

Grounded Situation Recognition with Transformers

Junhyeong Cho*¹
 junhyeong99@postech.ac.kr
 Youngseok Yoon*¹
 yys8646@postech.ac.kr
 Hyeonjun Lee*²
 hyeonjun1882@postech.ac.kr
 Suha Kwak^{1,2}
 suha.kwak@postech.ac.kr

¹ Department of CSE
 POSTECH
 Pohang, Republic of Korea
² Graduate School of AI
 POSTECH
 Pohang, Republic of Korea

Abstract

Grounded Situation Recognition (GSR) is the task that not only classifies a salient action (*verb*), but also predicts entities (*nouns*) associated with semantic roles and their locations in the given image. Inspired by the remarkable success of Transformers in vision tasks, we propose a GSR model based on a Transformer encoder-decoder architecture. The attention mechanism of our model enables accurate verb classification by capturing high-level semantic feature of an image effectively, and allows the model to flexibly deal with the complicated and image-dependent relations between entities for improved noun classification and localization. Our model is the first Transformer architecture for GSR, and achieves the state of the art in every evaluation metric on the SWiG benchmark. Our code is available at <https://github.com/jhcho99/gstrr>.

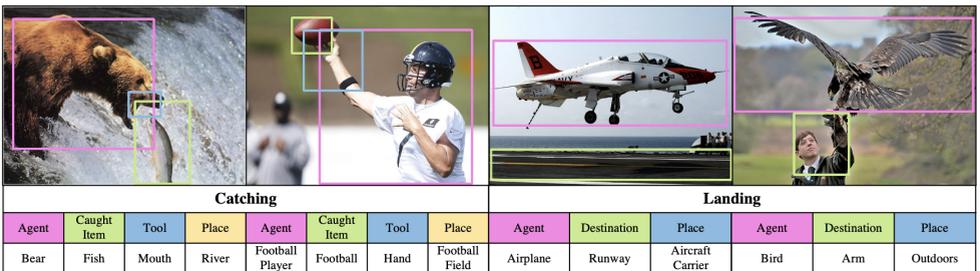


Figure 1: Predictions of our model on the SWiG dataset.

1 Introduction

Deep learning models have achieved or even surpassed human-level performance on basic vision tasks such as classification of objects [6, 15], actions [18, 51], and places [6, 12, 52].

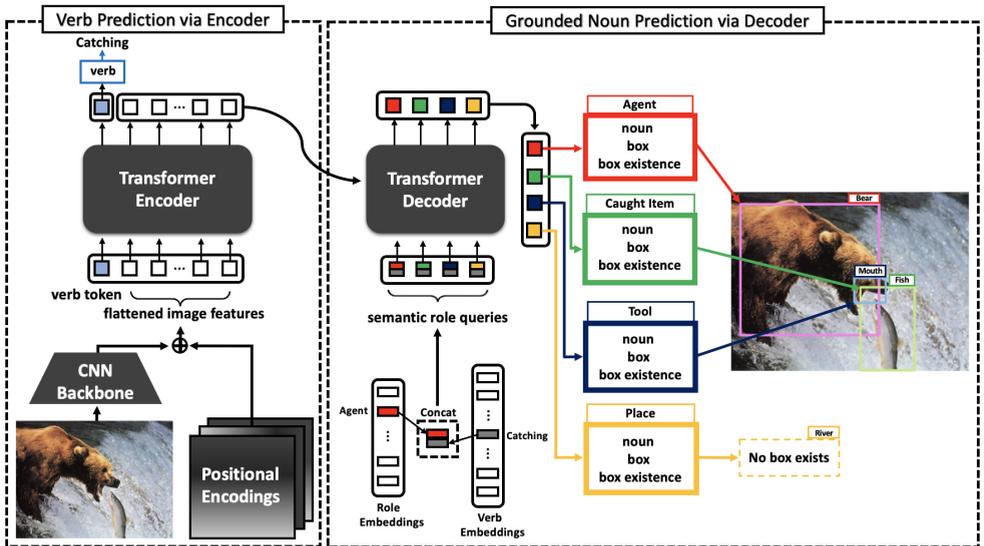


Figure 2: The overall architecture of our model (GSRTR). It mainly consists of two components: Transformer Encoder for verb prediction, and Transformer Decoder for grounded noun prediction. Diagram is best viewed in colored version.

However, it still remains challenging and less explored to expand such models for detailed and comprehensive understanding of natural scenes, *e.g.*, recognizing what happens and who are involved with which roles. Image captioning [1, 2, 29] and scene graph generation [11, 25, 26] have been studied in this context. These tasks aim at reasoning about image contents in detail and describing them through natural language captions or relation graphs of objects. However, quality evaluation of natural language captions is not straightforward, and scene graphs are limited in terms of expressive power as they represent an action only by a triplet of subject, predicate, and object.

