

Vision-Language Pre-Training with Triple Contrastive Learning

Jinyu Yang¹, Jiali Duan², Son Tran², Yi Xu², Sampath Chanda², Liqun Chen², Belinda Zeng²,
Trishul Chilimbi², and Junzhou Huang¹
¹University Of Texas at Arlington, ²Amazon

Abstract

Vision-language representation learning largely benefits from image-text alignment through contrastive losses (e.g., InfoNCE loss). The success of this alignment strategy is attributed to its capability in maximizing the mutual information (MI) between an image and its matched text. However, simply performing cross-modal alignment (CMA) ignores data potential within each modality, which may result in degraded representations. For instance, although CMA-based models are able to map image-text pairs close together in the embedding space, they fail to ensure that similar inputs from the same modality stay close by. This problem can get even worse when the pre-training data is noisy. In this paper, we propose triple contrastive learning (TCL) for vision-language pre-training by leveraging both cross-modal and intra-modal self-supervision. Besides CMA, TCL introduces an intra-modal contrastive objective to provide complementary benefits in representation learning. To take advantage of localized and structural information from image and text input, TCL further maximizes the average MI between local regions of image/text and their global summary. To the best of our knowledge, ours is the first work that takes into account local structure information for multi-modality representation learning. Experimental evaluations show that our approach is competitive and achieve the new state of the art on various common downstream vision-language tasks such as image-text retrieval and visual question answering.

1. Introduction

Self-supervision is an active research topic both in vision and language representation learning. Numerous methods have been proposed with an impressive performance on challenging tasks [4, 6, 9, 14, 16, 36]. A typical approach is to pre-train a model on massive amounts of unlabeled data in a self-supervised manner, then fine-tune it for downstream tasks (e.g., zero-shot learning and transfer learning) of interest. In the vision, self-supervision can be carried out us-

ing exemplars [12], predicting the relative position between two random patches [10] or via solving jigsaw [28]. In language, masked language modeling (MLM) is widely used as the method of choice for self-supervision.

Inspired by the success of self-supervision in individual modalities, there is a surging interest in self-supervised vision-language pre-training (VLP), which is essential for multi-modal tasks such as visual question answering (VQA), image-text retrieval, and visual entailment. These tasks heavily rely on joint multi-modal embeddings which are typically obtained by modeling interactions between vision and language features. To achieve this goal, various VLP frameworks are proposed by exploiting massive image-text pairs in the past few years [7, 13, 23, 24], where the key insight is to apply a fusion encoder to the concatenation of vision and language features to learn joint representations. Although simple and effective, this strategy suffers from the problem that vision and language features lie in different embedding spaces, which makes the feature fusion quite challenging [22]. To mitigate this problem, the most recent state of the art [22] disentangles the learning process into two stages: i) first align cross-modal features by using a contrastive loss (i.e., InfoNCE [29]) to pull the embeddings of matched image-text pairs together while pushing those of non-matched pairs apart; then ii) apply a fusion encoder to the aligned image and text representations to learn joint embeddings. Specifically, stage 1 aims to maximize the mutual information (MI) between matched image-text pair (I, T) through InfoNCE loss, which is spurred by the fact that I and T represent two "views" of the same semantic. However, the limitation of stage 1 lies in that: simply performing cross-modal alignment (CMA) cannot fully guarantee the expressiveness of the learned features that is essential for joint multi-modal representation learning. The main reason is that I and T are unable to fully describe each other. For instance, the text in (Figure 1 A) only focus on salient objects in the paired image, while overlooking other detailed and fine-grained information. To align I and T , only co-occurring features are captured by CMA. This is also evidenced by [19], where the performance of CMA-based features on image-text retrieval is far greater than on intra-modal retrieval (image-image and text-text). Further-

¹<https://github.com/uta-smile/TCL>

more, the pre-training datasets are usually collected from the web and are inherently noisy. This leads to learning degraded representations, where cross-modal features fail to capture certain key concepts.

As transformer became increasingly popular in both vision and language tasks, existing VLP methods adopted transformer architecture for extracting visual and linguistic features. Specifically, [CLS] tokens from the vision transformer (e.g., ViT [11]) and the text transformer (e.g., BERT [9]) are used to represent the global information of the input. For instance, ALBEF [22] maximizes MI between vision [CLS] and text [CLS]. However, global MI maximization fails to consider localized and structural information in the input [1, 17]. One potential side-effect is that it encourages the encoder to mainly extract information from certain unrelated/noisy image patches or text tokens that dominate MI.

In this paper, we introduce a novel VLP framework called triple contrastive learning (or TCL for short). **The key idea is to learn desirable representations by leveraging both cross-modal and intra-modal self-supervision, aiming to make it easier for the fusion encoder to learn multi-modal interactions.** To achieve this goal, TCL introduces three contrastive modules: cross-modal alignment (CMA), intra-modal contrastive (IMC), and local MI maximization (LMI), all of which rely on MI maximization. Specifically, i) CMA pulls the embeddings of matched image-text pairs together while pushing those of non-matched pairs apart by maximizing global MI between matched image and text; ii) complementary to CMA, IMC maximizes agreement between differently augmented views of the same data example through maximizing their global MI; iii) LMI encourages high MI between the global representation and every local region (e.g., image patches and text tokens) of the input, which is designed to remedy the side-effects that are introduced by the global MI maximization. The combination of these three modules allows us to i) learn representations that are semantically meaningful not only for cross-modal image-text pairs but also for intra-modal inputs; ii) capture the structural and localized information by extracting relevant features that are shared across local patches/tokens.

