

Scene Graph Modification as Incremental Structure Expanding

Xuming Hu^{1*}, Zhijiang Guo^{2*}, Yu Fu¹, Lijie Wen^{1†}, Philip S. Yu^{1,3}

¹Tsinghua University

²University of Cambridge

³University of Illinois at Chicago

¹{hxm19, fy20}@mails.tsinghua.edu.cn

²zg283@cam.ac.uk ¹wenlj@tsinghua.edu.cn ³psyu@uic.edu

Abstract

A scene graph is a semantic representation that expresses the objects, attributes, and relationships between objects in a scene. Scene graphs play an important role in many cross modality tasks, as they are able to capture the interactions between images and texts. In this paper, we focus on scene graph modification (SGM), where the system is required to learn how to update an existing scene graph based on a natural language query. Unlike previous approaches that rebuilt the entire scene graph, we frame SGM as a graph expansion task by introducing the incremental structure expanding (ISE). ISE constructs the target graph by incrementally expanding the source graph without changing the unmodified structure. Based on ISE, we further propose a model that iterates between nodes prediction and edges prediction, inferring more accurate and harmonious expansion decisions progressively. In addition, we construct a challenging dataset that contains more complicated queries and larger scene graphs than existing datasets. Experiments on four benchmarks demonstrate the effectiveness of our approach, which surpasses the previous state-of-the-art model by large margins. Source code and data are available¹.

1 Introduction

A scene graph is a structural representation that captures the semantics of visual scenes by encoding object instances, attributes of objects, and relationships between objects. (Johnson et al., 2015). As shown in Figure 1, the scene graph encodes objects (e.g. “Boy”, “Racket”), attributes (e.g. “Girl is standing”), and relations (“Boy holding racket”). Scene graphs are able to capture the interactions between text and images by associating objects in the graph with regions of an image and modeling the relations between objects. Therefore, it has




Figure 1: Example images and their corresponding scene graphs. Given the query, the original scene graph (left) is modified to be the target scene graph (right).

been used in the cross modality task such as image retrieval, image captioning, and visual question answering (Schuster et al., 2015; Shi et al., 2019; Yang et al., 2019; Wang et al., 2020b).

Recently, modifying the scene graph based on the input becomes an emerging research direction as cross-modal systems may need to resort to an interactive process through multiple iterations (Ramnath et al., 2019; He et al., 2020). Take text-based image retrieval as an example, users start with a query describing the main objects or topics they are looking for, then modify the query to add more constraints or provide additional information based on previous search results. Instead of directly manipulating images, scene graphs can be used to convert the image-editing problem into a graph-editing problem, conditioned on the textual query. As shown in Figure 1, given a retrieved image from the last turn, if the user wants to see a girl rather than a boy holding a racket, he will enter the query “I would like to see a girl holding racket” to the system. According to the query, the object “Boy” in the original scene graph will be substituted with the object “Girl”. The target image can be retrieved

¹<https://github.com/THU-BPM/SGM>

* Equally Contributed.

† Corresponding Author.

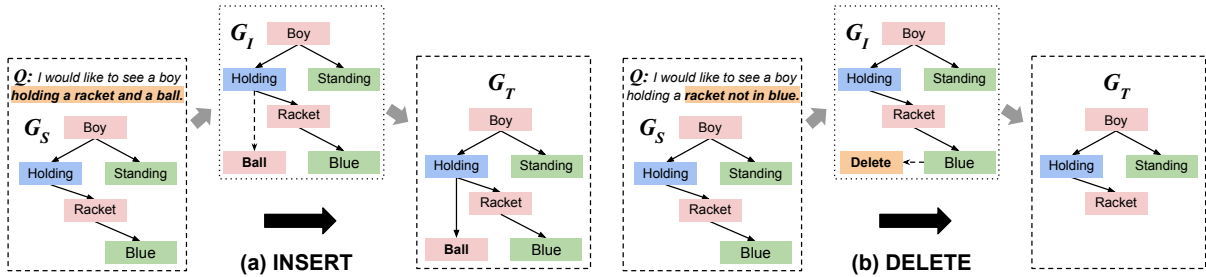


Figure 2: Examples of basic operations INSERT and DELETE for scene graph modification. Q denotes the textual query, G_S denotes the source scene graph, G_T denotes the target scene graph and G_I is the extended graph.

given the updated scene graph. The key challenge in this process is how to modify the corresponding partial structure in the original scene graph based on understanding the natural language query.

Prior effort framed this scene graph modification (SGM) task as conditional graph generation (He et al., 2020), where the scene graph is generated from the scratch condition on the original graph and query (You et al., 2018; Guo et al., 2019; Cai and Lam, 2020b). However, rebuilding the entire scene graph may not be an optimal solution, as the model has to generate the partial structure of the original graph that should be unmodified. Moreover, nodes and edges of the scene graph are constructed separately in their proposed framework, which generates all the nodes first then attaches edges between generated nodes in the second pass. Such an approach may lead to the lack of the modeling capability of interactions between node prediction and edge prediction.

Instead of rebuilding the whole scene graph, we introduce a novel formulation for SGM – incremental structure expanding (ISE), which is able to build the target graph by gradually expanding the original structure. At each step, ISE generates the connecting edges between the existing nodes and the newly generated node, upon which the type of the new node is jointly decided. Based on the formalism, our proposed model is able to iterate between finding the relevant part in the query and reading the partially constructed scene graph, inferring more accurate and harmonious expansion decisions progressively. Experiments on three SGM benchmarks demonstrate the effectiveness of the proposed approach, which is able to outperform previous state-of-the-art models by large margins. To test the ability of a model under a complex scenario, we further construct a more challenging dataset from the remote sensing domain (Lu et al., 2017), which has much more modification operations based on

the more complicated queries compared with the existing scene graph modification datasets. Our key contributions are summarized as follows:

- We propose a novel formulation for scene graph modification, allowing incremental expansion of the source scene graph rather than the regeneration of the target graph.
- We further construct a challenging dataset that contains more complicated queries and larger scene graphs. Extensive experiments on four SGM datasets show the effectiveness of our proposed approach.
- Experiments on four benchmarks demonstrate the effectiveness of our approach, which surpasses the previous state-of-the-art model by large margins.

