

An Empirical Study of Training End-to-End Vision-and-Language Transformers

Zi-Yi Dou^{1*}, Yichong Xu², Zhe Gan², Jianfeng Wang², Shuohang Wang², Lijuan Wang²,
Chenguang Zhu², Pengchuan Zhang², Lu Yuan², Nanyun Peng¹, Zicheng Liu², Michael Zeng²

¹University of California, Los Angeles, ²Microsoft Corporation

{zdou, violetpeng}@cs.ucla.edu

{yicxu, zhgan, jianfw, shuowa, lijuanw, chezhu, penzhan, luyuan, zliu, nzeng}@microsoft.com

Abstract

Vision-and-language (VL) pre-training has proven to be highly effective on various VL downstream tasks. While recent work has shown that fully transformer-based VL models can be more efficient than previous region-feature-based methods, their performance on downstream tasks often degrades significantly. In this paper, we present **METER**, a **Multimodal End-to-end TransformER** framework, through which we investigate how to design and pre-train a fully transformer-based VL model in an end-to-end manner. Specifically, we dissect the model designs along multiple dimensions: vision encoders (e.g., CLIP-ViT, Swin transformer), text encoders (e.g., RoBERTa, DeBERTa), multimodal fusion module (e.g., merged attention vs. co-attention), architectural design (e.g., encoder-only vs. encoder-decoder), and pre-training objectives (e.g., masked image modeling). We conduct comprehensive experiments and provide insights on how to train a performant VL transformer. **METER** achieves an accuracy of 77.64% on the VQA_{v2} test-std set using only 4M images for pre-training, surpassing the state-of-the-art region-feature-based model by 1.04%, and outperforming the previous best fully transformer-based model by 1.6%. Notably, when further scaled up, our best VQA model achieves an accuracy of 80.54%. Code and pre-trained models are released at <https://github.com/zdou0830/METER>.

1. Introduction

Vision-and-language (VL) tasks, such as visual question answering (VQA) [1] and image-text retrieval [27, 34], require an AI system to comprehend both the input image and text contents. Vision-and-language pre-training (VLP) has now become the *de facto* practice to tackle these tasks [5, 24, 26, 32, 42, 44]. Specifically, large amounts of image-caption pairs are fed into a model that consumes both images and texts to pretrain representations that contain rich

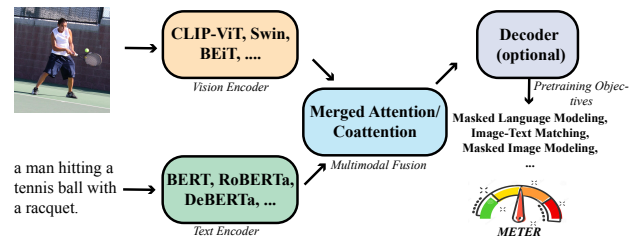


Figure 1. An overview of the proposed **METER** framework. We systematically investigate how to train a performant vision-and-language transformer, and dissect the model designs along multiple dimensions: vision encoder, text encoder, multimodal fusion module, architectural design (encoder-only vs. encoder-decoder), and pre-training objectives.

multimodal knowledge and is helpful for downstream tasks.

Transformers [48] are prevalent in natural language processing and have recently shown promising performance in computer vision [10, 30]. While almost all the existing VLP models adopt transformers for their multimodal fusion module, most of them [5, 24, 26, 32, 42, 44] use pre-trained object detectors (e.g., Faster RCNN [38]) on the vision side to extract *region* features from images. This can lead to several problems: first, the object detectors are not perfect, but are usually kept frozen during VLP, which limits the capacity of the VLP models; second, it is time-consuming to extract region features [20]. On the other hand, vision transformers (ViTs) have been an increasingly active research topic in computer vision and have shown great potential in vision feature extraction. Therefore, a natural question arises: *can we train a fully transformer-based VLP model with ViTs as the image encoder?*

Recent works [20, 23, 53] that tried to adopt vision transformers have not shown satisfactory performance and typically underperform state-of-the-art region-feature-based VLP models (e.g., VinVL [57]). To close the performance gap, we present **METER**, a **Multimodal End-to-end TransformER** framework, through which we thoroughly investigate how to design and pre-train a fully transformer-

* Work was done when the author interned at Microsoft.