Our main contributions can be summarized as

- We leverage both cross-modal and intra-modal self-supervision to provide complementary benefits in representation learning, which facilitates modeling of better joint multi-modal features in the fusion encoder;
- Rather than simply rely on global information for multi-modal contrastive learning, we propose to take advantage of localized and structural information in both image and text input by maximizing local MI maximization between local regions and their global summary;

Comprehensive empirical studies demonstrate that TCL

achieves a new state of the art on a wide range of vision+language benchmarks, such as image-text retrieval and VQA. Specifically, on zero-shot image-text retrieval tasks, our method achieves significant improvement than ALIGN [19] (a mean recall of 79.5% vs 70.9% on MSCOCO). It is noteworthy that ALIGN is pre-trained on 1.8B image-text pairs, which is approximately 350× larger than TCL (5M). By pre-training TCL on a large-scale dataset with 14M image-text pairs, we observe a significant performance boost, implying its potential for further improvement with larger datasets. To investigate the effectiveness of each component in TCL, comprehensive ablation studies are also carried out with detailed analyses.

2. Related Work

Vision-Language Pre-training (VLP) Inspired by the success of self-supervised learning in intra-modal tasks, there is a surging interest in developing pre-training objectives for tasks with multiple modalities (e.g., vision and language). For instance, to leverage a much broader source of supervision from text, pioneering work CLIP [35] predicts which text goes with which image, resulting in a task-agnostic model that is even competitive with task-specific supervised models. ALIGN [19] further scales up CLIP by leveraging a noisy dataset that covers more than one billion image alt-text pairs. Despite these advancements, CLIP and ALIGN are mainly designed for vision-based downstream tasks and ignore the interaction between multiple modalities during pre-training. To fit in vision+language tasks (such as VQA [15] and visual reasoning), recent studies propose to learn joint multi-modal representations of image content and natural language. Among them, OSCAR [24], UNIMO [23], VILLA [13], and UNITER [7] use an object detector (e.g., Faster R-CNN [38]) to capture vision features first, then a multi-layer transformer [41] is applied to the concatenation of the extracted vision features and text features to learn joint embeddings. However, such kind of strategy suffers from limitations such as i) extracting region features using an object detector is computationally inefficient and ii) the quality of visual features is largely limited by the predefined visual vocabulary in pre-trained object detectors. To address this issue, rather than rely on region-based visual features, SOHO [18] takes a whole image as input and extracts compact image features through a visual dictionary, which favors 10 times faster inference time than region-based methods. ViLT [20] totally discards convolutional visual features and adopt vision transformer [11] to model long-range dependencies over a sequence of fixed-size non-overlapping image patches. Although these aforementioned methods achieve remarkable performance, they fail to conduct image-text alignment before fusion, which makes it challenging to learn the interaction between different modalities.

To remedy this, ALBEF [22] applies a contrastive loss

to align image and text features before modeling their joint representations, which delivers the state-of-the-art performance. Our method shares similar spirits with ALBEF, but with clear differences as follows: i) instead of only performing cross-modal alignment (CMA), we propose to leverage both cross-modal and intra-modal self-supervision to enforce the learned representations are semantic meaningful. The rationale is that cross-modal alignment alone may result in feature degeneration problem: although features from different modalities are well-separated, those from the same modality fall within a narrow cone and have high similarity. ii) we introduce local alignment to the cross-modal scenario by maximizing mutual information (MI) between local regions and global representations. Compared with the global alignment strategy used in ALBEF, maximizing local MI encourages our model to learn features that are shared across image patches/text tokens. Furthermore, local alignment prevents simply capturing noise or unrelated features.

Mutual Information Maximization Mutual information (MI) aims to measure the relationship between random variables or determine the amount of shared information. MI is widely used in unsupervised feature learning, where the key idea is to maximize MI between the input and output [26]. However, the MI estimation of high dimensional random variables is quite difficult and intractable [31], especially for deep neural networks. To this end, MINE [2] exploits dual optimization to offer a general-purpose estimator of MI. Another alternative is InfoNCE [29], which is a categorical cross-entropy loss that identifies the positive sample amongst a set of negative samples. InfoNCE is proved to be a lower bound of MI [29], such that minimizing InfoNCE loss can indirectly maximize MI. However, existing studies simply maximize MI between the complete input and output (i.e., global MI maximization), which is proven to be insufficient for meaningful representation learning [17]. DIM [17] addresses this limitation by introducing local MI, i.e., maximizing the average MI between local regions of image input and the encoder output. AMDIM [1] further extends DIM to maximize MI between features extracted from independent augmentations of the same image.

Both DIM and AMDIM are conducted in intra-modal tasks. In contrast, we introduce local MI maximization for multi-modal problems to benefit cross-modal representation learning. Specifically, we encourage high MI between the global representation and every local region (e.g., image patches and text tokens) of the input. This enables more transferable representations which are evidenced by our empirical studies. Furthermore, rather than using a CNN-based network, our local MI is built upon transformer architecture. Therefore, the sequential patch tokens in the transformer actually give us free access to local features without the need to pull local information from intermediate layers. Our experiments show that patch embeddings from the

last layer outperform intermediate-layer patches with transformer backbone.

