

# ERNIE-ViL: Knowledge Enhanced Vision-Language Representations through Scene Graphs

Fei Yu\*, Jiji Tang\*, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, Haifeng Wang

Baidu Inc., Beijing, China

{yufei07, tangjiji, yinweichong, sunyu02, tianhao, wu\_hua, wanghaifeng}@baidu.com

## Abstract

We propose a knowledge-enhanced approach, ERNIE-ViL, which incorporates structured knowledge obtained from scene graphs to learn joint representations of vision-language. ERNIE-ViL tries to build the detailed semantic connections (objects, attributes of objects and relationships between objects) across vision and language, which are essential to vision-language cross-modal tasks. Utilizing scene graphs of visual scenes, ERNIE-ViL constructs Scene Graph Prediction tasks, i.e., Object Prediction, Attribute Prediction and Relationship Prediction tasks in the pre-training phase. Specifically, these prediction tasks are implemented by predicting nodes of different types in the scene graph parsed from the sentence. Thus, ERNIE-ViL can learn the joint representations characterizing the alignments of the detailed semantics across vision and language. After pre-training on large scale image-text aligned datasets, we validate the effectiveness of ERNIE-ViL on 5 cross-modal downstream tasks. ERNIE-ViL achieves state-of-the-art performances on all these tasks and ranks the first place on the VCR leaderboard with an absolute improvement of 3.7%.

## Introduction

Motivated by pre-trained models like BERT (Devlin et al. 2018) and GPT (Radford et al. 2018) which have significantly improved the performance of many NLP tasks, researchers (Lu et al. 2019; Li et al. 2019a; Su et al. 2019; Li et al. 2019b; Chen et al. 2019) have noticed the importance of pre-training for vision-language tasks, e.g., Visual Question Answering (VQA) (Antol et al. 2015) and Visual Commonsense Reasoning (VCR) (Zellers et al. 2019).

Existing vision-language pre-training methods attempt to learn joint representations through visual grounding tasks on large image-text datasets, including Masked Language Modelling based on randomly-masked sub-words, Masked Region Prediction and Image-Text Matching at the image/text-level. However, based on randomly-masking and predicting the sub-words, current models did not distinguish common words and words describing the detailed semantics (Johnson et al. 2015), e.g., objects (“man”, “boat”), attributes of objects (“boat is white”), relationships between objects (“man

standing on boat”). These methods neglect the importance of constructing detailed semantic alignments across vision and language, therefore the trained models can not well represent fine-grained semantics required by some real scenes. As shown in Figure 1, the detailed semantics are essential to distinguish the listed scenes which mainly differ in objects, attributes and relationships. Hence, better joint vision-language representations should characterize detailed semantic alignments across the modalities.

Inspired by the knowledge masking strategy of ERNIE (Sun et al. 2019), which aims at learning more structured knowledge by masking phrases and named entities rather than individual sub-words, we propose ERNIE-ViL, that incorporates knowledge obtained from scene graphs (Johnson et al. 2015) to construct better representations for vision-language joint modelling. Through constructing Scene Graph Prediction tasks, ERNIE-ViL puts more emphasis on detailed semantic alignments across vision and language. Concretely, we implement these pre-training tasks by masking and predicting different types of nodes in the scene graph parsed from the sentence. By concentrating on understanding detailed semantic words rather than common words, these Scene Graph Prediction tasks force the model to extract object/attribute/relationship information from the visual modality, thus establish semantic connections between vision and language. Pre-training with the Scene Graph Prediction tasks, ERNIE-ViL learns the vision-language detailed semantic alignments.

We pre-train ERNIE-ViL on two large commonly-used image-text out-of-domain datasets, namely Conceptual Captions (Sharma et al. 2018) and SBU Captions (Ordonez, Kulkarni, and Berg 2011). To evaluate the performance of ERNIE-ViL, we conduct experiments on various vision-language tasks, (1) Visual Question Answering (VQA 2.0) (Antol et al. 2015), (2) Visual Commonsense Reasoning (VCR) (Zellers et al. 2019), (3) Region-to-Phrase Grounding (RefCOCO+) (Kazemzadeh et al. 2014), (4, 5) Image-text Retrieval / Text-image Retrieval (Flickr30K) (Young et al. 2014). On all these tasks, ERNIE-ViL obtains significant improvements compared to those models pretrained on the same datasets. Especially, on the Region-to-Phrase grounding task that relies more heavily on detailed semantic alignments, we achieve an improvement of 2.4% on both test-sets. To compare with the models pretrained on both out-

\*indicates equal contribution.