Grounded Situation Recognition (GSR) [16] is a comprehensive scene understanding task that resolves the above limitations. It originates from Situation Recognition (SR) [27], the task of predicting a salient action, entities taking part of the action, and their roles altogether given an image. In SR, an action and entities are called *verb* and *nouns*, respectively, and the set of semantic *roles* of the entities in an action is termed *frame*; a frame is defined for each verb as prior knowledge by FrameNet [9], a lexical database of English. Then SR is typically done by predicting a verb then assigning a noun to each role given by the frame of the verb. GSR has been introduced to further address localization (*i.e.*, bounding box estimation) of the nouns in the image, which is missing in SR. It is thus more challenging yet enables more detailed scene understanding in comparison with SR.

The major challenge in GSR is two-fold. The first is the difficulty of verb prediction. This is caused by the fact that a verb is a high-level concept embodied by multiple entities; as illustrated in Fig. 1, images of the same verb often vary significantly due to different entities interacting in different ways. The second is the difficulty of modeling complicated relations between entities. Since an action (*i.e.*, verb) is performed by multiple entities (*i.e.*, nouns) related to each other, individual noun recognition per role is definitely suboptimal; relations between nouns have to be considered for improved noun prediction and localiza-

tion. However, modeling such relations is challenging since they are latent and depending on an input image.

Inspired by the recent success of Transformers [1, 3, 21], we present in this paper a new model, dubbed GSRTR, that addresses the aforementioned challenges through the attention mechanism. As illustrated in Fig. 2, it has an encoder-decoder architecture based on Transformer. The encoder takes as input a verb token and image features from a CNN backbone. The token then goes through self-attention blocks in the encoder and is finally processed by a verb classifier on top. Thanks to the self-attention with the image features, the encoder can capture rich and high-level semantic information for accurate verb prediction. Meanwhile, the decoder predicts a grounded noun per role, where target roles are determined by the frame of the target verb. It thus takes as input *semantic role queries* of target roles as well as image features given by the encoder; a semantic role query is obtained by a concatenation of two embedding vectors, one for its role and the other for a verb, which are learnable parameters dedicated to the role and verb, respectively. Each semantic role query is converted to a feature vector through attention blocks, then used to predict a noun class, a box coordinate and a box existence probability of its role. The attention blocks in our decoder allow to capture complicated and image-dependent relations among roles effectively and flexibly.

Contributions: Our GSRTR is the first Transformer architecture dedicated to GSR. Furthermore, its encoder-decoder architecture is carefully designed to address major challenges of the task. The efficacy of GSRTR is validated on the SWiG dataset [16], the standard benchmark for GSR, where it clearly outperforms existing models [16] in every evaluation metric. We also provide in-depth analysis on behaviors of GSRTR, which demonstrates that it has the capability of drawing attentions on local areas relevant to verb and grounded nouns.

2 Related Work

Situation Recognition: Situation Recognition (SR) is the task of predicting a salient action (*verb*) and entities (*nouns*) taking part of the action. Yatskar *et al.* [27] present the *imSitu* dataset as benchmark of Situation Recognition and propose Conditional Random Field (CRF) model. Their following work [28] figures out that sparsity of training examples compared to large output space could be problematic, and alleviates it through tensor-composition function. Since then, there have been attempts to model the relations among semantic roles. Inspired by image captioning task, Mallya and Lazebnik [14] adopt a Recurrent Neural Network (RNN) architecture to model the relations in the predefined order. Li *et al.* [8] use a Gated Graph Neural Network (GGNN) [9] to capture relations among roles, and Suhail and Sigal [19] propose a modified GGNN to learn context-aware relations among roles depending on the content of the image. Cooray *et al.* [2] formulate the relation modeling as an interdependent query based visual reasoning problem.

Grounded Situation Recognition: Recently, Grounded Situation Recognition (GSR) has been introduced by Pratt *et al.* [16] to further address localization of entities, which is missing in SR. They propose the Situation With Groundings (SWiG) dataset that provides bounding box annotations in addition to the *imSitu* dataset. They also propose Joint Situation Localizer (JSL) model which consists of a verb classifier and a RNN based object detector. The object detector sequentially produces noun and its bounding box prediction via the predefined role order. Compared with JSL, our GSRTR can flexibly capture the relations among the semantic roles rather than the predefined order. Furthermore, the verb prediction process in our model can capture long-range interactions of semantic concepts via a Transformer encoder.

Transformer in Vision Tasks: Dosovitskiy *et al.* [9] propose a standard Transformer encoder architecture [20] for image classification task. This model, called ViT, takes image patches flattened, linearly transformed, and combined with positional encodings as input together with a class token. On the other hand, the encoder of GSRTR takes image features from a CNN backbone as input, and is combined with a decoder for grounded noun prediction. Carion *et al.* [10] view object detection as a direct set prediction and bipartite matching problem, and propose a Transformer encoder-decoder architecture for object detection accordingly. Their model, called DETR, introduces learnable embeddings called *object queries* as inputs of the decoder, each of which is in charge of a certain image region and a set of bounding box candidates. Instead of the object queries, GSRTR uses *semantic role queries*, each of which focuses on entities taking part of a specified action with a specific role.