3. Method

In this section, we first describe the model architecture of our method (Figure 1), followed by the uni-modal representation learning. After that, we detail the proposed triple contrastive learning modules: cross-modal alignment (CMA), Intra Modal Contrastive (IMC), and Local MI maximization (LMI). In the end, we brief two pre-training objectives, i.e., image-text matching (ITM) and masked language modeling (MLM).

3.1. Model Architecture

An overview of our method is shown in Figure 1, which contains a vision encoder $g(\cdot)$ for learning visual features from the image input, a text encoder $h(\cdot)$ for learning linguistic features from the text input, and a fusion encoder for learning multi-modal interactions. All of these encoders adopt transformer-based architecture [41], which are detailed in section 4.3. For each encoder, we maintain a paired momentum encoder that is implemented by momentum-based moving average strategy by following the same setting in [16]. Formally, $\theta_{\hat{g}} = m\theta_{\hat{g}} + (1 - m)\theta_g$, where $\hat{g}(\cdot)$ is momentum vision encoder, $m \in [0, 1]$ is a momentum coefficient. Similarly, we use $\hat{h}(\cdot)$ to denote the momentum text encoder. Uni-modal encoders $g(\cdot)$ and $h(\cdot)$ are used to learn robust visual and linguistic features from the given input, then an alignment module is applied to the learned features to align both cross-modal and intra-modal representations before fusion. We detail each component in the following sections.

3.2. Uni-modal Representation Learning

Given an image-text pair (I, T) , two separate augmentations are applied to the image to obtain two correlated “views”, i.e., I_1 and I_2 . Following [4, 16], we consider two random “views” of the same image under random data augmentation as a positive pair. Each augmented image is split into fixed-size patches which are then linearly mapped and embedded with positional information [11]. Similar to BERT [9], a class token [CLS] is prepended to the image patches, serving as the representation of the whole image. The obtained sequential embeddings of I_1 are finally fed into $g(\cdot)$ to learn desired visual representations $\{v_{cls}, v_1, \dots, v_M\}$, where M is the total number of image patches. For I_2 , we use $\hat{g}(\cdot)$ to learn its representations $\{\hat{v}_{cls}, \hat{v}_1, \dots, \hat{v}_M\}$. For the text input T , we follow [9] to obtain $\{t_{cls}, t_1, \dots, t_N\}$ and $\{\hat{t}_{cls}, \hat{t}_1, \dots, \hat{t}_N\}$ by $h(T)$ and $\hat{h}(T_+)$, where N is the length of text tokens, $T_+ = T$.

To model the interaction between image and text features, previous VLP work directly apply a fusion encoder to the concatenation of $\{v_{cls}, v_1, \dots, v_M\}$ and $\{t_{cls}, t_1, \dots, t_N\}$




Figure 1. (A): An overview of our framework which consists of a vision encoder, a text encoder, and a fusion encoder. Each encoder has a paired momentum encoder updated by the momentum-based moving average. For the image input, we apply two separate data augmentation operators (a and b) which are sampled from the same family of augmentations. The alignment module contains three contrastive objectives (i.e., CMA, IMC, and LMI) for both cross-modal and intra-modal representation learning (make it easier for the fusion encoder to learn joint multi-modal embeddings). (B): The motivation of leveraging both cross-modal and intra-modal supervision. The original image (pink) is augmented to two different views (green). For CMA only, the middle image only has a positive text example (green) and treats other texts (red) as negatives. Its embedding (blue circle) would be close to its positive text example. By incorporating IMC, it has two positive examples (one text and one image) and two sets of negative examples (one from text and one from image) and tends to learn more reasonable embeddings (blue square).

to learn joint multi-modal embeddings. However, the most obvious drawback of this strategy is that visual and linguistic features lie in different embedding spaces, which is challenging for the fusion encoder to learn their interactions [22]. To alleviate this limitation, we propose an alignment module that is applied to the learned visual and linguistic features before fusion. Specifically, our alignment module contains three contrastive learning objectives, i.e., CMA, IMC, and LMI. We discuss each objective below and show that they play a complementary role in the feature alignment and can benefit multi-modal feature fusion.

3.3. Cross-Modal Alignment (CMA)

The goal of CMA is to pull embeddings of the matched image-text pair (sampled from the joint distribution) together while pushing those of unmatched pairs apart (sampled from the product of marginal distributions). In other words, CMA aims to maximize the MI between the image and text that are matched, which are assumed to describe the same semantic meaning. For instance, the text in Figure 1 (A) describes high-level information (e.g., the occurrence of certain events or presence of certain objects) in the paired image. Since direct maximization of MI for continuous and high-dimensional variables is intractable [2], we instead minimize InfoNCE loss [29] which represents the lower bound of MI. Formally, the InfoNCE loss for image-to-text is defined as:

$$\mathcal{L}_{nce}(I_1, T_+, \tilde{T}) = -\mathbb{E}_{p(I, T)} \left[\log \frac{e^{\text{sim}(I_1, T_+)/\tau}}{\sum_{k=1}^K e^{\text{sim}(I_1, \tilde{T}_k)/\tau}} \right] \quad (1)$$

where τ is a temperature hyper-parameter, $\tilde{T} = \{\tilde{T}_1, \dots, \tilde{T}_K\}$ is a set of negative text examples that are not matched to I_1 , $\text{sim}(I_1, T_+) = f_v(v_{cls})^T \hat{f}_t(\hat{t}_{cls})$, where $f_v(\cdot)$ and $\hat{f}_t(\cdot)$ are two projection heads that map representations to the space where InfoNCE loss is applied. To maintain the negative text samples \tilde{T} , following [22], we use a large queue that keeps the most recent K projected representations $\hat{f}_t(\hat{t}_{cls})$. Similarly, the loss of text-to-image is formulated by:

$$\mathcal{L}_{nce}(T, I_2, \tilde{I}) = -\mathbb{E}_{p(I, T)} \left[\log \frac{e^{\text{sim}(T, I_2)/\tau}}{\sum_{k=1}^K e^{\text{sim}(T, \tilde{I}_k)/\tau}} \right] \quad (2)$$

where $\text{sim}(T, I_2) = f_t(t_{cls})^T \hat{f}_v(\hat{v}_{cls})$, $f_t(\cdot)$ and $\hat{f}_v(\cdot)$ are two projection heads. $\tilde{I} = \{\tilde{I}_1, \dots, \tilde{I}_K\}$ is a queue of negative image examples which store the most recent K projected features $\hat{f}_v(\hat{v}_{cls})$. Taken together, we define the loss of CMA as:

$$\mathcal{L}_{cma} = \frac{1}{2} [\mathcal{L}_{nce}(I_1, T_+, \tilde{T}) + \mathcal{L}_{nce}(T, I_2, \tilde{I})] \quad (3)$$

Intuitively, by minimizing \mathcal{L}_{cma} , we encourage the visual features and linguistic features to be aligned well in the embedding space and in turn ease the feature fusion.

ALBEF [22] applies a special case of \mathcal{L}_{cma} by setting $I_1 = I_2$ and achieves the state-of-the-art performance. However, CMA loss ignores the self-supervision within each modality, thus failing to guarantee the desirable expressiveness of learned features. The reason is that i) text usually cannot fully describe the paired image. For instance, although the text in Figure 1 (A) captures most of the salient

objects in the image, it overlooks detailed features of each object, such as the cloth of the man. Therefore, simply pulling embeddings of image-text pair together results in degraded representations (Figure 1 B); and ii) image-text pairs used for pre-training are inherently noisy, which makes the problem in i) even worse. To mitigate these limitations, we propose to further make use of intra-modal self-supervision by introducing Intra-Modal Contrastive (IMC) objective as follows.

3.4. Intra-Modal Contrastive (IMC)

Different from CMA, IMC attempts to learn the semantic difference between positive and negative samples within the same modality. For the visual modality, we consider two random “views” (I_1, I_2) of the same image I under random data augmentation as a positive pair. Following [4, 16], we maximize agreement between (I_1, I_2) by using the contrastive loss $\mathcal{L}_{nce}(I_1, I_2, \tilde{I})$. Similar to Equation 2, we define $\text{sim}(I_1, I_2) = f_v(v_{cls})^T \hat{f}_v(\hat{v}_{cls})$.

For the text input, we follow [14] to take a text and predict itself in a contrastive objective. This is achieved by considering standard dropout as minimal data augmentation for the text, and applying independently sampled dropout masks for identical positive pairs, i.e., $T_+ = T$. Different from [14] that uses in-batch negatives, we use the same negative text queue \tilde{T} in Equation 1 instead. The contrastive objective can be described by $\mathcal{L}_{nce}(T, T_+, \tilde{T})$, where $\text{sim}(T, T_+) = f_t(t_{cls})^T \hat{f}_t(\hat{t}_{cls})$. Overall, we minimize the following objective to guarantee reasonable intra-modal representation learning.

$$\mathcal{L}_{imc} = \frac{1}{2}[\mathcal{L}_{nce}(T, T_+, \tilde{T}) + \mathcal{L}_{nce}(I_1, I_2, \tilde{I})] \quad (4)$$

Specifically, our model is encouraged to learn representations that keep alignment between semantically-related positive pairs within a modality. Most importantly, \mathcal{L}_{imc} enforces the uniformity of the whole representation space of image and text such that the embeddings are uniformly distributed [42]. Therefore, CMA and IMC are designed to play a complementary role in the representation learning: i) CMA maps matched image-text pair close in the embedding space, and ii) IMC maximizes agreement between differently augmented views of the same data example. Combining them together improves the quality of the learned representations (Figure 1 B) and can further facilitate joint multi-modal learning in the fusion encoder.

One limitation of IMC is that it simply performs the contrastive objective on [CLS] tokens of vision encoders and text encoders, where [CLS] tokens are assumed to represent the global information of the input. In other words, IMC maximizes the global MI between differently augmented views. However, the drawbacks of global MI maximization lie in that: i) it ignores the localized and structural information in the input [1, 17]; ii) certain unrelated local regions may dominate the MI, resulting in the model

that is biased to learning unrelated features. For instance, noisy patches can represent a larger “quantity” of information than semantic-meaningful patches that occur repeatedly [17]. To remedy this issue, we introduce local MI maximization into multi-modal representation learning as detailed below.

