

GRIT-VLP: Grouped Mini-batch Sampling for Efficient Vision and Language Pre-training

Jaeseok Byun^{1*}, Taebaek Hwang^{2*}, Jianlong Fu³, and Taesup Moon^{1**}

¹ Department of ECE/ASRI/IPAI, Seoul National University

² Department of ECE, Sungkyunkwan University

³ Microsoft Research Asia

Abstract. Most of the currently existing vision and language pre-training (VLP) methods have mainly focused on how to extract and align vision and text features. In contrast to the mainstream VLP methods, we highlight that two routinely applied steps during pre-training have crucial impact on the performance of the pre-trained model: *in-batch* hard negative sampling for image-text matching (ITM) and assigning the large masking probability for the masked language modeling (MLM). After empirically showing the unexpected effectiveness of above two steps, we systematically devise our GRIT-VLP, which adaptively samples mini-batches for more effective mining of hard negative samples for ITM while maintaining the computational cost for pre-training. Our method consists of three components: 1) GGrouped mIni-baTch sampling (GRIT) strategy that collects similar examples in a mini-batch, 2) ITC consistency loss for improving the mining ability, and 3) enlarged masking probability for MLM. Consequently, we show our GRIT-VLP achieves a new state-of-the-art performance on various downstream tasks with much less computational cost. Furthermore, we demonstrate that our model is essentially in par with ALBEF, the previous state-of-the-art, only with one-third of training epochs on the same training data. Code is available at <https://github.com/jaeseokbyun/GRIT-VLP>.

Keywords: Efficient vision and language pre-training, hard negative sampling, batch-sampling strategy, shuffling

1 Introduction

Recently, the pre-training and fine-tuning approach of the Transformer [35] based models have made exciting progress in natural-language-processing (NLP) [6] and vision tasks [8]. Particularly, despite the huge computational cost, vision and language pre-training (VLP) [20,33,13,27,4,17,22,31,21,9], which aims to learn cross-modal representations from large-scale image-text pairs, enabled to achieve the state-of-the-art results in various vision and language downstream tasks, *e.g.*,

* Equal contribution. This work was performed when Jaeseok Byun did an internship at Microsoft Research Asia.

** Corresponding author (tmoon@snu.ac.kr)

Anchor	Positive	Negative		
Text	Image	Random	Hard (in-batch)	Hard (Ours)
a zebra is walking in a field in some grass				
Image	Text	Random	Hard (in-batch)	Hard (Ours)
	two surfers are riding the waves in the ocean	a large long train on a steel track	a boat full of people in the ocean	two surfers carry their boards into the ocean

(a)
(b)
(c)

Fig. 1: A comparison of negative samples for ITM task selected by (a) Random, (b) *In-batch* sampling by ALBEF [20], and (c) Our GRIT strategy.

image-text retrieval (IRTR), natural language for visual reasoning (NLVR) [32], and visual question answering (VQA) [1], etc. For the joint understanding of image and text, a multi-modal encoder used in VLP is typically trained with the self-supervised learning objectives, such as image-text matching (ITM) and masked language modeling (MLM).

Majority of the existing VLP methods have focused on how to make the vision features to align with those of the text. The first popular approach [4,33,24,22] is to utilize the salient region-based features extracted from a pre-trained object detector. However, these region feature based VLP methods suffer from severe computational inefficiency and heavy dependency on the pre-trained object detectors. In order to overcome such drawbacks, recent approaches have replaced the object detectors with CNN backbones [14,13] or linear embedding inspired by the recently developed Vision Transformer (ViT) [8], which enables efficient end-to-end training of the vision-language representation learning.

Recently, ALBEF [20] was proposed as another attempt to lift the dependency on the object detectors. They designed a novel VLP architecture to integrate the uni-modal encoder for each modality (*i.e.*, an object-detector-free vision encoder and a text encoder) by employing a multi-modal Transformer encoder that fuses features from them. Additionally, ALBEF employed the image-text contrastive (ITC) loss for uni-modal encoders to *pre-align* the features before fusing, the *in-batch* hard negative sampling strategy for the ITM, and a momentum distillation to further improve the performance. As a result, it achieved the state-of-the-art performance for the multiple vision and language downstream tasks.

While the main emphasis of [20] was on the pre-aligning stage via ITC, we double-check that proposition and carry out careful ablation analyses on ALBEF and identify that the two routinely applied sampling steps in fact have crucial impacts on the final downstream performance. Firstly, the *hard negative sampling* for the ITM task, of which effect was described as marginal compared to the pre-aligning in [20, Section 6], in fact turns out to be an essential component,

even more than the ITC itself, for efficient VLP. Namely, when training for the ITM task, rather than using the randomly selected negatives as shown in Fig. 1(a) (for the text and image anchor, respectively), selecting *hard* negative samples as in Fig. 1(b), which is sampled from the given *mini-batch* by using the contrastive image-text similarity scores already computed for the ITC, becomes much more useful for promoting a more fine-grained representation learning. Secondly, the *mask sampling probability* for the MLM task, which typically is naively set to 15%, also is shown to have a significant impact on the VLP performance. Namely, when the probability is enlarged up to 50%, the multi-modal encoder is enforced to use more visual context features for predicting the masked token, hence, a more consolidated multi-modal representations could be learned. A recent concurrent work [37] also suggests the enlarging of the masking probability for MLM; however, their focus was on the NLP domain, thus, they have not investigated the impact of enlarging the probability on the multi-modal encoder.

Motivated by above analyses, in this paper, we make the following three modifications on ALBEF to significantly improve the downstream performance and computational efficiency of VLP. First, we devise GRIT (GRouped mIni-baTch sampling) strategy that enables to select much more informative hard negative samples (as shown in Fig. 1(c)) than those in [20] (Fig. 1(b)), *without* introducing any significant memory and computational overhead. Note such improvement is far from being straightforward since a naive extension of previous approaches would require either additional GPU memory (when simply enlarging the batch size) or forward pass computation (when utilizing additional queues as in [36,12,40]). We elaborate on this point more in details in a later section. Second, we devise a consistency loss between the image-text similarity scores used for ITC such that the contrastive learning and pre-aligning become more effective and, as a result, enables our GRIT to sample more exquisite negative samples. Third, we use enlarged mask sampling probability (50%) for MLM such that the visual features can be further integrated with the text features when solving the downstream tasks.

