

Learning Relation Alignment for Calibrated Cross-modal Retrieval

Shuhuai Ren¹, Junyang Lin³, Guangxiang Zhao¹, Rui Men³, An Yang³,
Jingren Zhou³, Xu Sun^{1,2*}, Hongxia Yang³

¹MOE Key Lab of Computational Linguistics, School of EECS, Peking University

²Center for Data Science, Peking University

³Alibaba Group, China

shuhuai_ren@stu.pku.edu.cn, {zhaoguangxiang, xusun}@pku.edu.cn
{junyang.ljy, menrui.mr, ya235025}@alibaba-inc.com
{jingren.zhou, yang.yhx}@alibaba-inc.com

Abstract

Despite the achievements of large-scale multimodal pre-training approaches, cross-modal retrieval, e.g., image-text retrieval, remains a challenging task. To bridge the semantic gap between the two modalities, previous studies mainly focus on word-region alignment at the object level, lacking the matching between the linguistic relation among the words and the visual relation among the regions. The neglect of such relation consistency impairs the contextualized representation of image-text pairs and hinders the model performance and the interpretability. In this paper, we first propose a novel metric, **Intra-modal Self-attention Distance (ISD)**, to quantify the relation consistency by measuring the semantic distance between linguistic and visual relations. In response, we present **Inter-modal Alignment on Intra-modal Self-attentions (IAIS)**, a regularized training method to optimize the ISD and calibrate intra-modal self-attentions from the two modalities mutually via inter-modal alignment. The IAIS regularizer boosts the performance of prevailing models on Flickr30k and MS COCO datasets by a considerable margin, which demonstrates the superiority of our approach.¹

1 Introduction

Cross-modal retrieval, including image-text retrieval, video-text retrieval, etc., has long been an important downstream task in cross-modal representation learning. Image-Text Retrieval (ITR) aims at modeling the similarity of image-text pairs and recalling the most relevant one. It remains quite challenging due to the heterogeneity of the data and the semantic gap between two different modalities. To bridge this gap, neural networks are responsible for learning global representations of images

*Corresponding Author

¹Our code is available at <https://github.com/lancopku/IAIS>

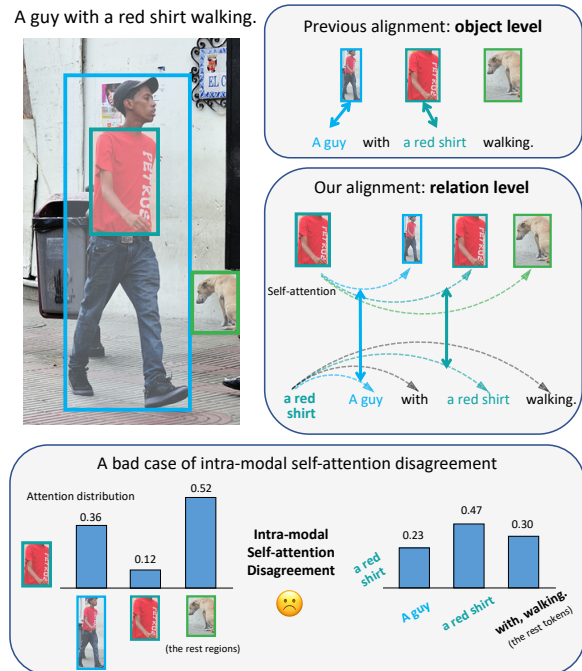


Figure 1: The upper part shows a comparison of previous object-level alignment and our relation-level alignment. The symbol \leftrightarrow denotes alignment and \dashrightarrow denotes the self-attention stems from a query. The lower panel gives a bad case of inconsistent textual and visual relations. The region of “a red shirt” pays considerable attention to the region of the dog, which does not benefit the matching and is inconsistent with the self-attention of the corresponding phrase.

and texts in a joint semantic space and aligning the images and texts with the same semantics (Faghri et al., 2018; Kiros et al., 2014). A straightforward way to enhance the alignment is to enforce the local matching between the object-oriented words and the corresponding image regions, and then leverage the object co-occurrence statistics (Liu et al., 2020; Zhang et al., 2020a) in the pairs for inference. Previous studies incorporate auxiliary knowledge source like scene graphs (Yu et al., 2020) or object tags (Li et al., 2020) to explicitly indicate the cross-

modal mapping. Other researches try to establish fine-grained interaction on cross-modal attention to reinforce the focus from words to their most relevant regions, and vice versa (Chen et al., 2020; Wang et al., 2019; Messina et al., 2020; Lee et al., 2018; Zhang et al., 2020b; Yang et al., 2020).

However, such word-region alignment at object level serves only as the basis because it mainly focuses on the local semantics but lacks the matching of global features like the **intra-modal relation**. The intra-modal relation refers to the correlation of items within a textual or visual sequence. More specifically, given a sentence and an image that describe the same scene and are highly matched, the correlation of the items in the textual sequence should also agree with the correlation of the corresponding items in the visual sequence. But such constraint of relation consistency is neglected in previous works, which hinders performance and interpretability of the models. To corroborate this, we conduct a case study on Flickr30k Entities dataset (Plummer et al., 2015) to probe the agreement of relation-level semantics in pre-trained models like UNITER (Chen et al., 2020). We utilize the self-attention distribution as a representation of the intra-modal relations (Clark et al., 2019; Htut et al., 2019; Kovaleva et al., 2019).

As shown in Figure 1, the attention distributions grouped by the annotated object of the given text and image are in disagreement with each other. Specifically, the attention distribution in the linguistic modality is reasonable. However, in the visual modality, the region “a red shirt” pays inappropriate attention to the region of the dog that doesn’t appear in the text, which impairs the representation of this visual item, i.e., “a red shirt” under the condition of the corresponding text. Such mismatched attention distributions suggest that the model represents the same concept with inconsistent semantics, which misleads the model to reduce the estimated similarity of the positive pairs and further leads to the wrong predictions that they are unmatched. What’s even worse is that in practice, the input regions of the existing methods are extracted by a pre-trained object detector like Faster R-CNN (Ren et al., 2015). The visual features are much noisier due to over-sampling (Li et al., 2020; Anderson et al., 2018), which necessitates a stronger regularizer to guide the alignment of the intra-modal relations.