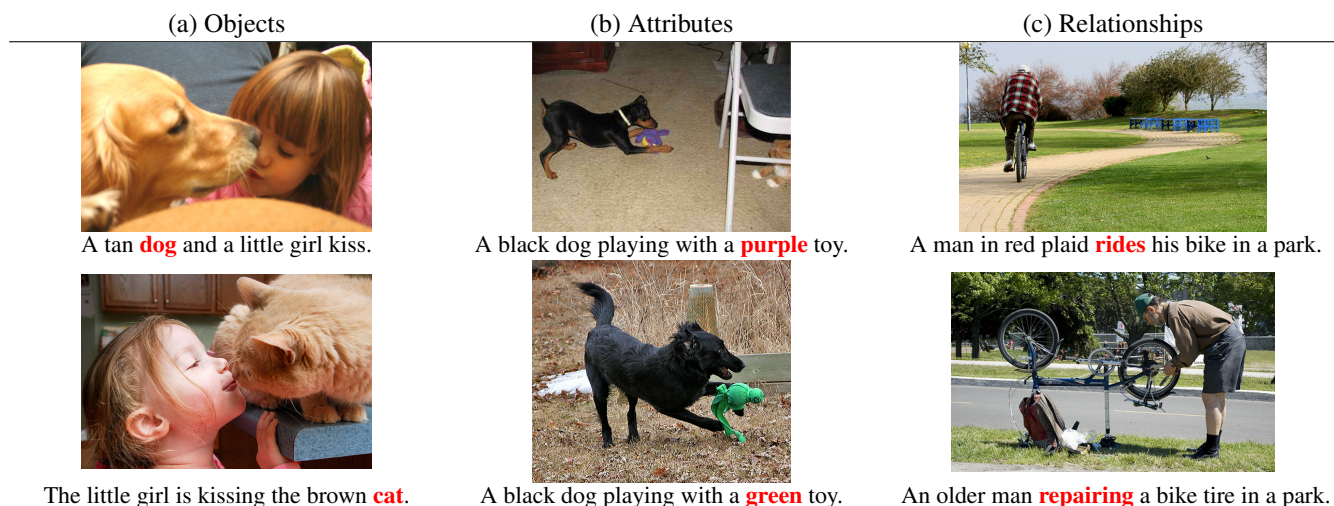


Figure 1: Similar scene pairs from the Flickr30K datasets (Young et al. 2014). It is the detailed semantics that determine the interpretation of the scenes, objects (dog, cat) in scene pair (a), attributes (purple, green) in scene pair (b) and relationships (rides, repairing) in scene pair (c).

of-domain and in-domain datasets, we continually pre-train ERNIE-ViL on MS-COCO (Lin et al. 2014) and Visual-Genome (Krishna et al. 2017) (in-domain datasets for downstream tasks). ERNIE-ViL achieves the state-of-the-art performances on all downstream tasks. Also ERNIE-ViL obtains the best single model performance and ranks the first place on the leaderboard with an absolute improvement of 3.7% on the Q→AR task compared to the state-of-the-art performance. And our code and pre-trained models are scheduled to be public.

Overall, the contributions of our method are three-folds:

- To the best of our knowledge, ERNIE-ViL is the first work that has introduced structured knowledge to enhance vision-language pre-training.
- ERNIE-ViL constructs Scene Graph Prediction tasks during the pre-training of vision-language joint representations, putting more emphasis on the cross-modal detailed semantics alignments.
- ERNIE-ViL achieves state-of-the-art performances on 5 downstream cross-modal tasks and ranks the first place on the VCR leaderboard.

## Related Works

### Cross-modal Pre-training

Inspired by text pre-training models (Devlin et al. 2018), many cross-modal pre-training models for vision-language have been proposed. These researchers put their efforts mainly on three aspects, which are model architecture, pre-training tasks and pre-training data.

- **Model Architecture** Current works are based on different variables of Transformers (Vaswani et al. 2017). Most of them (Li et al. 2019a; Su et al. 2019; Zhou et al. 2019;

Li et al. 2019b; Huang et al. 2020) used a uniform cross-modal Transformer modelling both image and text representations, while the others like ViLBERT (Lu et al. 2019) and LXMERT (Tan and Bansal 2019) were based on two-stream cross-modal Transformers, which brings more specific representations for images and texts.

- **Pre-training Tasks** Inspired by the pre-training tasks in text models, Masked Language Model and similar Masked Region Prediction tasks (Lu et al. 2019) are utilized in cross-modal pre-training. And similar to Next-Sentence Prediction, Image-Text Matching (Lu et al. 2019; Su et al. 2019; Chen et al. 2019) task is also widely used. However, based on randomly masking and predicting sub-words, these methods did not distinguish the common words and words describing the detailed semantics. Hence, the cross-modal fine-grained semantic alignments cannot be well characterized in those learned joint representations.
- **Pre-training Data** Unlike text pre-training models that can leverage tremendous natural language data, vision-language tasks require high-quality aligned image-text data that are hard to obtain. Conceptual Captions (Sharma et al. 2018) and SBU Captions (Ordonez, Kulkarni, and Berg 2011) are two widely-used datasets for image-text pre-training, with 3.0M and 1.0M image-description pairs respectively. These two datasets are out-of-domain for vision-language downstream tasks, while some existing works (Chen et al. 2019; Huang et al. 2020) incorporate in-domain datasets, such as MS-COCO and Visual-Genome, that are highly correlated with downstream tasks.