Our final method that combines above modifications is dubbed as GRIT-VLP, and we show that it can significantly improve the efficiency of VLP compared to ALBEF. Namely, trained on the exact same training data, GRIT-VLP significantly outperforms ALBEF on all of the downstream tasks we tested with 33% fewer number of epochs, 21% less training time per epoch. Furthermore, our thorough analyses show that GRIT-VLP is model agnostic and can be easily applied to existing VLP with different model architectures and objectives, which demonstrates the potential of our method being an essential tool for VLP.

2 Preliminaries and Related Work

[Vision-language pre-training] Existing VLP methods, which can be categorized into three frameworks, have mainly focused on the development of objectives and architectures to learn multi-modal representations. The first approach is to adopt dual uni-modal encoders which are composed of separate image and text

encoder. CLIP [27] and ALIGN [15] pre-trained with contrastive learning have been shown to be effective for IRTR, without object detectors. However, they suffer from the performance degradation in other downstream tasks (*e.g.*, VQA, NLVR). The second approach [33,24,31,22,19] mainly utilizes a single multi-modal encoder where concatenated text and image representations are used as input. In contrast to the former approach, these works consistently show promising results on various downstream tasks. However, these methods heavily depend on the pre-trained object detectors which are computationally inefficient. Thus, recent works [14,13,17,44] have struggled to replace object detectors with more efficient ones. The last category [20,45] offsets the shortcomings of the previous approaches by combining them, and achieves state-of-the-art performance. ALBEF [20] combines them by adding pre-alignment before fusing. Our method is built upon this ALBEF [20], but, deviating from the mainstream of VLP, our attention is on the *sampling strategy* for efficient pre-training.

[Hard negative mining] Most prior works on negative mining [38,11,29,42,43] point out that hard negatives can help a training model to converge faster. Recent approaches [28,5,39,3,46] mainly focus on the unsupervised contrastive learning setting where true dissimilarity of pairs are not available. However, these methods can not be applied to the VLP methods (second, third categories in the previous paragraph) due to the inherent architecture and input of multi-modal encoder.

2.1 ALign BEfore Fuse (ALBEF) [20]

Since ALBEF is the base model on which we build our method, we review it in details here. It consists of an image encoder f_v , a text encoder f_t , and a multi-modal encoder h , all of which are based on the Transformer architecture. Each input image V and sentence T is encoded into respective embedding sequences: $f_v(V) = \{v^{cls}, v^1, v^2, \dots, v^{S_V}\}$ and $f_t(T) = \{t^{cls}, t^1, t^2, \dots, t^{S_T}\}$, in which v^{cls} and t^{cls} denote the embedding of the [CLS] token for each modality, and S_V and S_T denote the sequence length of image and text, respectively. Then, vision and text representations are fused by a cross-attention module in the multi-modal encoder which requires both vision and text features as input (*i.e.*, $h(f_v(V), f_t(T)) = \{w^{cls}, w^1, w^2, \dots, w^{S_T}\}$). The three pre-training objectives of ALBEF are briefly introduced below⁴.

(a) Image-text contrastive learning (ITC) focuses on the pre-alignment of uni-modal representations before fusing them with a multi-modal encoder. Like conventional contrastive learning, it promotes positive image-text pairs to have similar representations and negative ones to be dissimilar. Inspired by MoCo [12], ALBEF utilizes two *queues* for storing recent [CLS] embeddings from the unimodal encoders, *i.e.*, v^{cls} and t^{cls} , and use them as extra negatives for the contrastive learning. More specifically, a similarity between V and T is defined as $s(V, T) = g_v(v^{cls})^T g_t(t^{cls})$ in which $g_v(\cdot)$ and $g_t(\cdot)$ are linear projections for mapping [CLS] embeddings to the normalized lower dimensional features. Then,

⁴ Note the Momentum Distillation (MD), which utilizes the soft outputs from an additional momentum model is omitted, since we do NOT use the momentum model.

for each V and T , the normalized image-to-text and text-to-image similarities for $j = 1, \dots, N$ are defined as:

$$p_j^{v2t}(V) = \frac{\exp(s(V, T_j)/\tau)}{\sum_{j=1}^N \exp(s(V, T_j)/\tau)}, \quad p_j^{t2v}(T) = \frac{\exp(s(V_j, T)/\tau)}{\sum_{j=1}^N \exp(s(V_j, T)/\tau)}, \quad (1)$$

in which τ is a learnable temperature, and N is the size of the queue. The ITC loss is then defined as:

$$\mathcal{L}_{\text{ITC}} = \frac{1}{2} \mathbb{E}_{(V, T) \sim D} [\text{CE}(\mathbf{y}^{v2t}(V), \mathbf{p}^{v2t}(V)) + \text{CE}(\mathbf{y}^{t2v}(T), \mathbf{p}^{t2v}(T))], \quad (2)$$

in which $\mathbf{y}^{v2t}(V)$ and $\mathbf{y}^{t2v}(T)$ denotes the ground-truth one-hot vector for the true pair sample for V and T , respectively. Now, in the pre-training, we do *not* use the queues for ITC but use the *in-batch* version, *i.e.*, N in (1) is the size of the mini-batch, to implement a lightweight version in terms of memory/computation. **(b) Image-text matching (ITM)** is a binary classification task that predicts whether a pair of image and text, (V, T) , is matched or not. The prediction probability of the classifier, $\mathbf{p}^{\text{ITM}}(V, T)$, is obtained by using the joint embedding feature of [CLS] token (w^{cls}) from the multi-modal encoder. Then, the ITM loss is defined as

$$\mathcal{L}_{\text{ITM}} = \mathbb{E}_{(V, T) \sim D} [\text{CE}(\mathbf{y}^{\text{ITM}}, \mathbf{p}^{\text{ITM}}(V, T))]. \quad (3)$$

in which \mathbf{y}^{ITM} is the ground truth one-hot vector, and $\text{CE}(\cdot, \cdot)$ stands for the cross-entropy between the two probability vectors. The effectiveness of ITM is determined by the quality of the negative pair, and, as outlined in the Introduction, ALBEF proposes the *in-batch* hard negative sampling (ITM_{hard}) by utilizing $\mathbf{p}^{v2t}(V)$ and $\mathbf{p}^{t2v}(T)$ defined in (1) for sampling text and image that has high similarity for given V and T , respectively, as a negative sample pair.

