

Dense Contrastive Visual-Linguistic Pretraining

Lei Shi¹, Kai Shuang¹, Shijie Geng², Peng Gao³, Zuohui Fu², Gerard de Melo⁴
Yunpeng Chen⁵, Sen Su¹

¹Beijing University of Posts and Telecommunications, China

²Rutgers University, USA

³Shanghai AI Laboratory, China

⁴Hasso Plattner Institute, University of Potsdam, Germany

⁵YITU Technology, China

{SLei,shuangk,susen}@bupt.edu.cn,{sg1309,zuohui.fu}@rutgers.edu
gaopeng@pjlab.org.cn,gdm@demelo.org,yunpeng.chen@yitu-inc.com

ABSTRACT

Inspired by the success of BERT, several multimodal representation learning approaches have been proposed that jointly represent image and text. These approaches achieve superior performance by capturing high-level semantic information from large-scale multimodal pretraining. In particular, LXMERT and UNITER adopt visual region feature regression and label classification as pretext tasks. However, they tend to suffer from the problems of noisy labels and sparse semantic annotations, based on the visual features having been pretrained on a crowdsourced dataset with limited and inconsistent semantic labeling. To overcome these issues, we propose unbiased Dense Contrastive Visual-Linguistic Pretraining (DCVLP), which replaces the region regression and classification with cross-modality region contrastive learning that requires no annotations. Two data augmentation strategies (Mask Perturbation and Intra-/Inter-Adversarial Perturbation) are developed to improve the quality of negative samples used in contrastive learning. Overall, DCVLP allows cross-modality dense region contrastive learning in a self-supervised setting independent of any object annotations. We compare our method against prior visual-linguistic pretraining frameworks to validate the superiority of dense contrastive learning on multimodal representation learning.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; **Computer vision tasks**; **Computer vision**;

KEYWORDS

Visual-Linguistic Pretraining; Contrastive Learning; Adversarial Learning; Multimodal Applications

^{*}Work done during an internship at Yitu Tech.

[†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475637>

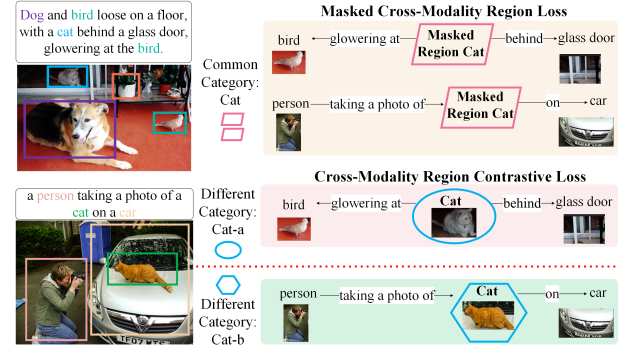


Figure 1: Different instances belonging to the same category will have different descriptions when in different scenes. However, the masked cross-modality region loss will classify them into the same category. Using the cross-modality region contrastive loss can help our model to distinguish different instances in the same category because it can fuse pertinent semantic information in different scenarios.

ACM Reference Format:

Lei Shi¹, Kai Shuang¹, Shijie Geng², Peng Gao³, Zuohui Fu², Gerard de Melo⁴, Yunpeng Chen⁵, Sen Su¹. 2021. Dense Contrastive Visual-Linguistic Pretraining. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3474085.3475637>

1 INTRODUCTION

Large-scale self-supervised pretraining of Transformers [8, 37] has had a transformative effect on Natural Language Processing (NLP). Models such as BERT [8] and RoBERTa [29] have led to strong gains across a wide swath of diverse NLP tasks. Inspired by this, Visual-Linguistic Pretraining (VLP) has been proposed to learn multimodal models covering both vision and language, by adding extra masked prediction self-supervised strategies for the visual branch [30, 46].

Among recent VLP approaches, LXMERT [46] and UNITER [5] are two prominent representatives. They perform feature regression or classification of masked visual regions as the pretext task of self-supervised learning.

However, we have identified several important drawbacks: (i) Noisy Labels: L_2 feature regression and classification suffer from noisy annotations in the Visual Genome dataset [21] used to train

these models. The visual area labels in the data come from Faster-RCNN, which generally has relatively low confidence, unlike text labels. (ii) Sparse Semantic Annotations: After establishing a relationship with the corresponding text caption information, it is often difficult to classify the visual region into a sparse label annotation. Especially when two different instances of the same category have divergent semantic information, a single classification is unable to capture the rich semantics, resulting in misleading gradients, as we can see in Figure 1. In addition, previous approaches use masked region labels for a training loss, which is not sufficiently robust with respect to domain shifts. For instance, region labels annotated on MSCOCO [27] and VG [21] have a large domain gap to NLVR2 [43], which consists of images collected online and thus has entirely different image manifolds compared with the sorts of images used for pretraining.

To overcome the aforementioned noisy label and sparse semantic annotation problems, we propose Dense Contrastive Visual-Linguistic Pretraining (DCVLP).

Specifically, DCVLP replaces the region regression and classification with contrastive learning, which simultaneously avoids both of the above problems. Contrastive learning [15] aims to discriminate between positive examples and negative ones, which does not require any annotation and thus sidesteps the noisy label and sparse semantic annotation issues. However, the strength of such comparative learning hinges on our ability to identify high-quality negative samples. In vision-language pre-training, the input visual features directly come from bottom-up features [1], so we cannot directly use the rotation, filter, and resize operations on the original images to generate positive and negative samples [3]. At the same time, classic forms of data augmentations in the field of visual contrastive learning tend to focus on natural attributes of the image itself. Thus, cross-modal visual language pre-training requires new data augmentation methods that can reflect high-level semantic changes in image context. To this end, we explore two new methods for generating positive and negative sample pairs, namely mask perturbation and intra-/inter-adversarial perturbation for vision region contrastive learning. Mask perturbation can make positive sample pairs remain similar in the absence of regions or texts. Intra-/inter-adversarial perturbation is able to make the positive sample pairs as similar as possible under the interference category. Our main contributions can be summarized as below:

- We propose a novel dense contrastive learning framework for visual-linguistic pretraining that avoids the sparse semantic annotation and noisy label problems encountered by previous visual-linguistic pretraining approaches such as LXMERT and UNITER.
- In order to discover fine-grained relationships between image objects and text descriptions, we first explore two new data augmentation methods, mask perturbation and intra-/inter-adversarial perturbation, to facilitate contrastive visual-linguistic pretraining.
- We carry out an extensive set of empirical studies over variants of DCVLP to validate our proposed approach. Our DCVLP pretraining achieves significant improvements over

strong baselines (LXMERT, UNITER) on downstream visual-linguistic tasks, especially when there is a domain gap between the pretraining and finetuning stages.

