Chum 2016), or has explicit tag labels (Gong et al. 2014).

However, performing image retrieval with complex images that have multiple objects and various relationships between them remains challenging for two reasons. First, deep convolutional neural networks (CNNs), on which most image retrieval methods rely heavily, tend to be overly sensitive to low-level and local visual features (Zheng, Yang, and Tian 2017; Zeiler and Fergus 2014; Chen et al. 2018). As shown in Figure 1, nearest-neighbor search on ResNet-152 penultimate layer feature space returns images that are superficially similar but have completely different content. Second, there is no publicly available labeled data to train and evaluate the image retrieval system for complex images, partly because quantifying similarity between images with multiple objects as label information is difficult. Furthermore, a similarity measure for such complex images is desired to reflect

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

\*Work done during an internship at Kakao Brain

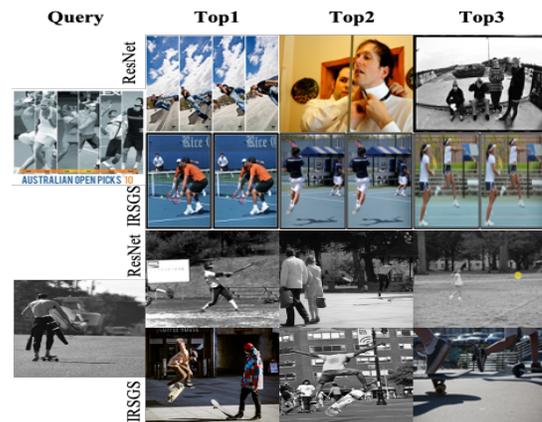


Figure 1: Image retrieval examples from ResNet and IRSGS. ResNet retrieves images with superficial similarity, e.g., grayscale or vertical lines, while IRSGS successfully returns images with correct context, such as playing tennis or skateboarding.

semantics of images, i.e., the context and relationship of entities in images.

In this paper, we address these challenges and build an image retrieval system capable of finding semantically similar images to a query from a complex scene image database. First of all, we propose a novel image retrieval framework, **Image Retrieval with Scene Graph Similarity (IRSGS)**, which retrieves images with a similar scene graph to the scene graph of a query. A scene graph represents an image as a set of objects, attributes, and relationships, summarizing the content of a complex image. Therefore, the scene graph similarity can be an effective tool to measure semantic similarity between images. IRSGS utilizes a graph neural networks to compute the similarity between two scene graphs, becoming more robust to confounding low-level features (Figure 1).

Also, we conduct a human experiment to collect human decisions on image similarity. In the experiment, annotators are given a query image along with two candidate images and asked to select which candidate image is more similar to the query than the other. With 29 annotators, we collect more than 10,000 annotations over more than 1,700 image triplets.

Thanks to the collected dataset, we can quantitatively evaluate the performance of image retrieval methods. Our dataset is available online<sup>1</sup>.

However, it is costly to collect enough ground truth annotation from humans to supervise the image retrieval algorithm for a large image dataset, because the number of pairwise relationships to be labeled grows in  $O(N^2)$  for the number of data  $N$ . Instead, we utilize human-annotated captions of images to define proxy image similarity, inspired by Gordo and Larlus (2017) which used term frequencies of captions to measure image similarity. As a caption tends to cover important objects, attributes, and relationships between objects in an image, the similarity between captions is likely to reflect the contextual similarity between two images. Also, obtaining captions is more feasible, as the number of the required captions grow in  $O(N)$ . We use the state-of-the-art sentence embedding (Reimers and Gurevych 2019) method to compute the similarity between captions. The computed similarity is used to train a graph neural network in IRSGS and evaluate the retrieval results.

Tested on real-world complex scene images, IRSGS show higher agreement with human judgment than other competitive baselines. The main contributions of this paper can be summarized as follows:

- We propose IRSGS, a novel image retrieval framework that utilizes the similarity between scene graphs computed from a graph neural network to retrieve semantically similar images;
- We collect more than 10,000 human annotations for semantic-based image retrieval methods and publish the dataset into the public;
- We propose to train the proposed retrieval framework with the surrogate relevance measure obtained from image captions and a pre-trained language model;
- We empirically evaluate the proposed method and demonstrate its effectiveness over other baselines.

## Related Work

### Image Retrieval

Conventional image retrieval methods use visual feature representations, object categories, or text descriptions (Zheng, Yang, and Tian 2017; Babenko et al. 2014; Chen, Davis, and Lim 2019; Wei et al. 2016; Zhen et al. 2019; Gu et al. 2018; Vo et al. 2019; Gordo et al. 2017). The activation of intermediate layers of CNN is shown to be effective as a representation of an image for image retrieval tasks. However, as shown in Figure 1, CNN often fails to capture semantic contents of images and is confounded by low-level visual features.

Image retrieval methods which reflects more semantic contents of images are investigated in Gordo and Larlus (2017); Johnson et al. (2015). Gordo and Larlus (2017) used term frequencies in regional captions to supervise CNN for image retrieval, but they did not utilize scene graphs. Johnson et al. (2015) proposed an algorithm retrieving images

given a scene graph query. However, their approach does not employ graph-to-graph comparison and is not scalable.