Motivated by the above observations, we pro-

mote the semantic alignment from object level to relation level. We leverage self-attention matrix to characterize the relation of items within one modality, and design **Intra-modal Self-attention Distance (ISD)**, a novel metric to measure the consistency between textual and visual relations. Our empirical analysis illustrates that the ISD and the model performance on image-text retrieval are highly correlated, which verifies our hypothesis and inspires us to minimize the semantic distance between intra-modal self-attentions in training. Accordingly, we propose a new regularized training method called **Inter-modal Alignment on Intra-modal Self-attentions (IAIS)** to calibrate two intra-modal attention distributions mutually via inter-modal alignment, which helps learn better contextualized representations for image-text pairs. The model performance of image-text retrieval on Flickr30k and MS COCO datasets is improved by a considerable margin with IAIS, which demonstrates the superiority of our proposal.

2 Measuring Semantic Distance between Intra-modal Relations

In this section, we present a formal definition of intra-modal relation alignment (Section 2.1). Such alignment requires extracting the visual and linguistic items corresponding to all objects and sorting them in the same order to make their self-attention distributions comparable. We first introduce the mechanism for multimodal attention calculation, and then present the method of attention weight extraction for constructing comparable intra-modal self-attentions (Section 2.2). Finally, we propose a metric named Intra-modal Self-attention Distance (ISD) to quantify the relation consistency. We conduct an empirical analysis on prevailing models to verify the correlation of the model performance and our metric (Section 2.3).

2.1 From Intra-modal Relation to Self-attention

Given a sequence $\mathbf{O} = [o_1, \dots, o_N]$ of N objects appeared in an image-text pair, the linguistic and visual representation of such object sequence can be written as $\mathbf{L} = [l_1, \dots, l_N]$ and $\mathbf{V} = [v_1, \dots, v_N]$, respectively. Each item l_i, v_i with the same index refers to the same object o_i .² For every object, its

²An object o_i may require one or more tokens in the text and one or more regions in the image to describe, such that the linguistic item l_i and the visual item v_i may refer to a collection of tokens and regions, respectively.

relation to the others is depicted in both the linguistic and the visual modality. From a linguistic view, we regard the following textual self-attention distribution as the relation \mathbf{R}_{l_i} stems from l_i :

$$\mathbf{R}_{l_i} = [a_{l_i \rightarrow l_1}, \dots, a_{l_i \rightarrow l_i}, \dots, a_{l_i \rightarrow l_N}], \quad (1)$$

where $a_{l_i \rightarrow l_j}$ is the attention weight from l_i to l_j . Similarly, the relation \mathbf{R}_{v_i} from the view of the visual modality can be written as

$$\mathbf{R}_{v_i} = [a_{v_i \rightarrow v_1}, \dots, a_{v_i \rightarrow v_i}, \dots, a_{v_i \rightarrow v_N}]. \quad (2)$$

Consequently, we can achieve relation-level alignment by narrowing the semantic distance, e.g., Kullback-Leibler Divergence, between the linguistic and visual self-attention distribution for all objects from $i = 1$ to N :

$$\min \sum_{i=1}^N \text{distance}(\mathbf{R}_{l_i}, \mathbf{R}_{v_i}). \quad (3)$$

In the original self-attention matrix, however, the attention weights of specific objects are scattered and disordered. We need to extract the target weights and reorder them to construct comparable attention distributions \mathbf{R}_{l_i} and \mathbf{R}_{v_i} .

2.2 Intra-modal Self-attention Reconstruction

In this subsection, we first introduce the vanilla multimodal attention mechanism and then present a specific way of attention weight extraction.

Consider models of single-stream Transformer-based architecture like UNITER (Chen et al., 2020). The model consists of a stack of Transformer layers with attention mechanism (Vaswani et al., 2017) and is responsible for encoding image-text pairs into feature representations. Given $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{N \times d}$, the matrix of N query, key and value vectors with dimension d , respectively, the attention function $\text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$ is defined as:

$$\text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \sigma(\mathbf{Q}\mathbf{K}^\top) \mathbf{V} = \sigma(\mathbf{S}) \mathbf{V}. \quad (4)$$

Here, σ is a row-wise, scaled softmax and \mathbf{S} is a matrix of attention scores that measure the similarity between every pair of query and key vectors. Let \mathcal{L} and \mathcal{V} denote the linguistic and the visual modality, respectively. Given a textual sequence $\mathbf{X}_{\mathcal{L}}$ of $N_{\mathcal{L}}$ tokens and a visual sequence $\mathbf{X}_{\mathcal{V}}$ of $N_{\mathcal{V}}$ regions, the input $\mathbf{X} = [\mathbf{X}_{\mathcal{L}} \parallel \mathbf{X}_{\mathcal{V}}]$ in the single-stream architecture is a concatenation of two sequences with length $N = N_{\mathcal{L}} + N_{\mathcal{V}}$. Accordingly, the query and

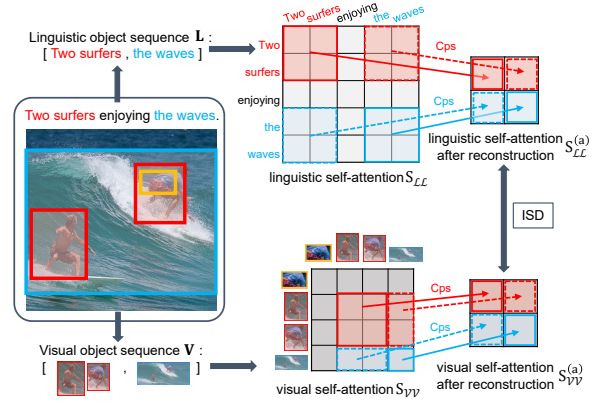


Figure 2: An example of calculating Intra-modal Self-attention Distance (ISDa) for a matched image-text pair. Two inputs in the pair both contain the object of “two surfers” and “the waves”. For self-attention matrix $\mathbf{S}_{\mathcal{L}\mathcal{L}}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}$ from each modality, we extract object-orientated patches according to the annotations and summarize it with the Cps operation (Eq. (7)) to synthesize new matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}$. Finally, we use our ISDa metric to measure their semantic distance.