3.5. Local MI Maximization (LMI)

The goal of local MI maximization is to encourage high MI between the global representation and every local region (e.g., image patches and text tokens) of the input. Rather than considering the [CLS] token pair (e.g., (v_{cls}, \hat{v}_{cls}) of (I_1, I_2)) as a positive pair, we pair [CLS] token from one augmented version, with patch embeddings in the other independently augmented version of the input. Without loss of generality, we take the vision input (I_1, I_2) as an example. Specifically, we consider $\{\hat{v}_i\}_{i=1}^M$ as positive examples of v_{cls} , while patch embeddings from other images in the same batch are used to build up negative examples. Similarly, $\{\hat{t}_j\}_{j=1}^N$ are considered as positive examples of t_{cls} , while text tokens from other in-batch texts are negative examples. We maximize the average MI between global and local regions by minimizing the following loss:

$$\mathcal{L}_{lmi} = \frac{1}{2} \left[\frac{1}{M} \sum_{i=1}^M \mathcal{L}_{nce}(I_1, I_2^i, \tilde{I}_l) + \frac{1}{N} \sum_{j=1}^N \mathcal{L}_{nce}(T, T_+^j, \tilde{T}_l) \right] \quad (5)$$

where $\text{sim}(I_1, I_2^i) = f_v(v_{cls})^T \hat{f}_v(\hat{v}_i)$, $\text{sim}(T, T_+^j) = f_t(t_{cls})^T \hat{f}_t(\hat{t}_j)$, \tilde{I}_l and \tilde{T}_l are in-batch negative image and text patch embeddings, respectively. Therefore, minimizing \mathcal{L}_{lmi} allows our model to encode the representations of data that are shared across all patches, rather than encoder representation from certain patches which dominate MI. Another perspective is that local MI maximization encourages the model to predict local from the global representation, which forces the model to also capture fine-grained information and in turn to benefit joint representation learning.

3.6. Image-Text Matching (ITM)

To fuse vision and language representations, we adopt ITM which is widely used in previous VLP studies. Given an image-text pair, ITM predicts whether they are matched (positive examples) or not (negative examples), which can be regarded as a binary classification problem. Following [22], the fusion encoder takes $\{v_{cls}, v_1, \dots, v_M\}$ and $\{t_{cls}, t_1, \dots, t_N\}$ as input. We use [CLS] token of the fusion encoder as the joint representation of the input image-text pair, which is then fed into a fully-connected layer to predict the matching probability $\phi(I, T)$. We assume that each image-text pair (I, T) sampled from the pre-training datasets is a positive example (with label 1) and construct negative examples (with label 0) through batch-sampling [22]. The ITM loss is defined as:

$$\mathcal{L}_{itm} = \mathbb{E}_{p(I, T)} H(\phi(I, T), y^{(I, T)}) \quad (6)$$

	COCO	VG	SBU	CC	CC12M
# images	113K	100K	859K	2.92M	10.97M
# text	567K	769K	859K	2.92M	10.97M

Table 1. Statistics of pre-training datasets.

where $H(\cdot)$ is the cross-entropy, $y^{(I,T)}$ denotes the label.

3.7. Masked Language Modeling (MLM)

We adopt MLM from BERT [9], which aims to predict the ground truth labels of masked text tokens T^{msk} . Specifically, we randomly mask out text tokens with a probability of 15%, and replace them with a special [MASK] token 80% of the time, and 10% with random words, and leave it unchanged for the remaining 10% of the time. Different from BERT, our MLM is conditioned on both surrounding text tokens of T^{msk} and image representations. The MLM loss is defined as:

$$\mathcal{L}_{mlm} = \mathbb{E}_{p(I, T^{msk})} H(\Phi(I, T^{msk}), y^{T^{msk}}) \quad (7)$$

where $\Phi(I, T^{msk})$ is the predicted probability of T^{msk} , and $y^{T^{msk}}$ is ground truth.

The overall training objective of our model is:

$$\mathcal{L} = \mathcal{L}_{cma} + \mathcal{L}_{imc} + \mathcal{L}_{lmi} + \mathcal{L}_{itm} + \mathcal{L}_{mlm} \quad (8)$$

4. Experiments

To verify the effectiveness of our method, we conduct comprehensive studies on commonly used benchmarks and present experimental comparisons against state-of-the-art VLP methods as shown below.

4.1. Pre-training Datasets

Following previous experimental protocols [7, 22], we use COCO [25], Visual Genome (VG) [21], Conceptual Captions (CC) [39], and SBU Captions [30] as the pre-training dataset in our study, where a total of 4.0M unique images and 5.1M image-text pairs are covered. We term this dataset as a 4M dataset in our study. To prove that our method is able to be applied to large-scale datasets, we further use CC12M [3]. Together with the 4M dataset, we, therefore, reach large-scale pre-training data with 14.97M unique images and 16M image-text pairs (Table 1).

