

# How to Understand “Support”? An Implicit-enhanced Causal Inference Approach for Weakly-supervised Phrase Grounding

Jiamin Luo\*, Jianing Zhao\*, Jingjing Wang†, Guodong Zhou

School of Computer Science and Technology, Soochow University, China  
No.1, Shizi Street, Suzhou City, Jiangsu Province, China  
{20204027003, jnzhao1106}@stu.suda.edu.cn, {djingwang, gdzhou}@suda.edu.cn

## Abstract

Weakly-supervised Phrase Grounding (WPG) is an emerging task of inferring the fine-grained phrase-region matching, while merely leveraging the coarse-grained sentence-image pairs for training. However, existing studies on WPG largely ignore the implicit phrase-region matching relations, which are crucial for evaluating the capability of models in understanding the deep multimodal semantics. To this end, this paper proposes an Implicit-Enhanced Causal Inference (IECI) approach to address the challenges of modeling the implicit relations and highlighting them beyond the explicit. Specifically, this approach leverages both the intervention and counterfactual techniques to tackle the above two challenges respectively. Furthermore, a high-quality implicit-enhanced dataset is annotated to evaluate IECI and detailed evaluations show the great advantages of IECI over the state-of-the-art baselines. Particularly, we observe an interesting finding that IECI outperforms the advanced multimodal LLMs by a large margin on this implicit-enhanced dataset, which may facilitate more research to evaluate the multimodal LLMs in this direction.

**Keywords:** Weakly-supervised Phrase Grounding, Implicit Phrase-Region Matching, Causal Inference

## 1. Introduction

Phrase Grounding (PG) (Wang et al., 2019), a fundamental task in the field of multimodal learning, aims to find all the regions within an image that correspond to various phrases present in a given sentence. This correspondence serves as a fundamental foundation for numerous vision-language tasks, including image captioning (Feng et al., 2019), vision question answering (Mun et al., 2018), and visual dialog (Guo et al., 2020). However, PG heavily relies on expensive annotations of linking phrases to the corresponding image regions, which is labor-intensive and time-consuming. Thus, existing studies on PG mainly seek to address this task in the Weakly-supervised Phrase Grounding (WPG) (Chen et al., 2018) setting which merely leverages coarse-grained sentence-image pairs during training while subsequently evaluates the performance on fine-grained phrase-region pairs, achieving substantial advancements.

Despite this, these studies fail to deeply explore the semantic nature of phrases that some phrases often exhibit the implicit and intricate semantics, rendering it arduous for models to establish correct connections with image regions. Take Figure 1 as an example, the phrase “support” necessitates the integration of the commonsense knowledge to precisely find its corresponding image region, i.e., red boxes. In this study, we refer to such correspondence as one type of implicit phrase-region match-

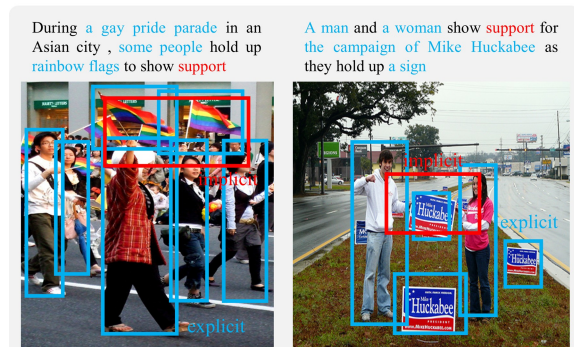


Figure 1: Two sentence-image pairs to illustrate the implicit (red phrases and boxes) and explicit (blue phrases and boxes) relations between phrases and regions.

ing relations<sup>1</sup>, which could be defined as phrases in the sentence that are not specific and explicit nominal phrases of the objects in the image, and the grasp of such implicit relations serves as a valuable evaluation of the model capability in understanding the deep multimodal semantics. Additionally, the weakly-supervised setting of WPG could bring the supervised noise problem (Xiao et al., 2017) which may lead to more difficulties in capturing such implicit information, making it urgent to address the implicit relations problem in WPG. In this study, we contend that capturing such relations at least faces two main challenges.

<sup>1</sup>In this paper, this commonsense-involved implicit relation is named as “commonsense understanding”. Besides, we also propose another three implicit relations as shown in Figure 3.

\*Equal Contribution.

†Corresponding Author: Jingjing Wang.

For one thing, we argue that modeling the implicit relations is challenging. Still take the implicit relation sample in Figure 1 as the example, the implicit phrase “support” actually only corresponds to a special small-scale red-box region (i.e., the people’s arms with the held flags or signs), which is easily confounded by the other common regions and thus difficult to be predicted. Fortunately, in the literature, a few recent studies on other tasks (e.g., Wang et al., 2020b) have also encountered with the similar confounding bias problem and suggested to leverage the causation-based approach for mitigating this confounding bias. Inspired by this, we believe that a well-behaved approach to WPG should take advantage of the causation-based approach (e.g., causal inference (Pearl and Mackenzie, 2018)) to model the implicit relations.

For another, we argue that highlighting the implicit relations beyond the explicit is rather challenging. As exemplified in Figure 1, it is obvious that the occurrence ratio of implicit relations (i.e., “support”) is significantly lower (about 1:9, see Section 4.1) compared to explicit relations (i.e., “some people”, “rainbow flags”, “a woman”, “a sign”), which could mislead the model to prefer capturing the explicit phrase-region relations instead of the implicit. Therefore, we believe that a better-behaved causation-based approach to WPG should further consider this imbalance for better aligning the implicit phrase-region pairs.

To tackle the aforementioned challenges, this paper proposes a causation-based approach namely **Implicit-Enhanced Causal Inference (IECI)** for WPG. Specifically, this approach first leverages the intervention technique (Pearl and Mackenzie, 2018) and proposes an implicit-aware deconfounded attention (IDA) block to model the implicit relations for mitigating the confounding bias. Furthermore, this approach leverages the counterfactual technique (Pearl and Mackenzie, 2018) and proposes an implicit-aware counterfactual inference (ICI) block to highlight the implicit relations beyond the explicit for better aligning the implicit phrase-region pairs. Particularly, a high-quality implicit-enhanced dataset is annotated for benchmarking our IECI approach. The main contributions of our work are summarized as follows:

- We are the first to address the implicit relations problem in WPG, and annotate a high-quality implicit-enhanced dataset to evaluate the ability of models in understanding deep multimodal semantics.
- We propose a new IECI approach, which integrates both the intervention and counterfactual techniques for addressing the implicit challenges inside WPG. Detailed evaluations on our implicit-enhanced dataset demonstrate the superiority of our IECI approach over the state-

of-the-art baselines.