2 RELATED WORK

Self-supervised Learning in Vision and Language. Deep Neural Networks (DNN) trained on ImageNet [6] have revolutionized automatic feature representation learning [22]. Compared to supervised training, which incurs a substantial cost for data annotation, self-supervised learning learns useful features automatically by constructing a loss from a pretext task, which does not require human annotation. In computer vision, context encoders [36] learn features by image in-painting. Jigsaw [34] learns features by predicting the position of permuted features. Kolesnikov et al. [20] carry out a large-scale study of previously proposed self-supervised learning methods and show that the performance of self-supervised tasks varies as the backbone changes. In NLP, large-scale pretraining with next-word prediction (GPT) [37], next sentence prediction, or masked word prediction (BERT) [8], typically trained with the Transformer architecture [48], is now the predominant approach for state-of-the-art results in many NLP tasks. Motivated by the success of transformer-based self-supervised learning in both vision and language, DFAF [10], QBN [42], MCAN [52] and STSGR [11] has shown that transformer can perform efficient multi-modality fusion. UNITER [5] and LXMERT [46] have shown that masked words and visual regions can yield a good visual-linguistic representation.

Contrastive Learning & Adversarial Learning. Contrastive learning is a sub-branch of self-supervised learning, employing a contrastive loss to learn a representation that is useful in downstream tasks. The contrastive loss encourages the encoded instance features to be similar to positive keys while keeping away from negative ones. Different contrastive learning methods adopt different strategies to generate positive and negative keys, which is an essential factor for the quality of the learned representation. Wu et al. [50] select the keys from a large memory bank that stores the instance features for the entire training dataset. Some works [3, 47] generate keys using the current mini-batch samples. MoCo V1, V2 [4, 16] proposes a momentum encoder to generate the keys on-the-fly and store them in a fixed-size queue. However, in each iteration, only part of the information in the negative sample queue is updated. Different from MoCo, our dense contrastive learning uses the original form of contrastive learning without utilizing a memory bank. BYOL [14] no longer uses negative samples, but rather attempts to learn using only positive samples. AdCo [17] generates negative samples directly through training. Recently, CLIP [28] has successfully trained visual-linguistic representation using contrastive learning. Different from CLIP which simply utilizes global information during pretraining, our DCVLP employs dense contrastive learning which can better model fine-grained information. Neural networks are susceptible to adversarial attacks in the form of small perturbations [12, 32, 45]. Adversarial learning aims at improving the robustness of neural network models to defend against adversarial attacks. Madry et al. [31] conducts effective adversarial training through Projected Gradient Descent (PGD). Shafahi et al. [40] proposes to make neural network more robust compared

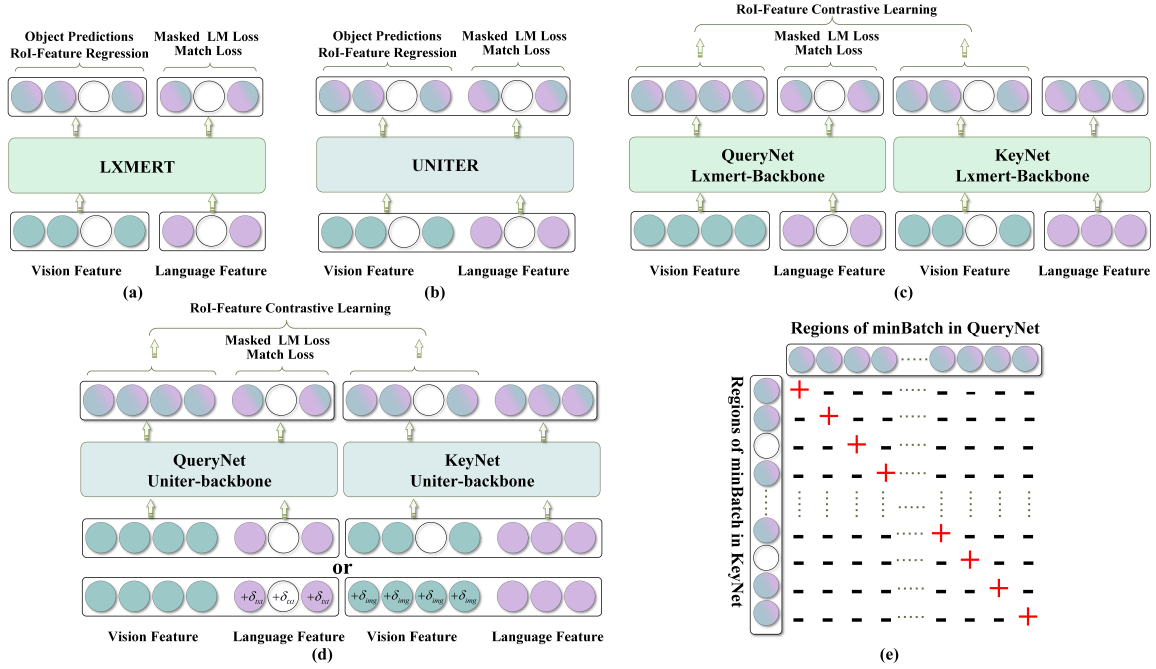


Figure 2: (a) represents the LXMERT model, which has dual-stream encoders. The visual and linguistic features are first processed in independent encoders and then use a cross-modal Transformer for multimodal information fusion. (b) stands for the UNITER model, a single-stream encoder, i.e., visual and linguistic features are directly cascaded and then encoded. (c), (d) describe the overall architecture of the proposed DCVLP approach with LXMERT-backbone and UNITER-backbone, respectively; The masked token is represented by white circles. DCVLP includes a Query Network and a Key Network. The entire model is trained with a combination of three cross-modality losses. Cross-Modality Contrastive loss is applied to all regions. δ_{img} , δ_{txt} in (d) denote adversarial perturbation. In (e), Regions in the same image position from QueryNet and KeyNet serve as positive pairs, while other regions are used to construct negative pairs. Since an image contains multiple object regions, the multiple positive pairs are from different image positions.