4.2. Downstream Tasks

Image-Text Retrieval includes two tasks: (1) image as query and text as targets (TR); (2) text as query and image as targets (IR). The pre-trained model is evaluated on Flickr30K [33] and COCO [25] by following both fine-tuning and zero-shot settings. For the fine-tuning setting, the pre-trained model is fine-tuned on the training data and evaluated on the validation/test data. For the zero-shot setting, the pre-trained model is directly evaluated on the test

data. In particular, for zero-shot retrieval on Flickr30K, we follow [22] to evaluate the model fine-tuned on COCO.

Visual Question Answering (VQA) [15] aims to predict the answer given an image and a question (in text format), which requires an understanding of vision, language, and commonsense knowledge to answer. We consider this task as a generation problem by following the same setting in [22]. Specifically, an answer decoder is fine-tuned to generate the answer from the 3,192 candidates.

Visual Entailment (SNLI-VE) [43] predicts whether a given image semantically entails a given text, which is a three-classes classification problem. Specifically, the class or relationship between any given image-text pair can be entailment, neutral, or contradictory. Compared with VQA, this task requires fine-grained reasoning.

Visual Reasoning (NLVR²) [40] determines whether a natural language caption is true about a pair of photographs. We evaluate our model on NLVR² dataset which contains 107,292 examples of human-written English sentences paired with web photographs. Since this task takes a text and two images as input, we extend our model by following [22].

4.3. Implementation Details

All of our experiments are performed on 8 NVIDIA A100 GPUs with PyTorch framework [32]. Our vision encoder is implemented by ViT-B/16 with 12 layers and 85.8M parameters. Both the text encoder and the fusion encoder are implemented by a 6-layer transformer and are initialized by the first 6 layers and the last 6 layers of BERT_{base} which contains 123.7M parameters. We set $K = 65,536$ and $m = 0.995$. For the pre-training stage, the model is trained for 30 epochs with a batch size of 512. We use mini-batch AdamW optimizer [27] with a weight decay of 0.02. The learning rate is initialized as $1e - 5$ and is warmed up to $1e - 4$ after 2,000 training iterations. We then decrease it by the cosine decay strategy to $1e - 5$. For data augmentation, a 256×256 -pixel crop is taken from a randomly resized image and then undergoes random color jittering, random grayscale conversion, random Gaussian Blur, random horizontal flip, and RandAugment [8]. During the fine-tuning stage, the image resolution is increased to 384×384 and the positional encoding is interpolated according to the number of image patches.

4.4. Evaluation on Image-Text Retrieval

To assess the generalization of the learned representations, the common practice is to perform the zero-shot transfer of the trained model to downstream tasks. We evaluate our model by benchmarking the zero-shot image-text retrieval tasks on Flickr30K and COCO datasets by following the standard evaluation protocol. As shown in Table 2, our

Method	#Images	MSCOCO (5K)						Flickr30K (1K)					
		Text Retrieval			Image Retrieval			Text Retrieval			Image Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
ImageBERT [34]	6M	44.0	71.2	80.4	32.3	59.0	70.2	70.7	90.2	94.0	54.3	79.6	87.5
UNITER [7]	4M	\times	\times	\times	\times	\times	\times	80.7	95.7	98.0	66.2	88.4	92.9
ViLT [20]	4M	56.5	82.6	89.6	40.4	70.0	81.1	73.2	93.6	96.5	55.0	82.5	89.8
CLIP [35]	400M	58.4	81.5	88.1	37.8	62.4	72.2	88.0	98.7	99.4	68.7	90.6	95.2
ALBEF [22]	4M	68.7	89.5	94.7	50.1	76.4	84.5	90.5	98.8	99.7	76.8	93.7	96.7
Ours	4M	71.4	90.8	95.4	53.5	79.0	87.1	93.0	99.1	99.6	79.6	95.1	97.4
ALIGN [19]	1.2B	58.6	83.0	89.7	45.6	69.8	78.6	88.6	98.7	99.7	75.7	93.8	96.8

Table 2. Performance comparison of zero-shot image-text retrieval on Flickr30K and COCO datasets. For completeness, we also provide the results of ALIGN [22] which uses 1.8B image-text pairs for pre-training. For text-retrieval (TR) and image-retrieval (IR), we report the average of R@1, R@5 and R@10.

Method	#Images	MSCOCO (5K)						Flickr30K (1K)					
		Text Retrieval			Image Retrieval			Text Retrieval			Image Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
ImageBERT [34]	6M	66.4	89.8	94.4	50.5	78.7	87.1	87.0	97.6	99.2	73.1	92.6	96.0
UNITER [7]	4M	65.7	88.6	93.8	52.9	79.9	88.0	87.3	98.0	99.2	75.6	94.1	96.8
VILLA [13]	4M	\times	\times	\times	\times	\times	\times	87.9	97.5	98.8	76.3	94.2	96.8
OSCAR [24]	4M	70.0	91.1	95.5	54.0	80.8	88.5	\times	\times	\times	\times	\times	\times
ViLT [20]	4M	61.5	86.3	92.7	42.7	72.9	83.1	83.5	96.7	98.6	64.4	88.7	93.8
UNIMO [23]	4M	\times	\times	\times	\times	\times	\times	89.7	98.4	99.1	74.7	93.47	96.1
SOHO [18]	200K	66.4	88.2	93.8	50.6	78.0	86.7	86.5	98.1	99.3	72.5	92.7	96.1
ALBEF [22]	4M	73.1	91.4	96.0	56.8	81.5	89.2	94.3	99.4	99.8	82.8	96.7	98.4
Ours	4M	75.6	92.8	96.7	59.0	83.2	89.9	94.9	99.5	99.8	84.0	96.7	98.5
ALIGN [19]	1.2B	77.0	93.5	96.9	59.9	83.3	89.8	95.3	99.8	100.0	84.9	97.4	98.6

Table 3. Performance comparison of fine-tuned image-text retrieval on Flickr30K and COCO datasets. For completeness, we also provide the results of ALIGN [22] which uses 1.8B image-text pairs for pre-training.