2 Incremental Structure Expanding

In scene graph modification, a node or multiple nodes can be inserted to, deleted from or replaced with other nodes in the scene graph. He et al. (2020) defined the scene graph modification task as a conditional graph generation problem. Formally, given the source scene graph G_S and the natural language query Q , the target scene graph G_T is generated by maximizing the conditional probability $p(G_T|G_S, Q)$.

Instead of generating the entire target graph G_T , we frame the task as an incremental structure expanding, which extends the source scene graph G_S one node at a time, as well as the edges associated with the node. Such a formulation does not require the model to rebuild the unmodified structure of the source scene graph.

Under this formulation, we first define two basic operations: INSERT and DELETE. Scene graph modification can be viewed as combining and applying these two operations multiple times. For-




Figure 3: Overview of the model architecture.

mally, given the query Q , a sequence of n operations a_1, a_2, \dots, a_n are selected from a set of graph modification operations $\mathcal{A} = \{\text{INSERT}, \text{DELETE}\}$. After applying the operations to the source scene graph G_s , the target scene graph G_t is derived. Each operation is defined as:

- **INSERT:** A new node o is added to G_s , and edges are attached between o and existing nodes in G_s . As shown in Figure 2 (a), the node “Ball” is added to G_s and an edge between “Ball” and “Holding” is attached, according to the query “holding a racket and a ball”.
- **DELETE:** As shown in Figure 2 (b). A node o is removed from G_s , as well as its associated edges. As shown in Figure 2 (b), the node “Blue” is removed from G_s and the edge between “Racket” and “Blue” is removed either, according to the query “a racket not a blue”.

Inspired by incremental parsing (Nivre, 2004; Dyer et al., 2015; Cai and Lam, 2020a; Zhang et al., 2021, 2022), we design a data structure called extended graph G_I , which can be used to model INSERT and DELETE under the graph expansion setting. As shown in Figure 2 (a), the extended graph G_T is identical with the target graph G_T after applying INSERT. As for DELETE, we introduce a dummy node “Delete”, which is attached to the node in the source graph G_S that should be removed. For example, the dummy node “Delete” is attached to the node “Blue” in G_T . In the post-processing stage, nodes attached with the dummy node “Delete” will be removed. Using this formulation, we are able to model scene graph modification by incrementally expanding the source graph G_S to

the extended graph G_I , which can be converted to the target graph G_T without any losses.

If the modification requires multiple operations, there will exist multiple node orderings. Take node substitution as an example, replacing a node o_i with o_j in G_s can be viewed as DELETE the node o_i first, then INSERT the node o_j , or vice versa. In practice, we impose that the DELETE operation always comes before INSERT, then the breadth-first search is used to define a deterministic node ordering.

3 Model Architecture

In this section, we will present the model based on the incremental structure expanding formulation. Figure 3 gives an overview of the proposed model, which consists of five components including query encoder, graph encoder, feature fusion, edge decoder and node decoder.

Query Encoder This module is used to encode the query Q by generating the representation of each token of it.

Graph Encoder This module is used to encode the graph by generating the representation of each node of it. Note that the representations of the graph are constructed incrementally during the expanding progresses based on the updated graph of the last time step. The graph is the source graph G_S at the first timestep.

Feature Fusion this module aims to combine the representations from query and graph encoder, then served as a writable memory, which is updated based on the information from edge and node decoder during the incremental expansion.

Edge Decoder this module is used to predict the edges between the newly generated node and existing nodes of the graph, then update the memory of the feature fusion module with edge information.

Node Decoder this module is used to generate a new node of the graph, then update the memory of the feature fusion module with node information.

3.1 Query Encoder & Graph Encoder

For fair comparisons with the previous work (He et al., 2020), our query encoder and graph encoder are based on the vanilla transformer (Vaswani et al., 2017), which consists of multi-head self attention (MSA) and position-wise feed-forward network (FFN) blocks. The FFN contains two layers with a ReLU non-linearity. Layer normalization (Ba et al. 2016) is applied before every block, and residual connections (He et al., 2016) after every block.

Formally, given an input query Q with n tokens, each token embedding is randomly initialized and positional encoding is added to the token embedding to retain positional information. The resulted embeddings are denoted as $\mathbf{x} = \{x_0, x_1, \dots, x_n\}$. Similar to BERT (Devlin et al., 2019), a special token is appended to the query as x_0 for sentence encoding. Transformations in the query encoder can be denoted as:

$$\mathbf{x}^l = LN(MSA(\mathbf{x}^{l-1}) + \mathbf{x}^{l-1}), \quad (1)$$

$$\mathbf{x}^l = LN(FFN(\mathbf{x}^l) + \mathbf{x}^l). \quad (2)$$

After stacking L blocks, we obtained the contextualized token representations from the query encoder, denoted as $\{x_0^L, x_1^L, \dots, x_n^L\}$. The first vector x_0 is treated as the sentence-level representation of the query and will be used as the initial state during expansion. For clarity, we denote the vectors as $\mathbf{x} \in \mathbb{R}^{(n+1) \times d}$, where d is the dimension.

As for the graph encoder, we treat the input graph as a sequence of nodes in the chronological order of when they are inserted into the graph as discussed in Section 2. Formally, given the graph G_t at the time step t , we take its node sequence $\{o_1, o_2, \dots, o_{t-1}\}$ as the input. A transformer architecture is also applied to obtain the contextualized node embeddings. Notice that the contextualized representation of the graph is constructed incrementally as the expanding progress. Therefore, we apply the vanilla transformer with masked self-attention as the graph encoder, which only allows each position in the node sequence to attend

to all positions up to and including that position. For brevity, we denoted the resulted contextualized node representations as $\mathbf{y} \in \mathbb{R}^{m \times d}$.

3.2 Feature Fusion

Unlike the conventional sequence-to-sequence model that only has one encoder, our model contains two encoders. Previous work (He et al., 2020) proposed to use gating mechanism and cross attention to combine the representations of resulted representations from query and graph encoders. We choose to use vanilla multi-head attention mechanism (Vaswani et al., 2017) to fuse the features from these encoders. Formally, at each time step t , the feature fusion component combines the query and graph representations for gradually locating and collecting the most relevant information for the next expansion:

$$z_t^l = LN(MSA(h_t^{l-1}, \mathbf{x}) + h_t^{l-1}), \quad (3)$$

$$z_t^l = LN(MSA(z_t^l, \mathbf{y}) + z_t^l), \quad (4)$$

$$h_t^l = LN(FFN(z_t^l) + z_t^l). \quad (5)$$

The initial expansion state of h_t^0 is initialized with x_0 . For clarity, we denote the last hidden state h_t^L as h_t , which is the expansion state at the time step t . We now proceed to present the details of each decision stage of one expansion step.