### Scene Graph

Scene graphs contain structured knowledge of visual scenes, including the present objects, attributes of objects, and relationships between objects. As a beneficial prior knowledge

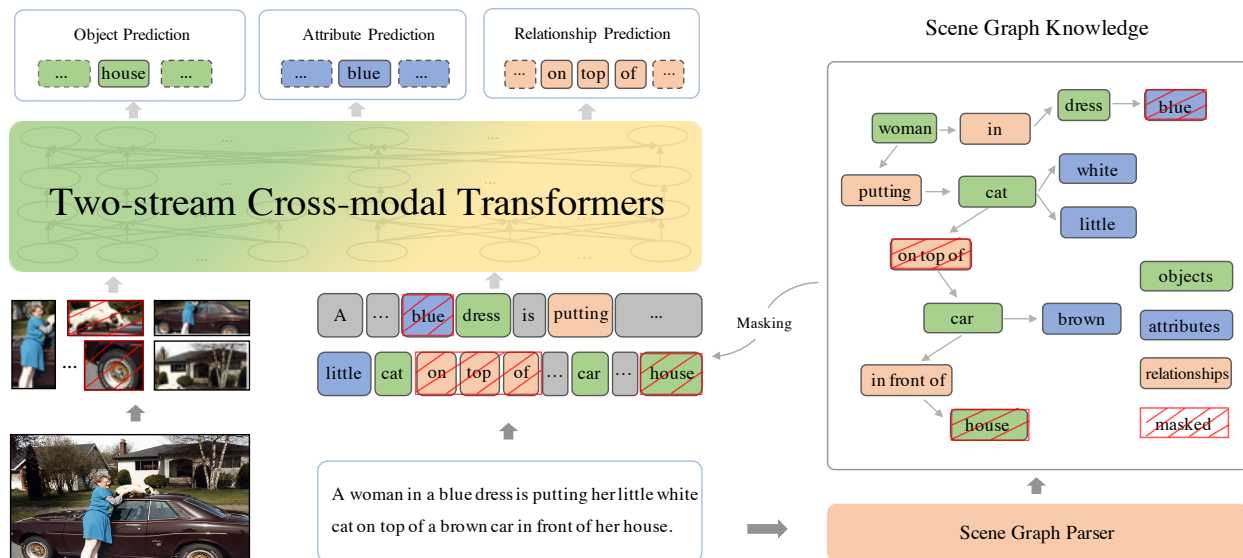


Figure 2: Illustration of Scene Graph Prediction tasks for ERNIE-ViL. Given detected regions of the image and token sequence of the text, ERNIE-ViL uses a two-stream cross-modal Transformers network to model the joint vision-language representations. Based on the scene graph parsed from the text using Scene Graph Parser, we construct Object Prediction, Attribute Prediction and Relationship Prediction tasks to learn cross-modal detailed semantics alignments.

describing the detailed semantics of images and captions, scene graphs have led to many state-of-the-art models in image captioning (Yang et al. 2019), image retrieval (Wu et al. 2019), VQA (Zhang, Chao, and Xuan 2019) and image generation (Johnson, Gupta, and Fei-Fei 2018).

### Approach

In this section, we first introduce the architecture of ERNIE-ViL. Then we illustrate our newly-proposed Scene Graph Prediction tasks. Finally, pre-training with Scene Graph Prediction tasks in ERNIE-ViL is introduced.

#### Model Architecture

The vision-language model aims at learning the joint representations that integrates information of both modalities and the alignments across the modalities. The inputs of ERNIE-ViL are a sentence and an image. Given a sequence of words and an image, we introduce the methods to embed the inputs to the feature space and the vision-language encoder.

**Sentence Embedding** We adopt the similar word pre-processing method as BERT. The input sentence is tokenized into sub-word tokens using WordPiece approach. Special tokens such as  $[CLS]$  and  $[SEP]$  are also added to the tokenized text sequence to form the text sequence as  $\{[CLS], w_1, \dots, w_T, [SEP]\}$ . The final embedding for each sub-word token is generated by combining its original word embedding, segment embedding and sequence position embedding.

**Image Embedding** For the image, we first use a pre-trained object detector to detect the salient image regions

from the image. The pooling features before multi-class classification layer are utilized as the region features. We also encode the location features for each region via a 5-dimensional vector  $(\frac{x_1}{W}, \frac{y_1}{H}, \frac{x_2}{W}, \frac{y_2}{H}, \frac{(y_2-y_1)(x_2-x_1)}{WH})$  for the region position and the fraction of image area covered, where  $(x_1, y_1)$  and  $(x_2, y_2)$  denote the coordinates of top-left and bottom-right corner while  $W$  and  $H$  are the width and height of the input image. The location vectors are projected to form the location features, which are then summed with the region visual features. We also add a special feature  $[IMG]$  that denotes the representation of the entire image (i.e. mean-pooled visual features with a spatial encoding corresponding to the entire image) to form the final region sequence  $\{[IMG], v_1, \dots, v_I\}$ .

**Vision-Language Encoder** Given the embedding of image regions and the words for the sentence  $\{[IMG], v_1, \dots, v_I, [CLS], w_1, \dots, w_T, [SEP]\}$ , we use two-stream cross-modal Transformers to joint model the intra-modal and inter-modal representations. Similar to ViLBERT (Lu et al. 2019), ERNIE-ViL consists of two parallel Transformer encoders for image and text segments, which are cross-attended with cross-modal Transformer blocks. The model outputs embeddings for each input of both the image and text. We take  $h_{[IMG]}$  and  $h_{[CLS]}$  as the holistic image and text representations.

