

Weakly Supervised Vision-and-Language Pre-training with Relative Representations

Chi Chen^{1,3}, Peng Li^{* 2,4}, Maosong Sun^{* 1,3}, Yang Liu^{1,2,3,4}

¹Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China

²Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

³Beijing National Research Center for Information Science and Technology

⁴Shanghai Artificial Intelligence Laboratory, Shanghai, China

Abstract

Weakly supervised vision-and-language pre-training (WVLP), which learns cross-modal representations with limited cross-modal supervision, has been shown to effectively reduce the data cost of pre-training while maintaining decent performance on downstream tasks. However, current WVLP methods use only local descriptions of images, i.e., object tags, as cross-modal anchors to construct weakly-aligned image-text pairs for pre-training. This affects the data quality and thus the effectiveness of pre-training. In this paper, we propose to directly take a small number of aligned image-text pairs as anchors, and represent each unaligned image and text by its similarities to these anchors, i.e., relative representations. We build a WVLP framework based on the relative representations, namely RELIT¹, which collects high-quality weakly-aligned image-text pairs from large-scale image-only and text-only data for pre-training through relative representation-based retrieval and generation. Experiments on four downstream tasks show that RELIT achieves new state-of-the-art results under the weakly supervised setting².

1 Introduction

Vision-and-language pre-training (VLP) (Chen et al., 2020; Zhang et al., 2021; Kim et al., 2021; Radford et al., 2021; Wang et al., 2022a) has received increasing attention in recent years for its great success on various vision-and-language tasks, such as visual question answering (Antol et al., 2015), cross-modal retrieval (Plummer et al., 2015), and image captioning (Lin et al., 2014). Different from other foundation models (Bommasani et al., 2021) such as BERT (Devlin et al.,

2018) and MAE (He et al., 2022) that only require single-modality data, VLP models rely on large-scale aligned image-text datasets (Lin et al., 2014; Sharma et al., 2018; Ordonez et al., 2011; Krishna et al., 2017) to bridge the gap between the two modalities, which requires either extensive manual annotations or heavy data cleaning processes (Lin et al., 2014; Sharma et al., 2018). The natural difficulty of obtaining paired data hinders the scale of cross-modal datasets, while the success of unimodal pre-trained models implies the potential to exploit the unlabeled data for pre-training. Therefore, besides collecting more paired data, it is a worthwhile direction to explore how to utilize low-cost unimodal data with limited cross-modal supervision, i.e., weakly supervised vision-and-language pre-training (WVLP).

The core challenge of WVLP is to establish the connection between the two modalities without using a large number of aligned image-text pairs. Existing works on WVLP (Li et al., 2021b; Zhou et al., 2022; Wang et al., 2022b; Chen et al., 2022) usually address this by taking object tags as anchors as they are in the form of text and cover the information of the image at the same time. They use tags to collect weakly-aligned image-text pairs from unaligned unimodal data for pre-training and achieve competitive results compared to standard VLP models, demonstrating that tags can effectively bridge the gap between the two modalities.

Despite its success, using object tags as anchors suffers from two limitations. First, tags are merely local descriptions instead of a complete representation of the whole image and text. Second, the vocabulary of tags only includes common concepts, making it difficult to represent images with complex semantics (Zhou et al., 2022). These limitations could deteriorate the quality of the weakly-aligned data (and possibly pre-trained models) based on the object tags. Therefore, to further improve the performance of WVLP, we need

*Corresponding authors: Peng Li (lipeng@air.tsinghua.edu.cn) and Maosong Sun (sms@tsinghua.edu.cn)

¹Relative rEpresentation-based Language-Image pre-Training

²Code is available at <https://github.com/THUNLP-MT/RELIT>.

to reconsider the choice of the cross-modal anchors and find a better approach to measure the alignment between an image and a text.

Recently *relative representation* has been proven to be effective in representation learning (Moschella et al., 2022) and zero-shot image classification (Norelli et al., 2022). The main idea is to represent a data point as its similarities to a set of selected data points (anchors). We argue that relative representations can be a good choice for WVLP because (1) they are built on the semantic similarities of well-trained neural network representations rather than on superficial human-designed features like tags and (2) they are modality-invariant by design because they reflect the intrinsic relationships between data points, which naturally enables communication between different modalities.

In this paper, we propose RELIT, a novel relative representation-based WVLP framework. Instead of object tags, we directly use a minuscule amount (compared to pre-training data) of available image-text pairs as anchors, and create a common relative representation space with respect to the anchors for unaligned images and text. This allows us to estimate the semantic similarity of any image-text pair by calculating their distance in the relative representation space. In addition, we design two relative representation-based data collection methods that can retrieve or generate weakly-aligned image-text pairs from unaligned unimodal corpora. Experimental results prove the effectiveness of relative representations in bridging the gap between image and text modalities. Moreover, our work reveals a promising research direction to establish cross-modal alignments by finding and aligning invariant data structures in different modalities, which may inspire future works on multimodal pre-training.

Our main contributions are as follows:

- We introduce the idea of relative representations in WVLP and demonstrate its superiority over object tags in effectively bridging the gap between different modalities.
- We propose a relative representation-based WVLP framework that can both retrieve and generate weakly-aligned image-text pairs for learning cross-modal representations.
- Extensive experiments on four diverse vision-and-language tasks show that our proposed

framework outperforms strong WVLP baselines and further closes the performance gap between WVLP and standard VLP.