Similar follow-ups to DETR: There have been attempts, including our GSRTR, to apply DETR to other domains such as video instance segmentation [23], video action recognition [30] and human-object-interaction detection [33]. Their models use latent queries for a Transformer decoder in the similar way, but GSRTR has notable differences. While their models employ a fixed number of latent queries in the decoder, GSRTR constructs a variable number of queries depending on a given image. Also, to the best of our knowledge, GSRTR is the first attempt to explicitly leverage the output of a Transformer encoder for building queries used in a Transformer decoder; semantic role queries use the verb embedding corresponding to the predicted verb from the encoder output at inference time.

3 Proposed Method

Inspired by ViT [9] and DETR [10], we propose a novel model called Grounded Situation Recognition TRansformer (GSRTR) to address the challenging GSR task; the architecture of GSRTR is illustrated in Fig. 2. This section first provides a formal definition of GSR, then describes details of our model architecture, training and inference procedures.

3.1 Task Definition

Let \mathcal{V} , \mathcal{R} , and \mathcal{N} denote the sets of verbs, roles, and nouns defined in the task, respectively. For each verb $v \in \mathcal{V}$, a set of semantic roles, denoted by $\mathcal{R}_v \subset \mathcal{R}$, is predefined as its frame by FrameNet [4]. For example, the frame of a verb *Catching* is a set of semantic roles $\mathcal{R}_{\text{Catching}} = \{\text{Agent, Caught Item, Tool, Place}\} \subset \mathcal{R}$. Also, a pair of a noun $n \in \mathcal{N}$ and its bounding box $\mathbf{b} \in \mathbb{R}^4$ is called a *grounded noun*. The goal of GSR is to predict a verb v of an input image and assign a grounded noun to each role in \mathcal{R}_v . Formally speaking, a prediction of GSR is in the form of $S = (v, \mathcal{F}_v)$, where $\mathcal{F}_v = \{(r, n_r, \mathbf{b}_r) \mid n_r \in \mathcal{N} \cup \{\emptyset_n\}, \mathbf{b}_r \in \mathbb{R}^4 \cup \{\emptyset_b\}\}$ for $r \in \mathcal{R}_v$; \emptyset_n and \emptyset_b mean *unknown* and *not grounded*, respectively. For example, the prediction for the leftmost image in Fig. 1 is given by $S = (\text{Catching}, \{(Agent, Bear, \square), (Caught Item, Fish, \square), (Tool, Mouth, \square), (Place, River, \emptyset_b)\})$.

3.2 Encoder for Verb Prediction

A CNN backbone first processes an input image to extract its feature map $X_{img} \in \mathbb{R}^{c \times h \times w}$, where c is the number of channels and $h \times w$ is the resolution of X_{img} . Then X_{img} is fed to a 1×1 convolution layer for reducing the channel size to d , and flattened, leading to flattened

image features $F_{img} \in \mathbb{R}^{d \times hw}$. Like the class token used in ViT [9], we append a learnable verb token $\mathbf{f}_v \in \mathbb{R}^d$ to F_{img} , forming an input of the encoder $F \in \mathbb{R}^{d \times (1+hw)}$.

The encoder is a stack of six layers, each of which consists of a Multi-Head Self-Attention (MHSA) block and a Feed Forward Network (FFN) block. Also, we apply Pre-Layer Normalization (Pre-LN) [24] before the MHSA and FFN blocks. Positional encodings are added to the input of each encoder layer. Please refer to the supplementary material for more details of the encoder.

The output of the encoder, denoted by $E \in \mathbb{R}^{d \times (1+hw)}$, is split into a verb feature $\mathbf{e}_v \in \mathbb{R}^d$ and hw image features $E_{img} \in \mathbb{R}^{d \times hw}$. The former is fed to the verb classifier, which in turn produces a logit vector $\mathbf{z}_v \in \mathbb{R}^{|\mathcal{V}|}$ as a result of verb classification. On the other hand, the latter will be used as observations for the decoder. Note that by exploiting the attention mechanism through the encoder layers, the verb token can effectively aggregate relevant semantic features of an image for accurate verb classification.

3.3 Decoder for Grounded Noun Prediction

In addition to the image features E_{img} given by the encoder, the decoder takes as input semantic role queries to predict corresponding nouns and their bounding boxes, inspired by the object queries in DETR [10]. To be specific, a semantic role query $\mathbf{w}_{(v,r)} \in \mathbb{R}^d$ is obtained by a concatenation of a verb embedding vector $\mathbf{w}_v \in \mathbb{R}^{d_v}$ and a role embedding vector $\mathbf{w}_r \in \mathbb{R}^{d_r}$ ($d = d_v + d_r$), both of which are learnable parameters; v is the ground-truth verb at training time and the predicted verb at inference time, while $r \in \mathcal{R}_v$. The number of semantic role queries fed to the decoder is thus $|\mathcal{R}_v|$.

The decoder is a stack of six layers, each of which consists of a MHSA block, a Multi-Head Attention (MHA) block, and a FFN block; Pre-LN is applied before each of the blocks. The first decoder layer input is set to zero. In each decoder layer, each semantic role query $\mathbf{w}_{(v,r)}$ is added to each key and query of the MHSA block and added to each query of the MHA block. The image features E_{img} serve as keys and values in the MHA block of each decoder layer. Through the MHSA block in each decoder layer, semantic role queries flexibly capture the role relations (Fig. 4). From the MHA block in each decoder layer, each semantic role query attends to image features considering image-dependent relations (Fig. 3).