(c) Masked language modeling (MLM) is a task to predict the randomly masked tokens in a text based on both contextual text and visual information. ALBEF uses the masking probability of 15% following [6], and by denoting the randomly masked text as \tilde{T} and the prediction probability for the masked tokens as $\mathbf{p}^{\text{mask}}(V, \tilde{T})$, the loss function of MLM becomes

$$\mathcal{L}_{\text{MLM}} = \mathbb{E}_{(V, \tilde{T}) \sim D} [\text{CE}(\tilde{\mathbf{y}}, \mathbf{p}^{\text{mask}}(V, \tilde{T}))], \quad (4)$$

in which $\tilde{\mathbf{y}}$ is a ground truth one-hot vector for the masked token.

3 Ablation Analyses on ALBEF

As mentioned in the Introduction, we carry out careful analyses on ALBEF to verify the true effect of the training objectives described in the previous section. To that end, we set the base model as “ALBEF_{Base}”, which mostly follows the model architecture and pre-training objectives of ALBEF, but does *not* use the additional momentum encoder and momentum distillation⁵. Highlighting the

⁵ We defer describing the detailed model architecture to Section 5.1.

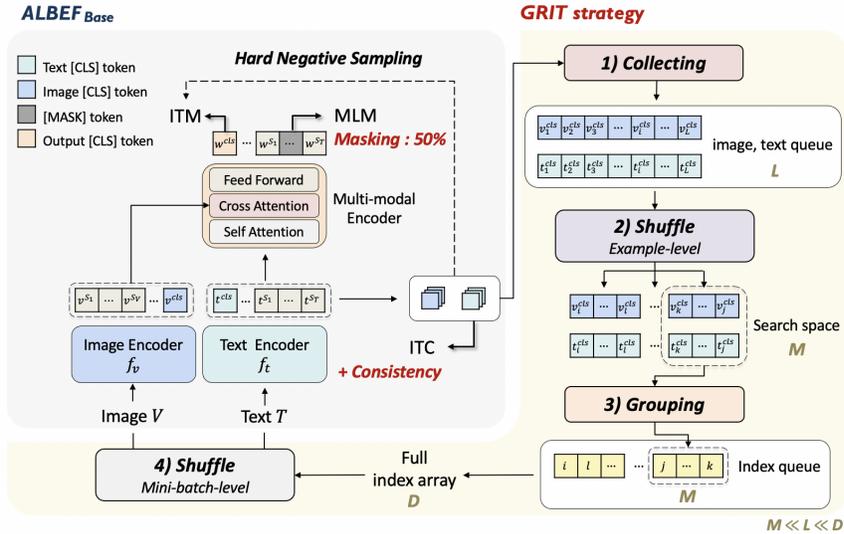


Fig. 2: The ALBEF_{Base} architecture and the overall process of GRIT-VLP.

training objectives, we can also denote “ALBEF_{Base}” by $\text{MLM} + \text{ITM}_{\text{hard}} + \text{ITC}$, and we ablate each of those components and evaluate the performance of the model variants on two downstream tasks (IRTR, NLVR). All models are pre-trained with 4M dataset, and evaluated with MS-COCO [23] and NLVR2 dataset [32]. Details on the tasks, datasets and additional results are described in Section 5 and the Supplementary Material (S.M.).

[Hard negative sampling on ITM] Table 1 compares the downstream task performance of models that have the fixed MLM objective (with masking probability 15%) but varying ITM and ITC objectives of ALBEF_{Base}. In the table, “MLM+ITM_{hard}” stands for the case in which only ITM_{hard} is carried out without the ITC objective — this case is missing in the analysis of the original ALBEF paper [20, Table 1], but we believe it is necessary for showing the effect of ITM_{hard} alone without the pre-aligning effect of ITC. The subtlety here is that, since ITM_{hard} utilizes the image-text similarity scores from ITC (1) for selecting the *in-batch* hard negative samples, we use the scores obtained from the uni-modal encoders of ALBEF (without the multi-modal encoder) that are *pre-trained* only with the ITC loss. Moreover, “ITM_{rand}” in Table 1 stands for the ITM loss with randomly selected negative samples.

The original ALBEF essentially focuses on the effect of ITC by mainly comparing “MLM+ITM_{rand}” and “MLM+ITM_{rand}+ITC” and argues that “ITM_{hard}” only gives a marginal improvement when it replaces ITM_{rand}. However, we observe a different story in Table 1. Namely, even without the pre-alignment of the representations via ITC, “MLM+ITM_{hard}” gives a significant performance boost over “MLM+ITM_{rand}”, which is substantially larger than the improvement we get by “MLM+ITM_{rand}+ITC”. Moreover, even with a shorter 10 epochs,

Table 1: Ablation study on ITM_{hard} and ITC for ALBEF_{Base}.