with naively trained neural networks by updating adversarial perturbations. FreeLB [54] introduces adversarial training to improve the robustness and generalization of pretrained language models on NLP tasks. Recently, VILLA [9] proposed a similar adversarial training strategy to benefit multimodal representation learning.

Semantically Dense Learning. Key tasks in image understanding include object detection, scene graph decomposition and image captioning. Single objects can contain dense semantic information when multiple objects appear in the image and information can flow between them. Recent research [2, 7, 38] investigates this aspect and shows that image captioning annotations allow us to learn better image representations than simple object category labels. Such labels also benefit several downstream tasks such as object detection and semantic segmentation with 10 times fewer annotations. In the VQA task [24], a novel implicit and explicit relation graph has been introduced to encode dense semantic information into a single object vector. Our proposed DCVLP framework can also infuse dense information into region and word vectors by drawing on contrastive learning.

3 DENSE CONTRASTIVE VISUAL-LINGUISTIC PRETRAINING

Figure 2 succinctly summarizes the original LXMERT and UNITER models in (a)–(b), the structure of the DCVLP model in (c)–(d), and the composition of the dense contrast loss in (e). As illustrated in Figure 2, (c)–(d), the architecture of DCVLP consists of a Query Network (QueryNet) and a Key Network (KeyNet). For LXMERT, they both contain a language Transformer encoder, a vision Transformer encoder and a multi-modal fusion Transformer. For UNITER, they only contain a single multi-modal self-attention Transformer encoder. LXMERT is a dual stream model, while UNITER is a single stream model. Our DCVLP is composed of QueryNet and KeyNet. We use either the LXMERT model or the UNITER model as the respective backbone of DCVLP in Figure 2, (c)–(d). KeyNet is copied from QueryNet with the same layers and parameters. In Figure 2, (a)–(b), both LXMERT and UNITER use mask region object predictions and mask region feature regression in the visual branch, which makes them susceptible to the influence of noisy labels. Thus, we use a new RoI-Feature Contrastive Loss to replace the original loss function of the visual branch in LXMERT and UNITER, and retain the classic Cross-Modality Match Loss, Masked Language Model Loss as later given in Eqn. (7) and (6).

3.1 Multi-modality Fusion

Given image–sentence pairs from a vision–language dataset, we first tokenize each sentence using the WordPieces technique [49] and map a token W_j to its corresponding embedding $h_{\text{emb}}(W_j) \in \mathbb{R}^{d_w}$, where $d_w = 768$. In addition, visual regions $B \in \mathbb{R}^{N \times 4}$ and their features $F \in \mathbb{R}^{N \times d_o}$ are extracted by a Faster-RNN [39] detector pretrained on Visual Genome [21] for each image I : $B, F = \text{RCNN}(I)$, where we detect $N = 36$ regions in the LXMERT-backbone (or $N = 10 \sim 100$ in UNITER-backbone) within each image and each such region is represented using a feature dimensionality of $d_o = 2048$. Then we can calculate the visual inputs v_i and textual inputs w_j of DCVLP as follows:

$$\begin{aligned} v_i &= \frac{g_F(F_i) + g_{P\text{-ROI}}(B_i)}{2} \\ w_j &= h_{\text{emb}}(W_j) + h_{P\text{-word}}(P_j) \end{aligned} \quad (1)$$

where g_F and $g_{P\text{-ROI}}$ are two fully-connected layers that map F_i and B_i , respectively, to the feature dimensionality d_w , while $h_{P\text{-word}}$ is a positional encoding function for the position P_j of token W_j .

Taking v_i and w_j as inputs, with mask perturbation, DCVLP adopts symmetrical masking for both QueryNet and KeyNet in Figure 2, (c)–(d). For the text modality, we uniformly choose 15% of the input textual tokens for replacement. The chosen tokens are replaced by the special *[MASK]* token. For the visual modality, we use a different masking strategy: the features of the chosen regions can either be set to zero or be replaced by region features from other images.

DCVLP contains two networks called QueryNet and KeyNet. QueryNet and KeyNet have exactly the same network structure. In the experiment, we can use the LXMERT or UNITER model as the backbone for QueryNet and KeyNet.

With LXMERT backbone for DCVLP in Figure 2-(c) the three encoders are implemented by 3 modules, namely, the visual self-attention, language self-attention and visual-linguistic co-attention modules. Visual self-attention performs information fusion between region features by using such features as both key, query and value in the attention model. We denote the key, query and value features for visual as K_v, Q_v, V_v , and for language as K_w, Q_w, V_w , respectively. Then the intra-modality information fusion for visual and language features can be denoted as:

$$\begin{aligned} \tilde{v} &= \text{Intra}_{v \leftarrow v}(Q_v, K_v, V_v) \\ \tilde{w} &= \text{Intra}_{w \leftarrow w}(Q_w, K_w, V_w) \end{aligned} \quad (2)$$

where the attention module of a Transformer layer can be expressed as follows:

$$\text{Attention}(Q, K, V) = \text{Softmax}(QK^T/\sqrt{d})V \quad (3)$$

After deploying intra-modality information flow for language and visual signals, we invoke an inter-modality fusion module to connect the signals in the language and visual features. The inter-modality fusion process is bi-directional, which includes information fusion from language to vision and vice versa:

$$\begin{aligned} \tilde{v} &= \text{Inter}_{v \leftarrow w}(Q_v, K_w, V_w) \\ \tilde{w} &= \text{Inter}_{w \leftarrow v}(Q_w, K_v, V_v) \end{aligned} \quad (4)$$

After intra-inter modality feature fusion, we can acquire a multi-modality contextual feature embedding for each word and visual

region. A contextual feature encodes the multimodal interactions in a compact feature vector. The contextual features are used by DCVLP for Modality Match Loss and Masked Cross-Modality Language Model Loss in the language branch and the RoI-Feature Contrastive Loss in the visual branch.