Through the decoder, each semantic role query $\mathbf{w}_{(v,r)}$ is converted to an output feature. The output feature of each role $r \in \mathcal{R}_v$ is in turn fed to three branches: One for noun classification, another for bounding box regression, and the other for predicting existence of its bounding box. The noun classifier produces a noun logit vector $\mathbf{z}_{n_r} \in \mathbb{R}^{|\mathcal{N} \cup \{0_n\}|}$. The bounding box regressor predicts $\hat{\mathbf{b}}'_r = (\hat{x}_x, \hat{x}_y, \hat{w}, \hat{h}) \in [0, 1]^4$, indicating the normalized center coordinate, height, and width of a box relative to the image size. This predicted box coordinate is transformed into top-left and bottom-right coordinate representation $\hat{\mathbf{b}}_r = (\hat{x}_1, \hat{y}_1, \hat{x}_2, \hat{y}_2) \in \mathbb{R}^4$. Finally, the box existence predictor produces a box existence probability $p_{b_r} \in [0, 1]$. Please refer to the supplementary material for more details of the decoder.

3.4 Training and Inference

The total loss for training GSRTR is a linear combination of five losses: A verb classification loss, a noun classification loss, a bounding box existence loss, a L_1 box regression loss, and a Generalized IoU (GIoU) [17] box regression loss. The verb classification loss \mathcal{L}_v is the cross entropy between the verb prediction probability $\mathbf{p}_v = \text{Softmax}(\mathbf{z}_v)$ and the ground-truth verb

distribution. The noun classification loss \mathcal{L}_n is formulated as the average of individual noun classification losses over the semantic roles, and is given by

$$\mathcal{L}_n = \frac{1}{|\mathcal{R}_v|} \sum_{r \in \mathcal{R}_v} \text{CrossEntropy}(\mathbf{p}_{n_r}, \mathbf{t}_{n_r}), \quad (1)$$

where \mathbf{p}_{n_r} denotes the noun prediction probability for each role r and \mathbf{t}_{n_r} indicates the ground-truth noun distribution for each role r . The bounding box existence loss \mathcal{L}_{exist} is the average of individual bounding box existence loss over the semantic roles, and is given by

$$\mathcal{L}_{exist} = \frac{1}{|\mathcal{R}_v|} \sum_{r \in \mathcal{R}_v} \text{CrossEntropy}(p_{b_r}, t_{b_r}), \quad (2)$$

where p_{b_r} denotes the bounding box existence probability for each role r and $t_{b_r} \in \{0, 1\}$ specifies the existence of the ground-truth bounding box for each role r (i.e., $t_{b_r} = 1$ when $\mathbf{b}_r \neq \emptyset_b$). The L_1 box regression loss \mathcal{L}_{L_1} is defined as the average of individual L_1 distances between predicted and ground-truth bounding boxes over semantic roles for which ground-truth bounding boxes exist, and are given by

$$\mathcal{L}_{L_1} = \frac{1}{|\tilde{\mathcal{R}}_v|} \sum_{r \in \tilde{\mathcal{R}}_v} \|\hat{\mathbf{b}}'_r - \mathbf{b}'_r\|_1, \quad (3)$$

where $\tilde{\mathcal{R}}_v = \{r \mid r \in \mathcal{R}_v \text{ and } \mathbf{b}_r \neq \emptyset_b\}$ is the set of roles associated with bounding boxes. Finally, the GIoU box regression loss \mathcal{L}_{GIoU} [14] is formulated as the average of individual GIoU losses over roles for which ground-truth bounding boxes exist, and are given by

$$\mathcal{L}_{GIoU} = \frac{1}{|\tilde{\mathcal{R}}_v|} \sum_{r \in \tilde{\mathcal{R}}_v} \left(1 - \left(\frac{|\mathbf{b}_r \cap \hat{\mathbf{b}}_r|}{|\mathbf{b}_r \cup \hat{\mathbf{b}}_r|} - \frac{|C(\mathbf{b}_r, \hat{\mathbf{b}}_r) \setminus \mathbf{b}_r \cup \hat{\mathbf{b}}_r|}{|C(\mathbf{b}_r, \hat{\mathbf{b}}_r)|} \right) \right), \quad (4)$$

where $C(\hat{\mathbf{b}}_r, \mathbf{b}_r)$ denotes the smallest box enclosing predicted box $\hat{\mathbf{b}}_r$ and ground-truth box \mathbf{b}_r for each role r . GIoU loss is a scale-invariant loss and it compensates for scale-variant L_1 loss. The total loss \mathcal{L}_{total} is formulated as $\mathcal{L}_{total} = \lambda_v \mathcal{L}_v + \lambda_n \mathcal{L}_n + \lambda_{exist} \mathcal{L}_{exist} + \lambda_{L_1} \mathcal{L}_{L_1} + \lambda_{GIoU} \mathcal{L}_{GIoU}$, where $\lambda_v, \lambda_n, \lambda_{exist}, \lambda_{L_1}, \lambda_{GIoU} > 0$ are hyperparameters.