2 Related Work

Relative Representations. The concept of relative representations is initially proposed by Moschella et al. (2022). They show that the relative representations obtained from the representation spaces of different models are similar, which enables comparison and alignment between latent embeddings of different learning models. Norelli et al. (2022) explore relative representations in a multimodal scenario to align images and text for zero-shot image classification tasks. Specifically, they use 1.6M image-text pairs to build the relative representation space, which is comparable to the size of the data used in pre-training. To the best of our knowledge, our work is the first that exploits relative representations for weakly supervised cross-modal pre-training.

Weakly Supervised Vision-and-Language Pre-training. Li et al. (2021b) first explore WVLP with unaligned image and text corpora and use object tags directly as pseudo captions for images to bridge the vision and language modalities. Zhou et al. (2022) use tags to retrieve weakly-aligned captions for each image and then apply multi-granular alignment tasks on this retrieved dataset. Wang et al. (2022b) propose the cross-modal CutMix to replace some grounded words with regions that have the same tags, and construct a multimodal view of the text-only sentences for pre-training. Chen et al. (2022) introduce an end-to-end framework with a referring expression matching task. Different from all of these WVLP works that utilize tags as anchors to provide object-level cross-modal alignment signals, our work uses relative representations to capture the overall semantic similarity between each image and text and demonstrates its effectiveness in WVLP.

Data Augmentation. Data augmentation has been extensively employed in various computer vision (Zhang et al., 2018; Cubuk et al., 2018) and natural language processing tasks (Sennrich et al., 2015; Guo et al., 2020). In the area of VLP, Li et al. (2022) augment the noisy web-crawled aligned data by filtering low quality image-text pairs and generating synthetic captions with an image captioner fine-tuned on clean image-text pairs. In this work,

we adopt a similar filter-and-generate process in the construction of weakly-aligned data for WVLP, but our relative representation-based pseudo caption generator is fine-tuned on the text-only dataset.

3 Method

3.1 Relative Representations

Figure 1a provides an illustration of relative representations. The basic idea is to represent a data point as its similarities to other data points (anchors). In this work, we consider the relative representations with cross-modal anchors, which has been shown its potential in zero-shot image classification (Norelli et al., 2022).

Formally, given a set of M cross-modal anchors $\mathcal{A} = \{a_1, a_2, \dots, a_M\}$ where $a_i = (\tilde{x}_i, \tilde{y}_i)$ is an image-text pair, \tilde{x}_i is the image and \tilde{y}_i is the text. For an image x , a pre-trained image encoder E_I is used to calculate the similarity between x and each anchor a_i as:

$$\text{sim}(x, a_i) = \cos(E_I(x), E_I(\tilde{x}_i)) \quad (1)$$

where $\cos(\cdot, \cdot)$ is the cosine similarity, and the relative representation of x is defined as:

$$\mathbf{r}_{\mathcal{A}}(x) = (\text{sim}(x, a_1), \dots, \text{sim}(x, a_M)) \quad (2)$$

Similarly, the relative representation of a text y is defined as $\mathbf{r}_{\mathcal{A}}(y)$ with a pre-trained text encoder E_T to compute $\text{sim}(y, a_i) = \cos(E_T(y), E_T(\tilde{y}_i))$.

Since the relationship between data points is objective, the relative representations obtained by different models should be similar, despite their independent representation spaces (Moschella et al., 2022). In other word, an image and its corresponding text should share similar relative representations. This allows us to leverage it to construct weakly-aligned image-text pairs from large-scale unpaired image and text datasets.

3.2 Weakly-Aligned Image-Text Pairs Retrieval

While there are no large-scale aligned image-text pairs available, having a joint input of image and text, even if they are not aligned, is still necessary for WVLP (Zhou et al., 2022; Wang et al., 2022b). To achieve this, inspired by previous work (Zhou et al., 2022), we construct a weakly-aligned image-text corpus from the unpaired unimodal corpora by retrieving semantically related sentences for each image based on the relative representations.

Figure 1b illustrates the process of our weakly-aligned image-text pairs retrieval method. First we collect a very small amount of image-text pairs as cross-model anchors (denoted by pairs of connected squares in the figure). Note that the number of anchors is negligible compared to the image-text pairs used in standard VLP, which keeps our method in a weakly supervised setting. Then, for all images and text we compute their relative representations with respect to the anchors, which only involves similarity computation within each modality using unimodal pre-trained encoders. We take the cosine distance between the relative representations of each image and text as their semantic relevance score and retrieve the best matching text with the highest score for each image to construct a weakly-aligned image-text pair.

Specifically, we randomly sample M image-text pairs as anchors \mathcal{A} from an aligned image-text dataset (e.g., Conceptual Captions (Sharma et al., 2018)) $\mathcal{D}_{\text{align}}$ ($M \ll |\mathcal{D}_{\text{align}}|$). Given unaligned image dataset \mathcal{D}_I and text dataset \mathcal{D}_T , we construct a retrieved weakly-aligned image-text pair dataset $\mathcal{D}_{\text{wa}} = \{(x_1, \hat{y}_1), \dots, (x_N, \hat{y}_N)\}$ where $N = |\mathcal{D}_I|$ and \hat{y}_i is the retrieved caption from \mathcal{D}_T for image x_i defined as:

$$\hat{y}_i = \underset{y \in \mathcal{D}_T}{\text{argmax}} \cos(\mathbf{r}_{\mathcal{A}}(x_i), \mathbf{r}_{\mathcal{A}}(y)) \quad (3)$$

We use the off-the-shelf ViT (Dosovitskiy et al., 2020) and Sentence-BERT (Reimers et al., 2019) to encode images and text, respectively.

Our retrieval method with relative representations can effectively improve the quality of retrieved weakly-aligned dataset compared to tag-based retrieval. This is because relative representations tend to capture the overall semantics while tags describe only local information of the image. As a result, our method can better measure the semantic similarities between images and text, especially in cases where tag-based retrieval fails to distinguish between images and text that have different semantics but share the same objects.