key matrix³ can be written as

$$\begin{aligned} \mathbf{Q} &= \mathbf{X}\mathbf{W}^Q = \begin{pmatrix} \mathbf{X}_{\mathcal{L}} \\ \mathbf{X}_{\mathcal{V}} \end{pmatrix} \mathbf{W}^Q = \begin{pmatrix} \mathbf{Q}_{\mathcal{L}} \\ \mathbf{Q}_{\mathcal{V}} \end{pmatrix} \\ \mathbf{K} &= \mathbf{X}\mathbf{W}^K = \begin{pmatrix} \mathbf{X}_{\mathcal{L}} \\ \mathbf{X}_{\mathcal{V}} \end{pmatrix} \mathbf{W}^K = \begin{pmatrix} \mathbf{K}_{\mathcal{L}} \\ \mathbf{K}_{\mathcal{V}} \end{pmatrix}, \end{aligned} \quad (5)$$

where \mathbf{W}^Q and \mathbf{W}^K are learnable parameters. Furthermore, the attention score matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$ can be organized into four submatrices (Bugliarello et al., 2020):

$$\begin{aligned} \mathbf{S} &= \mathbf{Q}\mathbf{K}^\top = \begin{pmatrix} \mathbf{Q}_{\mathcal{L}} \\ \mathbf{Q}_{\mathcal{V}} \end{pmatrix} \begin{pmatrix} \mathbf{K}_{\mathcal{L}}^\top & \mathbf{K}_{\mathcal{V}}^\top \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{Q}_{\mathcal{L}}\mathbf{K}_{\mathcal{L}}^\top & \mathbf{Q}_{\mathcal{L}}\mathbf{K}_{\mathcal{V}}^\top \\ \mathbf{Q}_{\mathcal{V}}\mathbf{K}_{\mathcal{L}}^\top & \mathbf{Q}_{\mathcal{V}}\mathbf{K}_{\mathcal{V}}^\top \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{S}_{\mathcal{L}\mathcal{L}} & \mathbf{S}_{\mathcal{L}\mathcal{V}} \\ \mathbf{S}_{\mathcal{V}\mathcal{L}} & \mathbf{S}_{\mathcal{V}\mathcal{V}} \end{pmatrix}. \end{aligned} \quad (6)$$

The matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}$ on the diagonal represent the linguistic and the visual **intra-modal** self-attention, respectively. $\mathbf{S}_{\mathcal{L}\mathcal{V}}$ and $\mathbf{S}_{\mathcal{V}\mathcal{L}}$ on back-diagonal represent the **inter-modal** attention scores from text to image, and the opposite. We regard the self-attention $\sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}})$ and $\sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}})$ as depictions of the intra-modal relations. Each row of the matrix represents the relation stemming from one linguistic or visual item to the others within the same modality.

To construct the comparable intra-modal self-attention matrices, we leverage the object annotations in the Flickr30k Entities dataset (Plummer

³The value matrix \mathbf{V} is omitted for brevity.

et al., 2015) to extract the tokens, regions, and attention weights with respect to the target objects. As shown in Figure 2, the text and the image both contain annotated objects of “two surfers” and “the waves”. The linguistic object sequence can be written as $\mathbf{L} = [l_1, l_2] = [\text{“two surfers”}, \text{“the waves”}]$. These two objects derive four intrinsic relations and can be described by four patches in the original linguistic self-attention matrix $\mathbf{S}_{\mathcal{L}\mathcal{L}}$. For clarity, we define an operation $\text{Ext}(\mathbf{S}, o_i, o_j)$ that extracts the patch of attention scores in matrix \mathbf{S} from the object o_i to o_j . Accordingly, the relation from “two surfers” to “the waves” can be denoted as $\text{Ext}(\mathbf{S}_{\mathcal{L}\mathcal{L}}, l_1, l_2)$. To describe the relation with a single value instead of a sub-matrix, we further construct an operation $\text{Cps}(\cdot)$ to summarize the attention patch $\mathbf{S} \in \mathbb{R}^{M \times N}$ to a scalar via column-wise sum and row-wise average:

$$\text{Cps}(\mathbf{S}) = \left(\sum_i^M \sum_j^N \mathbf{s}_{ij} \right) / M. \quad (7)$$

After the above processing, we complete the extraction of the linguistic self-attention $\mathbf{S}_{\mathcal{L}\mathcal{L}}$ through grouping the items by annotated object. The extraction of visual self-attention $\mathbf{S}_{\mathcal{V}\mathcal{V}}$ is similar and the final results are denoted as $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}$. As our processing for two intra-modal self-attentions follows the same order of object annotations, the matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}$ from two modalities are of the same dimension and comparable.

2.3 Intra-modal Self-attention Distance with Annotation (ISDa)

Given two comparable matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}$, we propose a metric called Intra-modal Self-attention Distance with annotation (ISDa) to quantify their semantic gap at the relation level. We define the following symmetric matrix-based Kullback-Leibler Divergence (m-KL) for measuring the distance between two matrices \mathbf{A} and \mathbf{B} :

$$\text{m-KL}(\mathbf{A}, \mathbf{B}) = \sum_i^N \text{KL}(\mathbf{A}_i \| \mathbf{B}_i) + \text{KL}(\mathbf{B}_i \| \mathbf{A}_i), \quad (8)$$

where $(\cdot)_i$ stands for the i^{th} row-vector in the matrix and KL denotes the Kullback-Leibler Divergence. Accordingly, the final ISDa metric for $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}$ is defined as:

$$\text{ISDa} = \text{m-KL} \left(\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}, \mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)} \right). \quad (9)$$

We present our algorithm for the calculation of ISDa in Algorithm 1.

Algorithm 1: Intra-modal Self-attention Distance with Annotation (ISDa)

Input: Intra-modal self-attention matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}, \mathbf{S}_{\mathcal{V}\mathcal{V}}$

Input: Linguistic object sequence \mathbf{L}

Input: Visual object sequence \mathbf{V}

for linguistic object l_i in \mathbf{L} **do**

for linguistic object l_j in \mathbf{L} **do**

$\mathbf{S}_{l_i \rightarrow l_j} \leftarrow \text{Ext}(\mathbf{S}_{\mathcal{L}\mathcal{L}}, l_i, l_j)$

$\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}[i, j] \leftarrow \text{Cps}(\mathbf{S}_{l_i \rightarrow l_j})$

for visual object v_i in \mathbf{V} **do**

for visual object v_j in \mathbf{V} **do**

$\mathbf{S}_{v_i \rightarrow v_j} \leftarrow \text{Ext}(\mathbf{S}_{\mathcal{V}\mathcal{V}}, v_i, v_j)$

$\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)}[i, j] \leftarrow \text{Cps}(\mathbf{S}_{v_i \rightarrow v_j})$

ISDa = m-KL $(\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(a)}, \mathbf{S}_{\mathcal{V}\mathcal{V}}^{(a)})$ // Eq. 9

return ISDa

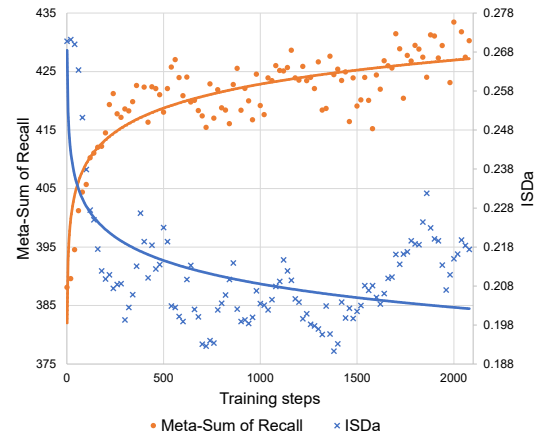


Figure 3: The ISDa (blue \times) and model performance (Meta-Sum of Recall, orange \bullet) with respect to the training steps. They are highly correlated with a Pearson’s correlation coefficient of -0.60.