- We observe an interesting finding that our IECI approach exhibits significant advantages compared to the advanced multimodal LLMs on the annotated implicit-enhanced dataset, which may further facilitate the evaluation of multimodal LLMs in this direction.

## 2. Related Work

### 2.1. Phrase Grounding

Phrase Grounding (PG) and Referring Expression Comprehension (REC) (Kazemzadeh et al., 2014) are two prevalent tasks within the field of Visual Grounding (VG) (Rohrbach et al., 2016). While both tasks involve establishing relations between sentence phrases and image regions, PG focuses on predicting regions for all phrases within a sentence-image pair, whereas REC pertains to identifying a single region in the image corresponding to the given sentence. PG can be broadly categorized into two forms, i.e., one-stage models (Yang et al., 2022; Deng et al., 2021) and two-stage models (Chen et al., 2021; Liu et al., 2020). Recognizing the expensive and difficult annotation of phrase-region, recent studies have predominantly shifted toward WPG. Early studies (Chen et al., 2018; Zhao et al., 2018) primarily focus on directly learning the phrase-region relations. Besides, some studies (Liu et al., 2019a,b, 2021) acknowledge the importance of context cues, which exploit linguistic contexts to enforce cross-modal consistency. Moreover, Datta et al. (2019) introduce a ranking-loss to minimize the distances between associated sentence-image and maximize the distance between irrelevant pairs, inspiring numerous studies (Gupta et al., 2020; Wang et al., 2021a, 2020a; Chen et al., 2022) to employ contrastive learning techniques to predict phrase-region matching in the WPG task.

In summary, all the above studies always ignore the implicit phrase-region matching relations problem, which however holds significant potential for evaluating the ability of models in understanding the deep multimodal semantics.

### 2.2. Causal Inference

In recent years, causal inference has sparked significant interest across a range of areas, including scene graph generation (Tang et al., 2020b), semantic segmentation (Zhang et al., 2020a), vision-language tasks (Chen et al., 2020a), etc. Pearl and Mackenzie (2018) have defined three levels of causality, encompassing intervention and counterfactual that are frequently employed to mitigate confounding bias and achieve unbiased estimations. For example, Wang et al. (2020b) employ

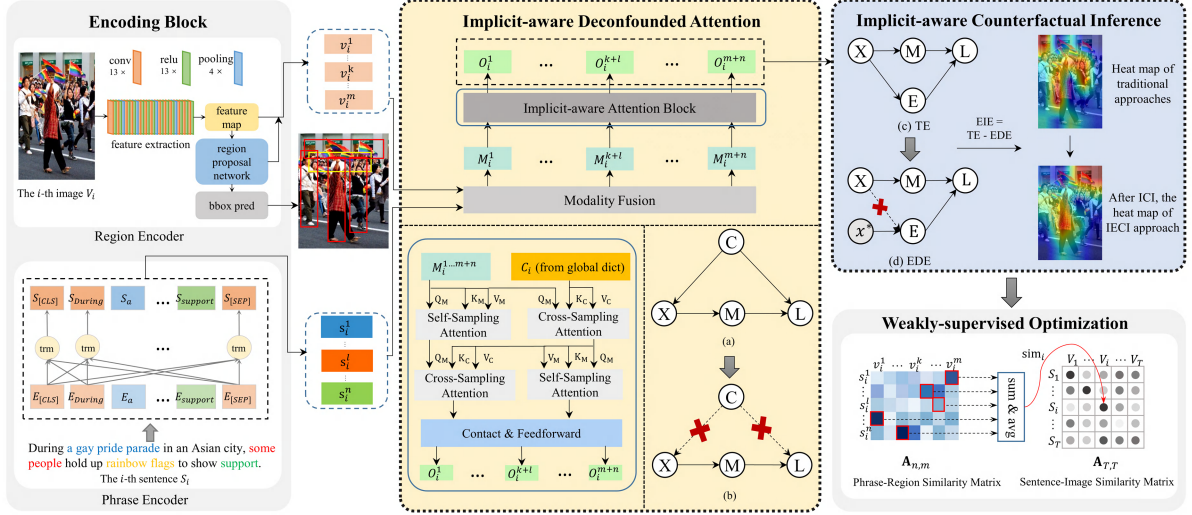


Figure 2: The overall framework of our proposed Implicit-Enhanced Causal Inference (IECI) approach. Wherein (a) and (b) are causal graphs for modeling the implicit relations (see Section 3.2), while (c) and (d) are those for highlighting the implicit relations beyond the explicit (see Section 3.3).

the causal intervention to deal with spurious correlation within datasets for visual common sense learning, and Zhang et al. (2020b) alleviate the spurious correlations between vision and language in vision-linguistic pre-training. For the strategy of intervention adjustment, Yang et al. (2021) employ front-door adjustment to realize the causal intervention, while Wang et al. (2021b) use back-door adjustment to self-annotate the confounder in an unsupervised way. Huang et al. (2022) propose a confounder-agnostic approach to remove the confounding bias between language and location. In addition, Tang et al. (2020a) and Niu et al. (2021) employ counterfactual inference to alleviate long-tailed categories bias in image classification and language bias in VQA, respectively.

Different from all above studies, our study is a pioneering effort in integrating both intervention and counterfactual techniques for multimodal tasks, where we leverage intervention to effectively mitigate the bias of confounding, and employ counterfactual to highlight the implicit matching relations.

