

PEVL: Position-enhanced Pre-training and Prompt Tuning for Vision-language Models

Yuan Yao^{1*}, Qianyu Chen^{1*}, Ao Zhang², Wei Ji²
Zhiyuan Liu^{1†}, Tat-Seng Chua², Maosong Sun^{1†}

¹Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China
Beijing National Research Center for Information Science and Technology
Innovation Center of Tsinghua University, Shanghai, China

²Sea-NExT Joint Lab, Singapore
School of Computing, National University of Singapore, Singapore
yaoyuanthu@163.com socqyc@gmail.com

Abstract

Vision-language pre-training (VLP) has shown impressive performance on a wide range of cross-modal tasks, where VLP models without reliance on object detectors are becoming the mainstream due to their superior computation efficiency and competitive performance. However, the removal of object detectors also deprives the capability of VLP models in explicit object modeling, which is essential to various position-sensitive vision-language (VL) tasks, such as referring expression comprehension and visual commonsense reasoning. To address the challenge, we introduce PEVL that enhances the pre-training and prompt tuning of VLP models with explicit object position modeling. Specifically, PEVL reformulates discretized object positions and language in a unified language modeling framework, which facilitates explicit VL alignment during pre-training, and also enables flexible prompt tuning for various downstream tasks. We show that PEVL enables state-of-the-art performance of detector-free VLP models on position-sensitive tasks such as referring expression comprehension and phrase grounding, and also improves the performance on position-insensitive tasks with grounded inputs. We make the data and code for this paper publicly available at <https://github.com/thunlp/PEVL>.

1 Introduction

Recent progress on self-supervised learning has led to powerful vision-language pre-training (VLP) models that achieve state-of-the-art performance on a wide range of cross-modal tasks (Lu et al., 2019;

Li et al., 2020c; Radford et al., 2021; Zhang et al., 2021; Kamath et al., 2021). Typically, VLP models are first pre-trained on large-scale image-text data to learn universal cross-modal representations, and then fine-tuned to adapt to downstream tasks (Bommasani et al., 2021). While most traditional VLP models heavily rely on external object detectors to obtain the visual inputs (Lu et al., 2019; Su et al., 2020; Li et al., 2020c; Zhang et al., 2021), recently there is a growing interest in VLP models that remove the reliance on object detectors due to their superior computation efficiency and competitive performance (Li et al., 2021; Kim et al., 2021; Radford et al., 2021; Kamath et al., 2021).

However, the removal of object detectors also deprives the capability of VLP models in explicit object modeling. The drawback hinders successful handling of vision-language (VL) tasks which are inherently object-centric, where deep understanding of objects and their interactions plays an essential role (Antol et al., 2015; Plummer et al., 2015; Krishna et al., 2017; Hudson and Manning, 2019). Therefore, it is typically difficult for detector-free¹ VLP models to handle various position-sensitive tasks (i.e., tasks that demand explicit object positions as input or output), such as visual commonsense reasoning (Zellers et al., 2019), visual relation detection (Krishna et al., 2017), referring expression comprehension (Yu et al., 2016) and phrase grounding (Plummer et al., 2015), which greatly undermines their generality and practicality as foundation models (Bommasani et al., 2021). For tasks that do not require explicit object modeling, such as visual question answering (Antol

* indicates equal contribution

† Corresponding authors: Z.Liu (liuzy@tsinghua.edu.cn), M.Sun (sms@tsinghua.edu.cn)

¹Note that by detector-free, we mean that no external object detector tools are required. However, object annotations may still be needed during pre-training and tuning.

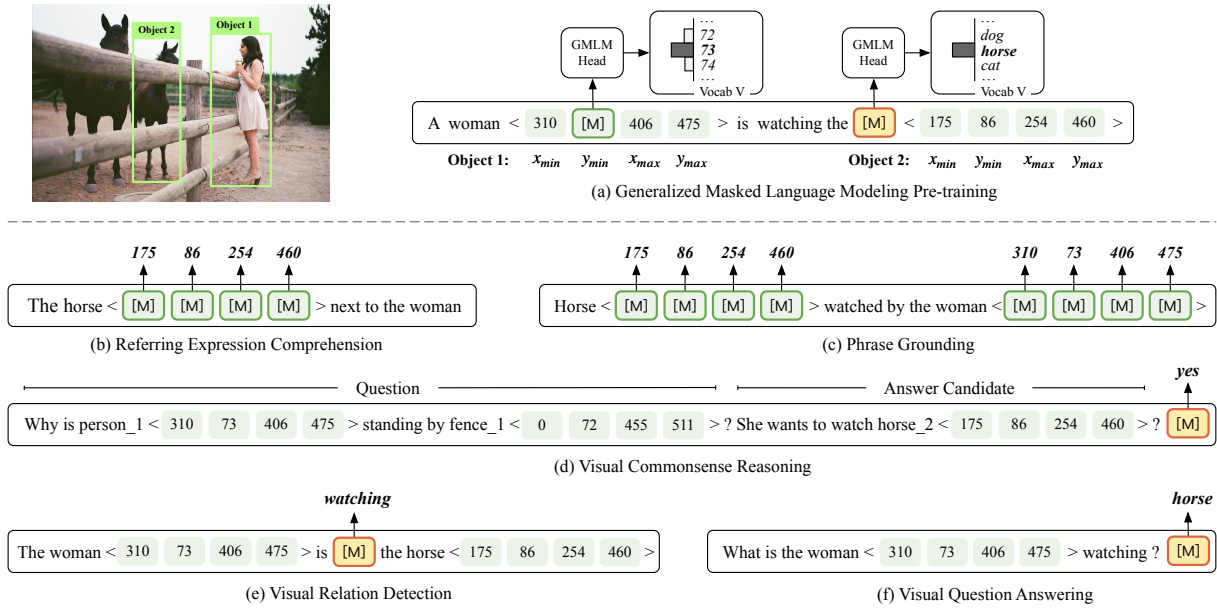


Figure 1: PEVL formulates positions and language into a unified language modeling framework. (a) During pre-training, PEVL recovers masked text and position tokens in a generalized masked language modeling (GMLM) task. (b) During prompt-tuning, PEVL reformulates various VL tasks into a fill-in-the-blank problem, which are addressed by the reused GMLM head.

et al., 2015), previous works have shown that introducing explicit grounding can also lead to better performance and robustness (Anderson et al., 2018; Huang et al., 2019), which can hardly be achieved in current detector-free VLP models.

In a preliminary exploration, MDETR (Kamath et al., 2021) proposes to enhance detector-free VLP models by regressing object positions with Transformer decoders, serving position-output tasks such as referring expression comprehension. However, it is still unknown how to deal with various position-input tasks, such as visual commonsense reasoning and visual relation detection. Moreover, during fine-tuning, task-specific classification heads are typically introduced, resulting in a significant gap between pre-training and fine-tuning, which hinders taking full advantage of pre-trained model capabilities in downstream tasks.