To study the correlation between the ISDa metric and the model performance,⁴ we conduct an empirical analysis on UNITER (Chen et al., 2020). As shown in Figure 3, the ISDa decreases during the training phase while the model performance continues to increase. They are strongly correlated with a Pearson’s correlation coefficient of -0.60. After the middle stage of training, the curve of the model performance and ISDa tends to be flat, suggesting that merely optimizing the task-oriented loss function while neglecting the constraint of relation consistency hinders the model from achieving better performance. To eliminate the bottleneck, we can minimize the ISD in the training phase as a regularization to induce further improvement for the ITR task and better the model interpretability.

⁴We use the Meta-Sum (Chen et al., 2020), sum of Recall@1, Recall@5, Recall@10 across the image and text retrieval as a metric for model performance.

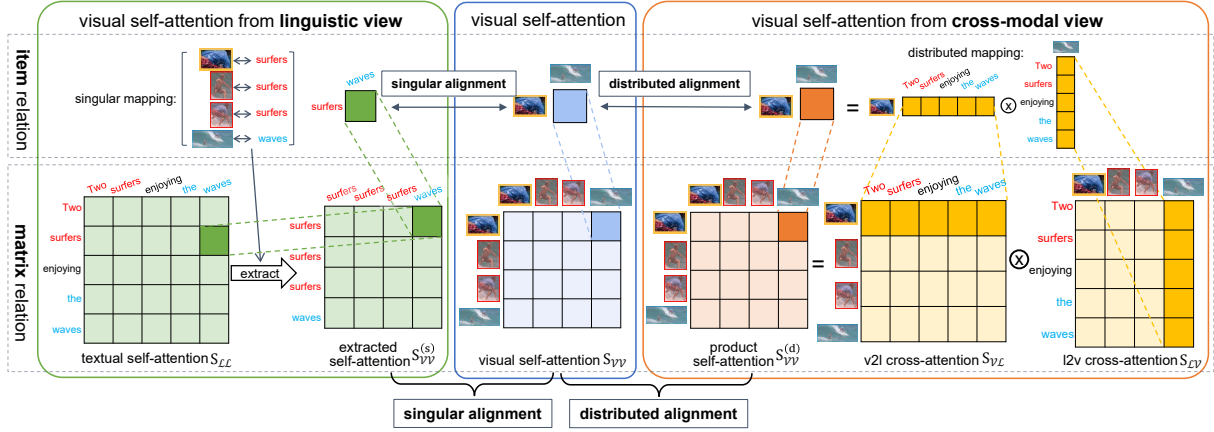


Figure 4: Singular alignment and distributed alignment. The image-text pair here is the same as that in Figure 2.

3 Inter-modal Alignment on Intra-modal Self-attentions (IAIS)

In this section, we propose a new regularized training method, Inter-modal Alignment on Intra-modal Self-attentions (IAIS), for image-text retrieval. Our goal is to enhance the semantic alignment of relations by minimizing the distance between two intra-modal self-attentions (ISD).

In practice, given the **original** visual and linguistic input sequence $\mathbf{V} = [v_1, \dots, v_{N_V}]$, $\mathbf{L} = [l_1, \dots, l_{N_L}]$ with the scattered items,⁵ there are no object annotations and the region features extracted by Faster R-CNN are much noisier (Li et al., 2020; Anderson et al., 2018), which results in difficulty in grouping the attention weights by ground-truth object. The ISDa thus cannot be used directly as the objective function to minimize.

To tackle this problem, we regard the input sequence from one modality (e.g., the visual sequence \mathbf{V}) as an anchor. For every item in the anchor sequence, we extract its corresponding representation from the other modality (e.g., one item or a collection of items in the linguistic sequence \mathbf{L}) to reconstruct a mirrored sequence. After that, the items and their relations within the anchor sequence have a one-to-one correspondence with the items and relations within the mirrored sequence, which makes the intra-modal self-attentions derived from the two sequences comparable. In the next two subsections, we propose two methods, **singular alignment** and **distributed alignment**, to accomplish the attention extraction and reconstruction. The former establishes a one-to-one mapping between

⁵As there are no object annotations in practice, each visual item now refers to only one region. Each linguistic item also refers to only one token, even if it is a sub-word.

linguistic and visual attention weight, while the latter establishes a distributed mapping. Besides, we design two losses $L_{IAIS}^{(s)}$ and $L_{IAIS}^{(d)}$ as a surrogate of the ISDa to measure the semantic distance between intra-modal self-attention matrices. Finally, we incorporate the surrogate loss minimization as a regularization to calibrate intra-modal self-attentions mutually and achieve the relation-level alignment.

3.1 Singular Alignment

For every item in the anchor sequence, singular alignment utilizes the inter-modal attention to find its most relevant item from the opposite modality. As the inter-modal attention score quantifies the similarity between the items from two modalities, the visual and the linguistic item with the highest score can be aligned with each other. For example, given the i^{th} visual item v_i and the inter-modal attention matrix \mathbf{S}_{V_L} , the similarities between v_i and all the linguistic items are depicted in $\mathbf{S}_{V_L}[i, :]$, i.e., the i^{th} row of the matrix. Hence the most relevant linguistic item for v_i can be denoted as l_{i^*} , where $i^* = \arg \max \mathbf{S}_{V_L}[i, :]$. Accordingly, for every weight $a_{v_i \rightarrow v_j}$ in the original visual self-attention matrix \mathbf{S}_{V_V} , its corresponding weight $a_{l_{i^*} \rightarrow l_{j^*}}$ in the linguistic self-attention matrix \mathbf{S}_{L_L} can be extracted by the following operation:⁶

$$\begin{aligned} a_{l_{i^*} \rightarrow l_{j^*}} &= \text{Ext}(\mathbf{S}_{L_L}, l_{i^*}, l_{j^*}), \\ i^* &= \arg \max \mathbf{S}_{V_L}[i, :], \\ j^* &= \arg \max \mathbf{S}_{V_L}[j, :], \end{aligned} \quad (10)$$

as a singular alignment. After all the extractions, we reconstruct a mirrored matrix $\mathbf{S}_{V_V}^{(s)}$ such that $\mathbf{S}_{V_V}^{(s)}[i, j] = a_{l_{i^*} \rightarrow l_{j^*}}$, which can be regarded as a

⁶Compared with Section 2.2, the Ext operation here extracts a singular attention weight instead of a patch.