### 3. Implicit-Enhanced Causal Inference Approach

In this section, we formulate the WPG task as follows. Given a collection of  $T$  sentence-image pairs  $(S, V)$ , where  $S = [S_1, \dots, S_T]$  and  $V = [V_1, \dots, V_T]$ . Each sentence  $S_i$  consists of multiple phrases  $S_i = [s_i^1, \dots, s_i^l, \dots, s_i^m]$ , while each image  $V_i$  comprises a set of regions  $V_i = [v_i^1, \dots, v_i^k, \dots, v_i^m]$ , where  $n$  and  $m$  represent the number of phrases and regions, respectively. The goal is to predict the region  $v_i^k$  from the set of  $m$  regions in image  $V_i$  that corresponds to the given phrase  $s_i^l$  in sentence  $S_i$ . However, under

the WPG setting, we only have access to coarse-grained sentence-image pair  $(S_i, V_i)$  for training, whereas fine-grained phrase-region pair  $(s_i^l, v_i^k)$  is available during inference.

In this paper, we propose an Implicit-Enhanced Causal Inference (IECI) approach to model the implicit relations. Figure 2 shows the overall architecture of the proposed IECI approach, consisting of three major components: 1) Encoding Block, 2) Implicit-aware Deconfounded Attention (IDA) Block, 3) Implicit-aware Counterfactual Inference (ICI) Block. Prior to delving into the intricacies of the core components within IECI, we provide an overview of the encoding block.

#### 3.1. Encoding Block

Given  $T$  pairs of sentence and image, following the setting by Gupta et al. (2020), BERT (Devlin et al., 2019) and Faster R-CNN (Ren et al., 2015) are adopted to encode phrases and regions.

**Phrase Encoder.** BERT-base model released by Devlin et al. (2019) is adopted as the phrase encoder, which is a light-weighting language encoding model. Specifically, the phrases are first extracted by following Plummer et al. (2015). Then, BERT is utilized to encode the sentence  $S_i$ , and finally the word vectors of all words in each phrase  $s_i^l$  are averaged as the phrase encoding.

**Region Encoder.** Faster R-CNN is adopted as the region encoder. Specifically, Faster R-CNN first utilizes the convolutional network (i.e., ResNet (He et al., 2016)) to compute a feature map of each image  $V_i$ . On this basis, a region proposal network is used to generate the encoding of the region  $v_i^k$  along with the corresponding bounding box.

### 3.2. Implicit-aware Deconfounded Attention Block

In this study, we take advantage of the intervention technique (Pearl and Mackenzie, 2018) and propose an Implicit-aware Deconfounded Attention (IDA) block to model the implicit relations for mitigating the confounding bias inside WPG. Specifically, we address two crucial questions: 1) how to mitigate the confounding bias through the front-door adjustment strategy (Pearl and Mackenzie, 2018); 2) how to implement the front-door adjustment strategy in the WPG task. We will provide comprehensive answers to these two questions in the subsequent section, formulated as follows.

**Deconfounded Causal Graph** is leveraged to answer the question 1). As illustrated in Figure 2 (a), we formulate the causation among sentence-image pairs  $X$ , multimodal knowledge  $M$ , phrase-region locations  $L$ , confounding factors  $C$ .  $X \rightarrow M \rightarrow L$  denotes the desired causal effect from sentence-image pairs  $X$  to phrase-region locations  $L$ , where multimodal knowledge  $M$  acts as a mediator.  $X \leftarrow C \rightarrow L$  denotes the causal effect from the invisible confounding factors  $C$  to sentence-image pairs  $X$  and phrase-region locations  $L$ .

We leverage *do*-operator (Pearl and Mackenzie, 2018) to mitigate the confounding bias ( $X, C$ ) present in the path  $M \rightarrow L$ . As shown in Figure 2 (b), we block the back-door path  $M \leftarrow X \leftarrow C \rightarrow L$  under the condition of  $X$ . Then, we leverage the front-door adjustment strategy to analyze the causal effect of  $X \rightarrow L$ , denoted as follows:

$$\begin{aligned} P(L = l | do(X = x)) \\ = \sum_m P(m|x) \sum_x P(x) [P(l|x, m)] \end{aligned} \quad (1)$$

**Implicit-aware Attention** is leveraged to answer the question 2). On the basis of the front-door adjustment strategy in Eq.(1), we consider implementing it through the utilization of attention mechanisms. Considering the expensive computation of network forward propagation for all samples, we introduce the Normalized Weighted Geometric Mean (NWGM) (Srivastava et al., 2014; Xu et al., 2015) approximation. Therefore, we can sample  $X, M$  and complete  $P(L | do(X))$  by feeding them into the network, and then leverage NWGM approximation to achieve the goal of Eq.(1), denoted as follows:

$$P(L | do(X)) \approx \text{softmax}[g(\hat{X}, \hat{M})] \quad (2)$$

where  $g(\cdot)$  is a network employed to parameterize the predictive distribution  $P(l|x, m)$ , which is followed by a softmax layer. Besides,  $\hat{M} = \sum_m P(M = m | h(X)) \mathbf{m}$  and  $\hat{X} = \sum_x P(X = x | f(X)) \mathbf{x}$  represent the estimations of self-sampling and cross-sampling, respectively. The variables  $m, x$  correspond to the embedding

vectors of  $m, x$ . The query embedding functions  $h(\cdot)$  and  $f(\cdot)$  are utilized to transform the input  $X$  into two distinct query sets, which can be parameterized as networks. Consequently, we leverage attention mechanisms to estimate the self-sampling  $\hat{M}$  and cross-sampling  $\hat{X}$  as shown in Figure 2:

$$\hat{M} = \begin{cases} \mathbf{V}_M \text{softmax}(\mathbf{Q}_M^\top \mathbf{K}_M) \\ \mathbf{V}_C \text{softmax}(\mathbf{Q}_C^\top \mathbf{K}_C) \end{cases} \quad (3)$$

$$\hat{X} = \mathbf{V}_C \text{softmax}(\mathbf{Q}_M^\top \mathbf{K}_C) \quad (4)$$

where Eq.(3) and Eq.(4) denote as the self-sampling attention and cross-sampling attention. Particularly, the upper formula of Eq.(3) calculates the self-sampling attention for multimodal knowledge  $M$ , while the lower formula of Eq.(3) calculates the self-sampling attention for confounding factors  $C$ . In the implementation,  $\mathbf{Q}_M$  and  $\mathbf{Q}_C$  are derived from  $h(X)$  and  $f(X)$ .  $\mathbf{K}_M$  and  $\mathbf{V}_M$  are obtained from the current input sample, while  $\mathbf{K}_C$  and  $\mathbf{V}_C$  come from other samples in the training set and serve as global dictionary compressed from the whole training dataset. Specifically, we initialize this dictionary by using K-means clustering (Hartigan and Wong, 1979) on all the embeddings of samples in the training set, such as region features.