Epochs	Training tasks	TR (COCO)			IR		NLVR		
		R@1	R@5	R@10	R@1	R@5	R@10	(val)	(test)
10	MLM + ITM _{rand}	61.6	86.1	92.5	47.8	75.4	84.8	77.02	78.44
	MLM + ITM _{rand} + ITC	66.8	88.8	94.5	51.1	78.4	86.8	76.59	78.69
	MLM + ITM _{hard}	68.6	89.4	94.9	52.1	79.0	87.1	79.18	79.32
	ALBEF _{Base}	72.3	91.3	96.0	55.1	81.0	88.5	79.21	79.78
20	MLM + ITM _{rand}	66.5	88.3	94.0	51.3	78.3	86.5	78.02	79.43
	MLM + ITM _{rand} + ITC	69.6	90.9	95.3	53.8	80.0	87.8	77.61	79.43
	MLM + ITM _{hard}	72.0	91.5	96.6	57.5	81.2	88.4	80.44	80.83
	ALBEF _{Base}	73.8	92.3	96.5	57.7	82.5	89.6	79.22	80.37

Table 2: Ablation study on the masking probability for MLM for ALBEF_{Base}.

Epochs	Training tasks	TR (COCO)			IR		NLVR		
		R@1	R@5	R@10	R@1	R@5	R@10	(val)	(test)
10	ALBEF _{Base}	72.3	91.3	96.0	55.1	81.0	88.5	79.21	79.78
	ALBEF _{Base50}	73.4	92.5	96.4	57.2	82.3	89.4	79.42	79.87
20	ALBEF _{Base}	73.8	92.3	96.5	57.7	82.5	89.6	79.22	80.37
	ALBEF _{Base50}	75.6	93.2	96.7	58.8	83.2	90.1	80.41	80.54

“MLM+ITM_{hard}” performs competitively or superior to “MLM+ITM_{rand}+ITC” trained for a longer 20 epochs. While the best performance is still obtained by using ITM_{hard} and ITC together, *i.e.*, ALBEF_{Base}, this result strongly motivates that further improving ITM_{hard} could be central in attaining efficient VLP.

[Mask sampling probability for MLM] Table 2 now focuses on MLM by varying the masking probability with fixed ITM and ITC. Namely, the original ALBEF_{Base} trains with the masking probability of 15%, and we also test the model with the probability 50%, dubbed as ALBEF_{Base50}. In the table, we observe that this simple change brings surprising performance gain; ALBEF_{Base50} always outperforms ALBEF_{Base} for the same epoch and becomes comparable to ALBEF_{Base} even when trained with significantly smaller number of epochs. This result clearly motivates using enlarged masking probability for MLM for VLP.

3.1 Motivation

The result in Table 1 suggests that improving the hard negative sampling strategy for ITM could bring further performance gain for VLP. An obvious way for such improvement is to enlarge the search space from which the negative samples are selected, hence, the sample that contains the nuanced difference with respect to the positive sample as in Figure 1(c) can be obtained. However, we note that such enlargement in a memory- and computation-efficient way is far from being straightforward, described as below.

The most naive way to enlarge the search space is to enlarge the size of the mini-batch during training. While conceptually simple, it clearly is limited by the GPU memory and high computational cost. An alternative is to utilize

the additional queues to store the compressed representations of the samples (*i.e.*, the [CLS] tokens $[v^{cls}, t^{cls}]$ from the uni-modal encoders), like MoCo [12] or MemoryBank [40], and include those representations in the search space for mining the hard negatives. While this queue-based solution is highly effective in the ordinary contrastive learning, it causes additional complication for VLP using the ITM loss. Namely, as described in Section 3, the ITM loss is calculated with the [CLS] token from the multi-modal encoder (w^{cls}), which needs the *entire* sequence embeddings ($f_v(V), f_t(T)$) to compute. Therefore, to employ the queue-based solution for ITM_{hard} , one should select between the following two options. One is to store the entire embedding sequences for both modalities in the queues, which is severely memory-inefficient due to the long sequence lengths (typically, $S_T : 30$ and $S_V : 200$). The other is to only store $[v^{cls}, t^{cls}]$ tokens from each modality to compute (1) and (2) for ITC, but carry out the additional forward passes for the samples that are not in the current mini-batch to compute w^{cls} and the ITM loss. Clearly, the second option would suffer from the additional computation cost required for the forward passes.

To overcome the limitations of the above naive solutions, we propose a new method that can enlarge the search space and select more informative negative samples for ITM *without* any significant overheads on the memory/computation.

4 Main Method: GRIT-VLP

4.1 GRIT: GGrouped mIni-baTch sampling (GRIT)

In this section, we describe our main contribution, the GGrouped mIni-baTch sampling (GRIT) strategy. The basic idea is to compose each mini-batch of size N with highly similar example groups such that the informative, hard negative samples are likely to be chosen by the *in-batch* sampling of ALBEF. In order to do that without significant memory/computation overhead, as described in Fig. 2 and Algorithm 1 in S.M, GRIT utilizes two additional queues of size L that store the [CLS] tokens from the uni-modal encoders, an index queue I of size M , and a full index array G of size D (the whole data size). The ordering of the sizes is $N \ll M \ll L \ll D$. Then, the procedure of constructing *grouped* mini-batches for the next epoch is performed concurrently with the loss calculation for pre-training at each epoch, and these grouped mini-batches are used for the ordinary mini-batch training in the following epoch.

A subtle challenge of above simultaneous creation process for the grouped mini-batches is that it removes the randomness of the mini-batches, which is an essential ingredient for the *stochastic* gradient descent based learning. Therefore, we add two-level shuffling phases for preserving the randomness among the grouped mini-batches. As a result, our GRIT is composed of the following four phases: 1) collecting, 2) example-level shuffling 3) grouping, and 4) mini-batch-level shuffling. We note the first three phases are repeated whenever the queue of size L is filled, and the last phase is repeated once every epoch.