Algorithm 2: Singular Alignment

Input: Intra-modal self-attention matrices $\mathbf{S}_{\mathcal{L}\mathcal{L}}, \mathbf{S}_{\mathcal{V}\mathcal{V}}$

for $i = 1$ **to** $N_{\mathcal{V}}$ **do**

$i^* \leftarrow \arg \max \mathbf{S}_{\mathcal{V}\mathcal{L}}[i, :]$

for $j = 1$ **to** $N_{\mathcal{V}}$ **do**

$j^* \leftarrow \arg \max \mathbf{S}_{\mathcal{V}\mathcal{L}}[j, :]$

$\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(s)}[i, j] \leftarrow \text{Ext}(\mathbf{S}_{\mathcal{L}\mathcal{L}}, l_{i^*}, l_{j^*})$

for $i = 1$ **to** $N_{\mathcal{L}}$ **do**

$i^* \leftarrow \arg \max \mathbf{S}_{\mathcal{L}\mathcal{V}}[i, :]$

for $j = 1$ **to** $N_{\mathcal{L}}$ **do**

$j^* \leftarrow \arg \max \mathbf{S}_{\mathcal{L}\mathcal{V}}[j, :]$

$\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(s)}[i, j] \leftarrow \text{Ext}(\mathbf{S}_{\mathcal{V}\mathcal{V}}, v_{i^*}, v_{j^*})$

$L_{\text{IAIS}}^{(s)} = \text{m-KL}(\sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}}), \sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(s)})) +$
 $\text{m-KL}(\sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}}), \sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(s)}))$

return $L_{\text{IAIS}}^{(s)}$

representation of the original visual self-attention $\mathbf{S}_{\mathcal{V}\mathcal{V}}$ from the linguistic view. The surrogate loss of ISDa between $\mathbf{S}_{\mathcal{V}\mathcal{V}}$ and $\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(s)}$ is denoted as $L_{\text{IAIS-V}}^{(s)}$ when taking vision as the anchor modality. The similar processing can also be performed when the linguistic sequence is the anchor. We can generate the matrix $\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(s)}$ as a visual representation of the linguistic self-attention $\mathbf{S}_{\mathcal{L}\mathcal{L}}$ and define a corresponding loss $L_{\text{IAIS-L}}^{(s)}$.

The detailed processing of singular alignment is illustrated in Algorithm 2 and Figure 4. The singular version of IAIS loss is defined as:

$$\begin{aligned} L_{\text{IAIS}}^{(s)} &= L_{\text{IAIS-V}}^{(s)} + L_{\text{IAIS-L}}^{(s)} \\ &= \text{m-KL}(\sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}}), \sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}}^{(s)})) + \\ &\quad \text{m-KL}(\sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}}), \sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}}^{(s)})). \end{aligned} \quad (11)$$

3.2 Distributed Alignment

As singular items from different modalities may not be able to give a full representation for each other, we further propose distributed alignment, which utilizes a collection of linguistic items as a representation of a visual item, and vice versa. Specifically, given two visual items v_i and v_j , we regard the inter-modal attentions $\sigma(\mathbf{S}_{\mathcal{V}\mathcal{L}}[i, :])$ ⁷ from v_i to all linguistic items and $\sigma(\mathbf{S}_{\mathcal{L}\mathcal{V}}[:, j])$ ⁸ from all linguistic items to v_j as a kind of features. Hence the original similarity $\mathbf{S}_{\mathcal{V}\mathcal{V}}[i, j] = a_{v_i \rightarrow v_j}$ between v_i and v_j can also be modeled as a dot-product of their distributed attention features from the cross-modal view: $\sigma(\mathbf{S}_{\mathcal{V}\mathcal{L}}[i, :]) \cdot \sigma(\mathbf{S}_{\mathcal{L}\mathcal{V}}[:, j])$. Such distributed

⁷The i^{th} row of $\mathbf{S}_{\mathcal{V}\mathcal{L}}$.

⁸The j^{th} column of $\mathbf{S}_{\mathcal{L}\mathcal{V}}$.

alignment leverages the language as a bridge to draw implicit connections within the visual modality, which can be intuitively regarded as the back-translation (Sennrich et al., 2016) for multimodal. As shown in Figure 4, the distributed version of mirrored self-attention matrix can be constructed by a matrix multiplication of two inter-modal attention matrices:

$$\begin{aligned} \mathbf{S}_{\mathcal{V}\mathcal{V}}^{(d)} &= \sigma(\mathbf{S}_{\mathcal{V}\mathcal{L}}) \sigma(\mathbf{S}_{\mathcal{L}\mathcal{V}}), \\ \mathbf{S}_{\mathcal{L}\mathcal{L}}^{(d)} &= \sigma(\mathbf{S}_{\mathcal{L}\mathcal{V}}) \sigma(\mathbf{S}_{\mathcal{V}\mathcal{L}}). \end{aligned} \quad (12)$$

Similar to the version of singular alignment, the distributed IAIS loss can be written as:

$$\begin{aligned} L_{\text{IAIS}}^{(d)} &= L_{\text{IAIS-V}}^{(d)} + L_{\text{IAIS-L}}^{(d)} \\ &= \text{m-KL}(\sigma(\mathbf{S}_{\mathcal{V}\mathcal{V}}), \mathbf{S}_{\mathcal{V}\mathcal{V}}^{(d)}) + \\ &\quad \text{m-KL}(\sigma(\mathbf{S}_{\mathcal{L}\mathcal{L}}), \mathbf{S}_{\mathcal{L}\mathcal{L}}^{(d)}). \end{aligned} \quad (13)$$