With UNITER-backbone for DCVLP, unlike LXMERT, UNITER is a single-stream model. Visual information and text information are cascaded, and then encoded by the self-attention module. Visual and language features can be denoted as:

$$\begin{aligned} \tilde{v} &= \text{Intra}_{v \leftarrow v+w}(Q_{v+w}, K_{v+w}, V_{v+w}) \\ \tilde{w} &= \text{Intra}_{w \leftarrow v+w}(Q_{v+w}, K_{v+w}, V_{v+w}) \end{aligned} \quad (5)$$

3.2 Cross-Modality Match Loss, Masked Language Model Loss for Language Branch

In the pretraining stage, DCVLP performs different pretext tasks compared with LXMERT and UNITER. DCVLP does not contain a supervised learning task and thus is independent of human-annotated labels. For the language branch, we keep masked language modeling and image–sentence matching prediction as two pretext tasks. The mask loss was first proposed by BERT [8]. Subsequent visual-linguistic BERT approaches [5, 25, 26, 30, 30, 46] add a visual feature mask loss besides the masked language modeling loss. This loss masks out the contextual representation obtained in Section 3.1 and predicts the masked feature using its contextual information. By optimizing the mask loss, the Transformer implicitly learns to encode contextual information, which facilitates the generalization on downstream tasks. In DCVLP, we only apply this mask loss for the textual inputs. Additionally, we also add a matching loss, which involves a binary Yes/No classification to predict whether the sentence matches the visual feature. The mask loss can be formalized as follows:

$$\mathcal{L}_{\text{MLM}} = -E_{w \sim D} \log P_{\theta}(w_m | \tilde{w}_m), \quad (6)$$

where θ denotes the parameters of the Language Transformer Encoder, w_m and \tilde{w}_m denote the masked token to be predicted and the contextual tokens that are not masked. The matching loss is defined as:

$$\begin{aligned} \mathcal{L}_{\text{MATCH}} &= -E_{w_{\text{CLS}} \sim D} [y \log P_{\theta}(\tilde{w}_{\text{CLS}}) \\ &\quad + (1 - y) \log (1 - P_{\theta}(\tilde{w}_{\text{CLS}}))], \end{aligned} \quad (7)$$

which is clearly a binary classification task. In the above equation, \tilde{w}_{CLS} stands for the *[CLS]* token, which encodes the visual-linguistic information for tackling the image–sentence matching pretext task.

3.3 RoI-Feature Dense Contrastive Loss for Visual Branch

Contrastive learning performs self-supervised representation learning by discriminating visually similar representation pairs from a group of negative ones. Given visual region features extracted by Faster-RCNN, we can obtain a positive query–key pair by feeding such features into both QueryNet and KeyNet. All region features from other images in the sampled batch and different positions serve as negative keys. Then we conduct contrastive learning by

updating network weights to minimize the following loss:

$$\mathcal{L}_C = -\log \frac{\exp(s^+/\tau)}{\exp(s^+/\tau) + \sum_{j=0}^K \exp(s_j^-/\tau)} \quad (8)$$

$$\begin{aligned} s^+ &= \widetilde{v_i^{\text{query}}} \cdot \widetilde{v_i^{\text{key+}}} \\ s^- &= \widetilde{v_i^{\text{query}}} \cdot \widetilde{v_j^{\text{key-}}} \end{aligned} \quad (9)$$

Here, τ is the temperature of the softmax operation, $\widetilde{v_i^{\text{key+}}}$ is a positive key of $\widetilde{v_i^{\text{query}}}$, and $\widetilde{v_j^{\text{key-}}}$ serves as negative examples for calculating $\mathcal{L}_{\text{CONTRAST}} (\mathcal{L}_C)$. In order to generate positive sample pairs, one simple strategy is to still use the mask operation, and another is to incorporate adversarial perturbation based on the PGD algorithm [31]. In the following, we outline these two methods.

3.3.1 Mask Perturbation Strategy for Visual Branch Dense Contrast Learning. In DCVLP, QueryNet and KeyNet have the same network structure, and we can use LXMERT or UNITER as the backbone for them. Specifically, in our experiments, we first consider a UNITER backbone as an example and finally also report the results of LXMERT as the backbone. In the mask perturbation DCVLP model, for QueryNet, we use a mask operation for textual information, while the visual information does not use any masking. Symmetrically, in KeyNet, we use masking for visual information, while the textual signals are not masked. The mask operation method is the same as in the UNITER model. The two sets of multimodal information processed by the mask are fed into QueryNet and KeyNet for encoding, respectively, and then the visual information is obtained from the last layer of QueryNet and KeyNet, respectively. The visual region information at the same position of QueryNet and KeyNet forms a pair of positive examples and the region information at different positions forms a negative example. We can consider different mask strategies as shown in Table 1. We empirically found that it is better to use the symmetric mask strategy for the multimodal information in QueryNet and KeyNet.