#### Scene Graph Prediction

Detailed semantics, includes objects, attributes of objects, and relationships between objects, are essential to the understanding of visual scenes (Johnson et al. 2015). As the scene shown in Figure 2, detailed semantics describes the vi-

sual scene from different aspects. The objects, such as “cat”, “car”, “woman” are the fundamental elements in the scene. And associated attributes, such as “little”, “brown”, “blue” characterize shape and color of objects. Relationships such as “on top of”, “putting” represent the spatial connections and actions between objects. Therefore detailed semantics are crucial in accurately understanding visual scenes. Since the goal of vision-language joint representations is to engrave the semantic connections across modalities, detailed semantic alignments are significantly important in cross-modal learning.

Scene graphs encode various fine-grained semantic information. Utilizing structured knowledge obtained from scene graphs, ERNIE-ViL learns the cross-modal detailed semantic alignments. As shown in Figure 2, according to the scene graph parsed from the text, we construct the corresponding Scene Graph Prediction tasks, including Object Prediction task, Attribute Prediction task, and Relationship Prediction task. These tasks force ERNIE-ViL to model the correlations of detailed semantics across modalities. For example, as the relationship words “on top of” is masked, based on the language context, the model may predict that the missing word is “under” or “into”. These words are grammatically fluent in the sentence, but are inconsistent with the scene “the cat is on top of the car”. Through training the Relationship Prediction task, the model obtains the spatial relation of the corresponding objects (“car”, “cat”) from the image, thus can accurately predict that the missing word is “on top of”. Through constructing Scene Graph Prediction tasks, ERNIE-ViL learns cross-modal detailed semantic alignments.

**Scene graph parsing** Given the text sentence  $\mathbf{w}$ , we parse it into a scene graph (Johnson et al. 2015), which denotes as  $G(\mathbf{w}) = \langle O(\mathbf{w}), E(\mathbf{w}), K(\mathbf{w}) \rangle$ , where  $O(\mathbf{w})$  is the set of objects mentioned in  $\mathbf{w}$ ,  $E(\mathbf{w}) \subseteq O(\mathbf{w}) \times R(\mathbf{w}) \times O(\mathbf{w})$  is the set of hyper-edges representing relationship triplets, and  $R(\mathbf{w})$  is the set of relationship nodes between object nodes.  $K(\mathbf{w}) \subseteq O(\mathbf{w}) \times A(\mathbf{w})$  is the set of attribute pairs, where  $A(\mathbf{w})$  is the set of attribute nodes associated with object nodes. Scene graphs describe the objects in more details with various associated attributes and relationships between objects. Thus integrating the knowledge of scene graphs can benefit learning more fine-grained joint representations for the vision-language. In this paper, the Scene Graph Parser provided by Anderson (Anderson et al. 2016) is adopted to parse texts to scene graphs. For a more intuitive understanding, we illustrate a specific case for the parsed scene graph from the text in Table 1.

**Object Prediction** Objects are the dominant elements of visual scenes, thus playing an important role in constructing the representations of semantic information. Predicting the objects forces the model to build the vision-language connections at object level.

Firstly, for all the object nodes in the scene graph, we randomly select 30% of them to mask. And for each selected object node  $O(\mathbf{w})$ , we replace it with the special token  $[MASK]$  in probability of 80%, another random token in

sentence: $\mathbf{w}$	A woman in blue dress is putting her little white cat on top of a brown car in front of her house.
objects: $O(\mathbf{w})$	dress, woman, cat, car, house
relationships: $R(\mathbf{w})$	in, putting, on-top-of, in-front-of
attributes: $A(\mathbf{w})$	blue, white, little, brown

Table 1: The scene graph parsed from the caption of the visual scene. For simplicity, we only list all the nodes leaving out the connections between them.

probability of 10%, and keep it in probability of 10%. Note that the objects actually correspond to the sub-sequences of text in the sentence, therefore the object masking are implemented by masking the corresponding sub-sequences in the text.

For Object Prediction, ERNIE-ViL recover these masked object tokens, denoted as  $\mathbf{w}_{o_i}$ , based on their surrounding words  $\mathbf{w}$  and all image regions  $\mathbf{v}$ , by minimizing the negative log-likelihood:

$$\mathcal{L}_{obj}(\theta) = -E_{(\mathbf{w}, \mathbf{v}) \sim D} \log(P(\mathbf{w}_{o_i} | \mathbf{w}_{\setminus \mathbf{w}_{o_i}}, \mathbf{v})) \quad (1)$$

**Attribute Prediction** Attributes characterize the specific information of the visual objects, such as color or shape of the objects, therefore representing the detailed information in the visual scenes in more fine-grained level.

Similarly, we randomly select 30% of the attribute pairs in the scene graph, and the mask strategy here is the same as that in Object Prediction. Since the attribute nodes in the scene graph are attached to objects, we keep the associated object while masking out the attribute node  $A(\mathbf{w})$  in each selected  $K(\mathbf{w}) \subseteq O(\mathbf{w}) \times A(\mathbf{w})$ .