3.3 Relation Alignment as Regularizer

With the IAIS loss, the surrogate of semantic distance between two intra-modal self-attentions, we present a new regularized training method to enhance the relation alignment for image-text retrieval. Our final loss is two-fold. The first is the task-orientated margin loss:

$$L_{\text{margin}} = \sum_{i=1}^{N_p} \left[\sum_{j=1}^{N_n} S_j - S_i + \alpha \right]_+, \quad (14)$$

where $[x]_+ = \max(0, x)$ and α is a preset margin. N_p and N_n denote the number of positive and negative pairs. S_i and S_j are the similarity scores of a positive and negative image-text pair, respectively. The second is the IAIS loss for all **positive** pairs that quantifies their relation distance. The IAIS loss is computed based on the attentions from the last Transformer-layer, and it can be either the singular alignment version (Eq. (11)) or the distributed alignment version (Eq. (13)). To summarize, our final final loss can be formalized as:

$$L = L_{\text{margin}} + \lambda_t L_{\text{IAIS}}, \quad (15)$$

where λ_t is a hyper-parameter w.r.t training steps t to balance two loss items. Since our relation-level alignment is based on mappings between linguistic and visual items, it is beneficial to focus on the item-level alignment at the previous training stage via the task-orientated loss. Accordingly, we utilize Training Signal Annealing (Xie et al., 2020) to gradually incorporate the signal of the IAIS loss and design the following exponential schedule:

$$\lambda_t = \exp((t/T - 1) \times 5). \quad (16)$$

Model	Flickr30k							MS COCO						
	Image Retrieval			Text Retrieval			Overall	Image Retrieval			Text Retrieval			Overall
	R@1	R@5	R@10	R@1	R@5	R@10	Meta-Sum	R@1	R@5	R@10	R@1	R@5	R@10	Meta-Sum
UNITER-base*	72.52	92.36	96.08	85.90	97.10	98.80	542.76	50.33	78.52	87.16	64.40	87.40	93.08	460.89
UNITER-base†	72.70	92.60	96.14	85.50	97.30	98.60	542.84	50.41	78.33	86.94	65.16	87.60	93.14	461.58
+ IAIS-singular	73.54	93.14	96.32	86.10	98.10	99.10	546.30	50.99	78.85	87.41	66.98	89.10	94.02	467.35
+ IAIS-distributed	73.66	92.88	96.28	87.10	97.90	99.20	547.02	51.10	78.70	87.09	66.88	88.90	94.10	466.77
UNITER-large*	73.56	94.08	96.76	87.30	98.00	99.20	548.90	52.93	79.93	87.95	65.68	88.56	93.76	468.81
UNITER-large†	75.98	93.40	96.68	85.80	97.80	98.80	548.46	52.57	79.76	88.00	64.24	88.00	93.62	466.19
+ IAIS-singular	76.86	93.30	95.72	88.30	98.40	99.40	551.98	53.17	80.07	87.94	67.78	89.70	94.48	473.14
+ IAIS-distributed	76.28	93.32	95.58	88.30	98.60	99.30	551.38	53.18	79.99	88.18	67.68	89.34	94.02	472.39

Table 1: Results of image and text retrieval on Flickr30k and MS COCO. R@K corresponds to whether the ground truth is recalled among top K results. * denotes the results of UNITER taken from Chen et al. (2020) and † denotes our reproduction. IAIS-singular and ISA-distributed denote the singular and distributed version of the proposed relation-level alignment, respectively.

Here T is the total training steps during fine-tuning phase and t is the current step. As a pluggable regularizer, our IAIS method does NOT incorporate any extra parameters and additional data collection yet empowers the models to capture the higher-level semantics of relation consistency efficiently.

4 Experimental Settings

4.1 Benchmark Datasets

We conduct experiments on the Flickr30k (Young et al., 2014) and MS COCO (Lin et al., 2014) datasets. Flickr30K contains 31K images collected from the Flickr website, with five textual descriptions per image. We follow Karpathy and Li (2015) to split the data into 30K/1K/1K training/validation/test splits. MS COCO consists of 123K images, each accompanied with five human-written captions. Following Karpathy and Li (2015), the data is divided into 82K/5K/5K training/validation/test images.

4.2 Fine-tuning Settings

Due to the limitation of computing resource, we only incorporate IAIS regularization in the phase of fine-tuning instead of pre-training. We use the base (12 layers) and the large (24 layers) version of UNITER (Chen et al., 2020), one of the most prevailing large-scale pre-trained models, as our baseline and backbone for IAIS. We follow the fine-tuning setting and hyper-parameter configuration of the original paper.⁹ The margin in Eq. (14) is 0.2. For each positive instance, 31 hard negative instances are sampled on the text and image side, respectively, and as each batch contains 8 different

positive instances, the batch size is 512. The learning rate is $5e-5$ and the training steps are 5000 for both base and large models. All experiments are run on 8 NVIDIA V100 GPUs.

5 Results and Analysis

5.1 Main Results

The main results of the UNITER performance with and without our IAIS regularization are reported in Table 1. Our methods of both singular and distributed version surpass the baseline by a considerable margin. The average improvement over all datasets and models is **4.49**.

There are also some interesting findings: (1) Compared with image retrieval, the model performance on text retrieval is boosted by IAIS more remarkably with an average improvement of 3.50. Note that each image in both datasets is paired with five ground-truth sentences, and our IAIS regularizer helps the model capture the common relations for the image and the corresponding texts so that more ground-truth texts can be successfully retrieved. (2) The improvement on UNITER-base is 17.2% higher than that on UNITER-large. A consistent result can be found in Table 2, which demonstrates various relation distance metrics of fine-tuned models. The ISDa of UNITER-large is smaller than that of UNITER-base, indicating UNITER-large learns more about the relation consistency due to its large capability while there is still room to improve the relation alignment with our IAIS method. (3) The relative improvement brought by the singular version of IAIS is 7.0%, higher than that of the distributed version. The ISDa and $L_{IAIS}^{(s)}$ are correlated with a Pearson’s cor-

⁹<https://github.com/ChenRocks/UNITER>

Model	ISDa	$L_{\text{IAIS}}^{(s)}$	$L_{\text{IAIS}}^{(d)}$
UNITER-base	0.26	0.59	0.36
+ IAIS-singular	0.18	1.31e-3	2.58e-3
+ IAIS-distributed	0.17	2.80e-3	2.72e-3
UNITER-large	0.23	0.40	0.16
+ IAIS-singular	0.18	2.27e-3	3.22e-3
+ IAIS-distributed	0.18	3.15e-3	3.70e-3

Table 2: Different relation distance metrics of each model after fine-tuning. Lower is better.