3.3 Edge Decoder

At the t -th time step, the edge decoder takes the expansion state h_t from the feature fusion module and the contextualized representation \mathbf{y} from the graph encoder as the inputs, and predicts which nodes in the current graph should be attached to the new node. Inspired by Cai and Lam (2019) and Cai and Lam (2020a), we leverage multi-head attention and take the maximum over different heads as the final edge probabilities. Formally, given h_t and \mathbf{y} , a set of attention weights can be obtained by using multi-head attention mechanism: $\{\alpha_t^{g_i}\}_{i=1}^k$, where k is the number of attention heads and $\alpha_t^{g_i}$ is the i -th probability vector. The probability of the edge between the new node and the node o_j is then computed by $\alpha_t^g = \max_i(\alpha_t^{g_i})$. Intuitively, each head is in charge of a set of possible relations (though not explicitly specified). The maximum pooling reflects that the edge should be built once one relation is activated.

Finally, the edge decoder passes the edge information to the feature fusion module by updating the expansion state h_{t+1} as follows:

$$h_{t+1} = LN(MSA(h_t, \mathbf{y}) + h_t). \quad (6)$$

3.4 Node Decoder

The node decoder needs to look at the input query and determine which tokens are the most important ones. This choice is a weighted matrix that gives an attention probability between each token in the query and generated nodes in the target graph. Concretely, a single-head attention α_t^s is computed based on the state h_t and the sentence representation $s_{1:n}$, where α_t^s denotes the attention weight of the word w_i in the current time step. This component then updates the parser state with the alignment information via the following equation:

$$h_{t+1} = LN(MSA(h_t, \mathbf{x}) + h_t). \quad (7)$$

We then compute the probability distribution of the new node through a hybrid of two channels. The new node can either be a DELETE node or a token copied from the input query. First, h_t is fed through a *softmax* to obtain a probability distribution over a pre-defined vocabulary, which contains the DELETE node and other dummy nodes such as EOS. The probability of the new node is calculated as $P^{vocab} = softmax(W^{vocab}h_t + b^{vocab})$.

Second, we used the attention scores α_t^s as the probability to copy a token from the input query as a node label similar to the copy mechanism (Gu et al., 2016; See et al., 2017). Therefore, the final prediction probability of a node o is defined as:

$$P(o) = p_{gen} \cdot P^{vocab}(o) + p_{copy} \cdot \sum_{i \in T(o)} \alpha_t^s[i], \quad (8)$$

where $[i]$ indexes the i -th element, and $T(o)$ are index sets of tokens respectively that have the surface form as o . $P(gen)$ and $P(copy)$ are the probabilities of generating and copying a node, respectively. They are computed by using a single layer neural network with *softmax* activation as:

$$[p_{gen}, p_{copy}] = softmax(W^{gate}h_t). \quad (9)$$

The whole expanding procedure is terminated if the newly generated node is the special node EOS.

Statistics	User Generated	MSCOCO	GCC	RSICD
Splits	30/1/1	196/2/2	400/7/7	8/1/1
Avg. Source Nodes	2.0	2.9	3.8	5.9
Avg. Target Nodes	2.0	2.9	3.7	5.9
Avg. Source Edges	1.0	1.9	2.8	3.7
Avg. Target Edges	1.0	1.9	2.8	3.6
OOV Nodes	10	4	3	12
OOV Edges	8	4	4	8

Table 1: Statistics of four SGM datasets.

4 Dataset Construction

Existing SGM datasets are synthetically constructed based on scene graphs from MSCOCO (Lin et al., 2014) and GCC (Sharma et al., 2018a), and via crowd sourcing. To construct scene graphs, He et al. (2020) used an in-house scene graph parser to parse a random subset of MSCOCO description data and GCC captions, thus the constructed scene graph is relatively simple. In Table 1, the average numbers of nodes and edges for each graph are limited to 2.9 and 1.9 respectively. GCC is more complicated than MSCOCO with a larger graph, but the percentage of nodes and edges from the development/test set that does not appear in the training set (OOV Nodes, OOV Edges) are still low, which will cause the model easily overfit to the dataset. To verify the generalization ability and the scalability of the model to handle more complex scene graphs, we constructed our own Scene Graph Modification dataset based on the Remote Sensing Image Captioning Dataset (RSICD) (Lu et al., 2017) in the remote sensing field for remote sensing image captioning task.

Inspired by the modification methods proposed by He et al. (2020). First, we adopt the parser (Schuster et al., 2015) to parse the caption for each graph and generate the original scene graph \mathbf{x} . Then we define three types of graph modification operations $\mathcal{A} = \{\text{INSERT}, \text{DELETE}, \text{SUBSTITUTE}\}$, and randomly apply them to the original scene graph to generate query (\mathbf{q}) and modified scene graph (\mathbf{y}). The data in RSICD consists of the triples $(\mathbf{x}, \mathbf{y}, \mathbf{q})$.²

Compared with the existing SGM dataset, each graph of RSICD has more nodes and edges, with an average of 5.9 and 3.7 on the training/development/test set, which is almost twice that of User Generated and MSCOCO. In addition, the dataset comes from the field of remote sensing. Due to the large

²We give three detailed operations and examples in the Appendix A.1.