Given object words  $w_{o_i}$  in attribute pair  $\langle w_{o_i}, w_{a_i} \rangle$ , Attribute Prediction is to recover the masked tokens  $w_{a_i}$  of attribute pairs. Based on the object tokens  $w_{o_i}$ , other surrounding words  $\mathbf{w}$  and all image regions  $\mathbf{v}$ , Attribute Prediction minimizes the negative log-likelihood:

$$\mathcal{L}_{attr}(\theta) = -E_{(\mathbf{w}, \mathbf{v}) \sim D} \log(P(\mathbf{w}_{a_i} | \mathbf{w}_{o_i}, \mathbf{w}_{\setminus \mathbf{w}_{a_i}}, \mathbf{v})) \quad (2)$$

**Relationship Prediction** Relationships describe the actions (semantic) or relative position (geometry) between the objects of the visual scenes, which contributes to distinguish scenes with same objects but different relationships.

Thus, ERNIE-ViL constructs the Relationship Prediction task to learn cross-modal relationships connections. When performing the mask strategy of selected relationship triplets  $E(\mathbf{w}) \subseteq O(\mathbf{w}) \times R(\mathbf{w}) \times O(\mathbf{w})$ , we keep the objects and mask out the relationship node  $R(\mathbf{w})$ . Specifically, given object tokens  $\mathbf{w}_{o_{i1}}, \mathbf{w}_{o_{i2}}$  in relationship triplet  $\langle \mathbf{w}_{o_{i1}}, \mathbf{w}_{r_i}, \mathbf{w}_{o_{i2}} \rangle$ , this task recovers the masked relationship tokens, predicting the probability for each masked relation tokens  $\mathbf{w}_{r_i}$ . Thus the context for the prediction is the given object tokens  $\mathbf{w}_{o_{i1}}, \mathbf{w}_{o_{i2}}$ , other surrounding words from the text and all image regions  $\mathbf{v}$ . The loss for this task is:

$$\mathcal{L}_{rel}(\theta) = -E_{(\mathbf{w}, \mathbf{v}) \sim D} \log(P(\mathbf{w}_{r_i} | \mathbf{w}_{o_{i1}}, \mathbf{w}_{o_{i2}}, \mathbf{w}_{\setminus \mathbf{w}_{r_i}}, \mathbf{v})) \quad (3)$$

## Pre-training with Scene Graph Prediction

Similar to ViLBER (Lu et al. 2019), ERNIE-ViL also adopts Masked Language Modelling (MLM) to capture the syntactic and lexical information in the text. Moreover, Masked Region Prediction and Image-text Matching are utilized for visual modality and cross-modality respectively. The losses for all these pre-training tasks are summed.

## Experiments

### Training ERNIE-ViL

**Pre-training Data** We use the Conceptual Captions (CC) dataset (Sharma et al. 2018) and SBU Captions (SBU) dataset (Ordonez, Kulkarni, and Berg 2011) as pre-training data. CC is a collection of 3.3 million image-caption pairs automatically scraped from alt-text enabled web images and SBU is a similar vision-language dataset which has 1.0 million image-caption pairs. Since some links have become broken, only about 3.0 million pairs for CC dataset and 0.8 million pairs for SBU dataset are available and utilized in our experiments. Note that CC and SBU are image-caption pairs automatically collected from the web and have no intersections with the down-stream task datasets, thus act as out-of-domain datasets for training vision-language models.

**Implementation Details** For each image-text pair in the training, the pre-processing is performed as follows. For the image, we adopt Faster R-CNN (Anderson et al. 2018) to select salient image regions and extract region features. Specifically, regions with class detection probability exceeds a confidence threshold of 0.2 are selected and 10 to 36 boxes are kept. And for each kept region, the mean-pooled convolutional representation is used as the region feature. For the text, we parse the scene graph from the sentence using the Scene Graph Parser and adopt WordPieces to tokenize the sentence similar to BERT.

For the masking strategies, we randomly mask 15% of tokens, 30% of scene graph nodes, and 15% of image regions. For the Image-text Matching task, we randomly select a image for each text to form the negative image-text pair. Note that only items in the positive pairs will be considered for token and region prediction tasks.

We train ERNIE-ViL on two scale settings: ERNIE-ViL-base and ERNIE-ViL-large, which mainly differ in model depth of the text stream. The detailed settings of text and visual streams are shown in Table 2. And similar to ViLBER (Lu et al. 2019), cross-transformers are used to co-attend the two streams. We initialize the text stream parameters from ERNIE 2.0 (Sun et al. 2019), and implement ERNIE-ViL via PaddlePaddle. After then, ERNIE-ViL is pre-trained on a total batch size of 512 for 700k steps on 8 V100 GPUs, using adam optimizer with initial learning rates of  $1e-4$  and Noam (Vaswani et al. 2017) as learning rate decay schedule.

### Downstream Tasks

**Visual Commonsense Reasoning (VCR)** The Visual Commonsense Reasoning (VCR) (Zellers et al. 2019) task contains two sub-tasks: visual question answering ( $Q \rightarrow A$ )

	Base				Large			
	$L$	$H$	$A$	$F$	$L$	$H$	$A$	$F$
Text	12	768	12	3072	24	1024	16	4096
Visual	6	1024	8	1024	6	1024	16	4096

Table 2: Settings for ERNIE-ViL model.  $L$ : number of layers,  $H$ : hidden size,  $A$ : number of self-attention heads,  $F$ : feed-forward/filter size.

and answer justification ( $QA \rightarrow R$ ), which are both multiple choice problems. The holistic setting ( $Q \rightarrow AR$ ) requires both the chosen answer and chosen rationale to be correct. In visual question answering ( $Q \rightarrow A$ ) task, we concatenate the question and each candidate answer for the language modality. We take dot product of final hidden state  $h_{[CLS]}$  and  $h_{[IMG]}$  to predict matching score with an additional FC layer. For the answer justification ( $QA \rightarrow R$ ) task, we concatenate the question, the answer and each candidate rationale as the input of the text stream. Similar with UNITER (Chen et al. 2019), a second-stage pre-training is adopted on VCR dataset. And then we fine-tune the model over 6 epochs with a batch size of 64 and adopt Adam optimizer with initial learning rate of  $1e-4$ .