3.3 Pseudo Caption Generation

Although relative representation-based retrieval can construct reasonable weakly-aligned image-text pairs for WVLP, there are still cases where non-relevant text are retrieved. This could happen especially when the unaligned unimodal corpora are collected individually and for some images there are no proper captions in the corpora.

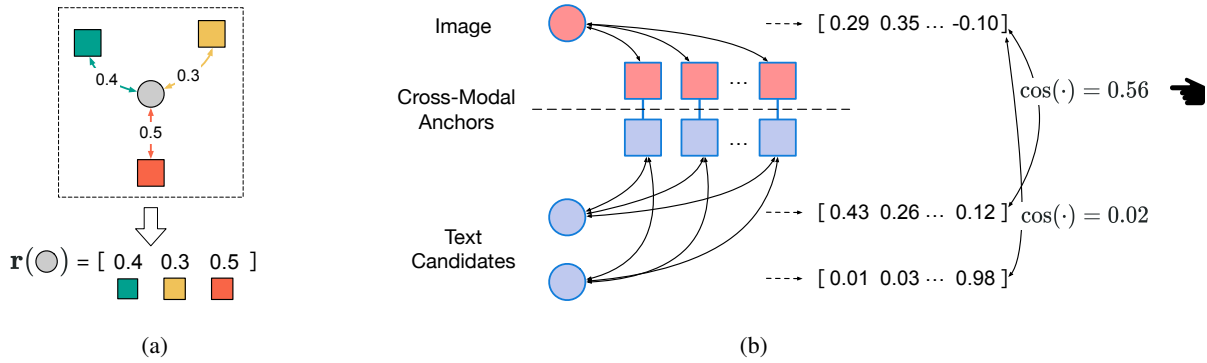


Figure 1: (a) Illustration of relative representations (Section 3.1), where three anchors (denoted by squares) are selected and the relative representation of the data point (denoted by circles) is a 3D vector with each dimension representing its similarity to the corresponding anchor. (b) Image-text retrieval based on the relative representations with cross-modal anchors (Section 3.2). Data of the same modality are represented by the same color.

To alleviate this problem, we propose to directly generate pseudo captions for these images. As shown in Figure 2, we first adapt a well-trained text generator to perform conditional text generation given relative representations. Then, since images and text share a common relative representation space, we can directly use this generator to predict the pseudo caption for an image based on its relative representation.

Specifically, given the text-only dataset \mathcal{D}_T , for each text $y \in \mathcal{D}_T$, we derive a prefix $\mathbf{P} \in \mathbb{R}^{M \times d}$ from its relative representations $\mathbf{r}_{\mathcal{A}}(y)$ as:

$$\mathbf{P} = [\mathbf{r}_{\mathcal{A}}(y)]^T \mathbf{W}_r + [E_T(\tilde{y}_1), \dots, E_T(\tilde{y}_M)] \mathbf{W}_e \quad (4)$$

where $E_T(\tilde{y}_i) \in \mathbb{R}^{d_T}$ is the encoder output of the text in the i -th anchor, $\mathbf{W}_r \in \mathbb{R}^{1 \times d}$ and $\mathbf{W}_e \in \mathbb{R}^{d_T \times d}$ are two learnable projection matrices. We fine-tune a pre-trained GPT-2 model (Radford et al., 2019) to learn to predict y given \mathbf{P} , and name the fine-tuned model as *Rel2Cap*. To further save computational cost, we only consider the entries in \mathbf{P} that correspond to the top K anchors with the highest similarities as the model input.

After training, the model can be used to predict the pseudo caption for an image x with low quality retrieved captions by constructing an input prefix \mathbf{P}' based on the relative representations of the image, i.e., $\mathbf{r}_{\mathcal{A}}(x)$. The definition of \mathbf{P}' is similar to Equation 4, except that $\mathbf{r}_{\mathcal{A}}(y)$ is replaced by $\mathbf{r}_{\mathcal{A}}(x)$. We define a quality score $s(x, \hat{y}) = \cos(\mathbf{r}_{\mathcal{A}}(x), \mathbf{r}_{\mathcal{A}}(\hat{y}))$ for each weakly-aligned image-text pair (x, \hat{y}) collected both by retrieval and generation, and replace the retrieved pair with the generated one if the latter has a higher quality score.

So far, we have discussed how we collect a

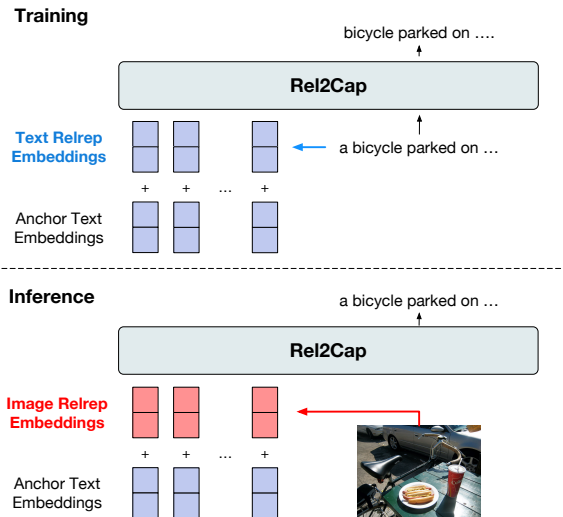


Figure 2: An illustration of the training and inference of the pseudo caption generator. In the training process, the model learns to generate text from its relative representation on the text-only dataset. During inference, the model is directly employed to predict the pseudo caption for an image from its relative representation.

weakly-aligned image-text dataset \mathcal{D}_{wa} from the unpaired unimodal corpora by relative representation-based retrieval and generation. Next, we describe how we use these data for WVLP.