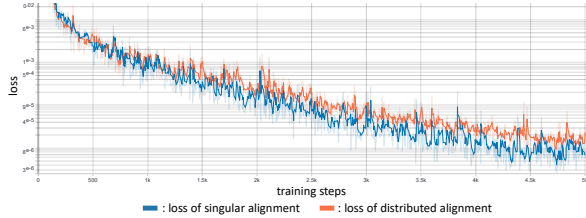


Figure 5: The singular and distributed version of IAIS loss with respect to the training steps.

relation coefficient of 0.779, which is also higher compared to $L_{\text{IAIS}}^{(d)}$ with 0.774. Besides, our empirical analysis in Figure 5 shows that it is slightly easier to optimize the $L_{\text{IAIS}}^{(s)}$, indicating it is a better surrogate of ISDa.

5.2 Effect of Anchor Modality

In Section 3.3, we leverage both the linguistic and the visual input as the anchor sequence to reconstruct the mirrored sequence from the opposite modalities. To study the impact of the anchor modality, we conduct an ablation study and the results are listed in Table 3. Compared to using language as the anchor modality, i.e., only $L_{\text{IAIS-L}}$ is incorporated, the overall model performance is 2.1 higher when vision is taken as the anchor. An explanation is that the description capability of visual regions is more concrete and powerful. However, introducing both $L_{\text{IAIS-V}} + L_{\text{IAIS-L}}$ to the final loss can achieve a further improvement of 2.22, which indicates the necessity of such combination.

5.3 Effect of Annealing Schedule

Besides the exp schedule in Eq. (16) for training signal annealing, we also try other schedules:

- log schedule: $\lambda_t = 1 - \exp(-t/T \times \gamma)$;
- linear schedule: $\lambda_t = t/T$;
- exp schedule: $\lambda_t = \exp((t/T - 1) \times \gamma)$,

where γ is chosen from $\{5, 10\}$. All the schedules are shown in Figure 6.

Model	Image Retrieval			Text Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
UNITER-base*	72.52	92.36	96.08	85.90	97.10	98.80
UNITER-base†	72.70	92.60	96.14	85.50	97.30	98.60
+ IAIS-singular-L	72.74	92.74	96.12	86.90	96.70	99.00
+ IAIS-singular-V	73.44	92.76	95.96	87.00	97.50	99.10
+ IAIS-distributed-L	72.48	92.96	96.26	86.90	97.10	99.10
+ IAIS-distributed-V	73.14	92.44	96.06	87.10	97.40	99.20

Table 3: Ablation study on the Flickr30k dataset. “-L” denotes that only $L_{\text{IAIS-L}}$ is incorporated, which regards language as the anchor modality. Similar for “-V”.

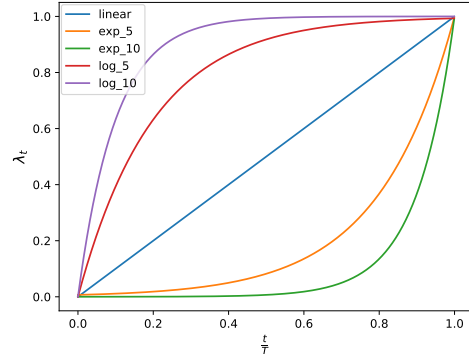


Figure 6: Schedules for IAIS signal annealing.

We compare the results of five schedules for IAIS signal annealing. The results in Figure 8 show that the exp schedule with scale $\gamma = 5$ achieves the best performance.

5.4 Effect of Layer to Apply IAIS

We also apply IAIS on different layers of UNITER-base. As illustrated in Figure 9, the optimal way is to apply IAIS on the last layer. We speculate that it is more important to learn relation alignment in the deeper layers because the attention in the deeper layers has a bigger impact on the final output, while the effect of the attention in shallow layers might fade away due to the normalization.

5.5 Case Study

We further discuss the advantage of our proposed relation-level alignment. Figure 7 shows two visualization examples of the intra-modal self-attentions from the Flickr30k Entities dataset. With IAIS regularization, the model is instructed to concentrate on the common relations within the linguistic and visual sequence, yielding more calibrated and consistent self-attention distributions.

6 Related Work

In this section, we introduce the task of image-text retrieval and review the representative studies of

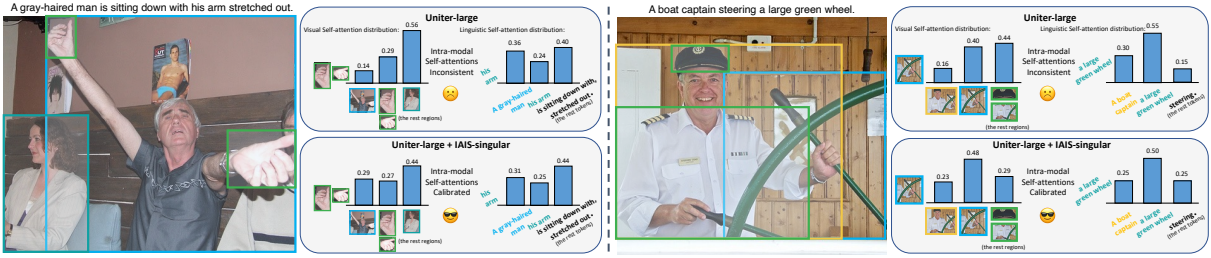


Figure 7: Visualization of intra-modal self-attentions with and without our IAIS method.

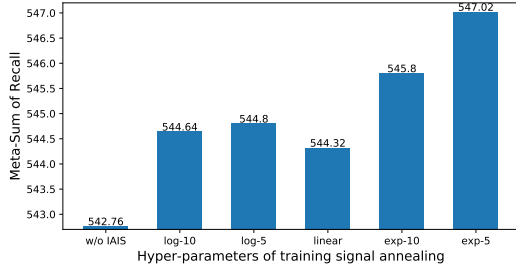


Figure 8: Comparison of the different hyper-parameters in training signal annealing. The exp schedule with scale $\gamma = 5$ achieves the best performance.

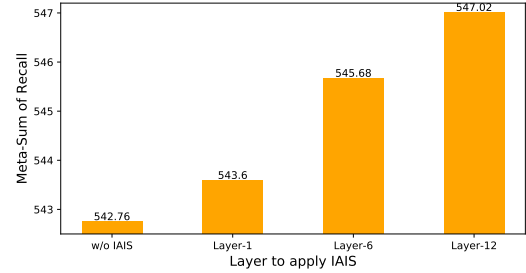


Figure 9: Comparison of the different layers to apply IAIS on the Flickr30k dataset. The IAIS applied on the last layer achieves the best performance.

large-scale multimodal pre-trained models.