**Visual Question Answering (VQA)** The VQA task requires answering natural language questions according to images. VQA 2.0 dataset (Antol et al. 2015) contains 204k images and 1.1M questions about these images. Also additional question-answer pairs from Visual Genome are used for data augmentation as in UNITER (Chen et al. 2019). We treat VQA as a multi-label classification task – assigning a soft target score to each answer based on its relevancy to the 10 human answer responses. We take dot product of final hidden state  $h_{[CLS]}$  and  $h_{[IMG]}$  to map this representation into 3,129 possible answers with an additional two-layer MLP. Fine-tuning of VQA model is performed over 12 epochs on batch size of 256 and using Adam optimizer with initial learning rate of  $1e-4$ .

**Grounding Referring Expressions** The referring expression task is to localize an image region given a natural language reference. We evaluate the task on RefCOCO+ dataset (Kazemzadeh et al. 2014). Bounding box proposals provided by Mattnet (Yu et al. 2018) are utilized. The representation for each region is denoted by its final hidden state  $h_{v_i}$  with an additional FC layer. Each region  $i$  is labelled as positive only when the IoU between it and the ground truth box is over 0.5. We fine-tune the model over 20 epochs with a batch size of 256 and adopt Adam optimizer with initial learning rate of  $1e-4$ .

**Image Retrieval & Text Retrieval** Caption-based image retrieval is a task of identifying an image from a pool based on a caption describing its content. Flickr30K (Young et al. 2014) contains 31,000 images and 5 captions for each image. Adopting the same split in ViLBER (Lu et al. 2019), we use each of 1,000 images for validation and for testing and the rest for training. We take dot product of final hidden state of  $h_{[CLS]}$  and  $h_{[IMG]}$  to predict matching score

Domains	Models	VCR			RefCOCO+		
		Q→A	QA→R	Q→AR	val	testA	testB
Out-of-domain	UNITER-base	-	-	-	72.78	-	-
	Unicoder-VL-base	72.6(73.4)	74.5(74.4)	54.4(54.9)	-	-	-
	ViLBERT-base	72.42(73.3)	74.47(74.6)	54.04(54.8)	72.34	78.52	62.61
	VLBERT-base	73.8(-)	74.4(-)	55.2(-)	71.60	77.72	60.99
	ERNIE-ViL-base	76.37(77.0)	79.65(80.3)	61.24(62.1)	74.02	80.33	<b>64.74</b>
	VLBERT-Large	75.5(75.8)	77.9(78.4)	58.9(59.7)	72.59	78.57	62.30
Out-of-domain + in-domain	ERNIE-ViL-Large	<b>78.52(79.2)</b>	<b>83.37(83.5)</b>	<b>65.81(66.3)</b>	<b>74.24</b>	<b>80.97</b>	64.70
	UNITER-large	77.22(77.3)	80.49(80.8)	62.59(62.8)	75.90	81.45	66.70
	VILLA-large	78.45(78.9)	82.57(82.8)	65.18(65.7)	<b>76.17</b>	81.54	66.84
	ERNIE-ViL-large	<b>78.98(-)</b>	<b>83.70(-)</b>	<b>66.44(-)</b>	75.89	<b>82.37</b>	<b>66.91</b>

Domains	Models	VQA		IR-Flickr30K			TR-Flickr30K		
		test-dev	test-std	R@1	R@5	R@10	R@1	R@5	R@10
Out-of-domain	UNITER-base	71.56	-	-	-	-	-	-	-
	Unicoder-VL-base	-	-	71.50	90.90	94.90	86.20	96.30	99.00
	VLBERT-base	71.16	-	-	-	-	-	-	-
	ViLBERT-base	70.55	70.92	58.20	84.90	91.52	-	-	-
	ERNIE-ViL-base	73.18	73.36	74.44	92.72	95.94	86.70	97.80	99.00
	VLBERT-large	71.79	72.22	-	-	-	-	-	-
Out-of-domain + in-domain	ERNIE-ViL-large	<b>73.78</b>	<b>73.96</b>	<b>75.10</b>	<b>93.42</b>	<b>96.26</b>	<b>88.70</b>	<b>97.30</b>	<b>99.10</b>
	UNITER-large	73.82	74.02	75.56	94.08	96.76	87.30	98.00	<b>99.20</b>
	OSCAR-large	73.61	73.82	-	-	-	-	-	-
	VILLA-large	74.69	74.87	76.26	<b>94.24</b>	<b>96.84</b>	87.90	97.50	98.80
	ERNIE-ViL-large	<b>74.95</b>	<b>75.10</b>	<b>76.66</b>	94.16	96.76	<b>89.20</b>	<b>98.50</b>	<b>99.20</b>

Table 3: Results of downstream vision-language tasks for ERNIE-ViL model, compared with previous state-of-the-art pre-trained models. IR: Image Retrieval. TR: Text Retrieval. For VCR task which has private test set, we only report the test results (in parentheses) for ERNIE-ViL models pre-trained on out-of-domain datasets.