At inference time, our method predicts a verb $\hat{v} = \arg \max_v \mathbf{p}_v$ then constructs corresponding semantic role queries $\mathbf{w}_{(\hat{v}, r)}$ for all $r \in \mathcal{R}_{\hat{v}}$. Each $\mathbf{w}_{(\hat{v}, r)}$ is used by the decoder to produce corresponding output noun logit \mathbf{z}_{n_r} , bounding box $\hat{\mathbf{b}}'_r$ and bounding box existence probability p_{b_r} . Note that if $p_{b_r} < 0.5$, the predicted bounding box $\hat{\mathbf{b}}'_r$ is ignored.

4 Experiments

4.1 Dataset and Metrics

SWiG [14] dataset is composed of 75k, 25k and 25k images for the train, development and test set respectively. There are $|\mathcal{V}| = 504$ verbs, $|\mathcal{R}| = 190$ roles, and $1 \leq |\mathcal{R}_v| \leq 6$ semantic roles per verb. We use about 10k nouns, the number of noun classes in the train set. The annotation for each image consists of a verb, a bounding box for each semantic role, and three nouns (from three annotators) for each semantic role.

Table 1: Requirements for each metric.

metric	requirement				
	correct verb	correct noun for a semantic role	correct nouns for all semantic roles	correct bounding box for a semantic role	correct bounding boxes for all semantic roles
<i>verb</i>	✓				
<i>value</i>	✓	✓			
<i>value-all</i>	✓	✓	✓		
<i>grounded-value</i>	✓	✓		✓	
<i>grounded-value-all</i>	✓	✓	✓	✓	✓

The predicted verb and grounded nouns are measured by five metrics: *verb*, *value*, *value-all*, *grounded-value*, and *grounded-value-all*. The *verb* metric denotes a verb prediction accuracy. The *value* metric denotes a noun prediction accuracy from its semantic role. The *value-all* metric denotes that all nouns corresponding to semantic roles are correctly predicted. The *grounded-value* metric denotes a grounded noun prediction accuracy for its semantic role. Note that the grounded noun prediction is considered correct if it correctly predicts noun and bounding box. The bounding box prediction is considered correct if it correctly predicts bounding box existence and the predicted box has an Intersection-over-Union (IoU) value of at least 0.5 with the ground-truth box. The *grounded-value-all* metric denotes that all grounded nouns corresponding to semantic roles are correctly predicted. The requirements for each metric are summarized in Table 1. Because the number of roles per verb is different and the number of images per verb could be different, all above metrics are calculated for each verb and then averaged over them.

Since these metrics depend heavily on the verb accuracy, the metrics are reported in 3 settings: **top-1 predicted verb**, **top-5 predicted verbs** and **ground-truth verb**. In **top-1 predicted verb** setting, five metrics are reported: a top-1 predicted verb accuracy, two noun metrics and two grounded noun metrics. If the top-1 predicted verb is incorrect, the noun and grounded noun metrics are considered incorrect. In **top-5 predicted verbs** setting, five metrics are reported: a top-5 predicted verbs accuracy, two noun metrics and two grounded noun metrics. If the ground-truth verb is not included in the top-5 predicted verbs, the noun and grounded noun metrics are considered incorrect, too. In **ground-truth verb** setting, four metrics are reported: two noun metrics and two grounded noun metrics. From the ground-truth verb assumed to be known, noun and grounded noun predictions are taken from the model by conditioning on the ground-truth verb.

4.2 Implementation Details

Following previous work [14], we use ImageNet-pretrained ResNet-50 backbone [6] except Feature Pyramid Network (FPN) [10]. The ResNet-50 backbone produces the image features $X_{img} \in \mathbb{R}^{c \times h \times w}$ from the input image where $c = 2048$. The hidden dimensions of each semantic role query, verb token and image feature are 512 ($d = 512$). The verb embedding dimension and role embedding dimension are 256 ($d_v = d_r = 256$). We use learnable 2D embeddings for the positional encodings. The number of heads for all MHSA and MHA blocks is 8. We use 2 fully connected layers with ReLU activation function for the four followings: the FFN blocks in the encoder and decoder, the verb classifier, the noun classifier, and the bounding box existence predictor. The size of hidden dimensions are 2048, $2d$, $2d$, and $2d$, respectively. The dropout rates are 0.15, 0.3, 0.3, and 0.2, respectively. The bounding box regressor is 3 fully connected layers with ReLU activation function and $2d$ hidden dimensions, using 0.2 dropout rate. The label smoothing regularization [20] is used

for the target verb and noun labels with label smoothing factor 0.3 and 0.2, respectively. We use AdamW [L3] optimizer with the learning rate 10^{-4} (10^{-5} for the backbone), weight decay 10^{-4} , $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We set the max gradient clipping value to 0.1 and train the BatchNorm layers in the backbone. The training epoch is 40 with batch size 16 per GPU on four 12GB TITAN Xp GPUs, which takes about 20 hours. The loss coefficients are $\lambda_v = \lambda_n = 1$ and $\lambda_{exist} = \lambda_{L_1} = \lambda_{GIoU} = 5$.