### 3.3. Implicit-aware Counterfactual Inference Block

In this study, we take advantage of the counterfactual technique (Pearl and Mackenzie, 2018) and propose an Implicit-aware Counterfactual Inference (ICI) block to highlight the implicit relations beyond the explicit, thereby addressing the imbalance problem between them. Specifically, we treat the explicit relations as the direct effect in the counterfactual technique, and then reduce such direct effect to achieve the goal of reducing the importance of the explicit relations while highlighting the implicit relations. Therefore, there are also two questions to be answered: 1) how to analyze the direct effect of the explicit relations; 2) how to reduce such direct effect to improve the alignment of implicit phrase-region pairs. Next, we will answer the two questions, formulated as follows.

**Counterfactual Causal Graph** is leveraged to answer the question 1). As illustrated in Figure 2 (c), we formulate the causation between explicit and implicit relations through the path  $X \rightarrow E \rightarrow L$ , which denotes the causal effect from the imbalanced explicit relations  $E$  to phrase-region locations  $L$ . On this basis, we leverage counterfactual to analyze the causal effects. Following the counterfactual annotations in Pearl and Mackenzie (2018) and Niu et al. (2021), we denote  $L_{x,e} = L(X = x, E = e)$  as the *total effect* (TE). We block the path  $X \rightarrow E$  to obtain the *explicit direct effect* (EDE) on  $L$  as shown in









Implicit Relations	Ratios (%)	Sentence-Image Examples with Implicit Relations
Commonsense Understanding (CU)	34.5	During a gay pride parade in an Asian city , some people hold up rainbow flags to show <b>support</b> .  
Complex Context Understanding (CCU)	26.9	A man aboard a red train helps a small child onto it while <b>another person</b> tries to get on.  
Spatial Relations Understanding (SRU)	23.5	A man is playing a guitar <b>next to another man</b> who is sitting behind a green cart wearing a mask.  
Numerical Understanding (NU)	15.1	Six ladies at the dining table and <b>three of them</b> are knitting.  

Figure 3: Four main types of the implicit phrase-region matching relations together with their corresponding ratios within the implicit phrase-region pairs.

Figure 2 (d), denoted as  $L_{x^*,e} = L(X = x^*, E = e)$ , which represents the value of  $L$  when we set  $x$  to  $x^*$ . Note that only in the counterfactual world,  $X$  can be simultaneously set to different values  $x$  and  $x^*$ . To reduce the EDE from TE, we aim to derive the *explicit indirect effect* (EIE), denoted as follows:

$$\text{EIE} = \text{TE} - \text{EDE} = L_{x,e} - L_{x^*,e} \quad (5)$$

**Implicit-aware Inference** is leveraged to answer the question 2). Upon obtaining the output representations  $o_i$  from IDA block, we utilize ICI to reduce the direct effect of explicit relations in Eq.(5). In this context, the representations  $o_i$  can be seen as value  $x$  in ICI. For the value of  $x^*$ , we assume that the model will randomly guess with equal probability, denoted as follows:

$$x = o_i, x^* = r \quad (6)$$

where  $r$  denotes a learnable parameter. Finally, we can compute the EIE representation (i.e., similarity matrix  $\mathbf{A}_{n,m}$ ) in Eq.(5), which are then employed to calculate similarities for the WPG task.

### 3.4. Weakly-supervised Optimization

During the training stage, our access is limited to coarse-grained sentence-image pairs  $(S, V)$ , from which we derive the ground-truth label  $y$  based on the matching relations of each sentence-image pair. Therefore, following [Chen et al. \(2018\)](#), we convert the phrase-region similarity matrix  $\mathbf{A}_{n,m}$  to sentence-image similarity matrix  $\mathbf{A}_{T,T}$  for weakly-supervised optimization. Specifically, we first compute the similarity value  $\text{sim}_i$  for the sentence-image pair  $(S_i, V_i)$ , i.e.,  $\text{sim}_i = \frac{1}{n} \sum_n \max_m \mathbf{A}_{n,m}$ . On this basis, we then calculate the similarity values between the current sentence  $S_i$  and all images  $V$ . Finally, we obtain the sentence-image similarity matrix  $\mathbf{A}_{T,T}$  encompassing all sentence-image pairs. After the argmax operation, we obtain the predicted label  $\hat{y}$ . To train our IECI approach end-to-end, we

leverage the cross-entropy loss function, denoted as follows:

$$\mathcal{L}_{wpg} = - \sum_{i=1}^T y_i \log \hat{y}_i \quad (7)$$

Besides, we incorporate a learnable parameter  $r$  in Eq.(6), which controls the sharpness of the distribution of  $L(x^*, e)$ . If  $r$  is inappropriate, EIE would be guided by TE or EDE. To avoid this, we leverage a Kullback-Leibler divergence to update  $r$  via back propagation:

$$\mathcal{L}_{kl} = \sum -p(y|x, e) \log \frac{p(y|x^*, e)}{p(y|x, e)} \quad (8)$$

Consequently, our optimization objective comprises both  $\mathcal{L}_{wpg}$  and  $\mathcal{L}_{kl}$ , denoted as  $\mathcal{L}_{total} = \mathcal{L}_{wpg} + \alpha \mathcal{L}_{kl}$ , where  $\alpha$  serves as a hyper-parameter responsible for balancing the loss between the WPG task and ICI block.