### Scene Graphs

A *scene graph* (Johnson et al. 2015) represents the content of an image in the form of a graph nodes of which represent objects, their attributes, and the relationships between them. After a large-scale real-world scene graph dataset manually annotated by humans in Visual Genome dataset (Krishna et al. 2017) was published, a number of applications such as image captioning (Wu et al. 2017; Lu et al. 2018; Milewski, Moens, and Calixto 2020) visual question answering (Teney, Liu, and van den Hengel 2017), and image-grounded dialog (Das et al. 2017) have shown the effectiveness of the scene graphs. Furthermore, various works, such as GQA(Hudson and Manning 2019), VRD(Lu et al. 2016), and VrR-VG(Liang et al. 2019) provided the human-annotated scene graph datasets. Also, recent researches (Yang et al. 2018; Xu et al. 2017; Li et al. 2017) have suggested methods to generate scene graphs automatically. Detailed discussion on scene graph generation will be made in Experimental Setup Section.

### Graph Similarity Learning

Many algorithms have been proposed for solving the isomorphism test or (sub-)graph matching task between two graphs. However, such methods are often not scalable to huge graphs or not applicable in the setting where node features are provided. Here, we review several state-of-the-art algorithms that are related to our application, image retrieval by graph matching. For the graph pooling perspective, we focus on two recent algorithms, the Graph Convolutional Network (GCN;Kipf and Welling (2016)) and the Graph Isomorphism Network (GIN;(Xu et al. 2018)). GCN utilized neural network-based spectral convolutions in the Fourier domain to perform the convolution operation on a graph. GIN used injective aggregation and graph-level readout functions. The learned graph representations, then, can be used to get the similarity of two graphs. Both networks transforms a graph into a fixed-length vector, enabling distance computation between two graphs in the vector space. Other studies viewed the graph similarity learning problem as the optimal transport problem (Solomon et al. 2016; Maretic et al. 2019; Alvarez-Melis and Jaakkola 2018; Xu, Luo, and Carin 2019; Xu et al. 2019; Titouan et al. 2019). Especially in Gromov Wasserstein Learning (GWL;(Xu et al. 2019)), node embeddings were learned from associated node labels. Thus the method can reflect not only a graph structure but also node features at the same time. Graph Matching Network (GMN;(Li et al. 2019)) used the cross-graph attention mechanism, which yields different node representations for different pairs of graphs.

### Image Retrieval with Scene Graph Similarity

In this section, we describe our framework, Image Retrieval with Scene Graph Similarity (IRSGS). Given a query image, IRSGS first generates a query scene graph from the image

<sup>1</sup><https://github.com/swyoon/aaai2021-scene-graph-img-retr>

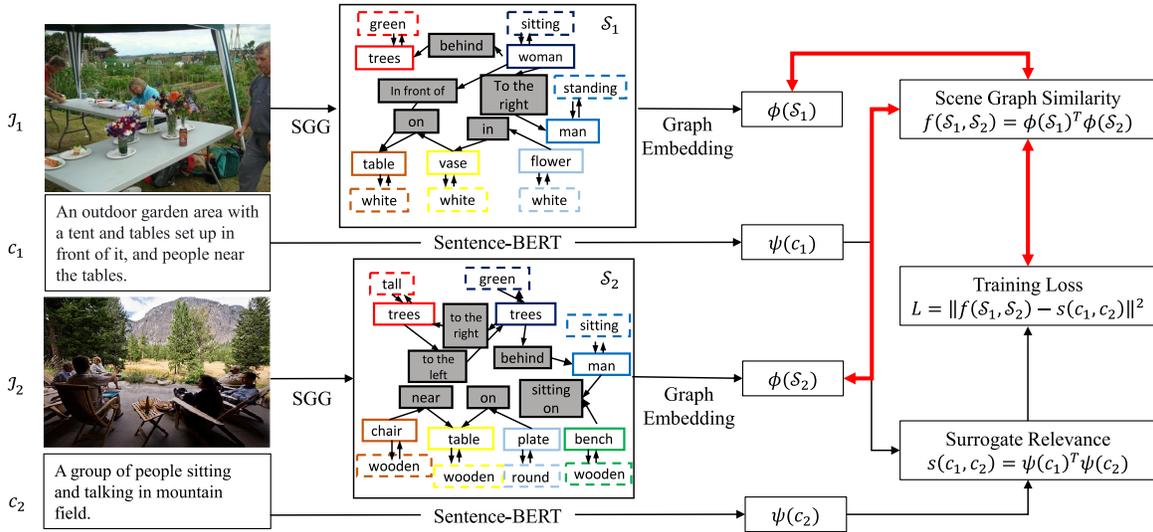


Figure 2: An overview of IRSGS. Images  $\mathcal{I}_1, \mathcal{I}_2$  are converted into vector representations  $\phi(\mathcal{S}_1), \phi(\mathcal{S}_2)$  through scene graph generation (SGG) and graph embedding. The graph embedding function is learned to minimize mean squared error to surrogate relevance, i.e., the similarity between captions. The bold red bidirectional arrows indicate trainable parts. For retrieval, the learned scene graph similarity function is used to rank relevant images.

and then retrieves images with a scene graph highly similar to the query scene graph. Figure 2 illustrates the retrieval process. The similarity between scene graphs is computed through a graph neural network trained with surrogate relevance measure as a supervision signal.