3.4 Pre-training

Model Overview. We use the same model architecture as Chen et al. (2022) that consists of a vision and a multimodal encoder. For each weakly-aligned image-text pair, the image is encoded with the vision encoder and the outputs are fed to the multimodal encoder along with the text embeddings to obtain a multimodal representation. Such an end-to-end framework has been proven

to be more effective compared to others that use region features from external object detectors both in standard VLP and WVLP. We apply three pre-training objectives to learn multimodal representations from the collected weakly-aligned image-text pairs: masked tag prediction (MTP), masked language modeling (MLM) and image text matching (ITM).

Masked Tag Prediction. This objective aims to learn object-level cross-modal alignment from the image-only data and their detected object tags. Following previous works (Li et al., 2021b; Chen et al., 2022), we randomly mask out the tags with a probability of 15%, and then predict the masked tags conditioned on the image and other unmasked tags. Formally, given the image $x \in \mathcal{D}_I$ and its detected object tags t , the MTP objective is defined as:

$$\mathcal{L}_{\text{MTP}} = -\mathbb{E}_{x \in \mathcal{D}_I} \log P(t_m | t_{\setminus m}, x) \quad (5)$$

where t_m and $t_{\setminus m}$ represents masked and unmasked object tags, respectively.

Masked Language Modeling. To better fuse between the two modalities, the masked language modeling objective is adopted to learn from the joint image-text inputs from the weakly-aligned corpora. Since the weakly-aligned pairs may contain noise in the retrieved or generated text, we only mask out and predict the noun phrases in the text inspired by (Zhou et al., 2022). The MLM loss is formulated as:

$$\mathcal{L}_{\text{MLM}} = -\mathbb{E}_{(x, \hat{y}) \in \mathcal{D}_{\text{wa}}} \log P(\hat{y}_m | \hat{y}_{\setminus m}, x) \quad (6)$$

where \hat{y}_m and $\hat{y}_{\setminus m}$ are masked and unmasked text.

Image Text Matching. ITM is a commonly used objective for learning instance-level cross-modal alignment in VLP, which aims to distinguish whether an image-text pair is matched semantically. We random replace the text in half of the image-text pairs with another text to form training input, and define the label of each pair as $l \in \{0, 1\}$ where 1 indicates the pair is a match. The ITM objective is to minimize the binary cross-entropy loss:

$$\mathcal{L}_{\text{ITM}} = -\mathbb{E}_{(x, \hat{y}) \in \mathcal{D}'_{\text{wa}}} \log P(l | x, \hat{y}) \quad (7)$$

where \mathcal{D}'_{wa} is the dataset after random replacement.

Relative Representation-Guided Training. To further reduce the impact of the noisy image-text pairs in the weakly-aligned dataset, we apply the

quality score $s(x, \hat{y})$ of each pair described in Section 3.3 to \mathcal{L}_{MLM} and \mathcal{L}_{ITM} to guide the training to learn more from high-quality data:

$$\mathcal{L}_{\text{MLM}} = -\mathbb{E}_{(x, \hat{y}) \in \mathcal{D}_{\text{wa}}} s(x, \hat{y}) \log P(\hat{y}_m | \hat{y}_{\setminus m}, x) \quad (8)$$

$$\mathcal{L}_{\text{ITM}} = -\mathbb{E}_{(x, \hat{y}) \in \mathcal{D}'_{\text{wa}}} s(x, \hat{y}) \log P(l | x, \hat{y}) \quad (9)$$

4 Experiments

4.1 Datasets

We follow previous WVLP works (Li et al., 2021b; Zhou et al., 2022; Wang et al., 2022b; Chen et al., 2022) and conduct experiments in two different settings, each containing an image-only dataset and a text-only dataset. The first setting treats images and text from Conceptual Captions (CC) (Sharma et al., 2018) as individually collected unimodal dataset without the alignment information. The second setting uses images from CC and text from BookCorpus (Zhu et al., 2015), which is a more realistic scenario where images and text are gathered separately from different sources.

4.2 Implementation Details

Relative Representations. We randomly select 8,192 aligned image-text pairs from CC as anchors, yielding relative representations as vectors of 8,192 dimensions. To save computational cost, inspired by Norelli et al. (2022), we only keep the highest 50 dimensions and set the others to 0.

Weakly-Aligned Data Construction. We implement the retrieval system with the faiss (Johnson et al., 2019) library. For each image we only retrieve the text with the best match score. For Rel-Cap, we fine-tune GPT-2 with a learning rate of $5e-5$ and a batch size of 1,024 for 5 epochs on the text-only dataset. We generate 5 pseudo-captions for each image using nucleus sampling with $p = 0.9$ which proved effective in synthetic caption generation (Li et al., 2022), and rank the results with the quality scores. We also include the weakly-aligned dataset based on tag-based retrieval in the pre-training, as described in Zhou et al. (2022).