## 4. Experimental Settings

### 4.1. Data Annotation and Settings

To evaluate the effectiveness of our IECI approach in WPG, we construct an implicit-enhanced dataset based on the Flickr30K-Entities ([Plummer et al., 2015](#)) dataset. Specifically, for the annotation of the implicit-enhanced dataset, we first summarize four main types of implicit relations through preliminary annotation, and estimate their ratios by analyzing 100 randomly-selected annotated samples within Flickr30K-Entities. Specifically, the four types of implicit relations (as illustrated in Figure 3) are introduced as follows:

- **CU** denotes the model needs extra knowledge to understand commonsense, e.g., “*support*” needs to understand the commonsense of *holding a flag* to precisely predict the region.
- **CCU** denotes the model needs the context of sentences to understand the meaning of

Approach	Flickr30K						COCO					
	Implicit		Explicit		Full		Implicit		Explicit		Full	
	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5
KAC-Net	31.99	51.14	49.11	63.50	38.71 <sup>#</sup>	-	39.51	56.85	54.81	73.72	45.88	66.27
ARN	33.81	49.87	46.13	64.50	34.87	50.42	40.35	57.68	57.07	74.98	41.93	58.27
KPRN	30.95	44.50	47.31	62.93	33.41	47.33	35.31	55.22	55.77	76.51	38.30	57.00
InfoGround	44.72	74.79	55.07	80.87	47.88 <sup>†</sup>	76.63 <sup>†</sup>	45.66	74.08	61.17	84.67	51.67 <sup>†</sup>	77.69 <sup>†</sup>
ALBEF	56.40	78.27	69.37	85.03	57.64	77.56	52.00	76.23	66.19	84.06	54.22	76.34
CL&KD	50.33	73.75	62.50	82.00	53.10 <sup>‡</sup>	-	50.37	75.02	64.42	83.78	51.36	74.98
ReIR	57.98	<b>78.72</b>	69.65	85.33	59.27 <sup>§</sup>	-	54.01	77.60	66.77	84.45	55.26	76.72
BLIP	20.31	41.51	26.62	62.62	23.30	57.07	26.99	63.19	34.14	70.02	31.15	67.13
<b>IECI</b>	<b>61.32</b>	78.36	<b>72.37</b>	<b>86.27</b>	<b>62.29</b>	<b>79.28</b>	<b>56.32</b>	<b>78.01</b>	<b>68.62</b>	<b>85.25</b>	<b>56.92</b>	<b>78.31</b>
w/o IDA	57.72	77.16	71.10	84.96	59.07	77.87	53.96	76.88	67.19	84.13	54.92	77.25
w/o ICI	58.09	77.79	71.87	85.33	59.50	78.51	54.44	77.12	67.61	84.71	55.05	77.17
w/o Both	56.05	76.48	69.32	83.89	57.87	77.78	52.20	76.07	66.05	83.57	54.17	76.54

Table 1: Comparison of several state-of-the-art baselines and our approaches on both the Flickr30K and COCO datasets (training sets). *Implicit*, *Explicit* and *Full* represent the evaluation results on our annotated implicit and explicit datasets, the original Flickr30K-Entities dataset, respectively. The result with symbol <sup>#</sup> is retrieved from [Chen et al. \(2018\)](#); those with <sup>†</sup> are from [Gupta et al. \(2020\)](#); this with <sup>‡</sup> is from [Wang et al. \(2021a\)](#) and this with <sup>§</sup> is from [Liu et al. \(2021\)](#). The symbol - denotes the results are not reported by these papers.

phrases, e.g., “another person” can hardly correspond to the exact region without context.

- **SRU** represents the position relation between two target objects in the sentence, e.g., the region of “next to another man” contains the position relation between *two men*.
- **NU** represents a phrase that may correspond to multiple regions, e.g., “three of them” corresponds to three regions in the image.

For annotation, we assign two annotators to tag each phrase-region pair and the *Kappa* consistency check value of the annotation is 0.85. When two annotators cannot reach an agreement, an expert will make the final decision, ensuring the quality of data annotations. Furthermore, we annotate 2K coarse-grained sentence-image samples, including 15K fine-grained phrase-region pairs in the original Flickr30K-Entities dataset. After the manual annotation, we finally obtain 1.4K implicit phrase-region pairs. The ratio between the implicit and explicit phrase-region pairs is about 1:9, which shows an extreme imbalance between them, thereby encouraging us to address the challenge of highlighting the implicit relations beyond the explicit.

For training WPG, following [Gupta et al. \(2020\)](#), we leverage the training sets from Flickr30K and COCO datasets. For inference, we maintain the same setting of validation and test splits as the original Flickr30K-Entities (*Full*) dataset, and obtain the *Implicit*, *Explicit* datasets based on these splits.

## 4.2. Baselines

We choose several state-of-the-art baselines for WPG to compare performance with our IECI approach, described as follows.

- **KAC-Net** ([Chen et al., 2018](#)) explores the consistency in visual and language as complementary external knowledge.
- **ARN** ([Liu et al., 2019a](#)) builds the correspondence between image region and query in an adaptive manner.
- **KPRN** ([Liu et al., 2019b](#)) models the relation between target and contextual entities.
- **InfoGround** ([Gupta et al., 2020](#)) leverages contrastive learning to maximize a lower bound between region features and contextualized word representations.
- **ALBEF** ([Li et al., 2021](#)) introduces a contrastive loss to align the image and text representations before fusing them into cross-modal attention.
- **CL&KD** ([Wang et al., 2021a](#)) focuses on distilling knowledge from a generic object detector under the framework of contrastive learning.
- **ReIR** ([Liu et al., 2021](#)) incorporates coarse-to-fine object refinement and entity relation modeling into a two-stage deep network.

Moreover, an advanced multimodal pre-training model **BLIP** ([Li et al., 2022](#)) is leveraged to compare with IECI. Here, it should be noted that BLIP can not directly fine-tune the WPG task, since BLIP is pre-trained with the coarse-grained sentence-image

pairs. Thus, we treat the phrase-region pairs as the coarse-grained pairs to fine-tune the WPG task.

### 4.3. Implementation Details and Metrics

In our experiments, we leverage open-source codes to obtain experimental results of all the baselines on both *Implicit* and *Explicit* datasets, and re-implement the results of several baselines by their open-source codes on *Full* datasets. The hyper-parameters of these baselines reported by their public papers maintain the same setting, and the others are tuned according to the validation set. Specifically, for phrase encoder, the parameters of BERT are following Devlin et al. (2019). For region encoder, we utilize a pre-trained Faster R-CNN on the Visual Genome (Krishna et al., 2017) dataset to generate a maximum number of 100 object regions. For IECI, we employ the Adam optimizer with an initial learning rate of  $1e-5$  for training. The regularization weight of parameters is  $1e-4$ . The hyper-parameter  $\alpha$  is set to be 0.1 and the batch size is set to be 64. The layers of the self-sampling attention and cross-sampling attention in IDA are both set to be 6. The overall training parameters of our approach are 0.16B. All the baselines are reproduced using PyTorch on a machine equipped with an NVIDIA GeForce RTX 3090, an Intel(R) Xeon(R) E5-2650 v4 CPU (2.20 GHz), CUDA version 11.7, and the PyTorch 1.7.1 library with Python 3.6.13, running on Ubuntu 20.04.1 LTS. To facilitate the corresponding research in this direction, all codes alongside our implicit-enhanced dataset are released<sup>2</sup>.