Models	User Generated			MSCOCO			GCC		
	Node F1	Edge F1	GAcc	Node F1	Edge F1	GAcc	Node F1	Edge F1	GAcc
CopyGraph (He et al., 2020)	66.17	31.42	—	78.41	64.62	—	79.46	66.32	—
Text2Text (He et al., 2020)	78.59	52.68	52.15	91.47	72.74	64.42	—	—	—
GRNN (You et al., 2018)	80.68	57.17	56.75	80.64	55.76	50.72	—	—	—
DCGCN (Guo et al., 2019)	79.05	54.23	52.67	89.08	72.47	68.89	—	—	—
GTran (Cai and Lam, 2020b)	81.47	59.43	58.23	91.21	75.68	71.38	—	—	—
STran (He et al., 2020)	83.69	62.10	60.90	95.40	86.52	82.97	93.84	57.68	52.50
EGraph (Weber et al., 2021)	97.62	88.26	87.60	99.52	98.40	96.15	98.62	91.64	75.01
ISE	98.74 \pm 0.12	91.37 \pm 0.14	89.41 \pm 0.47	99.68 \pm 0.14	98.96 \pm 0.21	97.26 \pm 0.37	99.53 \pm 0.13	93.06 \pm 0.22	76.34 \pm 0.47
ISE (w/o BERT)	94.39 \pm 0.11	79.53 \pm 0.18	75.72 \pm 0.47	98.17 \pm 0.13	97.25 \pm 0.14	89.61 \pm 0.45	96.91 \pm 0.16	85.50 \pm 0.21	58.40 \pm 0.56

Table 2: Results of User Generated, MSCOCO and GCC datasets. GAcc denotes the graph-level accuracy. Both of our models are statistically significantly outperform ($p < 0.0001$) previous best-reported model (Weber et al., 2021).

number of geographical terms, the OOV Nodes of the development/test sets compared with the training set reach 12%/11%, and the OOV Edges reach 8%/8%, which are much higher than the MSCOCO and GCC datasets. Considering the complexity of RSICD, we construct it apart from User Generated, MSCOCO and GCC to further analysis the generalization and scalability of ISE.

5 Experiments and Analyses

5.1 Data

We evaluated our model on four benchmarks, including User Generated, MSCOCO and GCC proposed by He et al. (2020), and RSICD dataset proposed in this work. MSCOCO, GCC and RSICD are constructed synthetically from publicly available datasets (Lin et al., 2014; Sharma et al., 2018b; Lu et al., 2017), while the User Generated dataset is created via crowd sourcing. Detailed statistics of datasets are shown in Table 1.

5.2 Setup

For fair comparisons, we used the same data splits for User Generated, MSCOCO and GCC datasets as in Weber et al. (2021). For RSICD, we randomly split the data into 8K/1K/1K for training/development/test. Following Weber et al. (2021), we use three automatic metrics for the evaluation, including node-level and edge-level F1 score, and graph-level accuracy. Graph-level accuracy is computed based on exact string match, which requires the generated scene graph to be identical to the target scene graph for a correct prediction. We reported the mean score and standard deviation by using 5 models from independent runs. We refer to the Appendix A.2 for the detailed implementation.

5.3 Baselines

For comprehensive comparisons, we include six baselines as follows. Except for the CopyGraph, all of them aim to rebuild the target scene graph.

CopyGraph This baseline directly copies the source scene graph as the target scene graph, which can be viewed as the lower bound.

Text2Text This baseline is introduced by He et al. (2020). They used the standard sequence-to-sequence architecture by linearizing the scene graph based on depth-first search.

GRNN Graph RNN (You et al., 2018) is used as the graph encoder and edge decoder. Specifically, the edges are represented by an adjacency matrix, which is then generated in an auto-regressive manner. Both the query encoder and node decoder are based on Gated Recurrent Units (Cho et al., 2014).

DCGCN Densely-Connected Graph Convolutional Networks (Guo et al., 2019) are used as the graph encoder. Other components are kept the same as the GRNN.

GTran Graph Transformer (Cai and Lam, 2020b) is used as the graph encoder, while other modules are the same as GRNN and DCGCN.

STran The sparsely-connected transformer (He et al., 2020) is used to encode the source graph. In addition, a cross-attention mechanism is applied to fuse the features from graph encoder and query encoder. Node decoder and edge decoder are the same as GRNN.

EGraph This is the state-of-the-art model on graph modification task. Concretely, Weber et al. (2021) considerably increases performance on the graph modification by phrasing it as a sequence labelling task.

Models	RSICD		
	Node F1	Edge F1	GAcc
CopyGraph	66.35	58.68	—
EGraph (Weber et al., 2021)	72.09±0.12	53.96±0.31	23.93±0.74
ISE	81.78±0.13	67.01±0.25	44.20±0.59

Table 3: Results of the RSICD dataset. Results of STran is reproduced from the implementation of Weber et al. (2021). Both of our models are statistically significantly outperform ($p<0.0001$) previous best-reported model (Weber et al., 2021).

5.4 Main Results

According to Table 2, our proposed approach (ISE) significantly outperforms the state-of-the-art model (Weber et al., 2021) on three datasets. Specifically, ISE outperforms EGraph 1.81, 1.11 and 1.33 percentage points in terms of graph accuracy on User Generated, MSCOCO and GCC datasets, respectively. We observe that the improvement is especially prominent on the User Generated dataset, which is more challenging than the other two synthetic datasets in terms of the diversity in graph semantics and natural language expressions. All baseline models suffer from performance degradation as it is much harder to rebuild the entire target scene graph on this dataset. On the other hand, ISE constructs the target scene graph by incrementally expanding the source scene graph without changing the unmodified structure. We believe this formulation is able to effectively cope with this difficulty.

We also observe that both EGraph and ISE achieve lower graph accuracy on the GCC dataset. The main reason is the difficulty of predicting the correct edges between generated nodes. For example, EGraph achieves 98.62 Node F1 score on GCC, higher than 97.62 Node F1 score on the User Generated dataset. However, EGraph only achieves 75.01 Edge F1 score on GCC, while it can attain 88.26 Edge F1 score on User Generated. Our proposed model has larger improvements upon EGraph in terms of Edge F1 score on the same dataset (93.06 vs. 91.64). We attribute this stronger improvement to iterations between nodes prediction and edge prediction, which allows more accurate and harmonious expansion decisions progressively. On the other hand, EGraph predicts nodes and edges at two independent stages. Such an approach may lead to the lack of the modeling capability of interactions between node prediction and edge prediction.