3.3.2 Intra Adversarial Perturbation Strategy For Visual Branch Dense Contrast Learning. We use the PGD attack algorithm to induce adversarial perturbation $\delta_{\text{img}}, \delta_{\text{txt}}$ for the visual region feature space and text embedding feature space. In other words, we aim to use the idea of contrastive learning to make the tokens in the same position as similar as possible when the adversarial perturbations $\delta_{\text{img}}, \delta_{\text{txt}}$ are injected. The intra-adversarial perturbation strategy generates $\delta_{\text{img}}, \delta_{\text{txt}}$ that directly maximize \mathcal{L}_C . Similar to the above mask strategy idea, we add adversarial perturbation symmetrically to QueryNet and KeyNet. The algorithmic details are formalized in Algorithm 1, where \mathcal{L}_C follows Eqn. (9). Symmetrically injecting adversarial perturbation in QueryNet and KeyNet, respectively, can alleviate the degree of simultaneous attacks on the network [19], which is naturally suitable for multimodal conditions. In the Intra-Adversarial Perturbation Strategy For Visual Branch Dense Contrast Learning (DCVLP_ADV-intra), we first use gradient ascent with regard to \mathcal{L}_C to generate adversarial perturbation in the adversarial training process, which may easily lead to inconsistent representations of tokens in the same position of

		QueryNet	KeyNet	VQA (Test dev)	NLVR2 (Test -P)
1	Text Image	Mask_1 Mask_a	Mask_1 Mask_a	72.45	77.30
2	Text Image	Mask_1 No_Mask	Mask_1 Mask_a	73.30	77.57
3	Text Image	Mask_1 No_Mask	Mask_2 Mask_a	73.14	78.45
4	Text Image	Mask_1 Mask_b	Mask_2 Mask_a	72.30	77.18
5	Text Image	Mask_1 No_mask	No_mask Mask_a	73.54	78.70

Table 1: QueryNet and KeyNet apply UNITER-backbone. For text modality, Mask-1, Mask-2 indicate that two Mask operations are performed on the text description information, and the tokens masked are different each time. For the image region feature, Mask-a, Mask-b indicate that two Mask operations are performed on the region contained in a picture, and the region masked is different each time. No-mask represents the use of original visual or textual information.

QueryNet and KeyNet. With multiple gradient ascent steps, we use the contrastive loss to minimize dissimilarities of the tokens in the same position. This process can track changes in the global sampled-batch signal across time and produce negative samples that are more difficult to distinguish.

3.3.3 Inter-Adversarial Perturbation Strategy For Visual Branch Dense Contrast Learning. The Intra-Adversarial Perturbation Strategy is a very direct method [17, 19], but its effect is found to be suboptimal in our experiments. As using unsupervised \mathcal{L}_C without a clear category for guidance to optimize $\delta_{\text{img}}, \delta_{\text{txt}}$ is very difficult, it is hard for \mathcal{L}_C to drive positive sample pairs to be similar with such attacks. Thus, we propose to instead use \mathcal{L}_{MLM} to generate $\delta_{\text{img}}, \delta_{\text{txt}}$. The algorithmic details are given in Algorithm 2. This form of perturbation attempts to increase the gradient of the masked language model, but we minimize \mathcal{L}_C to defend against such perturbation. Our approach is different from VILLA [9] because VILLA uses the same loss when generating perturbations and when defending against perturbation. However, that method will rely too much on the labels, and especially when there is an error or bias in the labels, the approach will not work well. Hence, we propose Inter-Adversarial Perturbation Strategy For Visual Branch Dense Contrast Learning (DCVLP_ADV-inter), which combines \mathcal{L}_{MLM} and \mathcal{L}_C to generate negative samples with moderate difficulty. It does not rely on any labels and at the same time promotes the joint optimization of both \mathcal{L}_{MLM} and \mathcal{L}_C .

4 EXPERIMENTS

In this section, we first introduce the implementation details of the proposed Dense Contrastive Visual-Linguistic Pretraining framework. Then we conduct extensive comparative studies to demonstrate the effectiveness of the proposed method. We apply the DCVLP algorithm to the LXMERT and UNITER frameworks, respectively. With LXMERT backbone for DCVLP, the pre-training data

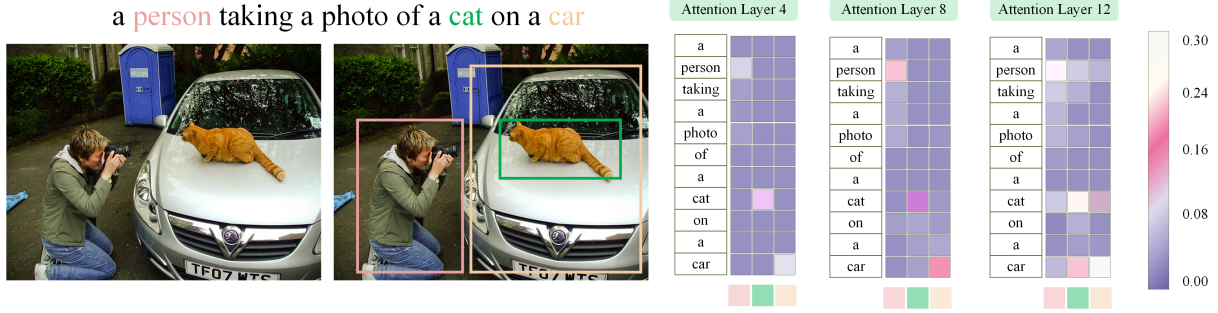


Figure 3: Illustration of attention weights in DCVLP with UNITER-backbone. The pink, green and orange in the right represents the people, cat and car bounding boxes in the left figure correspondingly. The lighter the color, the greater the weight of attention.

Algorithm 1 Multimodal Intra-Adversarial Training used in DCVLP

Input: \mathbf{v}, \mathbf{w} : clean image and text in sampled-batch; $\delta_{\text{img}}, \delta_{\text{txt}}$: adversarial perturbation in image and text; ϵ : perturbation bound; q, k represent *QueryNet* and *KeyNet* with parameters θ ; K : ascent steps; \mathcal{L}_C represent contrast loss.