### Scene Graphs and Their Generation

Formally, a scene graph  $\mathcal{S} = \{\mathcal{O}, \mathcal{A}, \mathcal{R}\}$  of an image  $\mathcal{I}$  is defined as a set of objects  $\mathcal{O}$ , attributes of objects  $\mathcal{A}$ , and relations on pairs of objects  $\mathcal{R}$ . All objects, attributes, and relations are associated with a word label, for example, "car", "red", and "in front of". We represent a scene graph as a set of nodes and edges, i.e., a form of a conventional graph. All objects, attributes, and relations are treated as nodes, and associations among them are represented as undirected edges. Word labels are converted into 300-dimensional GloVe vectors (Pennington, Socher, and Manning 2014) and treated as node features.

Generating a scene graph from an image is equivalent to detecting objects, attributes, and relationships in the image. We employ a recently proposed method (Anderson et al. 2018) in our IRSGS framework to generate scene graphs. While end-to-end training of scene graph generation module is possible in principle, a fixed pre-trained algorithm is used in our experiments to reduce the computational burden. We shall provide details of our generation process in Experimental Setup Section. Note that IRSGS is compatible with any scene graph generation algorithm and is not bound to the specific one we used in this paper.

### Retrieval via Scene Graph Similarity

Given a query image  $\mathcal{I}_q$ , an image retrieval system ranks candidate images  $\{\mathcal{I}_i\}_{i=1}^N$  according to the similarity to the query image  $\text{sim}(\mathcal{I}_i, \mathcal{I}_q)$ . IRSGS casts this image retrieval

task into a graph retrieval problem by defining the similarity between images as the similarity between corresponding scene graphs. Formally,

$$\text{sim}(\mathcal{I}_i, \mathcal{I}_j) = f(\mathcal{S}_i, \mathcal{S}_j) \quad (1)$$

where  $\mathcal{S}_i, \mathcal{S}_j$  are scene graphs for  $\mathcal{I}_i, \mathcal{I}_j$ , respectively. We shall refer  $f(\mathcal{S}_i, \mathcal{S}_j)$  as *scene graph similarity*.

We compute the scene graph similarity from the inner product of two representation vectors of scene graphs. With a scene graph, a graph neural network is applied, and the resulting node representations are pooled to generate a unit  $d$ -dimensional vector  $\phi = \phi(\mathcal{S}) \in \mathbb{R}^d$ . The scene graph similarity is then given as follows:

$$f(\mathcal{S}_1, \mathcal{S}_2) = \phi(\mathcal{S}_1)^\top \phi(\mathcal{S}_2). \quad (2)$$

We construct  $\phi$  by computing the forward pass of graph neural networks to obtain node representations and then apply average pooling. We implement  $\phi$  with either GCN or GIN, yielding two versions, IRSGS-GCN and IRSGS-GIN, respectively.

### Learning to Predict Surrogate Relevance

We define *surrogate relevance measure* between two images as the similarity between their captions. Let  $c_i$  and  $c_j$  are captions of image  $\mathcal{I}_i$  and  $\mathcal{I}_j$ . To compute the similarity between the captions, we first apply Sentence-BERT (SBERT; Reimers and Gurevych (2019))<sup>2</sup> and project the output to the surface of an unit sphere to obtain representation vectors  $\psi(c_i)$  and  $\psi(c_j)$ . The surrogate relevance measure  $s(c_i, c_j)$  is then given by their inner product:  $s(c_i, c_j) = \psi(c_i)^\top \psi(c_j)$ . When there is more than one caption for an

<sup>2</sup>We use the code and the pre-trained model (bert-large-mlm-mean-tokens) provided in <https://github.com/UKPLab/sentence-transformers>.

image, we compute the surrogate relevance of all caption pairs and take the average. With the surrogate relevance, we are able to compute a proxy score for any pair of images in the training set, given their human-annotated captions. To validate the proposed surrogate relevance measure, we collect human judgments of semantic similarity between images by conducting a human experiment (details in Human Annotation Collection Section).

We train the scene graph similarity  $f$  by directly minimizing mean squared error from the surrogate relevance measure, formulating the learning as a regression problem. The loss function for  $i$ -th and  $j$ -th images is given as  $L_{ij} = \|f(\mathcal{S}_i, \mathcal{S}_j) - s(c_i, c_j)\|^2$ . Other losses, such as triplet loss or contrastive loss, can be employed as well. However, we could not find clear performance gains with those losses and therefore adhere to the simplest solution.

## Human Annotation Collection

We collect semantic similarity annotations from humans to validate the proposed surrogate relevance measure and to evaluate image retrieval methods. Through our web-based annotation system, a human labeler is asked whether two candidate images are semantically similar to a given query image. The labeler may choose one of four answers: either of the two candidate images is more similar than the other, images in the triplet are semantically identical, or neither of the candidate images is relevant to the query. We collect 10,712 human annotations from 29 human labelers for 1,752 image triplets constructed from the test set of the VG-COCO, the dataset we shall define in Experimental Setup Section.