We further compare our model with EGraph on

Dataset/Model		Node F1	Edge F1	GAcc
User Generated	ISE	94.58	79.61	76.23
	ISE Rebuild	82.97	66.74	63.47
	ISE - Copy	88.38	74.95	71.64
MSCOCO	ISE	98.60	97.99	92.24
	ISE Rebuild	92.90	87.61	83.23
	ISE - Copy	95.58	91.87	88.09
GCC	ISE	96.87	85.50	58.90
	ISE Rebuild	89.48	62.59	51.67
	ISE - Copy	92.74	76.89	53.91

Table 4: An ablation study for ISE. Rebuild denotes that we regenerate the scene graph rather than extend it. - Copy denotes model without using the copy mechanism

the newly constructed dataset RSICD as shown in Table 3. ISE is able to achieve a graph accuracy of 44.20% and improves upon the EGraph model by 21 percentage points. However, the graph accuracy of all the models is much lower than the one attained on the previous three SGM datasets. One reason is that RSICD has more complex queries paired with larger scene graph, which brings a challenge to existing models. The RSICD dataset also suffers from the data sparsity issue where many words (39%) and nodes (42%) only appear once in the training data. Incorrect node prediction will further propagate the errors to edge prediction. Our iterative node and edge prediction paradigm help to alleviate this issue. Specifically, ISE only outperforms EGraph 9.69 percentage points on Node F1 score, while the improvement on Edge F1 score is 13.05%. Therefore, ISE is able to achieve a higher accuracy. In order to further address this data sparsity issue, one potential solution is transfer learning, where the model is pretrained on User Generated dataset first then fine-tuned on RSICD. However, this approach may suffer from a domain-shift problem, as RSICD is constructed based on the remote sensing domain. We leave this direction as future works.

5.5 Analysis and Discussion

In this section, we provided a fine-grained analysis of our proposed model. We reported all the results on the development set by using the ISE model without contextualized embeddings from BERT.

Ablation Study As shown in Table 4, we examine the contributions of two main components used in our model. The first one is the incremental structure expanding. We use the same model architecture but try to rebuild the target scene graph

% of Training Set		Node F1	Edge F1	GAcc
20%	STran	87.92	71.35	68.15
	ISE	95.46	92.22	79.12
40%	STran	93.94	81.11	78.55
	ISE	97.50	96.12	88.64
60%	STran	95.32	82.70	80.65
	ISE	98.09	97.17	89.29
80%	STran	95.92	86.36	83.90
	ISE	98.37	97.48	90.69
100%	STran	96.24	87.88	85.20
	ISE	98.60	97.99	92.24

Table 5: Comparison of STran and ISE against different training data sizes on the dev set of MSCOCO. Results of STran are reproduced from He et al. (2020).

Query Length		Node F1	Edge F1	Graph Acc
<5	STran	51.68	91.38	40.98
	ISE	96.02	92.36	57.38
5~10	STran	92.73	57.01	50.39
	ISE	97.32	86.84	60.18
≥10	STran	91.38	51.68	40.98
	ISE	98.42	84.62	58.64

Table 6: Comparison of STran and ISE against different lengths of queries.

similar to previous efforts. We can observe significant drops on three SGM datasets, which further confirms the effectiveness of the extending strategy. The second one is the copy mechanism, which directly copies the token from the query as nodes in the target scene graph. It plays a significant role in predicting nodes especially when the training data is limited (User Generated).

Performance against Training Data Size Table 5 shows the performance of STran and ISE against different training settings on MSCOCO dataset. We considered four training settings (20%, 40%, 60%, 80%, 100% training data). ISE consistently outperforms STran under the same amount of training data. When the size of training data decreases, we can observe that the performance gap becomes more obvious. Particularly, using 40% of the training data, ISE is able to achieve a graph accuracy of 88.64%, higher than STran trained on the whole dataset. These results demonstrate that our model is more effective in terms of using training resources and more robust when the training data is limited.

Graph Size		Node F1	Edge F1	Graph Acc
<5	STran	94.04	62.77	58.74
	ISE	96.73	81.25	62.54
5~10	STran	91.61	51.40	37.73
	ISE	97.22	88.82	49.06
≥10	STran	79.12	30.13	24.62
	ISE	95.44	90.95	35.38

Table 7: Comparison of STran and ISE against different target scene graph sizes.

Performance against Query Length Table 6 shows the results of STran and ISE under different query lengths on GCC dataset. We partitioned the sentence length into three classes (<5, [5, 10), ≥10). In general, ISE outperforms STran against various sentence lengths. When the length of the query increases, we can observe that the performance gap becomes more obvious in terms of graph accuracy. Intuitively, with the increase of the query length, it is more challenging for the model to comprehend the sentence. This suggests that ISE is able to handle more complex instructions.

Performance against Graph Size Table 7 shows the results of STran and ISE against different target scene graph sizes on GCC dataset. We partitioned the scene into three classes (<5, [5, 10), ≥10). Based on the formulation of extending the source scene graph, our model is required to deal with larger graphs. For example, deleting a node in the scene graph becomes adding a special “Delete” node in the extended graph. However, ISE is able to consistently outperform STran against various target graph sizes, even when the target scene graph is large. This result suggests the superiority of the proposed formulation.³

Case Study We give two cases in Figure 4. STran generates scene graph from the scratch conditioned on the original graph and query may lead to the lack of the modeling capability of interactions between node prediction and edge prediction. For example, in Figure 4 (a), STran omitted the attribute: “Velvet” during the node prediction. In addition, during the edge prediction, STran redundantly generated the relation: “Of” in Figure 4 (b). However, these structures do not need to be modified in the source scene graph. ISE can infer more accurate target graph by incrementally expanding the source graph without changing the unmodified structure.

³We give an error analysis in the Appendix A.4.