In this work, we propose PEVL that enhances the pre-training and prompt tuning of VLP models with explicit object position modeling. Inspired by the recent Pix2Seq (Chen et al., 2022) that casts object detection as a language modeling task, PEVL reformulates object positions as discrete tokens, and learns the joint distribution of object positions and language in a unified language modeling framework, as shown in Figure 1. Specifically, PEVL exploits explicit region-text alignments in existing VL datasets. Discretized position tokens are placed

after object text tokens to indicate the object locations in both pre-training and prompt tuning:

(1) During *pre-training*, PEVL learns explicit VL alignment based on a generalized masked language modeling (GMLM) task, where the model recovers masked text tokens and position tokens from cross-modal context. We note that although the discretization of positions enables their unified modeling with language, it also eliminates the ordering of positions as compared with traditional continuous regression methods (i.e., predicting nearby and faraway positions to ground-truth are equally punished). The problem is exacerbated by the inevitable small disturbances in the human annotation of bounding boxes. To address the challenge, we present a novel ordering-aware objective for masked position reconstruction, which assigns larger probabilistic soft labels for nearby position tokens, and therefore retains the ordering.

(2) During *prompt tuning*, PEVL can support various downstream VL tasks in a flexible prompt tuning framework, where VL tasks are addressed by the reused GMLM head in a fill-in-the-blank paradigm. In this way, PEVL maximally mitigates the gap between pre-training and tuning, and better stimulates the pre-trained model capabilities.

We conduct comprehensive experiments on five VL tasks, including position-output, input and insensitive tasks. Experimental results show that

through position enhancement, PEVL enables state-of-the-art performance of detector-free VLP models on position-sensitive tasks such as referring expression comprehension and phrase grounding, and also improves the performance on position-insensitive tasks with grounded inputs.

Our contributions are threefold: (1) We unify the modeling of positions and language in a language modeling framework, which enhances both pre-training and prompt tuning of VLP models. (2) We present a novel ordering-aware objective that retains the ordering of position tokens and avoids the influence of position annotation noise. (3) We conduct comprehensive experiments on five position-sensitive and insensitive VL tasks, which demonstrates the effectiveness of the proposed model.

2 Preliminary

In principle, the PEVL framework is orthogonal to VLP architectures and can be built on any VLP models to achieve position enhancement. In this work, without loss of generality, we adopt ALBEF (Li et al., 2021) as the model backbone, which is a representative detector-free VLP model that achieves state-of-the-art performance on many VL tasks. We briefly introduce the pre-training and fine-tuning procedure of ALBEF, and refer readers to the original paper for additional details.

Pre-training. The ALBEF architecture is composed of two unimodal encoders followed by a cross-modal encoder. Images and text are first encoded using a vision Transformer (Dosovitskiy et al., 2021) and a text Transformer (Vaswani et al., 2017) respectively, and then fused with a cross-modal Transformer. The model is pre-trained with three tasks, including masked language modeling, image-text contrastive learning and image-text matching. (1) Masked language modeling aims to recover masked text tokens from the cross-modal context. (2) Image-text contrastive learning aligns the intermediate unimodal representations of image-text pairs by a contrastive loss. (3) Image-text matching classifies whether an image-text pair is aligned based on the [CLS] token of the cross-modal Transformer. To alleviate the noise in the pre-training text, a momentum model is maintained based on the moving-average of model parameters to provide pseudo-targets as additional supervision.

Fine-tuning. During fine-tuning, ALBEF introduces new classification heads or decoders to handle VL tasks, which leads to significant gap from

pre-training. The gap hinders taking full advantage of pre-trained capabilities for downstream tasks. Moreover, since object positions cannot be explicitly modeled, detector-free VLP models typically struggle on position-sensitive tasks, which greatly undermines their generality and practicality.

3 Methodology

We introduce the PEVL framework, including the position reformulation for VL models, and position-enhanced VL pre-training and prompt tuning.

3.1 Reformulating Positions for VL Models

Cross-modal position modeling that explicitly connects image regions and text units underpins a broad range of VL tasks. To enable strong cross-modal position modeling capability of VLP models, a primary challenge is to find a good position formulation that can be (1) easily integrated and unified into mainstream VLP models, and can be (2) flexibly prompt-tuned in various downstream tasks with minimal gap from pre-training as well.

To this end, previous works attempt to indicate image regions by introducing region embeddings (Cho et al., 2021) or colors (Yao et al., 2021b) that correspond to pre-defined text tokens, which require pre-detected image regions from costly external object detectors. In contrast, we note that for the mainstream VLP models with vision Transformers (Dosovitskiy et al., 2021) as visual encoders, image patch positions are already well indicated by positional embeddings (Vaswani et al., 2017), and therefore no special treatments are in fact needed for visual position coordination.

To explicitly express visual positions in text, inspired by Pix2Seq (Chen et al., 2022) that casts object detection as a language modeling task, PEVL reformulates object bounding box coordinates as *discrete position tokens*. The position tokens can be easily unified with text tokens in a language modeling framework, where the vocabulary includes both text and position tokens, and can also be easily pre-trained with existing VLP techniques. In addition to the convenience in pre-training, another important advantage is that VLP models can be easily prompt-tuned to handle various position-sensitive and insensitive VL tasks with minimal gap from pre-training, as shown in Figure 1.

Specifically, given an image-text pair for pre-training (I, T) , we exploit the composing object texts and their bounding boxes $O = \{(c_i, b_i)\}_{i=1}^N$,

where c_i is the object text (e.g., *person*) in text T , and $b_i = (x_{\min}, y_{\min}, x_{\max}, y_{\max})$ is the coordinates of the corresponding bounding box. The bounding box coordinates are discretized into position tokens as $\lfloor Mx/w \rfloor$ and $\lfloor My/h \rfloor$, where w and h are the width and height of the image, and M is the total number of the position tokens. Intuitively, a larger number of position tokens will lead to a coordinate system with higher resolution, but will be more compute- and data- expensive to learn. Finally the position tokens are placed after the corresponding object text c_i in T to explicitly indicate the object position. Note that two special tokens “<” and “>” are introduced to indicate the start and end of position tokens, which are useful in prompting models to produce position tokens in position-output tasks.

3.2 Position-enhanced VL Pre-training

After unifying positions and text in a language modeling framework, PEVL can be easily integrated into existing VLP models. To effectively learn the position and text interactions, in addition to the image-text contrastive and image-text matching tasks (see Section 2), we present a novel generalized masked language modeling (GMLM) pre-training task, which recovers both masked text and position tokens based on a generalized vocabulary \mathcal{V} that includes both types of tokens. We introduce two main components of the GMLM task, including masking strategy and reconstruction objective.