[Phase 1: Collecting] To construct mini-batches containing similar samples, we first store the [CLS] tokens $[v^{cls}, t^{cls}]$ (from uni-modal encoders) in the two

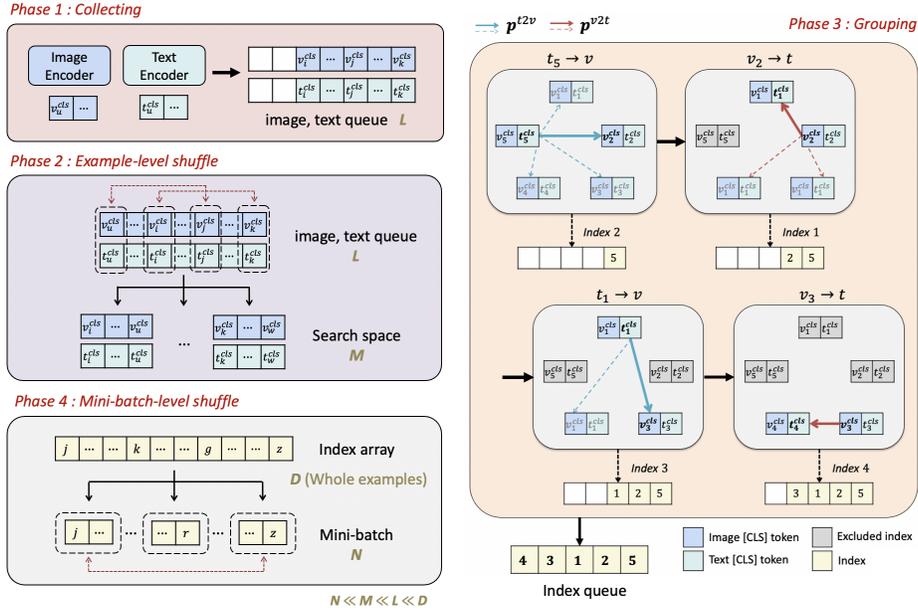


Fig. 3: Grouped mini-batch sampling (GRIT) strategy.

additional queues of size L , which is much larger than the size of mini-batches (N), until filled as shown in Figure 3.

[Phase 2: Example-level shuffle] Once the queues are filled, all the samples in the queues are randomly shuffled at the example-level, to secure the randomness mentioned above. After shuffling, each queue is divided into $\frac{L}{M}$ sub-queues of size M , which is the size of the enlarged search space for the hard negative samples for ITM. Then, the samples in each sub-queue are grouped based on similarity via the grouping phase below, which is sequentially performed for each sub-queue.

[Phase 3: Grouping] From the $[v^{cls}, t^{cls}]$ stored in the sub-queue, we can compute the image-to-text and text-to-image similarity scores, similarly as in (1), among the examples in the sub-queue. Accordingly, for each pair (V, T) , those scores can be denoted by $q^{v2t}(V) \in \Delta^M$ and $q^{t2v}(T) \in \Delta^M$, respectively.

Based on the computed similarities, we aim group *similar* (V, T) examples in the sub-queue to each mini-batch as much as possible. To that end, as described in Algorithm 2 in S.M, our grouping phase is summarized as: 1) randomly sample the first pair (V_1, T_1) from the sub-queue, then 2) iteratively find and store the index of the most similar example one by one until all examples inside the sub-queue are visited once, and finally, 3) the index queue $I \in \{1, \dots, M\}^M$ is generated. Note both the negative text for an anchor image and the negative image for an anchor text should be considered when constructing the negative samples for ITM. Thus, rather than using a one-way similarity score, two similarity scores are used alternatively; namely, as illustrated in Fig. 3 with a toy example of $M = 5$,

at the $(i + 1)$ -th iteration, given a pair (V_k, T_k) with index k , I_{i+1} is chosen as

$$I_{i+1} = \begin{cases} \operatorname{argmax}_{j \notin I} \mathbf{q}_j^{t2v}(T_k) & \text{if } I_i \text{ is chosen with } \mathbf{q}^{v2t} \\ \operatorname{argmax}_{j \notin I} \mathbf{q}_j^{v2t}(V_k) & \text{if } I_i \text{ is chosen with } \mathbf{q}^{t2v}. \end{cases} \quad (5)$$

Thus, during above *grouping* process for the sub-queue, half of the pairs are selected based on $(V \rightarrow T)$ direction, and the other half based on $(T \rightarrow V)$ direction. Whenever the index queue I is full, we convert the indices into the original data indices in $\{1, \dots, D\}$ and append those to the full index array G .

[Phase 4: Mini-batch-level shuffle] After each epoch, the full index array G , which is a permutation of $\{1, \dots, D\}$, is generated. Then, G is divided into multiple mini-batch-sized arrays, and these arrays are shuffled. Note this shuffling is done at the mini-batch level, not at the example level. Finally, these shuffled mini-batches are used for both training and GRIT for the next epoch.

Remark 1: We note the shuffling phases *Phase 2/4* in GRIT are important to secure the randomness among the mini-batches. Namely, since GRIT generates the indices during the previous epoch, it omits the conventional data re-shuffling performed at the start of each epoch. Hence, although the order of indices is continuously changed to some extent in *Phase 3*, such re-ordering happens only at the level of sub-queue of size M , hence the scope of shuffling is significantly limited. In Table 4 (Section 5), we verify that the performance of GRIT without shuffling is significantly degraded, justifying the proposed shuffling phases.

Remark 2: The naive implementation of GRIT would be to proceed *Phase 1/3* and training *separately*, not concurrently. To be specific, at the beginning of each epoch, the conventional re-shuffling of the whole data is done, followed by additional forward passes on the uni-modal encoders, and *Phase 1/3* are performed to generate grouped mini-batch indices. Then, the training begins with the generated indices. Since this naive version requires additional forward passes, it clearly has high computational cost and requires longer training time.

4.2 ITC consistency loss and increased masking probability for MLM

GRIT encourages similar examples to be grouped within each mini-batch, hence, the ITM_{hard} can become more effective since the mini-batch may contain informative, hard negative samples. However, when GRIT is combined with ITC, one potential drawback is that the representations for similar samples would move away from each other unexpectedly, since all negatives will be equally penalized during the contrastive learning regardless of the similarity.