Model	Vision Encoder	Text Encoder	Multimodal Fusion	Decoder	Pre-training Objectives
ViLBERT [32]	OD+Xformer	Xformer	Co-attn.		MLM+ITM+MIM
LXMERT [44]					MLM+ITM+MIM+VQA
VisualBERT [24]	OD	Emb.	Merged-attn.		MLM+ITM
VL-BERT [42]					MLM+MIM
UNITER [5]					MLM+ITM+MIM+WRA
OSCAR [26]					MLM+ITM
VinVL [57]					MLM+ITM
VL-T5 [6]					MLM+ITM+VQA+Grounding+Captioning
PixelBERT [16]					CNN
SOHO [15]	MLM+ITM+MIM				
CLIP-ViL [41]	MLM+ITM+VQA				
SimVLM [51]	PrefixLM				
MDETR [19]	OD+Token Prediction+Contrastive Alignment				
ViLT [20]	Patch Emb.	Emb.	Merged-attn.		MLM+ITM
Visual Parsing [53]					
ALBEF [23]	Xformer	Xformer	Co-attn.		MLM+ITM+MIM
METER (Ours)					MLM+ITM+ITC
CLIP [35]	CNN/Xformer	Xformer	None		ITC
ALIGN [17]	CNN				

Table 1. Glossary of representative VLP models. OD: objective detector. Xformer: transformer. Emb.: embedding. MLM/MIM: masked language/image modeling. ITM: image-text matching. WRA: word-region alignment. ITC: image-text contrastive learning.

based VLP model in an end-to-end manner. Specifically, as shown in Figure 1, we dissect the model designs along multiple dimensions, including vision encoders (e.g., CLIP-ViT [35], Swin transformer [30]), text encoders (e.g., RoBERTa [29], DeBERTa [13]), multimodal fusion module (e.g., merged attention vs. co-attention), architectural design (e.g., encoder-only vs. encoder-decoder), and pre-training objectives (e.g., masked image modeling [2]).

We perform the investigation by pre-training models under METER on four commonly used image-caption datasets: COCO [27], Conceptual Captions [40], SBU Captions [33], and Visual Genome [21]. We test them on visual question answering [1], visual reasoning [43], image-text retrieval [27, 34], and visual entailment [52] tasks. Through extensive analyses, we summarize our findings as follows:

- Vision transformer (ViT) plays a more vital role than language transformer in VLP, and the performance of transformers on pure vision or language tasks is not a good indicator for its performance on VL tasks.
- The inclusion of cross-attention benefits multimodal fusion, which results in better downstream performance than using self-attention alone.
- Under a fair comparison setup, the encoder-only VLP model performs better than the encoder-decoder model for VQA and zero-shot image-text retrieval tasks.
- Adding the masked image modeling loss in VLP will *not* improve downstream task performance in our settings.

These insights, combined with other useful tips and tricks detailed in later sections, enable us to train a strong model that achieves an accuracy of 77.64% on the VQAv2 test-std set, surpassing the previous best region-feature-based VinVL model [57] by 1.04% and outperforming the previously best ViT-based model (i.e., ALBEF [23]) by 1.6%.

Notably, when further scaled up, our best METER model achieves an accuracy of 80.54% on the VQAv2 test-std set.

2. Glossary of VLP Models

In this section, we provide an overview of representative VLP models, and divide them into three categories based on how they encode images, as summarized in Table 1.

OD-based Region Features. Most previous work use pre-trained object detectors (ODs) to extract visual features. Among them, ViLBERT [32] and LXMERT [44] use co-attention for multimodal fusion, where two transformers are applied independently to region and text features, and another transformer fuses the representations of the two modalities in a later stage. On the other hand, VisualBERT [24], VL-BERT [42], and UNITER [5] use a merged attention fusion module that feeds both region and text features together into a single transformer. OSCAR [26] and VinVL [57] feed additional image tags into the transformer model, and demonstrate state-of-the-art performance across VL tasks. However, extracting region features can be time-consuming, and the pre-trained ODs are usually frozen during pre-training, which limits the capacity of VLP models.