Masking Strategy. While the text tokens are usually masked with low ratios (e.g., 15%) in traditional MLM tasks (Devlin et al., 2019; Li et al., 2021), we find that the same masking strategy cannot well serve position modeling. The reason is that object positions are relatively low-level signals, and therefore models can easily reconstruct the masked position tokens when the masking ratio is low. For example, reconstructing a single masked position token (e.g., x_{\min}) given the other three unmasked ones will be largely equivalent to enclosing an object by moving a corner of the bounding box in a straight line, which does not require deep understanding of the VL semantics. Similar problems are also discussed in self-supervised learning on images (He et al., 2022).

To address the issue, we adopt *high masking ratios* for position tokens, and encourage masking a more complete subset of object positions. Specifically, for each object, we randomly mask n of its four position tokens with 0.25 probabil-

ity, where $n = 1, 2, 3, 4$. For example, for 25% of the time, the four object positions (i.e., $n = 4$) are completely masked for reconstruction. For text tokens, we follow the 15% masking strategy in previous works (Li et al., 2021).² In this way, models are forced to learn high-level semantic interactions among image regions, text and position tokens.

Reconstruction Objective. Traditional MLM tasks typically adopt a one-hot target for token reconstruction. However, we note that the one-hot target essentially eliminates the ordering of the positions: If the position prediction is not exactly correct, predicting nearby and faraway positions to ground-truth are equally punished. The problem is exacerbated by the inevitable small disturbances in the human annotation process of bounding boxes, which confuses models in discrete position learning. To address the problem, we present a novel *ordering-aware objective* for position reconstruction that assigns larger probabilistic soft labels for nearby position tokens. Specifically, given a masked position token, the unnormalized probabilistic label y_i for each position token p_i decreases exponentially with its distance to the ground-truth:

$$y_i = e^{-\alpha|p_i - p^*|}, \quad (1)$$

where $|p_i - p^*|$ is the distance between p_i and the ground-truth p^* , and α is a hyperparameter controlling the decay rate. The normalized probabilistic label \tilde{y}_i is then used to compute the ordering-aware objective for position tokens:

$$\mathcal{L}_p = - \sum_{p_i} \tilde{y}_i \log P([\text{MASK}] = p_i). \quad (2)$$

The probability of position tokens is given by the GMLM head as:

$$P([\text{MASK}] = p_i) = \frac{\exp(\mathbf{h}_{[\text{MASK}]}^\top \mathbf{p}_i)}{\sum_{p_j} \exp(\mathbf{h}_{[\text{MASK}]}^\top \mathbf{p}_j)}, \quad (3)$$

where $\mathbf{h}_{[\text{MASK}]}$ is the hidden representation of the [MASK] token, and \mathbf{p}_i is the representation of position token p_i in the GMLM head. In this way, the objective retains the ordering of position tokens and avoids the influence of position annotation noise. For the text token reconstruction loss \mathcal{L}_t , we follow the traditional implementation in ALBEF (Li et al., 2021). The final GMLM loss is the weighted sum of the loss for position token reconstruction and

²For the chosen text tokens, the replacements are 80% [MASK] tokens, 10% random tokens, and 10% unchanged.

text token reconstruction: $\mathcal{L}_{\text{GMLM}} = \lambda\mathcal{L}_p + \mathcal{L}_t$, where λ is a weighting hyperparameter.

Pre-training Corpora. PEVL exploits explicit object position annotation in VL datasets for position learning. The pre-training corpora consist of referring expressions (Yu et al., 2016; Mao et al., 2016), Flickr30k (Plummer et al., 2015), GQA (Hudson and Manning, 2019), VCR (Zellers et al., 2019) and Visual Genome (Krishna et al., 2017), with 4.7M image-text pairs in total. Following Chen et al. (2020), we remove the images in the downstream test and validation sets from the pre-training corpora.

3.3 Position-enhanced VL Prompt Tuning

To adapt VLP models to downstream tasks, previous works typically introduce new classification heads or even Transformer decoders (Kamath et al., 2021; Li et al., 2021; Chen et al., 2020), leading to significant gap from pre-training. Recent works in pre-trained language models have shown that a consistent tuning approach with pre-training (i.e., prompt tuning) can better stimulate the pre-trained capability in downstream tasks (Schick and Schütze, 2021; Gao et al., 2021; Liu et al., 2021). However, it is still unknown whether and how VLP models can be prompt tuned to support both position-sensitive and insensitive VL tasks.

In this context, a crucial advantage of unifying positions with language is that, VLP models can be easily prompt-tuned to handle various VL tasks based on the *reused GMLM head* with minimal gap from pre-training. We divide VL tasks according to the role of positions, including position-output tasks, position-input tasks, and position-insensitive tasks. Here we introduce the main prompt tuning procedure for position-sensitive tasks. In our experiments, we show that position-insensitive tasks can also benefit from well-grounded inputs in PEVL framework (see Section 4.1).

Position-output Tasks demand positions as task outputs (e.g., predicting the positions of objects described by text), such as referring expression comprehension and phrase grounding. To handle position-output tasks, we simply place four consecutive [MASK] tokens wrapped by “<” and “>” after object texts to be grounded for position prediction. (1) *Referring Expression Comprehension*. Since the task requires locating the head noun, we place the mask tokens after the first object text for position prediction. (2) *Phrase Grounding*. Since

the task requires locating all objects, mask tokens are placed after each object text. After placing mask tokens, the model is prompt-tuned to produce position tokens with reused GMLM head based on the ordering-aware objective as in Equation 2.

Position-input Tasks require a mixture of position and text (i.e., grounded text) as task inputs, such as visual commonsense reasoning and visual relation detection. To handle position-input tasks, PEVL first explicitly indicates the object positions in input text (see Section 3.1)³, and then produces answers in a fill-in-the-blank paradigm based on the reused GMLM head.

Visual Commonsense Reasoning. Given a question, models are asked to choose the answer sentence (and rationale) from multiple candidates. For answer selection, the question q and answer candidate a_i are put in a prompt template as: “ q a_i answer: [MASK]”. Then the model can be prompted to decide which token $t \in \{yes, no\}$ is more proper to reconstruct the [MASK] token. Another plausible alternative is to reuse the image-text matching head to discriminate whether the image is aligned with the concatenated question and answer. The intuition is that a question concatenated with the correct answer can better match the image content than concatenated with a wrong answer. In our experiments, we find that the latter approach yields better performance on VCR. Despite the essential equivalence of the two prompting approaches (i.e., classifying special tokens in the last layer into binary labels with reused pre-trained heads), image-text matching task focuses more on the holistic matching between cross-modal signals during pre-training, which better fits the VCR task containing typically long text answers.