Pre-training. We use the same architecture as Chen et al. (2022) which includes a 12-layer Swin-Transformer (Swin B-384/32) (Liu et al., 2021) as the vision encoder and a 12-layer Transformer initialized from BERT-base (Devlin et al., 2018) as the multimodal encoder. For object tags, we

Model	VQAv2	NLVR ²	VE	Flickr30k		
	Test-Dev	Test-P	Test	R@1	R@5	R@10
<i>Supervised (w/ Large-Scale Paired Image-Text Data)</i>						
VisualBERT (Li et al., 2019)	70.9	73.9	-	61.2	86.3	91.9
UNITER (Chen et al., 2020)	72.7	77.9	78.3	72.5	92.4	96.1
VinVL (Zhang et al., 2021)	76.0	83.1	-	-	-	-
ViLT (Kim et al., 2021)	71.3	76.1	-	66.4	88.7	93.8
ALBEF (Li et al., 2021a)	74.5	80.5	80.3	82.8	96.7	98.4
METER-CLIP-ViT _{BASE} (Dou et al., 2022)	77.7	83.0	81.2	82.2	96.3	98.3
<i>Weakly Supervised (w/o Large-Scale Paired Image-Text Data)</i>						
U-VisualBERT (Li et al., 2021b)	70.7	71.0	-	55.4	82.9	89.8
U-VisualBERT _{VinVL} (Zhou et al., 2022)	71.8	53.2	76.8	-	-	-
μ -VLA (Zhou et al., 2022)	72.1	73.4	77.3	-	-	-
VL Mixer (Wang et al., 2022b)	72.7	73.9	-	-	-	-
E2E-UVLP (Chen et al., 2022)	73.3	74.6	78.2	66.4	89.7	94.1
RELIT (Ours)	73.5	76.4	78.6	70.2	91.5	95.6

Table 1: Evaluation results on four V+L downstream tasks. All weakly-supervised models are pre-trained on non-parallel images and text from CC.

Method	VQAv2	NLVR ²	VE	Flickr30k		
	Test-Dev	Test-P	Test	R@1	R@5	R@10
U-VisualBERT	70.5	71.2	-	54.4	82.2	89.2
μ -VLA	71.2	67.1	77.1	-	-	-
E2E-UVLP	73.5	73.7	77.9	65.6	90.3	94.7
RELIT (Ours)	73.6	74.8	78.2	67.7	90.4	95.0

Table 2: Experimental results on downstream tasks of pre-training with images from CC and text from BookCorpus.

utilize the off-the-shelf object detector provided by VinVL (Zhang et al., 2021). We pre-train the model with a total training step of 150k and a batch size of 512. We use an AdamW optimizer (Kingma and Ba, 2014) with an initial learning rate of $3e-5$, and the warm-up ratio is set to 10%. The pre-training takes 3 days on 4 NVIDIA A100 GPUs.

Downstream Tasks. We follow previous works and test our pre-trained model on four downstream V+L tasks, including Visual Question Answering (VQAv2) (Goyal et al., 2017), Natural Language for Visual Reasoning (NLVR²) (Suhr et al., 2018), Visual Entailment (VE) (Xie et al., 2019) and image retrieval (Flickr30k) (Plummer et al., 2015). Details of the task settings and the fine-tuning strategies are in Appendix A.

4.3 Main Results

We first compare our proposed RELIT with previous methods pre-trained with unaligned images and text from CC. Note that these baselines only utilize object tags. Table 1 shows the experimental results on the downstream tasks. Our method outperforms previous WVLP methods on all downstream tasks.

Specifically, RELIT outperforms previous best results by 1.8% on NLVR² and by 3.8% on the image retrieval task (Flickr30k), both of which benefit from the instance-level cross-modal alignment capability of the pre-trained model (Chen et al., 2020; Zhou et al., 2022). This suggests that our relative representation-based method improves the alignment quality of weakly-aligned image-text pairs compared to previous tag-based approaches, resulting in improved cross-modal alignment capability of the pre-trained model.

When pre-trained with images from CC and text from BookCorpus, as shown in Table 2, our proposed RELIT also achieves the best results on all downstream tasks. This demonstrates that the proposed relative representation-based methods can effectively mine useful cross-modal alignment information for multimodal pre-training from image-only and text-only data, even if they are collected separately from different sources.

4.4 Ablation Study

We conduct an ablation study to verify the effectiveness of the proposed relative representation-based retrieval and generation. Table 3 shows the results.

Pre-training Data	VQAv2	NLVR ²	VE	Flickr30k		
	Test-Dev	Test-P	Test	R@1	R@5	R@10
Retrv (Tag)	73.2	74.5	77.8	66.3	89.3	94.2
Retrv (Relrep)	73.4	74.9	78.3	67.5	90.5	94.9
Retrv (Tag) + Retrv (Relrep)	73.5	75.3	78.4	67.3	90.4	94.6
Retrv (Tag) + Retrv (Relrep) + Rel2Cap	73.5	76.4	78.6	70.2	91.5	95.6

Table 3: Comparison of pre-training with different kinds of pseudo-aligned data.

Method	VQAv2	NLVR ²	VE	Flickr30k
	Test-Dev	Test-P	Test	R@1
RELIT	73.5	76.4	78.6	70.2
- Guided	73.2	76.1	78.5	70.0

Table 4: Ablation study on relative representation-guided training.

All models are pre-trained on weakly-aligned data derived from unaligned CC images and text. As we can see from the table, compared to tag-based retrieved data (Retrv (Tag)), pre-training with relative representation-based retrieved data (Retrv (Relrep)) performs better on downstream tasks. Besides, the model achieves the best results when the generated pseudo captions (Rel2Cap) are included during pre-training. We believe this is because the original CC dataset contains noisy captions, such as alt-texts that do not describe the image contents, which is suboptimal for VLP (Li et al., 2022). In summary, the experimental results demonstrate that both our retrieval and generation methods contribute to the performance of the pre-training.

We also compare the performance of the pre-trained models on downstream tasks with and without relative representation-guided training. As shown in Table 4, pre-training with guided training can consistently improve results across all downstream tasks, illustrating that relative representations can be used to detect noise in the weakly-aligned data and guide the model to learn from data with a higher level of alignment.