Besides, the WPG performance is evaluated using  $\text{Recall}@k$ ,  $k \in (1, 5)$ , which measures the fraction of phrases for which the ground-truth bounding box exhibits an  $\text{IoU} \geq 0.5$  with any of the top- $k$  predicted boxes. Moreover,  $t$ -test<sup>3</sup> is used to evaluate the significance of the performance difference by following Chen et al. (2020b).

## 5. Results and Discussions

### 5.1. Experimental Results

Table 1 shows the performance comparison of different approaches. From this table, we can see that: **1)** The performances of all approaches on the *Implicit* dataset are consistently lower than the *Explicit* dataset. This indicates that modeling the implicit relations is more challenging than the explicit, and encourages us to consider such implicit relations problem in WPG. Besides, we also observe that the performance of both *Implicit* and *Explicit* datasets is improved, and we analyze the reason

<sup>2</sup><https://github.com/Zhao-Jianing-SUDA/IECI>

<sup>3</sup><https://docs.scipy.org/doc/scipy/reference/stats.html>

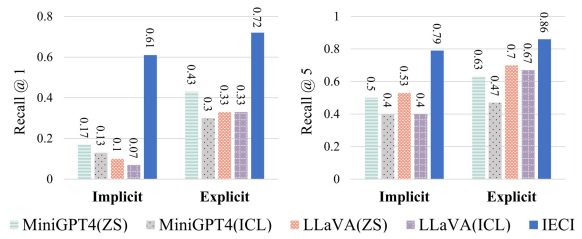


Figure 4: Comparison performance between multi-modal LLMs (i.e., MiniGPT4-13B, LLaVA-13B) and our IECI approach, where ZS and ICL represent zero-shot and in-context learning evaluation methods for LLMs.

is that there may be some overlap between the implicit and explicit relations in the current dataset, leading to performance improvements in explicit relations when employing causal inference. **2)** The performance improvements of our IECI approach on the *Implicit* dataset are larger than those on the *Explicit* dataset. This indicates that IECI can effectively address the challenge of highlighting the implicit relations beyond the explicit. Furthermore, our IECI approach consistently performs better than all other approaches. Impressively, compared to the best-performing ReIR approach, our IECI approach achieves the average  $\text{R}@1$  improvements of 2.83%, 2.29% and 2.34% on all three *Implicit*, *Explicit* and *Full* datasets, respectively. Significance test shows that all these improvements are significant ( $p$ -value  $< 0.01$ ). These demonstrate that IECI can better highlight the implicit relations and meantime maintain the performance of the explicit. Particularly, we observe that the performance of BLIP is quite worse than IECI. This is reasonable since BLIP is pre-trained on the coarse-grained sentence-images pairs, thus is not suitable for aligning the fine-grained phrase-region pairs. **3)** Our IECI approach trained on both the Flickr30K and COCO datasets still outperforms all other approaches. This justifies the robustness of IECI, and again encourages us to consider the important implicit relations in the WPG task.

### 5.2. Contributions of Causal Inference

To further investigate the influence of key causal components within our IECI approach, we conduct a series of ablation studies as shown in Table 1. From this table, we can see that: **1) w/o IDA** exhibits inferior performance compared to IECI, with an average decrease in  $\text{R}@1$  by 2.98% ( $p$ -value  $< 0.01$ ) and 1.35% ( $p$ -value  $< 0.05$ ) on the *Implicit* and *Explicit* datasets, respectively. This further justifies that our proposed IDA block can effectively model the implicit relations, encouraging us to leverage the intervention technique to mitigate the confounding bias. **2) w/o ICI** also shows inferior performance compared to IECI, with an average  $\text{R}@1$  decrease

of 2.56% ( $p$ -value  $< 0.01$ ) and 0.76% ( $p$ -value  $< 0.05$ ) on the *Implicit* and *Explicit* datasets, respectively. This further justifies that ICI block can effectively highlight the implicit relations beyond the explicit, encouraging us to leverage the counterfactual technique to address the imbalance problem. **3) w/o Both** yields a significant average decrease in R@1 by 4.70% ( $p$ -value  $< 0.01$ ) on the *Implicit* dataset, which excludes both IDA and ICI blocks. This again justifies the effectiveness of our IECI approach in modeling and highlighting the implicit phrase-region matching relations.

### 5.3. Comparison with Multimodal LLMs

Recently, the emergence of large language models (LLMs) has demonstrated their remarkable abilities across various fields and tasks. Since GPT-4 is closed-sourced and currently cannot evaluate the multimodal tasks, we choose two advanced open-sourced multimodal LLMs, i.e., MiniGPT4-13B (Zhu et al., 2023) and LLaVA-13B (Liu et al., 2023), to verify the effectiveness of IECI. Specifically, we first randomly select 30 sentence-image examples with the implicit and explicit relations for comparisons on the WPG task. Then, we select two well-studied evaluation methods for LLMs as follows: **1) Zero-Shot (ZS)**. Following the ZS setting for evaluating LLMs proposed by Bang et al. (2023), given a prompt including the task instructions, sentence-image pair and the bounding box list of each phrase, we utilize multimodal LLMs to “generate the bounding boxes from the list that can reflect the given phrase”. **2) In-Context Learning (ICL)**. Following the ICL setting for evaluating LLMs proposed by Li et al. (2023b), building upon the ZS, we introduce one additional demonstration example “*phrase, boxes;...;phrase, boxes*” as the prompt. As shown in Figure 4, we observe that IECI outperforms the multimodal LLMs by a large margin, indicating that the multimodal LLMs still face challenges in understanding the deep multimodal semantics. Moreover, we find that the performance of ZS is better than ICL, since the image encoders BLIP-2 and CLIP within MiniGPT4-13B and LLaVA-13B lack the ability of ICL as reported in Li et al. (2023a) and Radford et al. (2021).