Data Augmentation: Random Color Jittering, Random Gray Scaling, Random Scaling and Random Horizontal Flipping are used. The hue, saturate and brightness scale in random color jittering set to 0.1. The scale of random gray scaling sets to 0.3. The scales of random scaling set to 0.5, 0.75 and 1.0. The probability of random horizontal flipping sets to 0.5.

Final Noun Loss: In SWiG, three noun annotations exist per role. For each noun annotation, we calculate the loss (Eq. 1). The final noun loss is the summation of the three noun losses.

Batch Training: The number of semantic roles ranges from 1 to 6 depending on the frame of a verb. In GSRTR, the semantic role queries are constructed as much as the number of semantic roles. To ensure batch training, zero padding is used for each output of grounded noun prediction branches. We ignore the padded outputs in the loss computation.

4.3 Experiment Results

Quantitative Comparison with Previous Work: Table 2 quantitatively compares our model with previous work on the *dev* and *test* splits of SWiG dataset. In all evaluation metrics, GSRTR achieves the state-of-the-art accuracy. In the *dev* set, compared with JSL, GSRTR achieves the top-1 predicted verb and top-5 predicted verbs accuracies of 41.06% (+1.46%p) and 69.46% (+1.75%p), respectively. In ground-truth verb setting, GSRTR achieves the value and grounded-value accuracies 74.27% (+0.74%p) and 58.33% (+0.83%p), respectively. Note that previous work uses two ResNet-50 backbones and FPN, while our GSRTR only uses a single ResNet-50 backbone without FPN. Existing models in [L6] have about 108 million parameters, but our GSRTR only has about 83 million parameters. Although GSRTR has less backbone capacity and less parameters, it achieves the state-of-the-art accuracy in every evaluation metric. In addition, the reason for the small improvement by GSRTR in terms of grounded-value metrics is that these metrics require correct predictions of verb, noun and bounding box as shown in Table 1.

Existing models in [L6] are trained separately in terms of verb prediction part and grounded noun prediction part, while our GSRTR is trained in an end-to-end manner. For this reason, it is difficult to fairly compare the training time of ours with existing models. However, we can reasonably guess that GSRTR takes less training time than others. GSRTR takes about 20 hours with four 12GB TITAN Xp GPUs for whole training, but other models take about 20 hours with four 24GB TITAN RTX GPUs only for training of grounded noun prediction part. For the comparison of inference time, we compare GSRTR with JSL which was the previous state-of-the-art. We evaluate the models on the *test* set in the same environment with one 2080Ti GPU. GSRTR takes 21.69 ms (46.10 FPS) and JSL takes 80.00 ms (12.50 FPS) on the average of 10 trials.

Effect of Verb Embedding Concatenation: We also quantitatively show the effect of verb embedding concatenation in the semantic role query. If we do not concatenate the verb embedding (*i.e.*, $d_v = 0$ and $d_r = d$), the accuracies in the ground-truth verb setting decrease by around 1.3 ~ 2.3%p (GSRTR w/o VE in Table 2). It demonstrates that the verb embedding concatenation is helpful for grounded noun prediction.

Table 2: Quantitative evaluation on the SWiG dataset.

set	model	top-1 predicted verb					top-5 predicted verbs					ground-truth verb			
		verb	value	value-all	grnd value	grnd value-all	verb	value	value-all	grnd value	grnd value-all	value	value-all	grnd value	grnd value-all
dev	ISL []	38.83	30.47	18.23	22.47	7.64	65.74	50.29	28.59	36.90	11.66	72.77	37.49	52.92	15.00
	JSL []	39.60	31.18	18.85	25.03	10.16	67.71	52.06	29.73	41.25	15.07	73.53	38.32	57.50	19.29
	GSRTTR w/o VE (Ours)	40.81	32.05	19.31	25.64	10.31	69.33	53.09	29.78	42.01	15.36	72.55	37.07	57.00	18.93
	GSRTTR (Ours)	41.06	32.52	19.63	26.04	10.44	69.46	53.69	30.66	42.61	15.98	74.27	39.24	58.33	20.19
test	ISL []	39.36	30.09	18.62	22.73	7.72	65.51	50.16	28.47	36.60	11.56	72.42	37.10	52.19	14.58
	JSL []	39.94	31.44	18.87	24.86	9.66	67.60	51.88	29.39	40.60	14.72	73.21	37.82	56.57	18.45
	GSRTTR w/o VE (Ours)	40.61	31.87	19.01	25.21	9.69	69.75	53.25	29.67	41.65	14.93	72.32	36.75	56.03	18.02
	GSRTTR (Ours)	40.63	32.15	19.28	25.49	10.10	69.81	54.13	31.01	42.50	15.88	74.11	39.00	57.45	19.67