Output: Parameters θ in *QueryNet* and *KeyNet*;

- 1: **while** $t < \text{Iter}_{\max}$ **do**
 - 2: **for** $i = 1 \dots K$ **do**
 - 3: Accumulate gradient of θ given $\delta_{\text{img}, i-1}$ and $\delta_{\text{txt}, i-1}$
 - 4: $\delta_{\text{img}}, \delta_{\text{txt}} = \arg \max_{\|\delta_{\text{img}}\|_{\infty} \leq \epsilon, \|\delta_{\text{txt}}\|_{\infty} \leq \epsilon} \langle \mathcal{L}_C \rangle$
 - 5: $\mathcal{L}_C \stackrel{\text{def}}{=} q(\mathbf{v}, \mathbf{w} + \delta_{\text{txt}}; \theta) \circ k(\mathbf{v} + \delta_{\text{img}}, \mathbf{w}; \theta)$
 - 6: Update $\delta_{\text{img}}, \delta_{\text{txt}}$ via gradient ascent
 - 7: Update parameters θ to minimize \mathcal{L}_C ;
-

Algorithm 2 Multimodal Inter-Adversarial Training used in DCVLP

Input: \mathbf{v}, \mathbf{w} : clean image and text in sampled-batch; $\delta_{\text{img}}, \delta_{\text{txt}}$: adversarial perturbation in image and text; ϵ : perturbation bound; q, k represent *QueryNet* and *KeyNet* with parameters θ ; K : ascent steps; \mathcal{L}_C represents contrast loss; \mathcal{L}_{MLM} represents masked language model loss.

Output: Parameters θ in *QueryNet* and *KeyNet*;

- 1: **while** $t < \text{Iter}_{\max}$ **do**
 - 2: **for** $i = 1 \dots K$ **do**
 - 3: Accumulate gradient of θ given $\delta_{\text{img}, i-1}$ and $\delta_{\text{txt}, i-1}$
 - 4: $\delta_{\text{img}}, \delta_{\text{txt}} = \arg \max_{\|\delta_{\text{img}}\|_{\infty} \leq \epsilon, \|\delta_{\text{txt}}\|_{\infty} \leq \epsilon} \langle \mathcal{L}_{\text{MLM}} \rangle$
 - 5: $\mathcal{L}_C \stackrel{\text{def}}{=} q(\mathbf{v}, \mathbf{w} + \delta_{\text{txt}}; \theta) \circ k(\mathbf{v} + \delta_{\text{img}}, \mathbf{w}; \theta)$
 - 6: Update $\delta_{\text{img}}, \delta_{\text{txt}}$ via gradient ascent
 - 7: Update parameters θ to minimize \mathcal{L}_C ;
-

is the same as LXMERT, namely MSCOCO and Visual Genome. To assess the learned visual-linguistic features, we conduct finetuning experiments and compare LXMERT backbone for DCVLP against the original LXMERT on three downstream tasks, i.e., VQA v2 [13],

GQA [18] and NLVR2 [44], as considered in the LXMERT paper. With UNITER-backbone for DCVLP, the pre-training data is the same as for UNITER. The UNITER model uses a large-scale V+L dataset composed of four subsets: (i) COCO [27]; (ii) Visual Genome (VG) [21] [21]; (iii) Conceptual Captions (CC) [41]; and (iv) SBU Captions [35]. We conduct finetuning experiments and compare UNITER-backbone for DCVLP against the original UNITER on six downstream tasks, including: (i) VQA; (ii) Visual Commonsense Reasoning (VCR) [53]; (iii) NLVR2; (iv) Visual Entailment [51]; (v) Image-Text Retrieval (including zero-shot setting) [23]; and (vi) Referring Expression Comprehension [33].

Implementation Details. Following LXMERT and UNITER, we pretrain DCVLP on the same image-text pairs datasets from LXMERT and UNITER, respectively. In the pretraining stage, for a fair comparison, we retain the same basic parameters as LXMERT and UNITER respectively. For the dense contrastive loss for the visual branch in DCVLP, the temperature τ in the contrastive loss is set to 0.07. The embedding head dimensionality is 128. In the mask perturbation strategy for the visual branch, the masking rate of both textual and visual information is 15%, while in the adversarial perturbation strategy, the gradient ascent step number is $K = 3$, adversarial text and adversarial image perturbation learning rate is 10^{-3} . Due to limited compute resources, we only conducted the mask perturbation strategy for visual branch contrast learning experiment on the LXMERT model, naming this DCVLP-MASK with LXMERT backbone. However, we applied the mask perturbation strategy for visual branch contrast learning and intra- or inter-contrastive adversarial perturbation strategy for visual branch contrast learning based on the UNITER model, namely, DCVLP-MASK with UNITER-backbone and DCVLP-ADV-intra/inter with UNITER-backbone. Results are given in Tables 2, 3, and 4.

4.1 Comparison with State-of-The-Art VLP Methods

In Table 2, previous methods adopt masked visual region classification and regression on the visual modality. DCVLP, in contrast, only needs a region contrastive learning loss on the visual modality and shows a competitive effect. DCVLP-MASK with LXMERT backbone and LXMERT have sample training data. DCVLP-MASK

Method	VQA		VCR			GQA	NLVR2	SNLI-VE
	test-dev	test-std	Q-A	QA-R	Q-AR	dev	test-P	test
Label Supervision								
ViLBERT	70.55	70.92	72.42(73.3)	74.47(74.6)	54.04(54.8)	-	-	-
VisualBERT	70.80	71.00	70.8(71.6)	73.2(73.2)	52.2(52.4)	-	67.0	-
LXMERT	72.42	72.54	-	-	-	61.39	74.50	-
Unicoder-VL	-	-	72.6(73.4)	74.5(74.4)	54.4(54.9)	-	-	-
12-in-1	73.15	-	-	-	-	-	78.87	76.95
VL-BERT-Base	71.16	-	73.8(-)	74.4(-)	55.2(-)	-	-	-
Oscar-Base	73.16	73.44	-	-	-	61.58	78.36	-
VILLA-Pre	73.03	-	74.76	77.04	57.82	-	78.44	78.43
UNITER-Base	72.70	72.91	74.56(75.0)	77.03(77.2)	57.76(58.2)	-	77.85	78.28
No Label Supervision								
DCVLP_MASK (LXMERT backbone)	73.07	73.15	-	-	-	61.78	76.81	-
DCVLP_MASK (UNITER backbone)	73.54	73.73	74.90(75.56)	77.12(77.3)	57.94(58.4)	-	78.70	78.70
DCVLP_ADV-intra (UNITER-backbone)	73.15	73.34	74.46(74.38)	76.89(77.08)	57.59(58.10)	-	78.13	78.34
DCVLP_ADV-inter (UNITER-backbone)	73.33	73.49	74.72(74.70)	77.00(77.28)	57.70(58.3)	-	78.43	78.61

Table 2: Results on VQA, VCR, GQA, NLVR2, and SNLI-VE; _MASK and _ADV-intra/inter denotes Mask and Intra/Inter Adversarial Perturbation Strategy respectively.