$s(\mathbf{w}, \mathbf{v})$  for each image-text pair with an additional FC layer. We utilize circle loss (Sun et al. 2020) with 20 random negative samples for each image-text pair. We trained 40 epochs using Adam optimizer with a initial learning rate  $1e-5$ .

## Results

We compare ERNIE-ViL against other cross-modal pre-training models and the results are illustrated in Table 3.

Among the methods pre-trained on the same out-of-domain datasets (CC and SBU), ERNIE-ViL obtains the best performances on all 5 downstream tasks. For the visual reasoning tasks, ERNIE-ViL-large achieves a significant improvement of 6.60% on VCR (Q→AR) task and 1.74% on VQA (test-std) task compared with VLBERT-large. On visual grounding task, ERNIE-ViL-large obtains an improvement of 2.40% for both testA split and testB split on RefCOCO+ task compared to VLBERT-large. On the cross-modal retrieval tasks, where no large models pre-trained on out-of-domain datasets has released results, ERNIE-ViL-base achieves an improvement of 2.94% on R@1 for image retrieval and 0.50% on R@1 for text retrieval compared with Unicoder-VL-base.

For further comparison with those models pretrained with both out-of-domain and in-domain datasets, we pre-train

ERINE-ViL with all these datasets. As illustrated in Table 3, ERINE-ViL-large achieves state-of-the-art performances on these tasks compared to existing works, e.g., UNITER, OSCAR(Li et al. 2020) and VILLA(Gan et al. 2020).

## Analysis

**Effectiveness of Scene Graph Prediction tasks** To verify the effectiveness of Scene Graph Prediction (SGP) tasks, we first conduct experiments with ERNIE-ViL-base settings based on the text parameters initialized from BERT. As illustrated in Table 4, pre-training with SGP tasks in ERNIE-ViL brings significant improvements across all downstream tasks. Especially on Grounding Referring Expressions and Retrieval tasks, those require understanding detailed semantics alignments, SGP tasks make an improvement of 0.69% accuracy on RefCOCO+ and 2.22% of R@1 for image retrieval on Flickr30K.

Note that text parameter initialized from ERNIE 2.0 can lead to further improvements on all tasks and a relatively large improvement on VCR task. We consider that through continually learning on various pre-training tasks, ERNIE 2.0 learned more common sense knowledge which benefits the VCR task.

Overall, the SGP tasks significantly contribute to the



initialized text stream parameters	pre-training tasks	VCR	VQA	RefCOCO+	IR	TR
		Q→AR (dev)	dev	val	R@1 (dev)	R@1 (dev)
BERT	w/o <b>SGP</b>	59.06	72.38	72.81	70.74	85.00
BERT	w/ <b>SGP</b>	59.92	73.04	73.50	72.96	87.40
ERNIE-2.0	w/ <b>SGP</b>	61.24	73.18	74.02	73.58	87.80

Table 4: Results of downstream vision-language tasks for ERNIE-ViL pre-training with/without Scene Graph Prediction (**SGP**) tasks, and using different text stream parameters initialization. IR & TR: image retrieval & text retrieval on Flickr30K.


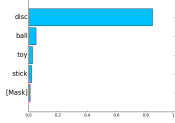
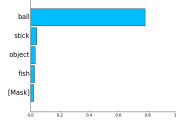

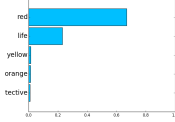
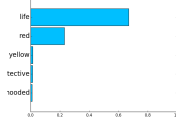

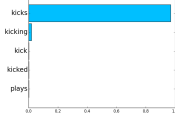
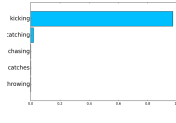
	Image	Text	with <b>SGP</b> task	without <b>SGP</b> task
1		a black dog about to catch a flying <b>disc</b> .		
2		two men wearing <b>red</b> jackets are looking out over some water and one man has yellow earphones on his ears .		
3		a little boy in a green shirt <b>kicks</b> a ball		

Table 5: Examples of cloze test predictions for ERNIE-ViL pre-training with and without **SGP** tasks. Masked token are colored in bold and red. The probabilities of the top 5 predictions, denoted as the light purple bars, are listed in the right columns.

Nodes	without <b>SGP</b> tasks		with <b>SGP</b> tasks	
	ACC@1	ACC@5	ACC@1	ACC@5
objects	57.14	79.22	<b>58.34</b>	<b>80.80</b>
attributes	44.32	67.58	<b>46.16</b>	<b>70.30</b>
relationships	47.57	68.10	<b>50.65</b>	<b>71.54</b>
overall	49.75	71.75	<b>51.75</b>	<b>74.31</b>

Table 6: Cloze test results for ERNIE-ViL. An improvement of 2.0% on overall ACC@1 between models with/without **SGP** tasks.

state-of-the-art results of ERNIE-ViL.