A query image of a triplet is randomly selected from the query set defined in the following section. Two candidate images are randomly selected from the rest of the test set, subjected to two constraints. First, the rank of a candidate image should be less than or equal to 100 when the whole test set is sorted according to cosine similarity in ResNet-152 representation to the query image. Second, the surrogate relevance of a query-candidate image pair in a triplet should be larger than the other, and the difference should be greater than 0.1. This selection criterion produces visually close yet semantically different image triplets.

We define the *human agreement score* to measure the agreement between decisions of an algorithm and that of the human annotators, in a similar manner presented in (Gordo and Larlus 2017). The score is an average portion of human annotators who made the same decision per each triplet. Formally, given a triplet, let  $s_1$  (or  $s_2$ ) be the number of human annotators who chose the first (or the second) candidate image is more semantically similar to the query,  $s_3$  be the number of annotators who answered that all three images are identical, and  $s_4$  be the number of annotators who marked the candidates as irrelevant. If an algorithm choose either one of candidate images is more relevant, the human agreement score for a triplet is  $\frac{s_i + 0.5s_3}{s_1 + s_2 + s_3 + s_4}$ , where  $i = 1$  if the algorithm determines that the first image is semantically closer and  $i = 2$  otherwise. The score is averaged over triplets with  $s_1 + s_2 \geq 2$ . Randomly selecting one of two

candidate images produces an average human agreement of 0.472 with a standard deviation of 0.01. Note that the agreement of random decision is lower than 0.5 due to the existence of the human choice of "both" ( $s_3$ ) and "neither" ( $s_4$ ).

The alignment between labelers is also measured with the human agreement score in a leave-one-out fashion. If a human answers that both candidate images are relevant, the score for the triplet is  $\frac{0.5s_1 + 0.5s_2 + s_3}{s_1 + s_2 + s_3 + s_4}$ , where  $s_1 \dots s_4$  are computed from the rest of annotators. If a human marks that neither of the candidates is relevant for a triplet, the triplet is not counted in the human agreement score. The mean human agreement score among those annotators is 0.727, and the standard deviation is 0.05. We will make the human annotation dataset public after the review.

## Experimental Setup

### Data

In experiments, we use two image datasets involving diverse semantics. The first dataset is the intersection of the Visual Genome (Krishna et al. 2017) and MS-COCO (Lin et al. 2014), which we will refer to as VG-COCO. In VG-COCO, each image has a scene graph annotation provided by Visual Genome and five captions provided by MS-COCO. We utilize the refined version of scene graphs provided by (Xu et al. 2017) and their train-test split. After removing the images with empty scene graphs, we obtain fully annotated 35,017 training images and 13,203 test images. We randomly select a fixed set of 1,000 images among the test set and define them as a *query set*. For each query image, a retrieval algorithm is asked to rank the other 13,202 images in the test set according to the semantic similarity. Besides the annotated scene graphs, we automatically generate scene graphs for all images and experiment with our approach to both human-labeled and machine-generated scene graphs.

The second dataset is Flickr30K (Plummer et al. 2017), where five captions are provided per an image. Flickr30K contains 30,000 training images, 1,000 validation images, and 1,000 testing images. For Flickr30k, the whole test set is the query set. During the evaluation, an algorithm ranks the other 999 images given a query image in a test set. Scene graphs are generated in the same manner as in the VG-COCO dataset.

### Scene Graph Generation Detail

Since we focus on learning graph embeddings when two scene graphs are given for the image-to-image retrieval task, we use the conventional scene graph generation process. Following the works (Anderson et al. 2018), objects in images are detected by Faster R-CNN method, and the name and attributes of the objects are predicted based on the ResNet-101 features from the detected bounding boxes. We keep up to 100 objects with a confidence threshold of 0.3. To predict relation labels between objects after extracting information about the objects, we used the frequency prior knowledge constructed from the GQA dataset that covers 309 kinds of relations.<sup>3</sup> For each pair of the detected objects,

<sup>3</sup>We have been tried to predict relation labels by using recently suggested SGG algorithms, such as (Yang et al. 2018; Xu et al.

Method	Data	nDCG						Human Agreement
		5	10	20	30	40	50	
Inter Human	-	-	-	-	-	-	-	$0.730 \pm 0.05$
Caption SBERT	Cap(HA)	1	1	1	1	1	1	0.700
Random	-	0.136	0.138	0.143	0.147	0.149	0.152	$0.472 \pm 0.01$
Gen. Cap. SBERT	Cap(Gen)	0.609	0.628	0.657	0.681	0.703	0.726	0.473
ResNet	I	0.687	0.689	0.691	0.692	0.693	0.693	0.494
ResNet-FT	I	0.642	0.656	0.682	0.703	0.724	0.745	0.478
Object Count	I+SG	0.736	0.749	0.770	0.788	0.804	0.819	0.587
GMN	I+SG	0.721	0.735	0.755	0.771	0.786	0.801	0.535
IRSGS-GIN	I+SG	0.751	0.768	0.790	0.808	0.824	0.839	0.576
IRSGS-GCN	I+SG	<b>0.784</b>	<b>0.795</b>	<b>0.814</b>	<b>0.829</b>	<b>0.844</b>	<b>0.856</b>	<b>0.602</b>