4.5 Data Quality

We evaluate the quality of different kinds of weakly-aligned data from unaligned CC images and text, and the results are listed in Table 5. We use CLIP-Score (Hessel et al., 2021) to measure the overall alignment of all weakly-aligned image-text pairs. As we can see from the table, the data quality of Retrv (Relrep) is significantly higher than that of Retrv (Tag), which again illustrates the superiority of relative representations as cross-modal anchors.

Data	CLIPScore
Retrv (Tag)	57.94
Retrv (Relrep)	63.31
+ Rel2Cap	65.23

Table 5: Data quality of different kinds of weakly-aligned data from unaligned CC images and text.

In addition, Rel2Cap further improves data quality by filtering and replacing low-quality pairs in Retrv (Relrep). The analysis of the data quality here is consistent with the analysis of pre-training results in Table 3, and again proves that our relative representation-based methods can produce high quality weakly-aligned data from unaligned unimodal data.

4.6 Effects of Anchor Selection

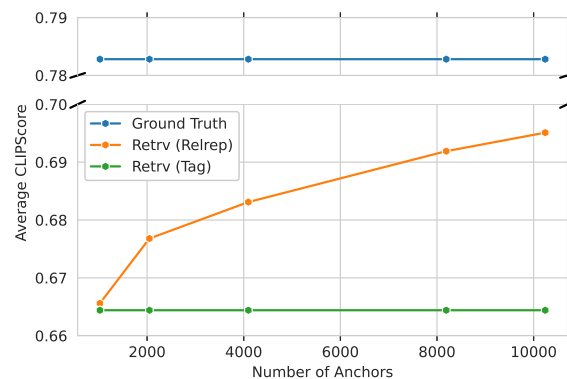


Figure 3: Data quality of the retrieved data using different number of anchors. Both the anchors and the images and text used for retrieval are from the COCO dataset.

The number of anchors has a significant influence on the effect of relative representations (Norelli et al., 2022). To verify its influence on the collected weakly-aligned image-text pairs, we test the quality of the data retrieved with different numbers of anchors on the COCO (Lin et al., 2014) dataset. From Figure 3, we can see that as the number of anchors increases, the quality of the retrieved data also improves. In addition, we evaluate

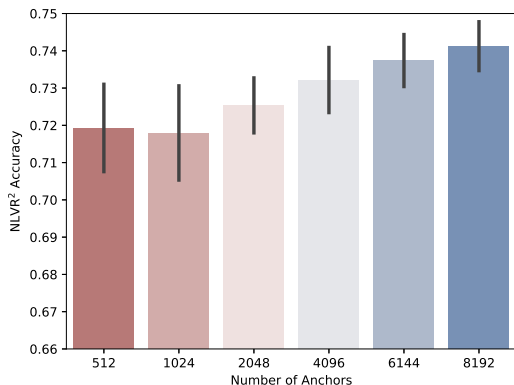


Figure 4: Fine-tuned NLVR² results of models pre-trained on data with different number of anchors.

Method	Number of Anchors			
	1024	2048	4096	8192
random	66.50	67.68	68.31	69.19
diverse	67.66	68.34	68.60	69.37
non-diverse	40.94	39.52	39.53	39.50

Table 6: The quality of data obtained from different anchor selection methods.

the downstream task accuracy using models pre-trained on data with varying numbers of anchors. Specifically, we generate 3 random sets of anchors for each size, and retrieve the weakly-aligned data with different sets of anchors. We pre-train models on each set of the retrieved data with the same hyperparameters, and fine-tune them on the NLVR² task. The results are shown in Figure 4. In general, the higher the number of anchors, the better the model performance. We use 8,192 anchors in our final experiments as a trade-off between representation capability and computational cost. However, using more anchors will almost certainly give better results due to better quality of the data, which indicates the scalability of our approach. We leave more exploration on this for future work.

We also conduct experiments to verify the impact of anchor diversity on data quality. Specifically, we considered three sampling methods on the COCO dataset: *random*, *diverse*, and *non-diverse*. The diverse sampling first performs K-means clustering on all the data, and selects one anchor from each cluster. The non-diverse sampling uses a greedy algorithm to select k anchors, at each step choosing the data closest to the average of the anchors already selected. Table 6 lists the data quality results obtained with different sam-

pling methods. In general, diverse anchors lead to better quality, while random anchors perform satisfactorily when the number of anchors is large enough. Non-diverse anchors can result in catastrophic data quality.

4.7 Case Study

To explore the reasons for the improvement in data quality, we show two examples of the comparisons between different weakly-aligned image-text pairs in Figure 6. In each example, we provide the ground truth caption of the image and the detected object tags, as well as three weakly-aligned captions. From these two examples, we can see that the captions retrieved by tags do have many of the same tags as the images (underlined in the figure), but are not good descriptions of the images. In contrast, our relative representation-based retrieval and generation methods are able to obtain captions that are more relevant to the overall semantics of the images. Specifically, in the example in Figure 6a, our proposed methods successfully identifies key information in the image such as “golfer”, which is difficult for tag-based retrieval since there are no such tag as “golfer”. The same thing happens to Retrv (Tag) in Figure 6b, which retrieves a caption related to “cat” instead of “lynx”. In this example, our retrieval method recognizes the animal in the image as “cheetah”, which is close but not exactly correct, while our generation method correctly generates a caption related to the correct concept “lynx”. This indicates that our generation method has the ability to generate pseudo captions of better quality when the retrieved ones are not good enough.