### 5.4. Visualization Study

We conduct a visualization analysis of IECI in our implicit-enhanced dataset. As illustrated in Figure 5, we can see that: **1)** The prediction of the region corresponding to the phrase “support” (red box in Figure 5 (a)) is challenging, which requires contextual and commonsense knowledge. This further indicates the difficulty in precisely aligning such implicit phrase-region pairs. **2)** We compare the performance of the best-performing baseline ReIR




Figure 5: A sentence-image example from our implicit-enhanced dataset, along with their ground-truth phrase-region pairs (a), predicted regions by best-performing baseline ReIR (b), and predicted regions by our IECI approach (c).

and IECI, as shown in Figure 5 (b) and (c). For the phrase “support”, ReIR fails to predict the corresponding region, while IECI successfully identifies the correct region (yellow box). This again justifies the effectiveness of IECI in precisely identifying such implicit phrase-region matching.

### 5.5. What Would We Do Next?

Our proposed IECI approach has demonstrated promising results in comprehensive experiments, showcasing its effectiveness in addressing the challenges of modeling the implicit relations and highlighting them beyond the explicit. However, we believe that there are still three potential directions to precisely predict implicit phrase-region matching relations. Despite this is not the focus of this paper, we believe that these three directions can not only further improve the performance of the WPG task, but also facilitate related research in this direction, as described below.

Firstly, for **Multimodal Representation**, our IECI approach has limitations in effectively representing underlying features. However, the existing multimodal pre-training models (e.g., CLIP (Radford et al., 2021), BLIP (Li et al., 2022)) are mostly pre-trained on coarse-grained sentence-image pairs, making them lack the ability of fine-grained multimodal semantics understanding, as reported in Section 5.3. In our future work, we would like to incorporate the multimodal LLMs (e.g., MiniGPT4 (Zhu et al., 2023), LLaVA (Liu et al., 2023)) to enhance the multimodal representation abilities of our approach for the WPG task.

Secondly, for **Knowledge Injection**, we observe that the ratio of commonsense understanding is 34.5% as illustrated in Figure 3, which further inspires us to consider integrating the external knowledge (e.g., multimodal knowledge graph (Zhu et al., 2022)) to assist in capturing the implicit relations.

Finally, for **Evaluation of Each Implicit Relation**, due to the imbalanced proportion of implicit relations in the whole dataset, the sample scale of



each implicit relation is relatively small (see Section 4.1), making it difficult to comprehensively evaluate the effectiveness of our IECI approach on each implicit relation. In our future work, we would like to leverage the multimodal LLMs to automatically annotate (e.g., the automatic annotation approach proposed by Pei et al. (2023)) different types of implicit relations, which may promote the further research in this direction.

## 6. Conclusion

In this paper, we propose an **Implicit-Enhanced Causal Inference** (IECI) approach to address the implicit phrase-region matching relations problem inside WPG. The key idea of IECI is to utilize the causal inference (i.e., intervention and counterfactual techniques) to effectively model the implicit relations and highlight them beyond the explicit. To comprehensively evaluate IECI, we construct a specialized implicit-enhanced dataset. Detailed experimental results on this dataset demonstrate the superior performance of IECI over several state-of-the-art baselines.

In our future work, we would like to introduce more information (e.g., multimodal knowledge graph (Zhu et al., 2022)) to assist in aligning implicit phrase-region pairs. Moreover, we would like to transfer IECI to other tasks which also have the implicit relation problems, such as Referring Expression Comprehension and video grounding.

## Ethics Statement

**Data Disclaimer.** We construct a high-quality implicit-enhanced dataset based on Flickr30K-Entities, a dataset which is widely used by other academics and are typically accessible to the public. Therefore, our proposed dataset does not involve any sensitive information that may harm others.

**Human Annotation.** When recruiting annotators for this study, we claim that all potential annotators are free to choose whether they want to participate, and they can withdraw from the study anytime without any negative repercussions. Thus, the establishment of our dataset is compliant with ethics.

## Acknowledgments

We thank our anonymous reviewers for their helpful comments. This work was supported by three NSFC grants, i.e., No.62006166, No.62076175 and No.62076176. This work was also supported by a Project Funded by the Priority Academic Program Development of Jiangsu Higher Education Institutions (PAPD).

## References

- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity](#). *CoRR*, abs/2302.04023.
- Kan Chen, Jiyang Gao, and Ram Nevatia. 2018. [Knowledge aided consistency for weakly supervised phrase grounding](#). In *Proceedings of CVPR 2018*, pages 4042–4050.
- Keqin Chen, Richong Zhang, Samuel Mensah, and Yongyi Mao. 2022. [Contrastive learning with expectation-maximization for weakly supervised phrase grounding](#). In *Proceedings of EMNLP 2022*, pages 8549–8559.
- Long Chen, Wenbo Ma, Jun Xiao, Hanwang Zhang, and Shih-Fu Chang. 2021. [Ref-nms: Breaking proposal bottlenecks in two-stage referring expression grounding](#). In *Proceedings of AAAI 2021*, pages 1036–1044.
- Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020a. [Counterfactual samples synthesizing for robust visual question answering](#). In *Proceedings of CVPR 2020*, pages 10797–10806.
- Xiao Chen, Changlong Sun, Jingjing Wang, Shoushan Li, Luo Si, Min Zhang, and Guodong Zhou. 2020b. [Aspect sentiment classification with document-level sentiment preference modeling](#). In *Proceedings of ACL 2020*, pages 3667–3677.
- Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. 2019. [Align2ground: Weakly supervised phrase grounding guided by image-caption alignment](#). In *Proceedings of ICCV 2019*, pages 2601–2610.
- Jiajun Deng, Zhengyuan Yang, Tianlang Chen, Wengang Zhou, and Houqiang Li. 2021. [Transvg: End-to-end visual grounding with transformers](#). In *Proceedings of ICCV 2021*, pages 1749–1759.
- 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 2019*, pages 4171–4186.
- Yang Feng, Lin Ma, Wei Liu, and Jiebo Luo. 2019. [Unsupervised image captioning](#). In *Proceedings of CVPR 2019*, pages 4125–4134.