Table 1: Image retrieval results on VG-COCO with human-annotated scene graphs. Data column indicates which data modalities are used. Cap(HA): human-annotated captions. Cap(Gen): machine-generated captions. I: image. SG: scene graphs.

relationships are predicted based on the frequency prior with confidence threshold 0.2. To give position-specific information, the coordinates of the detected bbox are used. Here, we should note that even though the suggested method to generate a scene graph is quite simple than other methods (Yang et al. 2018; Xu et al. 2017; Li et al. 2017), it outperforms all the others.

### Two-Step Retrieval using Visual Features

In information retrieval, it is a common practice to take a two-step approach (Wang et al. 2019; Bai and Bai 2016): retrieving roughly relevant items first and then sorting (or "re-ranking") the retrieved items according to the relevance. We also employ this approach in our experiment. For a query image, we first retrieve  $K$  images that are closest to the query in a ResNet-152 feature representation space formed by the 2048-dimension activation vector of the last hidden layer. The distance is measured in cosine similarity. This procedure generates a set of good candidate images which have a high probability of having strong semantic similarity. This approximate retrieval step can be further boosted by using an approximate nearest neighbor engine such as Faiss (Johnson, Douze, and Jégou 2017) and is critical if the following re-ranking step is computationally involved. We use this approximate pre-ranking for all experiments with  $K = 100$  unless otherwise mentioned. Although there is large flexibility of designing this step, we shall leave other possibilities for future exploration as the re-ranking step is our focus.

### Training Details

We use Adam optimizer with the initial learning rate of 0.0001. We multiply 0.9 to the learning rate every epoch. We set batch size as 32, and models are trained for 25 epochs. In each training step, a mini-batch of pairs is formed by randomly drawing samples. When drawing the second sample in a pair, we employ an oversampling scheme to reinforce the learning of pairs with large similarity values. With

2017; Li et al. 2017). However, we could not achieve any improvement in image retrieval tasks. The reasons might be that 1) small size vocabularies for object and relation are used for the conventional SGG setting (only 150/50 kinds of objects/relations), 2) the algorithms do not predict the attributes, and 3) the annotated scene graphs used for training the methods have very sparse relations.

Method	nDCG				Human Agreement
	5	10	20	40	
Inter Human	-	-	-	-	0.730
Caption SBERT	1	1	1	1	0.700
Random	0.136	0.138	0.143	0.149	0.472
Gen. Cap. SBERT	0.609	0.628	0.657	0.703	0.473
ResNet	0.687	0.689	0.691	0.693	0.494
ResNet-FT	0.642	0.656	0.682	0.724	0.478
Object Count	0.73	0.743	0.761	0.794	0.581
GWL	0.748	0.758	0.774	0.803	0.598
GMN	0.728	0.740	0.755	0.781	0.539
IRSGS-GIN	0.764	0.781	0.802	0.834	<b>0.612</b>
IRSGS-GCN	<b>0.771</b>	<b>0.784</b>	<b>0.805</b>	<b>0.836</b>	0.611

Table 2: Image retrieval results on VG-COCO with machine-generated scene graphs. Baselines which do not use scene graphs are identical to the corresponding rows of Table 1.

a probability of 0.5, the second sample in a pair is drawn from 100 most relevant samples with the largest surrogate relevance score to the first sample. Otherwise, we select the second sample from the whole training set. Oversampling improves both quantitative and qualitative results and is applied identically for all methods except for GWL where the scheme is not applicable.

## Experiments

### Evaluation

We benchmark IRSGS and other baselines with VG-COCO and Flickr30K. Images in the query set are presented as queries, and the relevance of the images ranked by an image retrieval algorithm is evaluated with two metrics. First, we compute normalized discounted cumulative gain (nDCG) with the surrogate relevance as gain. A larger nDCG value indicates stronger enrichment of relevant images in the retrieval result. In nDCG computation, surrogate relevance is clipped at zero to ensure its positivity. Second, the agreement between a retrieval algorithm and decision of human annotators is measured in a method described in Human Annotation Collection Section.

## Baseline Methods

**ResNet-152 Features** Image retrieval is performed based on the cosine similarity in the last hidden representation of ResNet-152 pre-trained on ImageNet.

**Generated Caption** To test whether machine-generated captions can be an effective means for semantic image retrieval, we generate captions of images by soft attention model (Xu et al. 2015) pretrained on Flickr30k dataset (Plummer et al. 2017). We obtain SBERT representations of generated captions, and their cosine similarity is used to perform image retrieval.

**Object Count (OC)** Ignoring relation information given in a scene graph, we transform a scene graph into a vector of object counts. Then, we compute the cosine similarity of object count vectors to perform image retrieval.

**ResNet Finetune (ResNet-FT)** We test whether a ResNet-152 can be fine-tuned to capture semantic similarity. Similarly to Siamese Network (Bromley et al. 1994), ResNet feature extractor is trained to produce cosine similarity between images close to their surrogate relevance measure.