Method	RefCOCO+						RefCOCO					
	val	testA	testB	val (d)	testA (d)	testB (d)	val	testA	testB	val (d)	testA (d)	testB (d)
Label Supervision												
ViLBERT	-	-	-	72.34	78.52	62.61	-	-	-	-	-	-
VL-BERT	79.88	82.40	75.01	71.60	77.72	60.99	-	-	-	-	-	-
VILLA-Pre	-	-	-	-	-	-	-	-	-	-	87.34	74.35
UNITER_Base	83.66	86.19	78.89	75.31	81.30	65.58	91.64	92.26	90.46	81.24	86.48	73.94
No Label Supervision												
DCVLP_MASK (UNITER-backbone)	83.83	86.69	79.26	75.45	81.58	65.82	91.84	92.38	90.66	81.45	86.85	74.03
DCVLP_ADV-intra (UNITER-backbone)	83.32	86.00	78.73	75.02	81.15	65.41	91.33	92.02	90.21	80.96	86.27	73.50
DCVLP_ADV-inter (UNITER-backbone)	83.77	86.63	79.01	75.23	81.34	65.76	91.67	92.30	90.55	81.30	86.65	73.89

Table 3: Results on RefCOCO+ and RefCOCO. The (d) denotes evaluation using detected proposals

Method	RefCOCOg				Flickr30 IR			Flickr30k TR		
	val	test	val-d	test-d	R @ 1	R @ 5	R @ 10	R @ 1	R @ 5	R @ 10
Label Supervision										
ViLBERT	-	-	-	-	58.20	84.90	91.52	-	-	-
Unicoder-VL	-	-	-	-	71.50	90.90	94.90	86.20	96.30	99.00
VILLA-Pre	-	-	-	-	73.76	93.02	96.28	-	-	-
UNITER_Base	86.52	86.52	74.31	74.51	72.52	92.36	96.08	85.90	97.10	98.80
No Label Supervision										
DCVLP_MASK (UNITER-backbone)	86.47	86.62	74.55	74.98	73.60	92.66	95.98	86.33	97.45	99.02
DCVLP_ADV-intra (UNITER-backbone)	86.48	86.13	74.10	74.33	73.01	92.08	95.49	85.50	96.87	98.43
DCVLP_ADV-inter (UNITER-backbone)	86.23	86.45	74.42	74.62	74.03	92.78	96.27	86.51	97.66	99.11

Table 4: Results on RefCOCOg and Flickr30k Image Retrieval (IR) and Text Retrieval (TR).

with LXMERT backbone outperforms LXMERT by +0.65 on VQA, +0.4 on GQA, and +2.31 on NLVR2. DCVLP-MASK with UNITER-backbone and UNITER-Base have sample training data. Notably, DCVLP-MASK with UNITER-backbone outperforms UNITER-Base by +0.84 on VQA. VILLA [9] is the first to apply large-scale adversarial training for vision-and-language and achieve outstanding performance. VILLA consists of VILLA-Pre and VILLA-Finetune. As in previous research on vision-language pretrained models, we focus on the effects of the pretraining phase. Comparing VILLA-Pre and DCVLP with UNITER-backbone, our model obtains fairly strong results. The core defense strategy of DCVLP-ADV against perturbation is different from VILLA-pre. VILLA-pre utilizes the same loss when generating perturbations and defending against the perturbation. DCVLP-ADV is a new idea for vision-language pretraining, which not only enjoys the advantages of adversarial training but also avoids the use of potentially noisy labels in the adversarial training process. Although the proxy task in DCVLP-ADV is more challenging, the introduction of adversarial training can improve the robustness of the model.

Methods	VQA	RefCOCO+	NLVR2
No Vision Task	71.55	72.75	75.98
Feature Regression	71.73	73.49	76.21
Feature Regression + Label	71.92	74.52	76.93
DCVLP-MASK	72.90	75.03	77.26

Table 5: Comparison of different loss compositions on test-dev. splits of VQA, RefCOCO+ and NLVR2 with UNITER in-domain data.

4.2 Ablation Studies and Analyses of DCVLP

Effects of Loss Composition. In Table 5, we perform an ablation study on different loss combinations. “No vision task” method performs visual-linguistic pretraining without adding masks on the visual features. After replacing feature regression and label classification loss with contrastive loss, we can achieve the best performance of the ablation study on all three downstream tasks. This consolidates our claim that contrastive learning is more powerful. Through this implicit method, the visual region can be integrated with richer textual information instead of merely being classified into a limited category. This is particularly useful when the gap between pretraining and finetuning is large.

Effects of Different Mask Strategies in UNITER-backbone for DCVLP. From Table 1, we can see that using a symmetric mask strategy yields the best results. Part of the text information in QueryNet is masked, while the complete original text signal is retained in KeyNet. Symmetrically, the complete visual region information is retained in the QueryNet, and part of the visual region information in the KeyNet is masked. We speculate that the advantage of a symmetric mask strategy is to ensure that the visual area can still learn appropriate similar representations under the condition of missing text information and visual information respectively, which promotes the learning ability of both intra-modality and inter-modality. When QueryNet and KeyNet visual

modalities use different mask strategies, the model adds too much noise perturbation, which impairs the learning of visual features. Although this strategy can be used to differentiate between positive and negative examples, the region features in this case are not relevant to the ground-truth region features; when the ground-truth feature is present, the feature representation of the positive sample is more similar to the related ground-truth feature, so this situation is more suitable for downstream tasks that do not have the mask region feature.