To address this issue, we add a consistency loss that can reflect the similarity among samples. Namely, when an image V and a text T form a positive pair (V, T) , it is natural to assume that they share a similar semantic. Hence, we would expect the similarity scores $\mathbf{p}^{v2t}(V)$ and $\mathbf{p}^{t2v}(T)$ to be similar to each other. To this end, we define the soft pseudo-target $\tilde{\mathbf{p}}^{t2v}(T)$ as $\text{sg}(\mathbf{p}^{t2v}(T))$ and $\tilde{\mathbf{p}}^{v2t}(V)$ as $\text{sg}(\mathbf{p}^{v2t}(V))$ for $\mathbf{p}^{v2t}(V)$ and $\mathbf{p}^{t2v}(T)$, respectively, in which $\text{sg}(\cdot)$ is

the stop-gradient operator. Then, our ITC with consistency loss is defined as

$$\mathcal{L}_{\text{ITC}_{\text{cons}}} = \mathcal{L}_{\text{ITC}} + \frac{\lambda_{\text{cons}}}{2} \mathbb{E}_{(V,T) \sim D} [KL(\tilde{\mathbf{p}}^{v2t}(V) || \mathbf{p}^{t2v}(T)) + KL(\tilde{\mathbf{p}}^{t2v}(T) || \mathbf{p}^{v2t}(V))], \quad (6)$$

in which λ_{cons} is the regularization parameter. We expect this loss refines similarity scores which affect the quality of the grouping phase of GRIT. We set λ_{cons} as 0.2 for all cases for simplicity.

Finally, our model, dubbed as GRIT-VLP and illustrated in Fig 2, is obtained as follows. We use ALBEF_{Base} as our base model architecture, and combine our GRIT, ITC consistency loss, and masking probability of 50% for MLM. Consequently, the pre-training objective of GRIT-VLP is

$$\mathcal{L} = \mathcal{L}_{\text{ITM}_{\text{hard}}} + \mathcal{L}_{\text{MLM}_{50}} + \mathcal{L}_{\text{ITC}_{\text{cons}}}, \quad (7)$$

in which the mini-batches are generated by the GRIT strategy. The pseudo-code for GRIT-VLP is given in Alg.1/2 in the S.M.

5 Experimental Results

5.1 Data and experimental settings

[Training data] Following ALBEF [20] and UNITER [4], we use four datasets (MS-COCO [23], Visual Genome [18], Conceptual Captions [30] and SBU Captions [25]) for training, which consist of 4M unique images and 5M image-text pairs.

[Implementation details] Here, we give the concrete model architecture of ALBEF_{Base}. We use a 12-layer vision transformer ViT-B/16 [8] with 86M parameters as the image encoder f_v and initialize it with the weights pre-trained on ImageNet-1k [34]. A 6-layer Transformer [35] is used for both the text encoder f_t and the multi-modal encoder h , which are initialized with the first 6 layers and the last 6 layers of BERT-base with 123.7M parameters[6], respectively. We use the same data augmentation technique of ALBEF, and our model is trained for 20 epochs. All experiments are performed on 4 NVIDIA A100 GPUs. Furthermore, unless otherwise noted, we set $N = 96$, $M = 960$, and $L = 48,000$. For all other hyper-parameter settings, we follow ALBEF [20]. More details on the dataset, software platform, training procedures, and hyper-parameters are in the S.M.

5.2 Downstream vision and language tasks

After the pre-training step, our model is fine-tuned on three well-established downstream vision and language tasks, including image-text retrieval (IRTR), visual question answering (VQA2 [10]), and natural language for visual reasoning (NLVR2 [32]). For IRTR, we use MS-COCO [23] and Flickr30K (F30K) [26] re-splited by [16]. We do not include SNLI-VE [41] in the evaluation, since the data set is known to be noisy according to [7]. We mostly follow the fine-tuning and evaluation process of ALBEF [20] except for using the momentum distillation. We compare our method with various VLP methods trained on the same 4M training set. More details on the downstream tasks including evaluation setting are given in S.M.

Table 3: Comparison with various methods on downstream vision-language tasks. **Bold** denotes the best result among models trained with 4M dataset.

Method	#Pre-train Images	Flickr R@1		COCO R@1		VQA		NLVR2	
		TR	IR	TR	IR	test-dev	test-std	dev	test-P
UNITER [4]	4M	87.3	75.6	65.7	52.9	72.70	72.91	77.18	77.85
VILLA [9]	4M	87.9	76.3	-	-	73.59	73.67	78.39	79.30
OSCAR [22]	4M	-	-	70.0	54.0	73.16	73.44	78.07	78.36
ViLT [17]	4M	83.5	64.4	61.5	42.7	71.26	-	75.70	76.13
ALBEF [20]	4M	94.3	82.8	73.1	56.8	74.54	74.70	80.24	80.50
GRIT-VLP_{E-10}	4M	94.7	82.0	74.9	58.1	74.72	74.74	79.98	80.11
GRIT-VLP	4M	96.0	83.8	77.1	59.5	75.11	75.26	80.73	81.60
ALBEF	14M	95.9	85.6	77.6	60.7	75.84	76.04	82.55	83.14

Table 4: Ablation study on the proposed method.

GRIT		λ_{cons}	Masking Prob(%)	TR		(COCO)		IR		NLVR	VQA	Time per epoch
Collecting	Shuffle			R@1	R@5	R@10	R@1	R@5	R@10	(test)	(test-std)	
\times	\times	0	15	73.8	92.3	96.5	57.7	82.5	89.6	80.37	74.70	2h 27m
\times	\times	0	50	75.6	93.2	96.7	58.8	83.2	90.1	80.54	75.07	2h 27m
\checkmark	\checkmark	0	50	76.4	93.6	96.7	59.6	83.3	90.1	81.32	75.14	2h 30m
\checkmark (naive)	\times	0	50	76.8	93.6	96.8	59.6	83.4	90.0	80.63	75.16	3h
\checkmark	\times	0	50	74.7	93.2	96.6	58.6	82.8	89.7	80.50	75.06	2h 30m
\checkmark	\checkmark	0.2	15	76.2	93.4	96.8	59.0	83.1	90.1	81.21	74.98	2h 30m
\checkmark	\checkmark	0.2	50	77.1	93.8	97.0	59.5	83.4	90.0	81.43	75.30	2h 30m

5.3 Comparison with the state-of-the-art VLP methods

Since we mainly build our method upon ALBEF, the previous state-of-the-art, we mainly compare our method with it. Table 3 reports the results of GRIT-VLP with $N = 128$ and $M = 1920$ on IRTR, VQA, and NLVR2. In S.M, we present additional results on these hyper-parameters showing the robustness of our method with respect to N .