**Gromov-Wasserstein Learning (GWL)** Based on Gromov-Wasserstein Learning (GWL) framework (Xu et al. 2019), we obtain a transport map using a proximal gradient method (Xie et al. 2018). A transport cost, a sum of Gromov-Wasserstein discrepancy and Wasserstein discrepancy, is calculated with the transport map and the cost matrix, and used for retrieval. The method is computationally demanding, and we only tested the method for VG-COCO with generated scene graphs setting in Table 2.

**Graph Matching Networks (GMN)** GMNs are implemented based on the publicly available code<sup>4</sup>. We use four propagation layers with shared weights. The propagation in the reverse direction is allowed, and the propagated representation is updated using the gated recurrent unit. Final node representations are aggregated by summation, resulting in a 128-dimensional vector which is then fed to a multi-layer perceptron to produce final scalar output. As GMN is capable of handling edge features, we leave relations as edges instead of transforming them as nodes. To indicate object-attribute connections, we append additional dimensionality to edge feature vectors and define a feature vector of an edge between an object and an attribute is a one-hot vector where only the last dimension is non-zero.

## Graph Embedding Methods in IRSGS

Here, we describe implementation details of graph neural networks used in IRSGS.

**IRSGS-GCN** A scene graph is applied with GCN and the final node representations are aggregated via mean pooling and scaled to the unit norm, yielding a representation vector  $\phi(\mathcal{S})$ . We use three graph convolution layers with 300 hidden neurons in each layer. The first two layers are followed by ReLU nonlinearity. Stacking more layers does not introduce clear improvement. We always symmetrize the adjacency matrix before applying GCN.

<sup>4</sup>[https://github.com/deepmind/deepmind-research/tree/master/graph\\_matching\\_networks](https://github.com/deepmind/deepmind-research/tree/master/graph_matching_networks)

Method	nDCG			
	5	10	20	40
Captions SBERT	1	1	1	1
Random	0.195	0.209	0.223	0.245
Gen. Cap. SBERT	0.556	0.576	0.610	0.659
Resnet	0.539	0.541	0.541	0.542
ResNet-FT	0.368	0.393	0.433	0.502
Object Count	0.511	0.530	0.560	0.615
IRSGS-GIN	0.564	0.584	0.618	<b>0.673</b>
IRSGS-GCN	<b>0.567</b>	<b>0.590</b>	<b>0.623</b>	0.672

Table 3: Image retrieval results on Flickr30K with machine-generated scene graphs.

**IRSGS-GIN** Similarly to GCN, we stack three GIN convolution layers with 300 hidden neurons in each layer. For multi-layer perceptrons required for each layer, we use one hidden layer with 512 neurons with ReLU nonlinearity. Other details are the same as that of the GCN case.

## Quantitative Results

From Table 1, Table 2, and Table 3, IRSGS shows larger nDCG score than baselines across datasets (VG-COCO and Flickr30K) and methods of obtaining scene graphs (human-annotated and machine-generated). IRSGS also achieves best agreement to human annotator’s perception on semantic similarity, as it can be seen from Table 1 and Table 2.

Comparing Table 1 and Table 2, we found that using machine-generated scene graphs instead of human-annotated ones does not deteriorate the retrieval performance. This result shows that IRSGS does not need human-annotated scene graphs to perform successful retrieval and can be applied to a dataset without scene graph annotation. In fact, Flickr30K is the dataset without scene graph annotation, and IRSGS still achieves excellent retrieval performance in Flickr30K with machine-generated scene graphs.

On the other hand, using machine-generated captions in retrieval results in significantly poor nDCG scores and human agreement scores. Unlike human-annotated captions, machine-generated captions are crude in quality and tend to miss important details of an image. We suspect that scene graph generation is more stable than caption generation since it can be done in a systematic manner, i.e., predicting objects, attributes, and relations in a sequential way.

While not showing the optimal performance, GWL and GMN also show competitive performance over other methods based on generated captions and ResNet. This overall tendency of competence of graph-based method is interesting and implies the effectiveness of scene graphs in capturing semantic similarity between images.

Note that in Caption SBERT, retrieval is performed with surrogate relevance, and their human agreement scores indicate the agreement between surrogate relevance and human annotations. With the highest human agreement score than any other algorithms, this result assures that the proposed surrogate relevance reflects the human perception of semantic similarity well.



Figure 3: Four most similar images retrieved by six algorithms. OC: Object Count, GIN: IRSGS-GIN, GCN: IRSGS-GCN. The visual genome ids for the query images are 2323522 and 2316427.

## Qualitative Results

Figure 1 and Figure 3 show the example images retrieved from the retrieval methods we test. Pitfalls of baseline methods that are not based on scene graphs can be noted. As mentioned in Introduction, retrieval with ResNet features often neglects the semantics and focuses on the superficial visual characteristics of images. On the contrary, OC only accounts for the presence of objects, yielding images with misleading context. For example, in the left panel of Figure 3, OC simply returns images with many windows. IRSGS could retrieve images containing similar objects with similar relations to the query image, for example, an airplane on the ground, or a person riding a horse.