Visual Relation Detection. Given an object pair (s, o) (e.g., *woman, horse*) in the image, models are required to classify their semantic relation r (e.g., *watching, riding*). We design the prompt template as: “The s is [MASK] the o ”. Then the model is prompted to produce the relational tokens from the relation set with reused GMLM head. To deal with relations that consist of different numbers of tokens, we pad relational tokens to a maximum length l , and place l consecutive masks in the template for relation prediction. We also include a special relation *no relation with* in the relation set, which indicates no relation between the object pair. Dur-

³We omit the position tokens in the task input in the following for simplicity.

Model	Object Detector	RefCOCO			RefCOCO+			RefCOCog		Flickr30k	
		val	testA	testB	val	testA	testB	val	test	val	test
MAttNet (Yu et al., 2018a)	w/	76.7	81.1	70.0	65.3	71.6	56.0	66.6	67.3	-	-
DDPN (Yu et al., 2018b)	w/	76.8	80.1	72.4	64.8	70.5	54.1	-	-	72.8	73.5
VL-T5 (Cho et al., 2021)	w/	-	-	-	-	-	-	71.2	71.3	-	-
ViLBERT (Lu et al., 2019)	w/	-	-	-	72.3	78.5	62.6	-	-	-	-
VL-BERT_L (Su et al., 2020)	w/	-	-	-	72.6	78.6	62.3	-	-	-	-
UNITER_L (Chen et al., 2020)	w/	81.4	87.0	74.2	75.9	81.5	66.7	74.9	75.8	-	-
VinVL_L (Zhang et al., 2021)	w/	81.8	87.2	74.3	74.5	80.8	64.3	74.6	75.7	-	-
VILLA_L (Gan et al., 2020)	w/	82.4	87.5	74.9	76.2	81.5	66.8	76.2	76.7	-	-
ERNIE-ViL_L (Yu et al., 2021)	w/	-	-	-	80.0	82.1	66.9	-	-	-	-
UniTAB (Yang et al., 2022)	w/o	88.6	91.1	83.8	81.0	85.4	71.6	84.6	84.7	-	79.6
MDETR (Kamath et al., 2021)	w/o	87.5	90.4	82.7	81.1	85.5	73.0	83.4	83.3	82.3	83.8
OFA (Wang et al., 2022a)	w/o	88.5	90.7	83.3	81.4	87.2	74.3	82.3	82.3	-	-
ALBEF [†] (Li et al., 2021)	w/o	-	-	-	58.5	65.9	46.3	-	-	-	-
PEVL	w/o	89.6	92.5	85.0	83.0	74.5	74.5	87.1	86.3	84.1	84.4
Δ	-	-	-	-	+24.5	+22.5	+28.2	-	-	-	-

Table 1: Experimental results on referring expression comprehension and phrase grounding. L: large size model. [†]: weakly supervised results, where only image-expression pairs are used due to the lack of explicit position modeling capability. Δ : improvements of PEVL over the ALBEF backbone.

Model	Q \rightarrow A	QA \rightarrow R	Q \rightarrow AR
R2C	63.8 (65.1)	67.2 (67.3)	43.1 (44.0)
TAB-VCR	69.9 (70.4)	72.2 (71.7)	50.6 (50.5)
VisualBERT	70.8 (71.6)	73.2 (73.2)	52.2 (52.4)
ViLBERT	72.4 (73.3)	74.5 (74.6)	54.0 (54.8)
Unicoder-VL	72.6 (73.4)	74.5 (74.4)	54.4 (54.9)
B2T2	73.2 (74.0)	77.1 (77.1)	56.6 (57.1)
UNITER	74.6 (75.0)	77.0 (77.2)	57.8 (58.2)
VL-BERT	73.8 (75.8)	74.4 (78.4)	55.2 (59.7)
ALBEF	71.9 (72.9)	74.5 (74.5)	54.1 (54.7)
PEVL	75.1 (76.0)	76.4 (76.7)	57.8 (58.6)
Δ	+3.2 (+3.1)	+1.9 (+2.2)	+3.7 (+3.9)

Table 2: Experimental results of visual commonsense reasoning on VCR validation (and test) sets.

Model	R@50	R@100	mR@50	mR@100
MSDN	64.6	66.6	15.9	17.5
VCTree	65.5	67.4	15.4	16.6
GPS-Net	65.2	67.1	15.2	16.6
Motif	66.0	67.9	14.6	15.8
VisualDS	64.4	66.4	16.1	17.5
Unbiased	47.2	51.6	25.4	28.7
IETrans	48.6	50.5	35.8	39.1
DT2-ACBS	23.3	25.6	35.9	39.7
ALBEF	57.6	63.5	12.2	15.4
PEVL	64.4	66.3	21.7	23.5
Δ	+6.8	+2.8	+9.5	+8.1

Table 3: Experimental results of visual relation detection on Visual Genome dataset.

ing inference, the score of relation r is given by the average log probability of non-padding tokens: $s_r = \frac{1}{|r|} \sum_{i=1}^{|r|} \log P([\text{MASK}]^{(i)} = r^{(i)})$, where $[\text{MASK}]^{(i)}$ is the i -th mask token, and $r^{(i)}$ is the i -th token of r . An important advantage of prompt tuning for the task is that, the large number of long-tail relations can be better learned thanks to the rich knowledge in VLP models.

4 Experiments

We evaluate PEVL on five popular VL tasks. The models are in base size unless otherwise specified.

4.1 Main Results

Referring Expression Comprehension. We adopt three popular datasets for the task, including RefCOCO, RefCOCO+ (Yu et al., 2016) and RefCOCog (Mao et al., 2016). We use accuracy@0.5

as the evaluation metric (Kamath et al., 2021). For baselines, we compare with state-of-the-art models for the task, and VLP models with large-size backbones. We report the weakly supervised results from ALBEF (Li et al., 2021), which uses GRAD-CAM (Selvaraju et al., 2017) heat map to rank the object candidates from external detectors.

From the experimental results in Table 1, we have the following observations: (1) PEVL outperforms all baseline models, achieving a new state-of-the-art on all three datasets for the task. Specifically, the base-size PEVL outperforms the state-of-the-art regression-based MDETR by 2.9 absolute points on the RefCOCO+ testA set, and large-size VLP models that use external detector feature inputs, such as ERNIE-ViL and VILLA. (2) PEVL significantly improves the ALBEF backbone by explicit object position modeling, effectively addressing the shortcoming in position-output tasks.

Models	LXMERT	BAN	CTI	CFR	ALBEF	PEVL [†]	Δ
Accuracy	59.8	61.5	61.7	73.6	64.8	77.0	+12.2

Table 4: Visual question answering results on GQA validation set. [†]: grounded inputs.