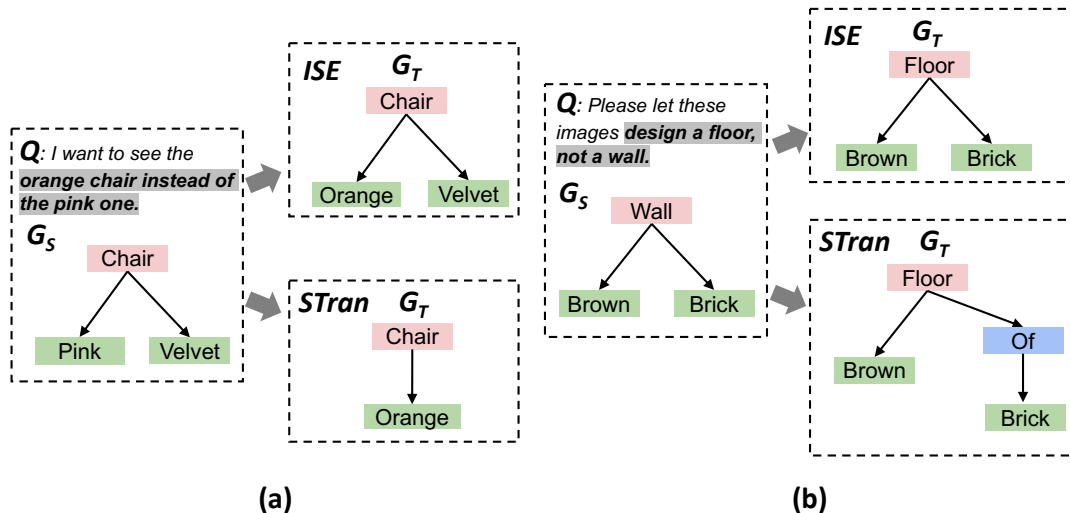


Figure 4: Two cases of STran and ISE for scene graph modification on User Generated. Q denotes the textual query, G_S denotes the source scene graph, G_T denotes the target scene graph generated by STran and ISE.

6 Related Work

We refer to the Appendix A.3 for the detailed related work of scene graph. Scene graph builds a bridge between image domain and text domain. Vision and natural language are all tremendously promoted by studying into scene graphs. Recently, scene graph modification becomes an emerging research direction. Chen et al. (2020) proposed a framework based on scene graph editing for text-based image retrieval. On the other hand, He et al. (2020) took the scene graph and the textual query as inputs and modified the source graph according to the query. They viewed the task as conditional graph generation, which is further decomposed into node prediction and edge prediction. For node prediction, all the nodes in the target scene graph is generated based on a graph-to-sequence model with dual encoder (Song et al., 2018; Beck et al., 2018; Zhang et al., 2020), then a graph RNN is adopted to predict the edges between generated nodes (You et al., 2018). More recently, Weber et al. (2021) developed an alternative formulation of this problem in which they model the modification as an auto-regressive sequence labelling task.

Instead of rebuilding the entire target graph, we framed the scene graph modification task as incremental graph expansion. This formulation is related to incremental parsing, where a sentence is scanned from left-to-right and the structured is built incrementally by inserting a node or attaching an edge. Incremental parsers are widely used in semantic parsing (Zhou et al., 2016; Cheng et al., 2017; Guo and Lu, 2018; Naseem et al., 2019; Liu

et al., 2022a) and syntactic parsing (Huang and Sagae, 2010; Dyer et al., 2015; Liu and Zhang, 2017), as they are computationally efficient, and can use machine learning to predict actions based on partially generated structures. Our feature fusion module can be viewed as the parser state as it carries the structural information and serves as a writable memory during the expansion step. Unlike Weber et al. (2021) linearize the scene graph and label it in an auto-regressive manner, our model iterates between finding the relevant part in the query and reading the partially constructed scene graph, inferring more accurate and harmonious expansion decisions progressively.

7 Conclusion

In this paper, we designed a novel formulation for scene graph modification, which allows us to incrementally expand the source scene graph instead of rebuilding the entire graph. Based on the formalism, we further propose a model that is able to leverage the mutual causalities between node prediction and edge prediction. Experiments on three SGM benchmarks demonstrate the effectiveness. To test our model under a complex scenario, we constructed a more challenging dataset from the remote sensing domain, which has more modification operations based on the more complicated queries compared with existing SGM datasets. For future work, we would like to explore how to integrate the model into the text-based image retrieval task.

8 Acknowledgement

We thank the reviewers for their valuable comments. The work was supported by the National Key Research and Development Program of China (No. 2019YFB1704003), the National Nature Science Foundation of China (No. 62021002 and No. 71690231), NSF under grants III-1763325, III-1909323, III-2106758, SaTC-1930941, Tsinghua BNRist and Beijing Key Laboratory of Industrial Bigdata System and Application.

References

- Iro Armeni, Zhi-Yang He, JunYoung Gwak, Amir R Zamir, Martin Fischer, Jitendra Malik, and Silvio Savarese. 2019. 3d scene graph: A structure for unified semantics, 3d space, and camera. In *ICCV*, pages 5664–5673.
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *CoRR*, abs/1607.06450.
- Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. Graph-to-sequence learning using gated graph neural networks. In *ACL*.
- Deng Cai and Wai Lam. 2019. [Core semantic first: A top-down approach for AMR parsing](#). In *EMNLP*.
- Deng Cai and Wai Lam. 2020a. Amr parsing via graph-sequence iterative inference. In *ACL*.
- Deng Cai and Wai Lam. 2020b. Graph transformer for graph-to-sequence learning. In *AAAI*, volume 34.
- Lichang Chen, Guosheng Lin, Shijie Wang, and Qingyao Wu. 2020. Graph edit distance reward: Learning to edit scene graph. In *ECCV*.
- Jianpeng Cheng, Siva Reddy, Vijay A. Saraswat, and Mirella Lapata. 2017. [Learning structured natural language representations for semantic parsing](#). In *ACL*.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](#). In *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *NAACL-HLT*.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. [Transition-based dependency parsing with stack long short-term memory](#). In *ACL*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *NLP-OSS*, Melbourne, Australia.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *ACL*. The Association for Computer Linguistics.
- Zhijiang Guo and Wei Lu. 2018. [Better transition-based AMR parsing with refined search space](#). In *EMNLP*.
- Zhijiang Guo, Yan Zhang, Zhiyang Teng, and Wei Lu. 2019. Densely connected graph convolutional networks for graph-to-sequence learning. *TACL*, 7:297–312.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *CVPR*.
- Xuanli He, Quan Hung Tran, Gholamreza Haffari, Walter Chang, Trung Bui, Zhe L. Lin, Franck Dernoncourt, and Nhan Dam. 2020. Scene graph modification based on natural language commands. In *EMNLP*.
- Xuming Hu, Fukun Ma, Chenyao Liu, Chenwei Zhang, Lijie Wen, and Philip S Yu. 2021a. Semi-supervised relation extraction via incremental meta self-training. In *EMNLP: Findings*.
- Xuming Hu, Lijie Wen, Yusong Xu, Chenwei Zhang, and S Yu Philip. 2020. Selfore: Self-supervised relational feature learning for open relation extraction. In *EMNLP*.
- Xuming Hu, Chenwei Zhang, Yawen Yang, Xiaohe Li, Li Lin, Lijie Wen, and S Yu Philip. 2021b. Gradient imitation reinforcement learning for low resource relation extraction. In *EMNLP*.
- Liang Huang and Kenji Sagae. 2010. [Dynamic programming for linear-time incremental parsing](#). In *ACL*.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, D. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. *CVPR*, pages 3668–3678.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *ICLR*.
- Shu’ang Li, Xuming Hu, Li Lin, and Lijie Wen. 2022. Pair-level supervised contrastive learning for natural language inference. In *ICASSP*.
- Xiaodan Liang, Lisa Lee, and Eric P Xing. 2017. Deep variation-structured reinforcement learning for visual relationship and attribute detection. In *CVPR*.
- Tsung-Yi Lin, M. Maire, Serge J. Belongie, James Hays, P. Perona, D. Ramanan, Piotr Dollár, and C. L. Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV*.