In Figure 5 we further visualize the relative representations of the image and two retrieved captions in Figure 6a, which helps understand the effectiveness of relative representations in aligning semantically related image-text pairs. From the figure we can see that the image and our retrieved caption Retrv (Relrep) activate the same group of anchors (i.e., have high similarities with these anchors), which makes them close in the relative representation space. On the other hand, Retrv (Tag) activates a completely different set of anchors, which leads to a large distance between it and the image in the relative representation space. These observations suggest that (1) relative representations are (almost) modality-invariant and (2) relative representations can be utilized to effec-



Figure 5: Comparison of the relative representations of the image and retrieved captions in Figure 6a. For simplicity, for each image and text on the left, we only display the anchors with the highest similarities on the right.



Ground Truth:
the couple share a love of the game .


Detected Tags:
collar, man, hair, nose, shirt, eye, ear, head ...

Retrv (Tag):
close up head shot of a small white fluffy long haired dog with a black nose , dark round eyes , black lips and a green ribbon.

Retrv (Relrep): golfer walks off the tee on the 18th hole with his caddie , during the first round .

Rel2Cap: the former star has been urging officials to closely track down the golfer.

(a)



Ground Truth:
lynx with a satisfied smile .

Detected Tags:
cat , ear , nose , eye , head, fur, leg, fence, spot , ground ...

Retrv (Tag):
a black and white cat with large green eyes and a spot next to nose

Retrv (Relrep):
cheetahs are the only big cat that can not roar .

Rel2Cap:
the lynx is a type of long - legged animal with muscular, bushy coat

(b)

Figure 6: Examples of different kinds of weakly-aligned data. We highlight in red the caption with the best quality and the words in it that match the key information of the image. Compared to Retrv (Tag) which focuses on tag matching (underlined), our proposed two methods Retrv (Tag) and Rel2Cap produce captions that are more semantically similar to the image.

tively estimate the cross-modal alignment of data in different modalities. These properties of the relative representations make it naturally suitable for WVLP, which is verified in this paper.

5 Conclusion

This paper introduces the idea of relative representations to weakly-supervised vision-and-language pre-training and demonstrates its effectiveness in bridging the gap between the two modalities. We propose a relative representation-based framework that can both retrieve and generate weakly-aligned image-text pairs for pre-training. Experimental results show that our method outperforms all previous tag-based approaches under the weakly-supervised setting. We hope our work will motivate future work in multimodal pre-training.

Limitations

As this work is mainly focused on weakly supervised vision-and-language pre-training, we do not fully explore the factors that may influence the performance of relative representations, such as the use of different unimodal encoders and the source of the anchors. Besides, we only validate the effectiveness of relative representations in a weakly supervised setting, while it remains to be explored whether it is also useful for standard VLP and multimodal learning in other modalities (e.g., audio and video). We will further exploit the potential of relative representations and validate it in more cross-modal learning scenarios in the future.

Acknowledgments

This work is supported by the National Key R&D Program of China (2022ZD0160502) and the National Natural Science Foundation of China (No. 61925601, 62276152, 62236011). We thank Ziyue Wang, Fuwen Luo, Rui Jiao and Zonghan Yang for their advices in paper writing.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (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*.
- Chi Chen, Peng Li, Maosong Sun, and Yang Liu. 2022. End-to-end unsupervised vision-and-language pre-training with referring expression matching. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10799–10810.
- 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 *ECCV*.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2018. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. 2022. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18166–18176.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Demi Guo, Yoon Kim, and Alexander M Rush. 2020. Sequence-level mixed sample data augmentation. *arXiv preprint arXiv:2011.09039*.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPScore: a reference-free evaluation metric for image captioning. In *EMNLP*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. ViLT: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- 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*, 123(1):32–73.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021a. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Liunian Harold Li, Haoxuan You, Zhecan Wang, Alireza Zareian, Shih-Fu Chang, and Kai-Wei Chang. 2021b. Unsupervised vision-and-language pre-training without parallel images and captions. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5339–5350.

- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin Transformer: Hierarchical vision Transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022.
- Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodolà. 2022. Relative representations enable zero-shot latent space communication. *arXiv preprint arXiv:2209.15430*.
- Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodolà, and Francesco Locatello. 2022. ASIF: Coupled data turns unimodal models to multimodal without training. *arXiv preprint arXiv:2210.01738*.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2Text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24.
- 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 the IEEE international conference on computer vision*, pages 2641–2649.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Nils Reimers, Iryna Gurevych, Nils Reimers, Iryna Gurevych, Nandan Thakur, Nils Reimers, Johannes Daxenberger, Iryna Gurevych, Nils Reimers, Iryna Gurevych, et al. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 671–688. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 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)*, pages 2556–2565.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2018. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*.
- 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 *International Conference on Machine Learning*, pages 23318–23340. PMLR.
- Teng Wang, Wenhao Jiang, Zhichao Lu, Feng Zheng, Ran Cheng, Chengguo Yin, and Ping Luo. 2022b. Vlmixer: Unpaired vision-language pre-training via cross-modal cutmix. In *International Conference on Machine Learning*, pages 22680–22690. PMLR.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*.
- Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6281–6290.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*.
- 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 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588.
- Mingyang Zhou, Licheng Yu, Amanpreet Singh, Mengjiao Wang, Zhou Yu, and Ning Zhang. 2022. Unsupervised vision-and-language pre-training via retrieval-based multi-granular alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16485–16494.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.

A Details of Downstream Tasks

Visual Question Answering (VQA) The task of VQA is to answer questions correctly according to the given images. We follow previous works (Yu et al., 2019; Chen et al., 2020) and formulate VQA as a classification task with 3,192 classes representing the most frequent answers in the dataset. We fine-tune the pre-trained model for 10 epochs with a batch size of 256. We use an AdamW optimizer with a peak learning rate of 5×10^{-5} .