Question Type	Relation	Attribute	Object	Category	Global	Overall
Percentage (%)	46.7	32.0	11.8	6.5	3.1	100.0
ALBEF	56.9	67.9	87.9	62.5	68.1	64.8
PEVL [†]	68.4	84.2	98.1	68.8	68.5	77.0
Δ	+11.5	+16.3	+10.2	+6.3	+0.4	+12.2

Table 5: Visual question answering results of different question types on GQA validation set. [†]: grounded inputs.

Model	RefCOCO+			Flickr30k		VCR			VG		GQA
	val	testA	testB	val	test	Q \rightarrow A	QA \rightarrow R	Q \rightarrow AR	R@50	mR@50	val
PEVL	83.1	88.4	74.5	84.1	84.4	75.1	76.4	57.8	64.4	21.7	77.0
w/o PT	-	-	-	-	-	75.1	76.2	57.6	61.7	14.2	77.0
w/o OAO	79.9	86.3	69.4	82.2	82.9	75.6	76.3	57.8	64.1	21.5	76.4
w/o Pos	-	-	-	-	-	71.9	74.5	54.1	61.5	18.9	66.6

Table 6: Ablation results. PT: prompt tuning, OAO: ordering-aware objective, Pos: position tokens.

Phrase Grounding. We perform experiments on the Flickr30k entities dataset (Plummer et al., 2015). Following MDETR, we adopt merged-box accuracy@0.5 as the evaluation metric, and compare our model with the state-of-the-art baselines for the task (Kamath et al., 2021; Yang et al., 2022). From Table 1 we observe that PEVL achieves a new state-of-the-art on the phrase grounding task in grounding multiple objects in text. The results show that PEVL can effectively integrate positions with language to achieve competitive performance for various position-output tasks.

Visual Commonsense Reasoning. We adopt the popular VCR benchmark (Zellers et al., 2019), which provides human-annotated positions for objects. We report the accuracy of predicting the answer (Q \rightarrow A), rationale (QA \rightarrow R) and both (Q \rightarrow AR). We compare with task-specific baselines and strong VLP models. For fair comparisons, we further pre-train ALBEF baseline on the same corpora as PEVL in all experiments. From the results in Table 2, we observe that PEVL significantly improves the ALBEF backbone (e.g., by 3.9 absolute points in Q \rightarrow AR), achieving comparable performance to strong UNITER equipped with external object detectors. While the results are not state-of-the-art on the VCR benchmark, they are quite reasonable considering the current literature. The results show that PEVL can effectively provide clues for complex reasoning through grounded inputs.

Visual Relation Detection. We evaluate PEVL on the widely used Visual Genome dataset (Krishna et al., 2017), which contains 50 visual relation types. Given the human-annotated positions and labels of an object pair, models are required to predict the relations. Following previous works (Tang et al., 2020; Lin et al., 2020), we report the recall@K (R@K) and mean recall@K (mR@K) as evaluation metrics. We compare PEVL with state-of-the-art baselines with detector feature inputs. From Table 3, we observe that: (1) Without task-specific designs or heuristics, PEVL achieves competitive performance in both R@K and mR@K. The results show that PEVL can effectively stimulate the knowledge in VLP models for both frequent and long-tail relations through prompt tuning. (2) ALBEF struggles on the visual relation detection task, since the positions of the target object pair cannot be informed. In contrast, PEVL can effectively integrate the object position information through simple position tokens for relation prediction.

Visual Question Answering. For position-insensitive tasks such as visual question answering, object positions are not required to be *explicitly* modeled. However, we argue that explicit object position modeling can provide fine-grained clues for complex question reasoning. Specifically, we are interested in the question: Can VLP models benefit from grounded text for answering complex questions in PEVL framework?



Figure 2: Case study on referring expression comprehension and phrase grounding tasks.

In principle, to achieve explicit position augmentation, VQA can be decomposed into two position-sensitive stages, including an object grounding (position-output) stage and a question answering (position-input) stage. However, in our experiments, we find that the unneglectable errors in current visual grounding models constitute a bottleneck for such a two-stage model. Therefore, we turn to an ideal experiment, where the ground-truth object positions are available. Specifically, we adopt the object position annotation provided by the GQA dataset (Hudson and Manning, 2019). We explicitly indicate the position of each object in a similar approach as in position-input tasks (see Section 3.3). Then the position-enhanced question is put into the prompt template: “*q* answer: [MASK]”. Finally models are asked to generate answer tokens from answer candidate set. We use the same approach as in visual relation detection to cope with multi-token answers.

From Table 4 we observe that with grounded inputs, PEVL significantly improves the performance of ALBEF backbone in compositional question answering. The results show that object grounding is still one of the key obstacles in VQA, and PEVL can effectively utilize grounded questions for VQA in a simple prompt tuning framework. To investigate which type of questions benefit from grounded inputs, we divide GQA validation set according to the question types from MDETR. From Table 5, we can see that high-quality grounding signals improve the performance on all question types. Interestingly, relation and attribute-based questions benefit more from grounded text than object-based questions, indicating the fundamental role of object modeling in reasoning over complex questions.

4.2 Experimental Analysis

Ablation Study. We ablate key components of PEVL, including prompt tuning, ordering-aware

Mask	val	testA	testB	Epochs
20%	89.3	92.3	84.3	5
40%	89.0	92.2	84.7	5
60%	89.4	92.4	83.4	7
Ours	89.6	92.5	85.0	3

Table 7: Performance and epochs of different masking strategies on the RefCOCO dataset.

objective, and position tokens to investigate their contribution. From the results in Table 6, we can see that all components contribute to the final performance. Specifically, position enhancement and prompt tuning are essential for PEVL to perform position-output tasks. Prompt tuning can also be helpful in learning long-tail relations for visual relation detection, which is consistent with the results from previous works (Yao et al., 2021b). The ordering-aware objective contributes more to position-output tasks than position-input tasks.

Influence of Masking Strategies. We investigate the influence of different masking strategies for position tokens during pre-training. Specifically, for baselines, the position tokens are independently chosen with a certain probability during pre-training, where the ratios of masked, replaced and unchanged tokens for the chosen token are kept identical to BERT (Devlin et al., 2019). We report the performance and the number of epochs required in the intermediate pre-training on the RefCOCO dataset. From the results in Table 7, we observe that our masking strategy achieves both better performance and faster convergence. The results show that a high masking ratio and a more complete subset of masked positions are both important for good position learning results.

Case Study. We visualize the position predictions on the validation sets of RefCOCO+, RefCOCOg and Flicker30k. Previous visual localization models are either based on continuous re-

gression (Kamath et al., 2021), or limited to non-language tasks (Chen et al., 2022). From Figure 2, we can see that discretized positions can be closely integrated with language in Transformers to achieve strong visual reasoning and localization results. Similar to regression-based models, the localization of small objects (e.g., *wheels* in the right figure) can also be challenging.