- Aiwei Liu, Xuming Hu, Li Lin, and Lijie Wen. 2022a. Semantic enhanced text-to-sql parsing via iteratively learning schema linking graph. In *KDD*.
- Jiangming Liu and Yue Zhang. 2017. [In-order transition-based constituent parsing](#). *TACL*, 5:413–424.
- Shuliang Liu, Xuming Hu, Chenwei Zhang, Shu’ang Li, Lijie Wen, and Philip S. Yu. 2022b. Hiure: Hierarchical exemplar contrastive learning for unsupervised relation extraction. In *NAACL*.
- Xiaochen Liu, Yu Bai, Jiawei Li, Yanan Hu, and Yang Gao. 2022c. Psp: Pre-trained soft prompts for few-shot abstractive summarization. *arXiv preprint arXiv:2204.04413*.
- Xiaoqiang Lu, Binqiang Wang, Xiangtao Zheng, and Xuelong Li. 2017. Exploring models and data for remote sensing image caption generation. *IEEE Transactions on Geoscience and Remote Sensing*, 56(4):2183–2195.
- Ramesh Manuvinakurike, Jacqueline Brixey, Trung Bui, Walter Chang, Doo Soon Kim, Ron Artstein, and Kallirroi Georgila. 2018. Edit me: A corpus and a framework for understanding natural language image editing. In *LREC*, Miyazaki, Japan.
- Tahira Naseem, Abhishek Shah, Hui Wan, Radu Florian, S. Roukos, and Miguel Ballesteros. 2019. Rewarding smatch: Transition-based amr parsing with reinforcement learning. In *ACL*.
- Joakim Nivre. 2004. Incrementality in deterministic dependency parsing.
- Hang Qi, Yuanlu Xu, Tao Yuan, Tianfu Wu, and Song-Chun Zhu. 2018. Scene-centric joint parsing of cross-view videos. In *AAAI*.
- Sahana Ramnath, Amrita Saha, Soumen Chakrabarti, and Mitesh M. Khapra. 2019. Scene graph based image retrieval - a case study on the clevr dataset. *ArXiv*, abs/1911.00850.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. 2021. A comprehensive survey of neural architecture search: Challenges and solutions. *ACM Computing Surveys (CSUR)*, 54(4):1–34.
- Brigit Schroeder and Subarna Tripathi. 2020. Structured query-based image retrieval using scene graphs. *CVPRW*, pages 680–684.
- Sebastian Schuster, Ranjay Krishna, Angel X. Chang, Li Fei-Fei, and Christopher D. Manning. 2015. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *VL@EMNLP*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *ACL*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018a. [Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *ACL*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018b. [Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *ACL*.
- Jiaxin Shi, Hanwang Zhang, and Juanzi Li. 2019. Explainable and explicit visual reasoning over scene graphs. In *CVPR*.
- Linfeng Song, Yue Zhang, Zhiguo Wang, and D. Gildea. 2018. A graph-to-sequence model for amr-to-text generation. In *ACL*.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Damien Teney, Lingqiao Liu, and Anton van den Hengel. 2017. [Graph-structured representations for visual question answering](#). In *CVPR*, pages 3233–3241. IEEE Computer Society.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NeurIPS*.
- Peng Wang, Qi Wu, Jiewei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. Neighbourhood watch: Referring expression comprehension via language-guided graph attention networks. In *CVPR*.
- Ruize Wang, Zhongyu Wei, Piji Li, Qi Zhang, and Xuanjing Huang. 2020a. Storytelling from an image stream using scene graphs. In *AAAI*, volume 34, pages 9185–9192.
- Sijin Wang, Ruiping Wang, Ziwei Yao, S. Shan, and Xilin Chen. 2020b. Cross-modal scene graph matching for relationship-aware image-text retrieval. *WACV*, pages 1497–1506.
- Yu-Siang Wang, Chenxi Liu, Xiaohui Zeng, and Alan L. Yuille. 2018. [Scene graph parsing as dependency parsing](#). In *NAACL-HLT*, pages 397–407. Association for Computational Linguistics.
- Leon Weber, Jannes Münchmeyer, Samuele Garda, and Ulf Leser. 2021. Extend, don’t rebuild: Phrasing conditional graph modification as autoregressive sequence labelling. In *EMNLP*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin

- Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *EMNLP*, Online.
- Danfei Xu, Yuke Zhu, Christopher B. Choy, and Li Fei-Fei. 2017. [Scene graph generation by iterative message passing](#). In *CVPR*, pages 3097–3106. IEEE Computer Society.
- Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. 2019. [Auto-encoding scene graphs for image captioning](#). In *CVPR*, pages 10685–10694. Computer Vision Foundation / IEEE.
- Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. 2018. Exploring visual relationship for image captioning. In *ECCV*.
- Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and J. Leskovec. 2018. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *ICML*.
- Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. 2018. [Neural motifs: Scene graph parsing with global context](#). In *CVPR*, pages 5831–5840. IEEE Computer Society.
- Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, Xiaobin Wang, and Min Zhang. 2022. Identifying chinese opinion expressions with extremely-noisy crowdsourcing annotations. In *ACL*.
- Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, and Pengjun Xie. 2021. Crowdsourcing learning as domain adaptation: A case study on named entity recognition. In *ACL*.
- Yan Zhang, Zhijiang Guo, Zhiyang Teng, Wei Lu, Shay B. Cohen, Zuozhu Liu, and Lidong Bing. 2020. [Lightweight, dynamic graph convolutional networks for amr-to-text generation](#). In *EMNLP*.
- Junsheng Zhou, Feiyu Xu, Hans Uszkoreit, Weiguang Qu, Ran Li, and Yanhui Gu. 2016. AMR parsing with an incremental joint model. In *EMNLP*.