Why Cross-Modality Region Contrastive Loss Works. In the pretraining process, QueryNet and KeyNet merge the corresponding text information for each region through an attention mechanism. Under normal circumstances, the corresponding positions of QueryNet and KeyNet for each visual region information should be as similar as possible. However, in our DCVLP model, the input visual information of KeyNet is applied with the mask or adversarial perturbation, so the original visual information is disturbed. In the QueryNet, each visual region feature establishes the intra-modality and inter-modality relationships through self-attention and co-attention; However, because of the mask or adversarial perturbation in the KeyNet, it is more difficult to establish the intra-modality and inter-modality relationships for each visual region feature. Therefore, the use of cross-modality region contrastive loss can encourage QueryNet and KeyNet to generate similar representations.

4.3 Visualizing DCVLP Encoder

In Figure 3, we visualize the attention weights in layers (i.e., the 4th, 8th, 12th layers) of DCVLP with UNITER-backbone. We can see that as the layer grows, the attention weight which indicates correct word-bounding box matching also increases gradually. The visual region not only gradually establishes a relationship with the target text category, but also establishes a certain relationship with related actions, directions, and category words in the text description, which well captures the densely semantic information of the visual region.

5 CONCLUSION

In this paper, we proposed dense contrastive visual-linguistic pretraining (DCVLP) in order to overcome the problems of noisy labels and sparse semantic annotations suffered by existing large-scale multimodal pretraining methods that utilize region feature regression and label classification. Moreover, end-to-end VLP will likely become a mainstream trend for VLP in the near future, but such an end-to-end training process has no opportunity to draw on classification labels for the image context. Hence, it appears natural and reasonable to combine our DCVLP method into end-to-end VLP training processes. Hence, we also conduct related basic experimental explorations.

ACKNOWLEDGMENTS

This work was supported by the Foundation for Innovative Research Groups of the National Natural Science Foundation of China (Grant No. 61921003).

REFERENCES

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6077–6086.
- [2] Mert Bulent Sariyildiz, Julien Perez, and Diane Larlus. 2020. Learning Visual Representations with Caption Annotations. *arXiv e-prints* (2020), arXiv–2008.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020).
- [4] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. 2020. Improved Baselines with Momentum Contrastive Learning. *arXiv preprint arXiv:2003.04297* (2020).
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740* (2019).
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [7] Karan Desai and Justin Johnson. 2020. VirTex: Learning Visual Representations from Textual Annotations. *arXiv preprint arXiv:2006.06666* (2020).
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [9] 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. *arXiv preprint arXiv:2006.06195* (2020).
- [10] Peng Gao, Zhengkai Jiang, Haoxuan You, Pan Lu, Steven CH Hoi, Xiaogang Wang, and Hongsheng Li. 2019. Dynamic fusion with intra- and inter-modality attention flow for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6639–6648.
- [11] Shijie Geng, Peng Gao, Moitrya Chatterjee, Chiori Hori, Jonathan Le Roux, Yongfeng Zhang, Hongsheng Li, and Anoop Cherian. 2021. Dynamic graph representation learning for video dialog via multi-modal shuffled transformers. In *AAAI Conference on Artificial Intelligence*.
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [13] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [14] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaoan Daniel Guo, Mohammad Gheshlaghi Azar, et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733* (2020).
- [15] Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, Vol. 2. IEEE, 1735–1742.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2019. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722* (2019).
- [17] Qianjiang Hu, Xiao Wang, Wei Hu, and Guo-Jun Qi. 2020. AdCo: Adversarial Contrast for Efficient Learning of Unsupervised Representations from Self-Trained Negative Adversaries. *arXiv preprint arXiv:2011.08435* (2020).
- [18] Drew A Hudson and Christopher D Manning. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. *Conference on Computer Vision and Pattern Recognition (CVPR)* (2019).
- [19] Ziyu Jiang, Tianlong Chen, Ting Chen, and Zhangyang Wang. 2020. Robust pre-training by adversarial contrastive learning. *arXiv preprint arXiv:2010.13337* (2020).
- [20] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. 2019. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 1920–1929.
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* 123, 1 (2017), 32–73.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [23] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked cross attention for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 201–216.
- [24] Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Relation-aware graph attention network for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*. 10313–10322.
- [25] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557* (2019).
- [26] Xijun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. *arXiv preprint arXiv:2004.06165* (2020).
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [28] Yang Liu, Alexandras Neophytou, Sunando Sengupta, and Eric Sommerlade. 2021. Cross-modal spectrum transformation network for acoustic scene classification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 830–834.
- [29] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [30] 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*. 13–23.
- [31] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rjlBfZAb>
- [32] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2017. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1765–1773.
- [33] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. 2016. Modeling context between objects for referring expression understanding. In *European Conference on Computer Vision*. Springer, 792–807.
- [34] Mehdi Noroozi and Paolo Favaro. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*. Springer, 69–84.
- [35] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems* 24 (2011), 1143–1151.
- [36] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2536–2544.
- [37] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- [38] Guanghui Ren, Lejian Ren, Yue Liao, Si Liu, Bo Li, Jizhong Han, and Shuicheng Yan. 2020. Scene graph generation with hierarchical context. *IEEE Transactions on Neural Networks and Learning Systems* 32, 2 (2020), 909–915.
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [40] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. 2019. Adversarial training for free! *arXiv preprint arXiv:1904.12843* (2019).
- [41] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2556–2565.
- [42] Lei Shi, Shijie Geng, Kai Shuang, Chiori Hori, Songxiang Liu, Peng Gao, and Sen Su. 2020. Multi-Layer Content Interaction Through Quaternion Product For Visual Question Answering. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4412–4416.
- [43] Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 217–223.
- [44] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2018. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491* (2018).
- [45] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [46] Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490* (2019).
- [47] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849* (2019).
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [49] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between

- human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [50] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3733–3742.
 - [51] Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706* (2019).
 - [52] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6281–6290.
 - [53] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6720–6731.
 - [54] Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. FreeLB: Enhanced Adversarial Training for Natural Language Understanding. In *ICLR*. <https://openreview.net/forum?id=BygzbyHFvB>