## Discussion

**Ablation Study** We also perform an ablation experiment for effectiveness of each scene graph component (Table 4). In this experiment, we ignore attributes or randomize relation information from IRSGS-GCN framework. In both cases, nDCG and Human agreement scores are higher than the Object Count that uses only object information. This indicates that both attributes and relation information are useful to improve the image retrieval performance of the graph matching-based algorithm. Further, randomizing relations drops performance more than ignoring attribute information, which means that relations are important for capturing the human perception of semantic similarity.

**Comparison to Johnson et al. (2015)** We exclude Johnson et al. (2015) from our experiment because the CRF-based algorithm from Johnson et al. (2015) is not feasible in a large-scale image retrieval problem. One of our goals is to tackle a large-scale retrieval problem where a query is compared against more than ten thousand images. Thus, we mainly consider methods that generate a compact vector representation of an image or a scene graph (Eq.(2)). However, the method in Johnson et al. (2015) requires object detection results to be additionally stored and extra computation for all query-candidate pairs to be done in the retrieval phase. Note that Johnson et al. (2015) only tested their algorithm on 1,000 test images, while we benchmark algorithms using 13,203 candidate images.

Method	nDCG				Human Agreement
	5	10	20	40	
IRSGS-GCN	0.771	0.784	0.805	0.836	0.611
No Attribute	0.767	0.782	0.803	0.834	0.606
Random Relation	0.764	0.777	0.797	0.828	0.604
Object Count	0.730	0.743	0.761	0.794	0.581

Table 4: Scene graph component ablation experiment results on VG-COCO. Machine-generated scene graphs are used.

**Effectiveness of Mean Pooling and Inner Product** One possible explanation for the competitive performance of IRSGS-GCN and IRSGS-GIN is that the mean pooling and inner product are particularly effective in capturing similarity between two sets. Given two sets of node representations  $\{a_1, \dots, a_N\}$  and  $\{b_1, \dots, b_M\}$ , the inner product of their means are given as  $\sum_{i,j} a_i^\top b_j / (NM)$ , the sum of the inner product between all pairs. This expression is proportional to the number of common elements in the two sets, especially when  $a_i^\top b_j$  is 1 if  $a_i = b_j$  and 0 otherwise, measuring the similarity between the two sets. If the inner product values are not binary, then the expression measures the set similarity in a "soft" way.

## Conclusion

In this paper, we tackle the image retrieval problem for complex scene images where multiple objects are present in various contexts. We propose IRSGS, a novel image retrieval framework, which leverages scene graph generation and a graph neural network to capture semantic similarity between complex images. IRSGS is trained to approximate surrogate relevance measure, which we define as a similarity between captions. By collecting real human data, we show that both surrogate relevance and IRSGS show high agreement to human perception on semantic similarity. Our results show that an effective image retrieval system can be built by using scene graphs with graph neural networks. As both scene graph generation and graph neural networks are techniques that are rapidly advancing, we believe that the proposed approach is a promising research direction to pursue.

## Acknowledgements

Sangwoong Yoon is partly supported by the National Research Foundation of Korea Grant (NRF/MSIT2017R1E1A1A03070945) and MSIT-IITP (No. 2019-0-01367, BabyMind).