Method	#Images	VQA		NLVR ²		SNLI-VE	
		test-dev	test-std	dev	test-P	val	test
OSCAR [24]	4M	73.16	73.44	78.07	78.36	\times	\times
UNITER [7]	4M	72.70	72.91	77.18	77.85	78.59	78.28
ViLT [20]	4M	71.26	\times	75.7	76.13	\times	\times
UNIMO [23]	4M	73.29	74.02	\times	\times	80.0	79.1
VILLA [13]	4M	73.59	73.67	78.39	79.30	79.47	79.03
ALBEF [22]	4M	74.54	74.70	80.24	80.50	80.14	80.30
Ours	4M	74.90	74.92	80.54	81.33	80.51	80.29

Table 4. Performance comparison on vision+language tasks.

approach achieves the best performance while outperforming the existing state-of-the-art by a large margin. Compared with ViLT [20] which directly uses a transformer encoder to model the interaction between word and image patch embeddings, we improve +9.5% (average) on COCO and +12.2% (average) on Flickr30K, revealing the necessity of conducting cross-modal alignment before fusion. ALBEF [22] is closely related to our work, which aligns image and text embeddings first, then uses a fusion encoder to learn joint representations. Furthermore, ALBEF shares the same pre-training datasets with our method, thus making them comparable. However, ALBEF ignores the intra-modal supervision, therefore the expressiveness of the

learned features cannot be guaranteed. Compared with ALBEF, our method brings +2.5% TR/R@1 boost and +2.8% IR/R@1 boost on Flickr30K dataset by explicitly leveraging the intra-modal information from both global and local perspectives. It is worth mentioning that our method demonstrates a significant improvement over ALIGN [19], i.e., a mean of 79.5% vs 70.9% on COCO and 94.0% vs 92.2% on Flickr30K. Note, ALIGN is pre-trained on 1.8B image-text pairs which is approximately 360 \times times more image-text pairs than our model. This observation suggests that our method is more data-efficient which is mainly attributed to the consideration of intra-modal supervision. Overall, the representations learned by our method are more general and transferable than existing baselines.

For fine-tuned experiments, we set up new benchmark results as shown in Table 3. On the medium-sized COCO dataset, we surpass ALBEF [22] by 2.5% absolute TR/R@1 and 2.2% absolute IR/R@1, revealing that our model can further benefit from fully-supervised training. We also compete favorably against prior baselines on the small-sized Flickr30K dataset. The only exception is ALIGN [19], which outperforms our method by +0.48% averaged (89.69% vs 89.21%) on COCO and Flickr30K, while at the expense of huge computational resources. This is especially

Module	Zero-Shot				Fine-Tune			
	MSCOCO		Flickr30K		MSCOCO		Flickr30K	
	TR	IR	TR	IR	TR	IR	TR	IR
CMA+ITM+MLM	68.7*	50.1*	90.5	76.8	73.1	56.8	94.3	82.8
+IMC (w/o aug)	71.1	52.2	92.0	78.6	75.0	58.6	94.5	82.9
+IMC	71.4	53.3	92.1	78.9	75.6	58.8	95.1	83.1
+IMC+LMI (Ours)	71.4	53.5	93.0	79.6	75.6	59.0	94.9	84.0

Table 5. Ablation study of each component on image-text retrieval tasks. The R@1 is reported. For CMA+ITM+MLM, we use the results in ALBEF [22]. Since the results of zero-shot on COCO is not provided, we use the pre-trained model from ALBEF to generate these numbers which are denoted by *.

Pooling	Intermediate	Zero-Shot				Fine-Tune			
		MSCOCO		Flickr30K		MSCOCO		Flickr30K	
		TR	IR	TR	IR	TR	IR	TR	IR
		71.5	52.9	92.4	79.1	75.7	58.6	94.6	83.3
	✓	71.4	52.9	91.5	77.9	75.7	58.6	94.4	82.3
✓	✓	71.8	53.2	93.2	79.2	75.6	58.7	94.8	82.8
✓		71.4	53.5	93.0	79.6	75.6	59.0	94.9	84.0

Table 6. Ablation study of image patch pooling and intermediate local feature on image-text retrieval. R@1 is reported.