Natural Language for Visual Reasoning (NLVR²) The objective of NLVR² is to decide if a natural language description is true for a given pair of images. We follow previous work (Chen et al., 2020) to form two image-text pairs as inputs, and concatenate the two [CLS] outputs of the model as the final representation for classification. We fine-tune the model for 10 epochs with a batch size of 128 and a peak learning rate of 2.5×10^{-5} .

Visual Entailment (VE) Given an image and a text hypothesis, the task of VE is to determine whether the image implies the hypothesis. This is formulated as a three-way classification task to predict whether the logical relationship between the image and the text is *entailment*, *neutral* or *contradiction*. For the VE task, we fine-tune the pre-trained model with a batch size of 64 and a peak learning rate of 1×10^{-5} for 5 epochs.

Image Retrieval (Flickr30k) We follow previous works (Li et al., 2021b; Chen et al., 2022) to conduct the image retrieval task on the Flickr30k (Plummer et al., 2015) dataset. We sample 15 negative image-text pairs for each positive pair by replacing its text with randomly sampled ones. The batch size is set to 512. We fine-tune the model with a peak learning rate of 2.5×10^{-5} for 10 epochs.

B Additional Examples

In Figure 7, we provide more examples of different kinds of weakly-aligned image-text pairs. From these examples, we can see that our relative representation-based approaches yield higher quality weakly-aligned image-text pairs compared to tag-based retrieval.



Ground Truth:
a hike into an ancient forest outside a city

Detected Tags:
tree, fence, path, flower, bush

Retrv (Tag):
a wooden fence and field full of flowers

Retrv (Relrep): **movement** along the **trail** in the pine **forest**

Rel2Cap: the removal of the temporary pole created an open natural pathway to the woods which connects us state to the southeast.




Ground Truth:
happy man with a dog

Detected Tags:
dog, glove, nose, tree, head, smile, coat, ...

Retrv (Tag):
black , white , silver , and gray christmas tree .

Retrv (Relrep): a shepherd dog briefly stops while playing and running

Rel2Cap: the young **skier** helps his **dog** remove its fur as part of the massive snowfall.




Ground Truth:
couple buying a bottle of red wine in a supermarket

Detected Tags:
hair, man, bottle, jacket, woman, bar, shelf, table, ...

Retrv (Tag):
photo of the shelves , with colourful bottles of drink

Retrv (Relrep): company of **people** at a bar : beautiful **women** drinking **wine** while bartender is pouring beer for a **man**

Rel2Cap: the conversation is always online with our purchase of wine by person.



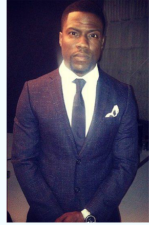
Ground Truth:
airliner being prepared for flight sitting on the tarmac

Detected Tags:
airplane, van, cloud, vehicle, person, sky, logo, tree ...

Retrv (Tag):
cargo and passenger white van on the road of fuzzy flowers and grass in front and trees, mountains and sky

Retrv (Relrep): aircraft handling at an airport

Rel2Cap: the **plane** is being **serviced** at an **airport**



Ground Truth:
comedian was all business in a sleek suit on the set of a photo shoot .

Detected Tags:
man, tie, watch, shirt, nose, ear, eye, suite, face, ...

Retrv (Tag):
watches to wear with a black , white and beige **outfit**

Retrv (Relrep):
center-of-attention : all eyes were on the rapper at the event

Rel2Cap:
the acclaimed composer shakes hands with his new partner after he is introduced to them on wednesday



Ground Truth:
person standing on the out - stretched arms of the crowd

Detected Tags:
hand, man, arm, leg, person, face, crowd ...

Retrv (Tag):
extended arm with the hand up on white background

Retrv (Relrep): cheerful young **man dancing** , moving **arms** up in the air , isolated on black background

Rel2Cap: the teenager skates around the rings to the live band at a member's retreat.




Ground Truth:
beautiful girl in hat riding a horse in prairie at sunset in slow motion

Detected Tags:
horse, tail, hat, girl, shirt, ear, head, grass, ...

Retrv (Tag):
colored hair and a bun on girl 's head

Retrv (Relrep): **person** enjoying time with the **horses**

Rel2Cap: the herd of horses is moving south while his faithful man is resting near the horizon.



Ground Truth:
hug : person hugs a statue

Detected Tags:
building, shirt, man, statue, bus, car, woman, ground, park, ...

Retrv (Tag):
colorful buses standing on the parking lot of a bus station

Retrv (Relrep):
the statue is torn down .

Rel2Cap:
the **man** is kissing a **statue**, passing the people near the town center.



Ground Truth:
scientist examining flora in a greenhouse

Detected Tags:
hair, hat, head, person, rabbit, ground, ...

Retrv (Tag):
rabbit or person head isolated on white background .

Retrv (Relrep): person is a professional who works on agriculture , including the cultivation of plants in gardens and fields

Rel2Cap: the senior **scientist** carrying out an **experiment** in a **greenhouse**



Ground Truth:
little boy , little girl , woman and man crop small rice to plan in another field on a sunny day .

Detected Tags:
hat, man, shirt, cloud, person, sky, tree, boy ...

Retrv (Tag):
plus size male standing wearing a blue shirt and black trousers with a tree as background and some colourful

Retrv (Relrep): woman cutting grass in between the rice plants in a paddy field .

Rel2Cap: the **woman** is carefully **cropping rice** in the **field** with a **team** of young hikers

Figure 7: Examples of different kinds of weakly-aligned data. We highlight in red the caption with the best quality and the words in it that match the key information of the image.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
6
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Left blank.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4

- B1. Did you cite the creators of artifacts you used?
4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
4
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
4
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
We only use publicly available datasets that are widely used in previous vision-and-language pre-training works. Relevant statistics have been fully discussed and can be easily found in these works.

C Did you run computational experiments?

4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.