On all downstream tasks (IRTR, VQA, NLVR2), GRIT-VLP outperforms other methods trained on the same 4M dataset, including the previous best model ALBEF (4M) by a large margin (+4% TR/R@1 on MS-COCO, +1.1% on NLVR test-P). Moreover, GRIT-VLP is even competitive with ALBEF (14M) on some metrics, while being trained on a much smaller dataset. Furthermore, “GRIT-VLP_{E-10}”, denoting GRIT-VLP trained for only 10 epochs, achieves competitive performance compared to ALBEF (4M) trained with 30 epochs, highlighting the efficiency of our method. We believe the performance gains in Table 3 clearly highlights effectiveness of GRIT-VLP.

5.4 Ablation studies on the proposed method

Table 4 shows the effectiveness of each proposed component: GRIT, ITC consistency loss, and enlarged masking probability (15% \rightarrow 50%) for MLM. First two rows indicate ALBEF_{Base} and ALBEF_{Base50} analyzed in Section 3, respectively. By integrating the ALBEF_{Base50} with GRIT-variants (row 3, 4), we can verify that the performance is significantly improved. However, in the case of “naive”

Table 5: Effect of GRIT on ITC

Training tasks	TR (COCO)			IR		
	R@1	R@5	R@10	R@1	R@5	R@10
ITC	59.7	84.8	91.9	43.2	72.2	82.0
Queue-based ITC	60.3	85.2	92.2	43.5	72.3	82.3
ITC + GRIT	63.0	87.1	93.2	45.2	72.9	82.1
ITC _{cons} + GRIT	64.3	87.4	93.2	46.3	73.7	82.6

Table 6: Results on top of TCL [45]

Method	TR (COCO)			IR		NLVR2	VQA
	R@1	R@5	R@10	R@1	R@5	R@10 (test-P)	(test-std)
TCL	75.6	92.8	96.7	59.0	83.2	89.9	81.33
TCL+ours	77.3	94.1	97.2	60.2	83.7	90.0	81.52

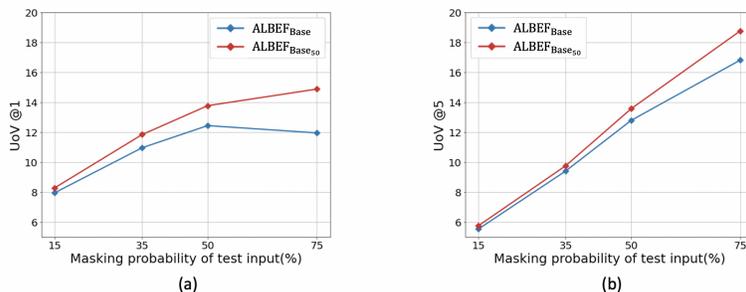


Fig. 4: UoV results on COCO validation dataset.

implementation version of GRIT described in Section 4 (row 4), the training time is significantly increased as expected. We believe that competitive results of row 3 and 4 clearly demonstrate the need of all components of GRIT. Moreover, if the both shufflings are removed from GRIT while collecting the mini-batches using the previous epoch (row 5), its performance is severely degraded due to the vanishing randomness. The last row denotes our final GRIT-VLP ($\text{ALBEF}_{\text{Base50}} + \text{GRIT} + \text{consistency}$); by adding the consistency loss from row 3, we verify that the overall performance is increased. Furthermore, the gains of the last two rows compared to the first two rows show the combined effect of “GRIT + consistency” at two different mask sampling probabilities.

5.5 Experiments on the model-agnostic property

To investigate the model-agnostic property of GRIT, we report the results when our method is integrated with different network architectures and objectives.

[Small model, ITC] Table 5 compares the IRTR results with several variants that are pre-trained and fine-tuned with only ITC loss and dual uni-modal encoders of ALBEF (without multi-modal encoder). “Queue-based ITC” denotes a model pre-trained with ITC and queues for leveraging the stored features from the momentum encoder (row 2), and the other two models denote models that adopt our approach. While the performance gain of the Queue-based ITC is marginal, our GRIT brings a significant gain when combined with ITC. Finally, $\text{ITC}_{\text{cons}} + \text{GRIT}$ achieves the best result, demonstrating the standalone effectiveness of GRIT and consistency loss on ITC. We believe this result shows that our method has a potential to be easily integrated with existing contrastive learning based models.

[Large model, more objectives] In Table 6, we additionally measure the gains of the recently proposed TCL [45] when it is combined with “ours” (GRIT and

enlarged masking probability). TCL introduces additional objectives and adopts almost the same but larger network architecture (additional momentum model) than ours. We omit the consistency loss since TCL already uses MD which has a similar role to it. Although the learning objectives and model sizes are different, we verify that the combination of our method and TCL again brings significant gains in Table 6, which clearly shows the model-agnostic property of our method. Details about this combined approach can be found in S.M.

5.6 Analysis on the masking probabilities

We believe the enlarged masking probability for MLM enables more usage of the visual features for predicting the masked token, resulting in a more effective multi-modal representation learning. To demonstrate this point, we introduce a *Usage of Vision* (UoV) metric like [2], which is defined as the difference between MLM accuracy of a pre-trained model with and without the image input (*Accuracy - Accuracy w/o image*). When evaluating the MLM *Accuracy* and *Accuracy w/o image* of the pre-trained model, test input sentences are masked with the same specific masking probability, and exactly the same tokens are masked for a fair comparison. Then, a high UoV value means that the pre-trained model is highly affected by visual information since it implies that the vision information is important for the model to correctly predict the masked token.