## References

- Alvarez-Melis, D.; and Jaakkola, T. 2018. Gromov-Wasserstein Alignment of Word Embedding Spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1881–1890.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6077–6086.
- Babenko, A.; Slesarev, A.; Chigorin, A.; and Lempitsky, V. 2014. Neural codes for image retrieval. In *European conference on computer vision*, 584–599. Springer.
- Bai, S.; and Bai, X. 2016. Sparse contextual activation for efficient visual re-ranking. *IEEE Transactions on Image Processing* 25(3): 1056–1069.
- Bromley, J.; Guyon, I.; LeCun, Y.; Säckinger, E.; and Shah, R. 1994. Signature verification using a “siamese” time delay neural network. In *Advances in neural information processing systems*, 737–744.
- Chen, B.-C.; Davis, L. S.; and Lim, S.-N. 2019. An Analysis of Object Embeddings for Image Retrieval. *arXiv preprint arXiv:1905.11903*.
- Chen, X.; Li, L.-J.; Fei-Fei, L.; and Gupta, A. 2018. Iterative visual reasoning beyond convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7239–7248.
- Das, A.; Kottur, S.; Gupta, K.; Singh, A.; Yadav, D.; Moura, J. M.; Parikh, D.; and Batra, D. 2017. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 326–335.
- Gong, Y.; Ke, Q.; Isard, M.; and Lazebnik, S. 2014. A multi-view embedding space for modeling internet images, tags, and their semantics. *International journal of computer vision* 106(2): 210–233.
- Gordo, A.; Almazán, J.; Revaud, J.; and Larlus, D. 2016. Deep image retrieval: Learning global representations for image search. In *European conference on computer vision*, 241–257. Springer.
- Gordo, A.; Almazán, J.; Revaud, J.; and Larlus, D. 2017. End-to-End Learning of Deep Visual Representations for Image Retrieval. *International Journal of Computer Vision* 124(2): 237–254.
- Gordo, A.; and Larlus, D. 2017. Beyond instance-level image retrieval: Leveraging captions to learn a global visual representation for semantic retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6589–6598.
- Gu, J.; Cai, J.; Joty, S.; Niu, L.; and Wang, G. 2018. Look, Imagine and Match: Improving Textual-Visual Cross-Modal Retrieval with Generative Models. *CVPR*.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Johnson, J.; Douze, M.; and Jégou, H. 2017. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734*.
- Johnson, J.; Krishna, R.; Stark, M.; Li, L.-J.; Shamma, D.; Bernstein, M.; and Fei-Fei, L. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3668–3678.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* 123(1): 32–73.
- Li, Y.; Gu, C.; Dullien, T.; Vinyals, O.; and Kohli, P. 2019. Graph Matching Networks for Learning the Similarity of Graph Structured Objects. *arXiv preprint arXiv:1904.12787*.
- Li, Y.; Ouyang, W.; Zhou, B.; Wang, K.; and Wang, X. 2017. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE International Conference on Computer Vision*, 1261–1270.
- Liang, Y.; Bai, Y.; Zhang, W.; Qian, X.; Zhu, L.; and Mei, T. 2019. VrR-VG: Refocusing Visually-Relevant Relationships. In *Proceedings of the IEEE International Conference on Computer Vision*, 10403–10412.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Lu, C.; Krishna, R.; Bernstein, M.; and Fei-Fei, L. 2016. Visual relationship detection with language priors. In *European Conference on Computer Vision*, 852–869. Springer.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2018. Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7219–7228.
- Maretic, H. P.; Gheche, M. E.; Chierchia, G.; and Frossard, P. 2019. GOT: An Optimal Transport framework for Graph comparison. *arXiv preprint arXiv:1906.02085*.
- Milewski, V.; Moens, M.-F.; and Calixto, I. 2020. Are scene graphs good enough to improve Image Captioning? *arXiv preprint arXiv:2009.12313*.
- Mohedano, E.; McGuinness, K.; O’Connor, N. E.; Salvador, A.; Marques, F.; and Giro-i Nieto, X. 2016. Bags of local convolutional features for scalable instance search. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, 327–331.

- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2017. Flickr30K Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models. *IJCV* 123(1): 74–93.
- Radenović, F.; Tolias, G.; and Chum, O. 2016. CNN image retrieval learns from BoW: Unsupervised fine-tuning with hard examples. In *European conference on computer vision*, 3–20. Springer.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *arXiv preprint arXiv:1908.10084*.
- Solomon, J.; Peyré, G.; Kim, V. G.; and Sra, S. 2016. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35(4): 72.
- Teney, D.; Liu, L.; and van den Hengel, A. 2017. Graph-structured representations for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1–9.
- Titouan, V.; Courty, N.; Tavenard, R.; Laetitia, C.; and Flamar, R. 2019. Optimal Transport for structured data with application on graphs. In *International Conference on Machine Learning*, 6275–6284.
- Vo, N.; Jiang, L.; Sun, C.; Murphy, K.; Li, L.-J.; Fei-Fei, L.; and Hays, J. 2019. Composing text and image for image retrieval-an empirical odyssey. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6439–6448.
- Wang, L.; Qian, X.; Zhang, Y.; Shen, J.; and Cao, X. 2019. Enhancing sketch-based image retrieval by cnn semantic re-ranking. *IEEE transactions on cybernetics*.
- Wei, Y.; Zhao, Y.; Lu, C.; Wei, S.; Liu, L.; Zhu, Z.; and Yan, S. 2016. Cross-modal retrieval with CNN visual features: A new baseline. *IEEE transactions on cybernetics* 47(2): 449–460.
- Wu, Q.; Shen, C.; Wang, P.; Dick, A.; and van den Hengel, A. 2017. Image captioning and visual question answering based on attributes and external knowledge. *IEEE transactions on pattern analysis and machine intelligence* 40(6): 1367–1381.
- Xie, Y.; Wang, X.; Wang, R.; and Zha, H. 2018. A Fast Proximal Point Method for Computing Exact Wasserstein Distance. *arXiv preprint arXiv:1802.04307*.
- Xu, D.; Zhu, Y.; Choy, C. B.; and Fei-Fei, L. 2017. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5410–5419.
- Xu, H.; Luo, D.; and Carin, L. 2019. Scalable Gromov-Wasserstein Learning for Graph Partitioning and Matching. *arXiv preprint arXiv:1905.07645*.
- Xu, H.; Luo, D.; Zha, H.; and Carin, L. 2019. Gromov-wasserstein learning for graph matching and node embedding. *arXiv preprint arXiv:1901.06003*.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, 2048–2057.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 670–685.
- Zeiler, M. D.; and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, 818–833. Springer.
- Zhen, L.; Hu, P.; Wang, X.; and Peng, D. 2019. Deep Supervised Cross-Modal Retrieval. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zheng, L.; Yang, Y.; and Tian, Q. 2017. SIFT meets CNN: A decade survey of instance retrieval. *IEEE transactions on pattern analysis and machine intelligence* 40(5): 1224–1244.