Image-Text Retrieval Image-Text Retrieval (ITR, Barnard et al., 2003; Barnard and Forsyth, 2001), also known as Image-Text Matching, is one of the popular and challenging Language-and-Vision (V+L) tasks. Given image-text pairs, the prevailing approaches project them into a joint representation space, on which cosine or dot-product similarities are defined, and recall the most relevant one according to the similarity.

Multimodal Pre-trained Models The development of the transformer-based large-scale pre-training paradigm sweeps across the area of multimodal learning and achieves many state-of-the-art results on V+L tasks like Image Captioning, Visual Question Answering, Visual Commonsense Reasoning, etc. Recent prevailing multimodal pre-trained models can be categorized into single-stream (Chen et al., 2020; Gan et al., 2020; Lin et al., 2020; Li et al., 2020; Su et al., 2020; Lin et al., 2021) and two-stream (Yu et al., 2020; Tan and Bansal, 2019; Lu et al., 2019) models. Given a piece of text and an image, the former architecture concatenates the features of tokens and regions and learns their joint representations with one transformer model, while the latter embeds the textual and the visual input separately with two independent intra-modal transformers and then utilizes an

inter-modal transformer to reinforce cross-modal interactions via cross-modal attention modules.

7 Conclusion

In this paper, we promote the semantic alignment for cross-modal retrieval from the object level to the relation level. We propose a surrogate metric to quantify the relation consistency by measuring the semantic distance between linguistic and visual relations. Furthermore, we present a regularized training method IAIS to calibrate intra-modal self-attentions mutually by minimizing the ISD metric. Our method improves both the performance and the interpretability of large-scale pre-trained models. Note that, without object annotation in practice, the singular and distributed version of the IAIS loss only provides a coarse-grained attention distribution alignment. We leave the elaborate design of ISDa proxy function for future work.

Acknowledgments

This work is partly supported by Beijing Academy of Artificial Intelligence (BAAI). We thank all the anonymous reviewers for their constructive comments, and Xuancheng Ren and Lei Li for their helpful suggestions in preparing the manuscript.

References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image captioning and visual question answering](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. IEEE Computer Society.
- Kobus Barnard, Pinar Duygulu, David A. Forsyth, Nando de Freitas, David M. Blei, and Michael I. Jordan. 2003. [Matching words and pictures](#). *J. Mach. Learn. Res.*, 3:1107–1135.
- Kobus Barnard and David A. Forsyth. 2001. [Learning the semantics of words and pictures](#). In *Proceedings of the Eighth International Conference On Computer Vision (ICCV-01), Vancouver, British Columbia, Canada, July 7-14, 2001 - Volume 2*, pages 408–415. IEEE Computer Society.
- Emanuele Bugliarelli, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. 2020. [Multimodal pretraining unmasked: Unifying the vision and language bert](#)s. *CoRR*, abs/2011.15124.
- 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 *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 104–120. Springer.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of bert’s attention](#). *CoRR*, abs/1906.04341.
- Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. [VSE++: improving visual-semantic embeddings with hard negatives](#). In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 12. BMVA Press.
- Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. 2020. [Large-scale adversarial training for vision-and-language representation learning](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R. Bowman. 2019. [Do attention heads in BERT track syntactic dependencies?](#) *CoRR*, abs/1911.12246.
- Andrej Karpathy and Fei-Fei Li. 2015. [Deep visual-semantic alignments for generating image descriptions](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3128–3137. IEEE Computer Society.
- Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. 2014. [Unifying visual-semantic embeddings with multimodal neural language models](#). *CoRR*, abs/1411.2539.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 4364–4373. Association for Computational Linguistics.
- Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. [Stacked cross attention for image-text matching](#). In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, volume 11208 of *Lecture Notes in Computer Science*, pages 212–228. Springer.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020. [Oscar: Object-semantics aligned pre-training for vision-language tasks](#). In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 121–137. Springer.
- Junyang Lin, Rui Men, An Yang, Chang Zhou, Ming Ding, Yichang Zhang, Peng Wang, Ang Wang, Le Jiang, Xianyan Jia, Jie Zhang, Jianwei Zhang, Xu Zou, Zhikang Li, Xiaodong Deng, Jie Liu, Jinbao Xue, Huiling Zhou, Jianxin Ma, Jin Yu, Yong Li, Wei Lin, Jingren Zhou, Jie Tang, and Hongxia Yang. 2021. [M6: A chinese multimodal pretrainer](#). *CoRR*, abs/2103.00823.
- Junyang Lin, An Yang, Yichang Zhang, Jie Liu, Jingren Zhou, and Hongxia Yang. 2020. [Interbert: Vision-and-language interaction for multi-modal pretraining](#). *CoRR*, abs/2003.13198.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft COCO: common objects in context](#). In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer.
- Chunxiao Liu, Zhendong Mao, Tianzhu Zhang, Hongtao Xie, Bin Wang, and Yongdong Zhang. 2020. [Graph structured network for image-text matching](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10918–10927. IEEE.

- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23.
- Nicola Messina, Giuseppe Amato, Andrea Esuli, Fabrizio Falchi, Claudio Gennaro, and Stéphane Marchand-Maillet. 2020. [Fine-grained visual textual alignment for cross-modal retrieval using transformer encoders](#). *CoRR*, abs/2008.05231.
- 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 *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2641–2649. IEEE Computer Society.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. [Faster R-CNN: towards real-time object detection with region proposal networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [VL-BERT: pre-training of generic visual-linguistic representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: learning cross-modality encoder representations from transformers](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. 2019. [CAMP: cross-modal adaptive message passing for text-image retrieval](#). In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5763–5772. IEEE.
- Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. 2020. [Unsupervised data augmentation for consistency training](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Pengcheng Yang, Boxing Chen, Pei Zhang, and Xu Sun. 2020. [Visual agreement regularized training for multi-modal machine translation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9418–9425. AAAI Press.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Trans. Assoc. Comput. Linguistics*, 2:67–78.
- Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2020. [Ernie-vil: Knowledge enhanced vision-language representations through scene graph](#). *CoRR*, abs/2006.16934.
- Bowen Zhang, Hexiang Hu, Vihan Jain, Eugene Ie, and Fei Sha. 2020a. [Learning to represent image and text with denotation graph](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 823–839. Association for Computational Linguistics.
- Zhihan Zhang, Zhiyi Yin, Shuhuai Ren, Xinhang Li, and Shicheng Li. 2020b. [DCA: diversified co-attention towards informative live video commenting](#). In *Natural Language Processing and Chinese Computing - 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14-18, 2020, Proceedings, Part II*, volume 12431 of *Lecture Notes in Computer Science*, pages 3–15. Springer.