In Fig. 4, two pre-trained networks ($ALBEF_{Base}$, $ALBEF_{Base50}$) are evaluated with test input sentences masked with various masking probability. We calculate the UoV@1 and UoV@5 by considering the top-1 and top-5 MLM accuracy, respectively. Fig. 4 shows that $ALBEF_{Base50}$ model always obtain higher UoV demonstrating the high usage of vision information. In particular, when the textual context almost disappears (75%), the difference in UoV becomes outright. As a result, we verify that enlarging the masking probability enriches the usage of visual information, which facilitates the alignment of image and text.

6 Concluding Remarks

We proposed GRIT-VLP, which effectively samples mini-batches for mining hard negatives while maintaining the computational overhead. We showed our method achieves state-of-the-art performance on various downstream tasks with much less computational overhead and can be easily integrated with existing VLP.

Acknowledgment

This work was supported in part by New Faculty Startup Fund from Seoul National University, NRF Mid Career Research Program [NRF-2021R1A2C2007884], IITP grants funded by the Korean government [No.2021-0-01696, High-Potential Individuals Global Training Program], [No.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)], [No.2021-0-02068, Artificial Intelligence Innovation Hub (Artificial Intelligence Institute, Seoul National University)], and SNU-NAVER Hyperscale AI Center.

References

1. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: ICCV (2015)
2. Bitton, Y., Stanovsky, G., Elhadad, M., Schwartz, R.: Data efficient masked language modeling for vision and language. arXiv preprint arXiv:2109.02040 (2021)
3. Chen, T.S., Hung, W.C., Tseng, H.Y., Chien, S.Y., Yang, M.H.: Incremental false negative detection for contrastive learning. arXiv preprint arXiv:2106.03719 (2021)
4. Chen, Y.C., Li, L., Yu, L., El Kholly, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: ECCV (2020)
5. Chuang, C.Y., Robinson, J., Lin, Y.C., Torralba, A., Jegelka, S.: Debiased contrastive learning. In: NeurIPS (2020)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
7. Do, V., Camburu, O.M., Akata, Z., Lukasiewicz, T.: e-snli-ve: Corrected visual-textual entailment with natural language explanations. arXiv preprint arXiv:2004.03744 (2020)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
9. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. In: NeurIPS (2020)
10. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In: CVPR (2017)
11. Harwood, B., Kumar BG, V., Carneiro, G., Reid, I., Drummond, T.: Smart mining for deep metric learning. In: ICCV (2017)
12. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
13. Huang, Z., Zeng, Z., Huang, Y., Liu, B., Fu, D., Fu, J.: Seeing out of the box: End-to-end pre-training for vision-language representation learning. In: CVPR (2021)
14. Huang, Z., Zeng, Z., Liu, B., Fu, D., Fu, J.: Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. arXiv preprint arXiv:2004.00849 (2020)
15. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. arXiv preprint arXiv:2102.05918 (2021)
16. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: CVPR (2015)
17. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. arXiv preprint arXiv:2102.03334 (2021)
18. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. IJCV (2017)
19. Li, G., Duan, N., Fang, Y., Gong, M., Jiang, D.: Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In: AAAI (2020)
20. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. In: NeurIPS (2021)

21. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 (2019)
22. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: ECCV (2020)
23. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
24. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. arXiv preprint arXiv:1908.02265 (2019)
25. Ordonez, V., Kulkarni, G., Berg, T.: Im2text: Describing images using 1 million captioned photographs. In: NeurIPS (2011)
26. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: ICCV (2015)
27. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021)
28. Robinson, J., Chuang, C.Y., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. arXiv preprint arXiv:2010.04592 (2020)
29. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR (2015)
30. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: ACL (2018)
31. Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., Dai, J.: Vi-bert: Pre-training of generic visual-linguistic representations. arXiv preprint arXiv:1908.08530 (2019)
32. Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., Artzi, Y.: A corpus for reasoning about natural language grounded in photographs. arXiv preprint arXiv:1811.00491 (2018)
33. Tan, H., Bansal, M.: Lxmert: Learning cross-modality encoder representations from transformers. arXiv preprint arXiv:1908.07490 (2019)
34. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML (2021)
35. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
36. Wang, X., Zhang, H., Huang, W., Scott, M.R.: Cross-batch memory for embedding learning. In: CVPR (2020)
37. Wettig, A., Gao, T., Zhong, Z., Chen, D.: Should you mask 15% in masked language modeling? arXiv preprint arXiv:2202.08005 (2022)
38. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: ICCV (2017)
39. Wu, M., Mosse, M., Zhuang, C., Yamins, D., Goodman, N.: Conditional negative sampling for contrastive learning of visual representations. arXiv preprint arXiv:2010.02037 (2020)
40. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018)
41. Xie, N., Lai, F., Doran, D., Kadav, A.: Visual entailment: A novel task for fine-grained image understanding. arXiv preprint arXiv:1901.06706 (2019)
42. Xuan, H., Stylianou, A., Liu, X., Pless, R.: Hard negative examples are hard, but useful. In: ECCV (2020)

43. Xuan, H., Stylianou, A., Pless, R.: Improved embeddings with easy positive triplet mining. In: WACV (2020)
44. Xue, H., Huang, Y., Liu, B., Peng, H., Fu, J., Li, H., Luo, J.: Probing inter-modality: Visual parsing with self-attention for vision-and-language pre-training. In: NeurIPS (2021)
45. Yang, J., Duan, J., Tran, S., Xu, Y., Chanda, S., Chen, L., Zeng, B., Chilimbi, T., Huang, J.: Vision-language pre-training with triple contrastive learning. In: CVPR (2022)
46. Zolfaghari, M., Zhu, Y., Gehler, P., Brox, T.: Crossclr: Cross-modal contrastive learning for multi-modal video representations. In: ICCV (2021)