**Cloze Test** To get a more intuitively understanding of the improvements brought by SGP tasks, we conduct the language cloze test conditioned on the visual modality. In the cloze test, language tokens represent detailed semantics (objects, attributes and relationships) are masked from the text and the model is required to infer them with the context from both the text and the image. To construct the dataset, we sampled 15,000 image-text pairs from Flickr30K dataset and 5,000 objects, attributes and relationships tokens each are selected. For the prediction, the top one accuracy (ACC@1) and top five accuracy (ACC@5) are adopted as the evaluation metric. The comparison of prediction results between two models, which are pre-trained models with **SGP** task

and without **SGP** task, are illustrated in Table 6. The text-stream parameters of both models are initialized from BERT. An absolute improvement of 1.20% for objects, 3.08% for relationships and 1.84% for attributes on ACC@1 demonstrates that ERNIE-ViL pre-trained with **SGP** tasks learns better cross-modal detailed semantics alignments.

Moreover, we illustrate some cases in Table 5, and the top 5 possible predictions are shown in the right columns. As in case 1-2, model pre-trained without **SGP** tasks cannot make the right predictions as it didn't learn accurate alignments of detailed semantics, without distinguishing common words and detailed semantics words while pre-training. While in case 3, the model can predict the reasonable tokens but with lower confidence compared with model pre-trained with **SGP** tasks.

## Conclusion

We proposed ERNIE-ViL to learn the joint representations of vision and language. In addition to conventional MLM for cross-modal pre-training, we introduce Scene graph Prediction tasks to characterize the cross-modal detailed semantic alignments. Experiment results on various downstream tasks demonstrate the improvements of incorporating structured knowledge obtained from scene graphs during cross-modal pre-training. For future work, scene graphs extracted from images could also be incorporated into cross-modal pre-training. Moreover, Graph Neural Networks that integrate more structured knowledge could be considered as well.

## References

- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, 382–398. Springer.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6077–6086.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Chen, Y.-C.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2019. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Gan, Z.; Chen, Y.-C.; Li, L.; Zhu, C.; Cheng, Y.; and Liu, J. 2020. Large-Scale Adversarial Training for Vision-and-Language Representation Learning. *arXiv preprint arXiv:2006.06195*.
- Huang, Z.; Zeng, Z.; Liu, B.; Fu, D.; and Fu, J. 2020. PixelBERT: Aligning Image Pixels with Text by Deep Multi-Modal Transformers. *arXiv preprint arXiv:2004.00849*.
- Johnson, J.; Gupta, A.; and Fei-Fei, L. 2018. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1219–1228.
- Johnson, J.; Krishna, R.; Stark, M.; Li, L.-J.; Shamma, D.; Bernstein, M.; and Fei-Fei, L. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3668–3678.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 787–798.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* 123(1): 32–73.
- Li, G.; Duan, N.; Fang, Y.; Jiang, D.; and Zhou, M. 2019a. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. *arXiv preprint arXiv:1908.06066*.
- Li, L. H.; Yatskar, M.; Yin, D.; Hsieh, C.-J.; and Chang, K.-W. 2019b. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Li, X.; Yin, X.; Li, C.; Hu, X.; Zhang, P.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; et al. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. *arXiv preprint arXiv:2004.06165*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, 13–23.
- Ordonez, V.; Kulkarni, G.; and Berg, T. L. 2011. Im2text: Describing images using 1 million captioned photographs. In *Advances in neural information processing systems*, 1143–1151.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. In URL [https://s3-us-west-2.amazonaws.com/openaiassets/research-covers/language\\_unsupervised/language\\_understanding\\_paper.pdf](https://s3-us-west-2.amazonaws.com/openaiassets/research-covers/language_unsupervised/language_understanding_paper.pdf).
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2556–2565.
- Su, W.; Zhu, X.; Cao, Y.; Li, B.; Lu, L.; Wei, F.; and Dai, J. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*.
- Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle loss: A unified perspective of pair similarity optimization. *arXiv preprint arXiv:2002.10857*.
- Sun, Y.; Wang, S.; Li, Y.; Feng, S.; Tian, H.; Wu, H.; and Wang, H. 2019. Ernie 2.0: A continual pre-training framework for language understanding. *arXiv preprint arXiv:1907.12412*.
- Tan, H.; and Bansal, M. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wu, H.; Mao, J.; Zhang, Y.; Jiang, Y.; Li, L.; Sun, W.; and Ma, W.-Y. 2019. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6609–6618.
- Yang, X.; Tang, K.; Zhang, H.; and Cai, J. 2019. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10685–10694.



Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics 2*: 67–78.

Yu, L.; Lin, Z.; Shen, X.; Yang, J.; Lu, X.; Bansal, M.; and Berg, T. L. 2018. Mattnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1307–1315.

Zellers, R.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6720–6731.

Zhang, C.; Chao, W.-L.; and Xuan, D. 2019. An empirical study on leveraging scene graphs for visual question answering. *arXiv preprint arXiv:1907.12133* .

Zhou, L.; Palangi, H.; Zhang, L.; Hu, H.; Corso, J. J.; and Gao, J. 2019. Unified vision-language pre-training for image captioning and vqa. *arXiv preprint arXiv:1909.11059* .