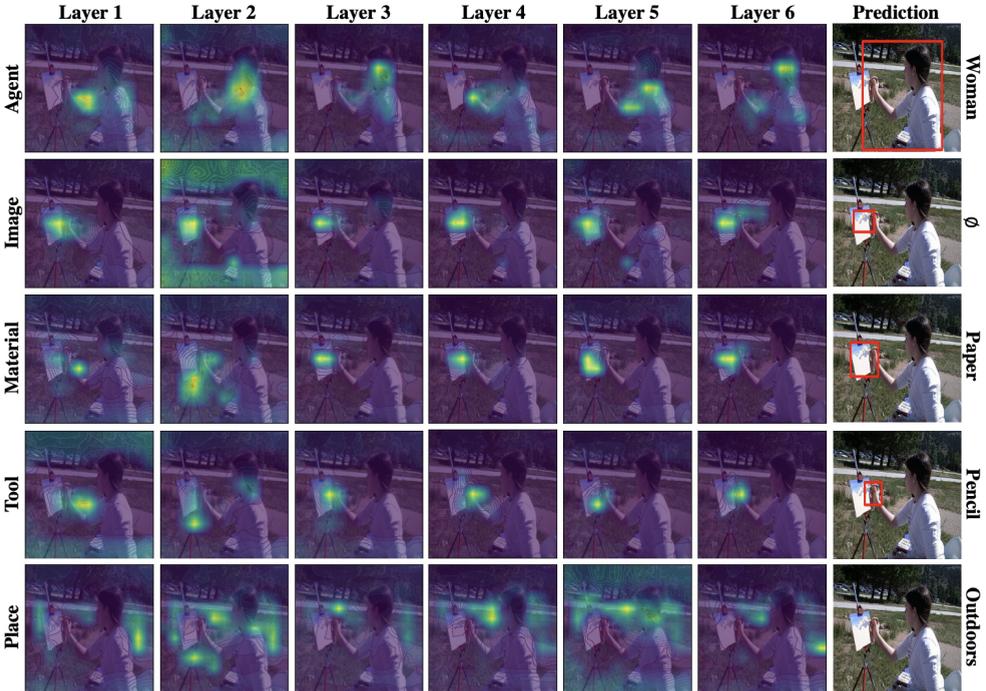


Figure 3: Role Attention Map on Image Features for a *Sketching* image from the MHA block in each decoder layer. The left labels are the semantic roles of the verb *Sketching*. The rightmost column images and labels are predicted bounding boxes and nouns of our model.

Role Attention Map on Image Features: In Figure 3, each column shows the difference of attention maps among semantic roles. For example, at Layer 6, the role *Agent* focuses on the woman, and the role *Place* focuses on the road and yard. Each row shows the transition of attention maps through the decoder layers. For example, in the role *Material*, the attention map gradually focuses on the paper in the image through the decoder layers. It shows that the semantic role queries can focus on the region related to them.

Visualization on Role Relations: In Figure 4, two images show different context for a verb *Swinging*. The role *Agent* and *Carrier* in Fig. 4(a) focus on the role *Place*, i.e., the forest (*Place*) is highly related to the monkey (*Agent*) and the vine (*Carrier*) given the verb *Swinging*. Meanwhile, the role *Place* in Fig. 4(b) focuses on the role *Carrier*, i.e., the golf club (*Carrier*) is highly related to the golf course (*Place*) given the verb *Swinging*. It shows that the relations among roles can be adaptively captured depending on the context of a given image.

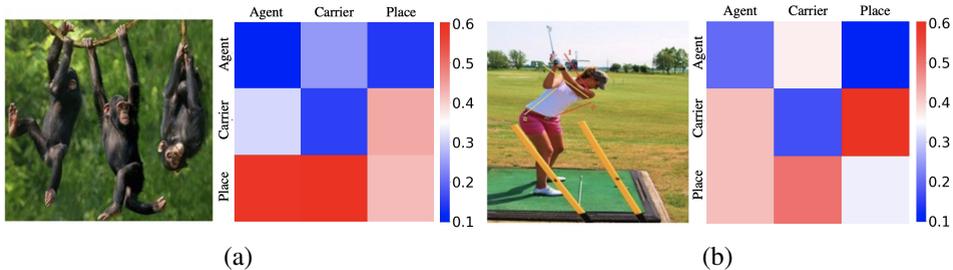


Figure 4: Visualization on Role Relations for two *Swinging* images. We visualize the attention scores between semantic role pairs computed in the MHSA block of the last decoder layer. Attention scores are represented as column-wise sum to 1.

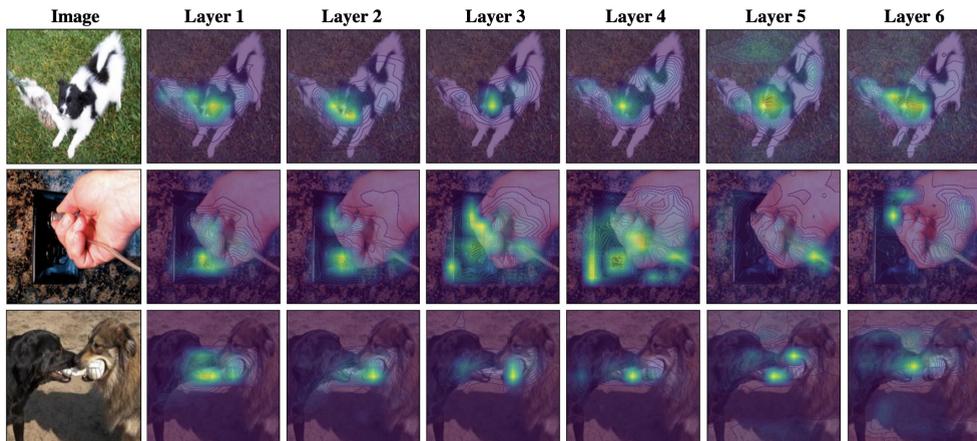


Figure 5: Verb Token Attention Map on Image Features for three *Tugging* images. Each row consists of an image and attention maps from the MHSA block in each encoder layer.