5 Related Work

Position Enhancement for VLP. Object position modeling underpins a wide range of VL tasks. To deal with position-output tasks, some works (Kamath et al., 2021; Gupta et al., 2022; Yang et al., 2022) propose to perform object position prediction using Transformer decoders, but are unable to handle various position-input tasks. To explicitly indicate position inputs for VLP models, previous works explored learning region embeddings (Cho et al., 2021), or color-based cross-modal coreferential markers (Yao et al., 2021b), but rely on external object detectors. MERLOT (Zellers et al., 2021) also proposes to highlight objects in images with colors for position-input tasks. X-VLM (Zeng et al., 2022) aligns multi-grained concepts in text and image for VLP models. In comparison, PEVL supports both position- input and output VL tasks in a unified prompt tuning framework.

Prompt Tuning. Prompt tuning for pre-trained language models is in rapid growth in natural language processing (Petroni et al., 2019; Raffel et al., 2020; Brown et al., 2020; Schick and Schütze, 2021; Gao et al., 2021; Qin and Eisner, 2021; Liu et al., 2021; Yao et al., 2022). Recently there is also growing interest in prompt tuning VLP models. Most existing works prompt tune contrastively pre-trained image-text matching models (Radford et al., 2021; Jia et al., 2021) for recognition tasks (Zhou et al., 2022; Rao et al., 2022; Wang et al., 2021; Gu et al., 2022; Xie and Zheng, 2021; Ju et al., 2022). CPT (Yao et al., 2021b) prompt tunes VLP models with color-based prompts, and achieves promising results for zero- and few-shot tasks. Some works perform pre-training and VL tasks using identical Transformer decoders in an auto-regressive fashion (Wang et al., 2022b; Cho et al., 2021; Yang et al., 2022; Tsimpoukelli et al., 2021), which avoids the gap between pre-training and tuning, but are typically limited in performance due to the unidirectional architecture.

6 Conclusion and Future Work

In this work, we present PEVL that enhances the pre-training and prompt-tuning of detector-free VLP models with unified position and language modeling. Comprehensive experimental results demonstrate the effectiveness of the proposed model. Future works include exploring weakly supervised signals for position and language learning without human annotation.

7 Acknowledgement

This work is supported by the National Key R&D Program of China (No. 2020AAA0106502), Institute Guo Qiang at Tsinghua University and NExT++ project from the National Research Foundation, Prime Minister’s Office, Singapore under its IRC@Singapore Funding Initiative.

Yuan Yao designed the framework and experiments, and wrote the paper. Qianyu Chen conducted the experiments. Ao Zhang and Wei Ji participated in the discussion. Zhiyuan Liu, Tat-Seng Chua and Maosong Sun advised the project.

8 Limitations

We identify several key limitations of PEVL that are promising for future explorations.

Computation Cost. To explicitly model object positions and text in a unified framework, we introduce four additional position tokens for each target object, which requires more computation cost. Improving the computation efficiency of additional position tokens is an important direction for future improvements.

Object Annotation. Similar to other explicit object position modeling VLP models, PEVL requires manual object annotation in multi-modal datasets. It will be promising to explore pre-training tasks that can learn position tokens with less supervision. For example, VLP models can be bootstrapped from small-scale human annotations and large-scale predicted annotations.

References

- Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. 2019. [Fusion of detected objects in text for visual question answering](#). In *Proceedings of EMNLP-IJCNLP*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image](#)

- captioning and visual question answering. In *Proceedings of CVPR*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. **VQA: Visual question answering**. In *Proceedings of ICCV*.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. **On the opportunities and risks of foundation models**. *arXiv preprint arXiv:2108.07258*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. **Language models are few-shot learners**. In *Proceedings of NeurIPS*.
- Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. 2022. **Pix2Seq: A language modeling framework for object detection**. In *Proceedings of ICLR*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. **UNITER: Universal image-text representation learning**. In *Proceedings of ECCV*.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. **Unifying vision-and-language tasks via text generation**. In *Proceedings of ICML*.
- Alakh Desai, Tz-Ying Wu, Subarna Tripathi, and Nuno Vasconcelos. 2021. **Learning of visual relations: The devil is in the tails**. In *Proceedings of ICCV*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of NAACL-HLT*.
- Tuong Do, Thanh-Toan Do, Huy Tran, Erman Tjiputra, and Quang D Tran. 2019. **Compact trilinear interaction for visual question answering**. In *Proceedings of ICCV*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. **An image is worth 16x16 words: Transformers for image recognition at scale**. In *Proceedings of ICLR*.
- Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. 2020. **Large-scale adversarial training for vision-and-language representation learning**. In *Proceedings of NeurIPS*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. **Making pre-trained language models better few-shot learners**. In *Proceedings of ACL*.
- Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. 2022. **Open-vocabulary object detection via vision and language knowledge distillation**. In *Proceedings of ICLR*.
- Tanmay Gupta, Amita Kamath, Aniruddha Kembhavi, and Derek Hoiem. 2022. **Towards general purpose vision systems: An end-to-end task-agnostic vision-language architecture**. In *Proceedings of CVPR*.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. **Masked autoencoders are scalable vision learners**. In *Proceedings of CVPR*.
- Pingping Huang, Jianhui Huang, Yuqing Guo, Min Qiao, and Yong Zhu. 2019. **Multi-grained attention with object-level grounding for visual question answering**. In *Proceedings of ACL*.
- Drew A Hudson and Christopher D Manning. 2019. **GQA: A new dataset for real-world visual reasoning and compositional question answering**. In *Proceedings of CVPR*.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. **Scaling up visual and vision-language representation learning with noisy text supervision**. In *Proceedings of ICML*.
- Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. 2022. **Prompting visual-language models for efficient video understanding**. In *Proceedings of ECCV*.
- Aishwarya Kamath, Mannat Singh, Yann LeCun, Ishan Misra, Gabriel Synnaeve, and Nicolas Carion. 2021. **MDETR: Modulated detection for end-to-end multimodal understanding**. In *Proceedings of ICCV*.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. **Bilinear attention networks**. In *Proceedings of NeurIPS*.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. **ViLT: Vision-and-language transformer without convolution or region supervision**. In *Proceedings of ICML*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. **Visual Genome: Connecting language and vision using crowdsourced dense image annotations**. *International Journal of Computer Vision*.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020a. **Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training**. In *Proceedings of AAAI*.
- Junnan Li, Ramprasaath R Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. 2021. **Align before fuse: Vision and language representation learning with momentum distillation**. In *Proceedings of NeurIPS*.

- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2020b. [VisualBERT: A simple and performant baseline for vision and language](#). In *Proceedings of ACL*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020c. [Oscar: Object-semantic aligned pre-training for vision-language tasks](#). In *Proceedings of ECCV*.
- J. Lin, U. Jain, and A. G. Schwing. 2019. [TAB-VCR: Tags and attributes based vcr baselines](#). In *Proceedings of NeurIPS*.
- Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. 2020. [GPS-Net: Graph property sensing network for scene graph generation](#). In *Proceedings of CVPR*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Computing Surveys*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [ViLBERT: Pretraining task-agnostic visual-linguistic representations for vision-and-language tasks](#). In *Proceedings of NeurIPS*.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. [Generation and comprehension of unambiguous object descriptions](#). In *Proceedings of CVPR*.
- Binh X Nguyen, Tuong Do, Huy Tran, Erman Tjiputra, Quang D Tran, and Anh Nguyen. 2022. [Coarse-to-fine reasoning for visual question answering](#). In *Proceedings of CVPR Workshops*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of EMNLP-IJCNLP*.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. [Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models](#). In *Proceedings of ICCV*.
- Guanghui Qin and Jason Eisner. 2021. [Learning how to ask: Querying LMs with mixtures of soft prompts](#). In *Proceedings of NAACL*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of ICML*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*.
- Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. 2022. [DenseCLIP: Language-guided dense prediction with context-aware prompting](#). In *Proceedings of CVPR*.
- Timo Schick and Hinrich Schütze. 2021. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of NAACL*.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. [GRAD-CAM: Visual explanations from deep networks via gradient-based localization](#). In *Proceedings of ICCV*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [VL-BERT: Pre-training of generic visual-linguistic representations](#). In *Proceedings of ICLR*.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: Learning cross-modality encoder representations from transformers](#). In *Proceedings of EMNLP-IJCNLP*.
- Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. 2020. [Unbiased scene graph generation from biased training](#). In *Proceedings of CVPR*.
- Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. 2019. [Learning to compose dynamic tree structures for visual contexts](#). In *Proceedings of CVPR*.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. 2021. [Multimodal few-shot learning with frozen language models](#). In *Proceedings of NeurIPS*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of NeurIPS*.
- Mengmeng Wang, Jiazheng Xing, and Yong Liu. 2021. [ActionCLIP: A new paradigm for video action recognition](#). In *Proceedings of CVPR*.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022a. [OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *Proceedings of ICML*.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2022b. [SimVLM: Simple visual language model pretraining with weak supervision](#). In *Proceedings of ICLR*.

- Johnathan Xie and Shuai Zheng. 2021. [ZSD-YOLO: Zero-shot YOLO detection using vision-language knowledge distillation](#). *arXiv preprint arXiv:2109.12066*.
- Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. [Scene graph generation by iterative message passing](#). In *Proceedings of CVPR*.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Faisal Ahmed, Zicheng Liu, Yumao Lu, and Lijuan Wang. 2022. [UniTAB: Unifying text and box outputs for grounded vision-language modeling](#). In *Proceedings of ECCV*.
- Yuan Yao, Bowen Dong, Ao Zhang, Zhengyan Zhang, Ruobing Xie, Zhiyuan Liu, Leyu Lin, Maosong Sun, and Jianyong Wang. 2022. [Prompt tuning for discriminative pre-trained language models](#). In *Findings of ACL*.
- Yuan Yao, Ao Zhang, Xu Han, Mengdi Li, Cornelius Weber, Zhiyuan Liu, Stefan Wernter, and Maosong Sun. 2021a. [Visual distant supervision for scene graph generation](#). In *Proceedings of ICCV*.
- Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2021b. [CPT: Colorful prompt tuning for pre-trained vision-language models](#). *arXiv preprint arXiv:2109.11797*.
- Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [ERNIE-ViL: Knowledge enhanced vision-language representations through scene graphs](#). In *Proceedings of AAAI*.
- Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. 2018a. [MAttNet: Modular attention network for referring expression comprehension](#). In *Proceedings of CVPR*.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. [Modeling context in referring expressions](#). In *Proceedings of ECCV*.
- Zhou Yu, Jun Yu, Chenchao Xiang, Zhou Zhao, Qi Tian, and Dacheng Tao. 2018b. [Rethinking diversified and discriminative proposal generation for visual grounding](#). In *Proceedings of IJCAI*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [From recognition to cognition: Visual commonsense reasoning](#). In *Proceedings of CVPR*.
- Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. 2021. [MERLOT: Multimodal neural script knowledge models](#). In *Proceedings of NeurIPS*.
- Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. 2018. [Neural motifs: Scene graph parsing with global context](#). In *Proceedings of CVPR*.
- Yan Zeng, Xinsong Zhang, and Hang Li. 2022. [Multi-grained vision language pre-training: Aligning texts with visual concepts](#). In *Proceedings of ICML*.
- Ao Zhang, Yuan Yao, Qianyu Chen, Wei Ji, Zhiyuan Liu, Maosong Sun, and Tat-Seng Chua. 2022. [Fine-grained scene graph generation with data transfer](#). In *Proceedings of ECCV*.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [VinVL: Revisiting visual representations in vision-language models](#). In *Proceedings of CVPR*.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022. [Learning to prompt for vision-language models](#). *International Journal of Computer Vision*.

A Pre-training Details

We provide pre-training details and statistics of the pre-training corpora.

Implementation details. Our backbone consists of a 6-layer text Transformer encoder, a ViT-B/16 visual encoder, and a 6-layer cross-modal Transformer encoder (commonly referred to as base size in the literature), with 209.5M parameters in total. The backbone is open-sourced for research usage. In pre-training, we initialize PEVL with pre-trained parameters from ALBEF for computation efficiency. PEVL is pre-trained with learning rate $8e-5$, batchsize 512 on 32 NVIDIA V100 GPUs for 5 epochs. The number of position tokens is 512, with decay rate $\alpha = 0.25$, and weighting hyperparameter $\lambda = 2$ in ordering-aware reconstruction. The hyperparameters are selected by grid search on the validation sets. For data augmentation, following MDETR (Kamath et al., 2021), we augment images with random size crop. We also follow Pix2Seq (Chen et al., 2022) to adopt horizontal flip to augment images, where “left” and “right” in text are swapped after flip to ensure the semantic correctness. Previous works suggest that an intermediate in-domain pre-training can better adapt VLP models to downstream tasks (Chen et al., 2020). We therefore conduct an intermediate pre-training before tuning on each downstream task.

Pre-training Corpora. The pre-training corpora consist of referring expressions (Yu et al., 2016; Mao et al., 2016), Flickr30k (Plummer et al., 2015), GQA (Hudson and Manning, 2019), VCR (Zellers et al., 2019) and Visual Genome dense captions (Krishna et al., 2017), with 4.7M image-text pairs and 210K images in total. We provide the detailed statistics of the datasets in Table 8.

Dataset	RefCOCO	RefCOCO+	RefCOCOg	Flickr	GQA	VCR	Visual Genome
# Image-text pairs	107K	107K	72K	148K	762K	1.7M	1.8M
# Images	15K	15K	20k	30K	63K	80K	46K

Table 8: Statistics of pre-training corpora. Numbers of image-text pairs and images are reported.