- Dan Guo, Hui Wang, Hanwang Zhang, Zheng-Jun Zha, and Meng Wang. 2020. [Iterative context-aware graph inference for visual dialog](#). In *Proceedings of CVPR 2020*, pages 10052–10061.
- Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. 2020. [Contrastive learning for weakly supervised phrase grounding](#). In *Proceedings of ECCV 2020*, pages 752–768.
- John A Hartigan and Manchek A Wong. 1979. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *Proceedings of CVPR 2016*, pages 770–778.
- Jianqiang Huang, Yu Qin, Jiabin Qi, Qianru Sun, and Hanwang Zhang. 2022. [Deconfounded visual grounding](#). In *Proceedings of AAAI 2022*, pages 998–1006.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Maten, and Tamara L. Berg. 2014. [Referitgame: Referring to objects in photographs of natural scenes](#). In *Proceedings of EMNLP 2014*, pages 787–798.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *Int. J. Comput. Vis.*, 123(1):32–73.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. [BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models](#). *CoRR*, abs/2301.12597.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. [BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *Proceedings of ICML 2022*, pages 12888–12900.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#). In *Proceedings of NeurIPS 2021*, pages 9694–9705.
- Mukai Li, Shansan Gong, Jiangtao Feng, Yiheng Xu, Jun Zhang, Zhiyong Wu, and Lingpeng Kong. 2023b. [In-context learning with many demonstration examples](#). *CoRR*, abs/2302.04931.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *CoRR*, abs/2304.08485.
- Xuejing Liu, Liang Li, Shuhui Wang, Zheng-Jun Zha, Dechao Meng, and Qingming Huang. 2019a. [Adaptive reconstruction network for weakly supervised referring expression grounding](#). In *Proceedings of ICCV 2019*, pages 2611–2620.
- Xuejing Liu, Liang Li, Shuhui Wang, Zheng-Jun Zha, Li Su, and Qingming Huang. 2019b. [Knowledge-guided pairwise reconstruction network for weakly supervised referring expression grounding](#). In *Proceedings of ACM MM 2019*, pages 539–547.
- Yongfei Liu, Bo Wan, Lin Ma, and Xuming He. 2021. [Relation-aware instance refinement for weakly supervised visual grounding](#). In *Proceedings of CVPR 2021*, pages 5612–5621.
- Yongfei Liu, Bo Wan, Xiaodan Zhu, and Xuming He. 2020. [Learning cross-modal context graph for visual grounding](#). In *Proceedings of AAAI 2020*, pages 11645–11652.
- Jonghwan Mun, Kimin Lee, Jinwoo Shin, and Bohyung Han. 2018. [Learning to specialize with knowledge distillation for visual question answering](#). In *Proceedings of NeurIPS 2018*, pages 8092–8102.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. [Counterfactual VQA: A cause-effect look at language bias](#). In *Proceedings of CVPR 2021*, pages 12700–12710.
- Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect*. Basic books.
- Xiaohuan Pei, Yanxi Li, and Chang Xu. 2023. [GPT self-supervision for a better data annotator](#). *CoRR*, abs/2306.04349.
- 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 2015*, 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 ICML 2021*, pages 8748–8763.

- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In *Proceedings of NeurIPS 2015*, pages 91–99.
- Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. 2016. [Grounding of textual phrases in images by reconstruction](#). In *Proceedings of ECCV 2016*, pages 817–834.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. 2020a. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Proceedings of NeurIPS 2020*.
- Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiabin Shi, and Hanwang Zhang. 2020b. [Unbiased scene graph generation from biased training](#). In *Proceedings of CVPR 2020*, pages 3713–3722.
- Liwei Wang, Jing Huang, Yin Li, Kun Xu, Zhengyuan Yang, and Dong Yu. 2021a. [Improving weakly supervised visual grounding by contrastive knowledge distillation](#). In *Proceedings of CVPR 2021*, pages 14090–14100.
- Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. 2019. [Learning two-branch neural networks for image-text matching tasks](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(2):394–407.
- Qinxin Wang, Hao Tan, Sheng Shen, Michael W. Mahoney, and Zhewei Yao. 2020a. [MAF: multimodal alignment framework for weakly-supervised phrase grounding](#). In *Proceedings of EMNLP 2020*, pages 2030–2038.
- Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. 2020b. [Visual commonsense R-CNN](#). In *Proceedings of CVPR 2020*, pages 10757–10767.
- Tan Wang, Chang Zhou, Qianru Sun, and Hanwang Zhang. 2021b. Causal attention for unbiased visual recognition. In *Proceedings of ICCV 2021*, pages 3091–3100.
- Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. 2017. [Weakly-supervised visual grounding of phrases with linguistic structures](#). In *Proceedings of CVPR 2017*, pages 5253–5262.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *Proceedings of ICML 2015*, pages 2048–2057.
- Li Yang, Yan Xu, Chunfeng Yuan, Wei Liu, Bing Li, and Weiming Hu. 2022. [Improving visual grounding with visual-linguistic verification and iterative reasoning](#). In *Proceedings of CVPR 2022*, pages 9489–9498.
- Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. 2021. [Causal attention for vision-language tasks](#). In *Proceedings of CVPR 2021*, pages 9847–9857.
- Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. 2020a. Causal intervention for weakly-supervised semantic segmentation. In *Proceedings of NeurIPS 2020*.
- Shengyu Zhang, Tan Jiang, Tan Wang, Kun Kuang, Zhou Zhao, Jianke Zhu, Jin Yu, Hongxia Yang, and Fei Wu. 2020b. [Devlbert: Learning deconfounded visio-linguistic representations](#). In *Proceedings of ACM MM 2020*, pages 4373–4382.
- Fang Zhao, Jianshu Li, Jian Zhao, and Jiashi Feng. 2018. [Weakly supervised phrase localization with multi-scale anchored transformer network](#). In *Proceedings of CVPR 2018*, pages 5696–5705.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. [Minigtpt-4: Enhancing vision-language understanding with advanced large language models](#). *CoRR*, abs/2304.10592.
- Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Penglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. 2022. [Multi-modal knowledge graph construction and application: A survey](#). *CoRR*, abs/2202.05786.