Graph Modification Operation: DELETE
Original Scene Graph: Some *trees* are in a medium residential area.
Query: Remove *trees*.
Modified Scene Graph: Some are in a medium residential area.

Graph Modification Operation: INSERT
Original Scene Graph: A bridge built on a river.
Query: Show me a *red* bridge.
Modified Scene Graph: A *red* bridge built on a river.

Graph Modification Operation: SUBSTITUTE
Original Scene Graph: Some gray and *green* mountains are together.
Query: I prefer *red* to *green*, modify *red* to *green*.
Modified Scene Graph: Some gray and *red* mountains are together.

Table 8: Examples on the three types of graph modification operations $\mathcal{A} = \{\text{INSERT, DELETE, SUBSTITUTE}\}$

A Appendix

A.1 Operations in RSICD

We introduce three operations in RSICD in details:

- **DELETE:** The original scene graph is \mathbf{x} . We randomly select a node \mathbf{o} in \mathbf{x} , and delete it both with related edges. The deleted graph is defined as \mathbf{y} . We choose a random sentence from the *DELETE Template* (Manuvinakurike et al., 2018), for example, “I do not want **.” We replace ** with \mathbf{o} to get modification operation \mathbf{q} .
- **INSERT:** It is the reverse process of **DELETE**. The graph before deleting the node is regarded as \mathbf{y} , and the corresponding graph after deletion is treated as \mathbf{x} . The modification operation is randomly selected from the *INSERT Template* (Manuvinakurike et al., 2018), for example, “Show me **.” We replace ** with \mathbf{o} to obtain query \mathbf{q} .
- **SUBSTITUTE:** We randomly select a node \mathbf{o} , use the AllenNLP toolkit (Gardner et al., 2018) to find the three most similar semantics nodes compared with \mathbf{o} . We randomly choose a node \mathbf{m} , and select a sentence from the *SUBSTITUTE Template* (Manuvinakurike et al., 2018), for example, “I prefer @@ to **, modify ** to @@.” We replace ** and @@ with \mathbf{o} and \mathbf{m} , and get modification operation \mathbf{q} . Note that SUBSTITUTE operation could be viewed as DELETE the node \mathbf{o} first and then INSERT the node \mathbf{m} , or vice versa.

In Table 8, we give the simple examples in RSICD to better understand three types of graph modification operations.

Embeddings	
concept	300
word	300
relation	100
Query Encoder	
transformer layers	4
Graph Encoder	
transformer layers	2
Feature Fusion	
heads	8
hidden size	512
feed-forward hidden size	1024
Node Decoder/ Edge Decoder	
heads	8
feed-forward hidden size	1024

Table 9: Hyper-parameters settings for ISE.

A.2 Implementation Details

Hyper-parameters of the model are tuned on the development set. All transformer (Vaswani et al., 2017) layers share the same hyper-parameter settings. Following He et al. (2020), we randomly initialized the word and node embeddings. We also report results with contextualized embeddings from BERT (Devlin et al., 2019). Specifically, we used the BERT-base-uncased implemented by (Wolf et al., 2020). The parameters in BERT are fixed during training. To mitigate over-fitting, we apply dropout (Srivastava et al., 2014) with the drop rate 0.2 between different layers. Following Cai and Lam (2020a), we use a special UNK token to replace the out-of-vocabulary lemmas of the input query and remove the UNK token in the generated graph. Parameter optimization is performed with the ADAM optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The learning rate schedule is similar to that in Vaswani et al. (2017), where warm-up steps being set to 2K. We used early stopping on the development set for choosing the best model. Please refer to Table 9 for the detailed hyper-parameters settings for ISE.

A.3 Scene Graph and Application

Deep learning has significantly promoted the advancement of computer vision (Liang et al., 2017; Ren et al., 2021). Simple visual understanding tasks such as object detection and recognition are no longer sufficient. To depict the relationship between objects in the scene as a driving force,




Figure 5: Two errors of ISE for scene graph modification on User Generated. Q denotes the textual query, G_S denotes the source scene graph, G_T denotes the target scene graph generated by ISE. G_G denotes the gold target scene graph.

higher-level visual understanding and reasoning skills are frequently necessary. Scene graphs were created specifically to address this issue. Scene graph was first proposed by Johnson et al. (2015) for image retrieval, which describes objects, their attributes, and relationships in images with a graph. A complete scene graph could represent the semantics of a dataset’s scenes, not just a single image or video; additionally, it contains powerful representations that encode 2D/3D images (Johnson et al., 2015; Armeni et al., 2019), and videos (Qi et al., 2018; Wang et al., 2020a) into their abstract semantic elements. Scene graph is beneficial for various downstream tasks, such as information extraction (Hu et al., 2020, 2021a,b; Liu et al., 2022b), natural language summarization (Liu et al., 2022c), and natural language inference (Li et al., 2022).

Following the graph representation paradigm, different methods have been proposed to generate scene graphs from images (Xu et al., 2017; Wang et al., 2018; Zellers et al., 2018). Many cross-modal tasks that require understanding and reasoning on image and text are able to benefit from incorporating scene graphs, such as visual question answering (Teney et al., 2017; Shi et al., 2019), grounding referring expressions (Wang et al.), image captioning (Yang et al., 2019; Yao et al., 2018), and image retrieval (Wang et al., 2020b; Schroeder and Tripathi, 2020).

A.4 Error Analysis

We give two wrong scene graphs generated by ISE in Figure 5. We can observe in Figure 5 (a) that although ISE successfully predicts the need to insert a relation between object “Plants” and attribute “Surface”, since the User Generated dataset contains a total of 2078 relations and the relations have serious long-tail effects. It is difficult for ISE to learn sparseness relations with few occurrences, leading to incorrectly predicting relation “in growing over” as “on”. We attempt to address the long-tail effects of relations in future work. Since a node can be attached to multiple nodes, when Edge Decoder determines which nodes in the current graph should be attached to the new node, a common error is predicting the wrong node that needs to be attached. As shown in Figure 5 (b), ISE incorrectly connects relation “behind” between “Giraffe” and “Tree” instead of “Head” and “Tree”.