B Downstream Tasks

We provide details of dataset, prompt tuning and baseline models for each downstream task.

B.1 Referring Expression Comprehension

Datasets. RefCOCO (Yu et al., 2016) is collected from a referential game between two players. The dataset is split into train, validation, testA and testB sets, containing 120,624, 10,834, 5,657 and 5,095 expression-object pairs respectively. RefCOCO+ (Yu et al., 2016) is also constructed in an interactive fashion, and contains 120,191, 10,758, 5,726 and 4,889 expression-object pairs in train, validation, testA and testB sets respectively. RefCOCOg (Mao et al., 2016) is built in a non-interactive way, and contains 80,512, 4,896 and 9,602 expression-object pairs in train, validation and test sets respectively.

Prompt Tuning. We tune the model with learning rate $1e-5$, weight decay 0.02, and batchsize 32 for 10 epochs. Following previous works (Dosovitskiy et al., 2021; Li et al., 2021), we use a higher image resolution of 512 in downstream tuning. The hyperparameters are selected by grid search on the validation set for all experiments. During inference, we select the position token with the largest reconstruction score for each of the four masked tokens.

Baselines. We compare with state-of-the-art baselines, including MAttNet (Yu et al., 2018a), DDPN (Yu et al., 2018b), VL-T5 (Cho et al., 2021), ViLBERT (Lu et al., 2019), UNITER (Chen et al., 2020), VL-BERT (Su et al., 2020), VinVL (Zhang et al., 2021), VILLA (Gan et al., 2020), ERNIE-ViL (Yu et al., 2021), MDETR (Kamath et al., 2021), and ALBEF (Li et al., 2021). We also compare with two concurrent works that achieve competitive performance on visual grounding tasks, including UniTAB (Yang et al., 2022) and OFA (Wang et al., 2022a). We adopt accuracy@0.5 as the evaluation metrics, where an expression is considered correctly grounded if the intersection over union between the top prediction and ground truth is greater than 0.5

B.2 Phrase Grounding

Datasets. Flickr30k entities dataset (Plummer et al., 2015) is collected through annotating 276K entities in the 158K captions from Flickr30k with object bounding boxes. The dataset is split into train, validation and test sets, with 148,915, 14,433, 14,481 noun phrases respectively.

Prompt Tuning. We tune the model with learning rate $1e-5$, weight decay 0.02, and batchsize 128 for 10 epochs. Following previous works (Kamath et al., 2021; Yang et al., 2022), we evaluate our model under the merged-boxes protocol, where the boxes of a phrase (e.g., *crowd*) referring to multiple objects are merged by their union. We use resolution 512 during downstream tuning. During tuning and inference, our model predicts the bounding box of each object separately.

Baselines. We compare with state-of-the-art baselines, including DDPN (Yu et al., 2018b), UniTAB (Yang et al., 2022) and MDETR (Kamath et al., 2021). UniTAB (Yang et al., 2022) performs multi-task fine-tuning with several downstream task datasets. We adopt accuracy@0.5 as the evaluation metrics.

B.3 Visual Relation Detection

Datasets. We use Visual Genome (Krishna et al., 2017) dataset for the evaluation of this task. The dataset is split into train, validation and test sets, with 65,651, 5,000, 32,422 images respectively. The of object categories and relation categories in the dataset are 150 and 50 respectively.

Prompt Tuning. We tune the model with learning rate $2e-5$, weight decay 0.02, and batchsize 256 for 5 epochs. The resolution of images is 512. The ratio of negative samples (i.e., no relations between the object pair) and positive samples is 3:1.

Baselines. We compare with strong baselines, including MotifNet (Zellers et al., 2018), Unbiased (Tang et al., 2020), GPS-Net (Lin et al., 2020), MSDN (Xu et al., 2017), VCTree (Tang et al., 2019), DT2-ACBS (Desai et al., 2021), VisualDS (Yao et al., 2021a), and IETrans (Zhang

et al., 2022). For evaluation metrics, we adopt Recall@K(R@K), which is the ratio of correct relationship in the top K confident relationship predictions, and mean Recall@K(mR@K), which is the average recall upon all predicate classes.

B.4 Visual Commonsense Reasoning

Datasets. VCR dataset (Zellers et al., 2019) is collected through creating questions requiring commonsense reasoning by workers for given images from 110K movie scenes. The dataset is split into train, validation, test sets with 212,923, 26,534, 25,263 questions respectively.

Prompt Tuning. We tune the model with learning rate $1e-5$, weight decay 0.02, and batchsize 4,096 for 5 epochs. We use resolution 512 in downstream tuning. PEVL predicts binary labels indicating the whether candidate is correct, given the text of question concatenated with answer, or question, answer concatenated with rationale. In $Q \rightarrow AR$, we first predict an answer from four answer candidates, and then pick a rationale from four rationale candidates based on the predicted answer.

Baselines. We compare with strong baselines, including R2C (Zellers et al., 2019), TAB-VCR (Lin et al., 2019) and strong VLP models including VisualBERT (Li et al., 2020b), ViLBERT (Lu et al., 2019), Unicoder-VL (Li et al., 2020a), VL-BERT (Su et al., 2020), B2T2 (Alberti et al., 2019) and UNITER (Chen et al., 2020). We report the accuracy of predicting the answer ($Q \rightarrow A$), rationale ($QA \rightarrow R$) and both ($Q \rightarrow AR$).

B.5 Visual Question Answering

Datasets. GQA dataset (Hudson and Manning, 2019) is collected through automatically generating questions and answers with functional programs based on the scene graphs in Visual Genome. The dataset is split into train, validation, test-dev, and test sets, with 14,305,356, 2,011,853, 172,174 and 1,340,048 questions respectively.

Prompt Tuning. We tune the model with learning rate $1e-5$, weight decay 0.02, batchsize 256 for 5 epochs. Following MDETR (Kamath et al., 2021), the intermediate pre-training is conducted on the unbalanced train set, and prompt tuning on the balanced train set. The resolution of image is 384. We infer the answers based on the 1,853 candidates from Kamath et al. (2021).

Baselines. We compare with existing methods reported on the GQA balanced validation dataset, including LXMERT (Tan and Bansal, 2019), BAN (Kim et al., 2018), CTI (Do et al., 2019) and CFR (Nguyen et al., 2022).

C Ethical Considerations

Potential risks of this work lie in (1) privacy issues of the pre-training images and text from the Web, (2) misuse of the model (e.g., visual relation detection for monitoring human activity), and (3) toxic model outputs. The initial version of this paper is released at <https://arxiv.org/abs/2205.11169>. The picture in Figure 1 is obtained from the RefCOCO dataset.