Verb Token Attention Map on Image Features: In Figure 5, the rightmost column shows the semantic regions where the verb token focuses on are similar. The verb token can capture the key feature (*e.g.*, tugged item) to infer the salient action. Each row shows the transition of attention maps through the encoder layers, *e.g.*, focusing on the tugged item gradually.

5 Conclusion

We propose the first Transformer architecture for GSR, which achieves the state-of-the-art accuracy on every evaluation metric. Our model, GSRTR, can capture high-level semantic feature, and flexibly deal with the complicated and image-dependent role relations. We perform extensive experiments and qualitatively illustrate the effectiveness of our method.

Acknowledgement: This work was supported by the NRF grant and the IITP grant funded by Ministry of Science and ICT, Korea (No.2019-0-01906 Artificial Intelligence Graduate School Program–POSTECH, NRF-2021R1A2C3012728–50%, IITP-2020-0-00842–50%).

References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–229, 2020.
- [2] Thilini Cooray, Ngai-Man Cheung, and Wei Lu. Attention-Based Context Aware Reasoning for Situation Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4736–4745, 2020.
- [3] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*, 2021.
- [4] Charles J. Fillmore, Christopher R. Johnson, and Miriam R.L. Petruck. Background to Framenet. *International Journal of Lexicography*, 16(3):235–250, 2003.
- [5] Yunchao Gong, Liwei Wang, Ruiqi Guo, and Svetlana Lazebnik. Multi-scale Orderless Pooling of Deep Convolutional Activation Features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 392–407, 2014.
- [6] He, Kaiming and Zhang, Xiangyu and Ren, Shaoqing and Sun, Jian. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [7] Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on Attention for Image Captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4634–4643, 2019.
- [8] Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. Situation Recognition with Graph Neural Network. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 4173–4182, 2017.
- [9] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated Graph Sequence Neural Networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- [10] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2125, 2017.
- [11] Hengyue Liu, Ning Yan, Masood Mortazavi, and Bir Bhanu. Fully Convolutional Scene Graph Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11546–11556, 2021.
- [12] Shaopeng Liu, Guohui Tian, and Yuan Xu. A novel scene classification model combining ResNet based transfer learning and data augmentation with a filter. *Neurocomputing*, 338:191–206, 2019.

- [13] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*, 2019.
- [14] Arun Mallya and Svetlana Lazebnik. Recurrent Models for Situation Recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 455–463, 2017.
- [15] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V. Le. Meta Pseudo Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11557–11568, 2021.
- [16] Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. Grounded Situation Recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 314–332, 2020.
- [17] RezaTofighi, Hamid and Tsoi, Nathan and Gwak, JunYoung and Sadeghian, Amir and Reid, Ian and Savarese, Silvio. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 658–666, 2019.
- [18] Marjaneh Safaei and Hassan Foroosh. Still Image Action Recognition by Predicting Spatial-Temporal Pixel Evolution. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 111–120, 2019. doi: 10.1109/WACV.2019.00019.
- [19] Mohammed Suhail and Leonid Sigal. Mixture-Kernel Graph Attention Network for Situation Recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10363–10372, 2019.
- [20] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [22] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and Tell: A Neural Image Caption Generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3156–3164, 2015.
- [23] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-End Video Instance Segmentation With Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8741–8750, 2021.
- [24] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tiejun Liu. On Layer Normalization in the Transformer Architecture. In *International Conference on Machine Learning (ICML)*, pages 10524–10533. PMLR, 2020.

- [25] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5410–5419, 2017.
- [26] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph R-CNN for Scene Graph Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 670–685, 2018.
- [27] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. Situation Recognition: Visual Semantic Role Labeling for Image Understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5534–5542, 2016.
- [28] Mark Yatskar, Vicente Ordonez, Luke Zettlemoyer, and Ali Farhadi. Commonly Uncommon: Semantic Sparsity in Situation Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7196–7205, 2017.
- [29] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image Captioning with Semantic Attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4651–4659, 2016.
- [30] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Temporal Query Networks for Fine-grained Video Understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4486–4496, 2021.
- [31] Zhichen Zhao, Huimin Ma, and Shaoqi You. Single Image Action Recognition using Semantic Body Part Actions. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3391–3399, 2017.
- [32] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning Deep Features for Scene Recognition using Places Database. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [33] Cheng Zou, Bohan Wang, Yue Hu, Junqi Liu, Qian Wu, Yu Zhao, Boxun Li, Chengguang Zhang, Chi Zhang, Yichen Wei, et al. End-to-End Human Object Interaction Detection with HOI Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11825–11834, 2021.