Module	Zero-Shot				Fine-Tune			
	MSCOCO		Flickr30K		MSCOCO		Flickr30K	
	TR	IR	TR	IR	TR	IR	TR	IR
+IMC (w/o aug) (4M)	71.1	52.2	92.0	78.6	75.0	58.6	94.5	82.9
+IMC (w/o aug) (14M)	72.7	54.1	94.6	83.6	77.9	60.9	96.2	86.0

Table 7. Ablation study of the size of pre-training datasets. R@1 is reported.

problematic for the scenario/researchers with limited budgets. We believe that our method can also largely benefit from a much larger pre-training dataset, which is evidenced in section 4.6.

4.5. VQA, VE, and NLVR²

Table 4 shows the performance comparison on VQA, VE, and NLVR² which requires image+text as inputs. In other words, to make succeed in these tasks, the model is supposed to have the capability in learning joint multi-modal embeddings. Among five out of six criteria, we deliver state-of-the-art results, implying that explicitly considering cross-modal alignment and intra-modal supervision contribute to the feature fusion.

4.6. Ablation Study

To learn the effectiveness of the newly proposed modules (i.e., IMC and LMI) in improving the multi-modal representation learning, we perform ablation studies on image-text retrieval tasks shown in Table 5. Since ALBEF [22] is implemented by using loss function CMA+ITM+MLM, we

thus use the results in ALBEF as the baseline. We investigate two choices of IMC, i) IMC (w/o aug): only random crop, random horizontal flip, and RandomAugment are applied to the input image and set $I_1 = I_2$ by following ALBEF; ii) IMC: I_1 and I_2 are two augmented views of the input image with stronger data augmentation as discussed in section 4.3. Both strategies can improve the performance by a large margin, while stronger augmentation works better that is consistent with previous studies [4, 5]. The performance is further improved by incorporating LMI, indicating the importance of localized and structural information in representation learning.

During the pre-training stage, each image is split into 256 patches with a size 16×16 . To rule out the probability that small image patches may not contain enough information for local MI maximization, we apply global average pooling to the last-layer patch embeddings $\{\hat{v}_1, \dots, \hat{v}_M\}$, resulting in $M = 16$ patches for LMI in Equation 5. To maintain the spatial relationship among patches, we reshape $\{\hat{v}_1, \dots, \hat{v}_M\}$ to the original 3D image space then apply the pooling operation. Notably, different from [17] which uses feature maps from the intermediate layer as local information, we use patch embeddings pulled from the last layer. We examine these two choices on image-text retrieval tasks as shown in Table 6 and observe the importance of image patch pooling. In addition, the performance of using last-layer patches is comparable to, if not better than, using patches from intermediate layers (i.e., 9th layer in $\hat{g}(\cdot)$ and 4th layer in $\hat{h}(\cdot)$). We suspect that the difference between features learned by CNNs and vision transformers lead to this observation [37].

To study the impact of training on larger-scale datasets, we perform an ablation study on 14M datasets by using +IMC (w/o aug) as shown in Table 7. We could clearly see that the larger scale dataset gave a significant boost in performance. We hypothesize that our model has the potential for further improvement if pre-trained on further large datasets.

5. Conclusion

In this paper, we propose a new vision-language pre-training framework named TCL (short for triple contrastive learning). Different from previous studies that simply align image and text representations through a cross-modal contrastive loss, TCL further considers intra-modal supervision to guarantee that the learned representations are also meaningful within each modality, and in turn benefits cross-modal alignment and joint multi-modal embedding learning. To incorporate the localized and structural information in representation learning, TCL further introduces the local MI which maximizes the mutual information between the global representation and the local information from image patches or text tokens. Experimental results on widely used benchmarks show that TCL outperforms existing SOTA methods by a large margin.

References

- [1] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019. **2, 3, 5**
- [2] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018. **3, 4**
- [3] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. **6**
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. **1, 3, 5, 8**
- [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. **8**
- [6] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. **1**
- [7] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. **1, 2, 6, 7**
- [8] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. **6**
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **1, 2, 3, 6**
- [10] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, 2015. **1**
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **2, 3**
- [12] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence (PAMI)*, 2015. **1**
- [13] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *arXiv preprint arXiv:2006.06195*, 2020. **1, 2, 7**
- [14] Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*, 2021. **1, 5**
- [15] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913, 2017. **2, 6**
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. **1, 3, 5**
- [17] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018. **2, 3, 5, 8**
- [18] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12976–12985, 2021. **2, 7**
- [19] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918*, 2021. **1, 2, 7**
- [20] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. *arXiv preprint arXiv:2102.03334*, 2021. **2, 7**
- [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. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. **6**
- [22] Junnan Li, Ramprasaath R Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *arXiv preprint arXiv:2107.07651*, 2021. **1, 2, 4, 5, 6, 7, 8**
- [23] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. *arXiv preprint arXiv:2012.15409*, 2020. **1, 2, 7**
- [24] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer, 2020. **1, 2, 7**
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. **6**

- [26] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988. 3
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [28] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision (ECCV)*, 2016. 1
- [29] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 1, 3, 4
- [30] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24:1143–1151, 2011. 6
- [31] Liam Paninski. Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253, 2003. 3
- [32] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [33] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 6
- [34] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*, 2020. 7
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 2, 7
- [36] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018. 1
- [37] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *arXiv preprint arXiv:2108.08810*, 4, 2021. 8
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 2
- [39] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 6
- [40] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Hua-jun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018. 6
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 3
- [42] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020. 5
- [43] Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*, 2019. 6