CNN-based Grid Features. To tackle the above two issues, researchers have tried different ways to pre-train VL models in an end-to-end fashion. Among them, PixelBERT [16] and CLIP-ViL [41] propose to feed grid features from convolutional neural networks (CNNs) and text directly into a transformer. SOHO [15] proposes to first discretize grid features using a learned vision dictionary, then feed the discretized features into their cross-modal module. While using grid features directly can be efficient, inconsistent optimizers are typically used for CNN and transformer. For example, PixelBERT [16] and CLIP-ViL [41]

use AdamW [31] for transformer and SGD for CNN. Recent work on vision transformers (ViTs) has also shown that CNN can achieve slightly worse accuracy/FLOPs trade-offs than their ViT counterparts [30], motivating researchers to develop ViT-based VLP models.

ViT-based Patch Features. ViLT [20] directly feeds image patch features and text token embeddings into a pre-trained ViT model, and fine-tunes the model on image-caption datasets. More recently, visual parsing [53] and ALBEF [23] use ViT as the image encoder and design different pre-training objectives for ViT-based VLP models. However, all these models lag behind the state-of-the-art performance on downstream tasks such as visual question answering. In this paper, we investigate how to pre-train a ViT-based model in an end-to-end manner that closes the performance gap while maintaining fast inference speed.

3. The METER Framework

Based on the previous work, we identify several important modules in VLP models as in Figure 1. In this section, we first illustrate our METER framework, then our default settings, which paves the way for our analyses hereinafter.

Overview. Given a text sentence l and an image v , a VLP model first extracts both text features $l = \langle l_1, \dots, l_N \rangle$ and visual features $v = \langle v_1, \dots, v_M \rangle$ via a *text encoder* and a *vision encoder*. The text and visual features are then fed into a *multimodal fusion module* to produce cross-modal representations, which are then optionally fed into a *decoder* before generating the final outputs.

3.1. Model Architecture

Vision Encoder. In this paper, we focus on patch features, and study the use of vision transformers (ViTs) [10] for vision encoder. Specifically, in ViT, an image is first segmented into patches, and then the patches are fed into a transformer model. ViT has become a popular research topic recently [2, 10, 30, 45, 45, 46, 56], and has been introduced into VLP [20, 23, 53]. However, all these models only achieve inferior performance compared to state-of-the-art region-feature-based models (e.g., VinVL [57]). Also, different pre-trained ViTs are used, lacking a systematic study of which ViTs are the best for VLP. In this work, we compare the original ViT [10], DeiT [45], Distilled-DeiT [45], CaiT [46], VOLO [56], BEiT [2], Swin Transformer [30] and CLIP-ViT [35], to provide a comprehensive analysis on the role of vision transformers.

Text Encoder. Following BERT [9] and RoBERTa [29], VLP models [5, 24, 26, 32, 42, 44] first segment the input sentence into a sequence of subwords [39], then insert two special tokens at the beginning and the end of the sentence to generate the input text sequence. After we obtain the text

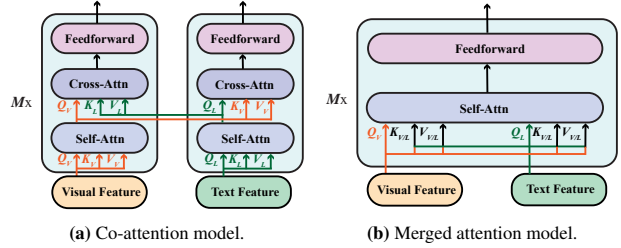


Figure 2. Illustration of two types of multimodal fusion modules: (a) co-attention, and (b) merged attention.

embeddings, existing works either feed them directly to the multimodal fusion module [5, 24], or to several text-specific layers [32, 44] before the fusion. For the former, the fusion module is typically initialized with BERT, and the role of text encoding and multimodal fusion is therefore entangled and absorbed in a single BERT model. Here, we aim to disentangle the two modules, and use a text encoder first before sending the features into the fusion module.

Language model (LM) pre-training has demonstrated impressive performance across tasks and different pre-trained LMs have been proposed; however, most VLP models still only use BERT for initialization [5]. In this work, we study the use of BERT [9], RoBERTa [29], ELECTRA [7], ALBERT [22], and DeBERTa [13] for text encoding. Besides, we also experiment on only using a simple word embedding look-up layer initialized with the BERT embedding layer as used in many previous works [5, 57].

Multimodal Fusion. We study two types of fusion modules, namely, *merged attention* and *co-attention* [14], as illustrated in Figure 2. In the *merged attention* module, the text and visual features are simply concatenated together, then fed into a single transformer block. In the *co-attention* module, on the other hand, the text and visual features are fed into different transformer blocks independently, and techniques such as cross-attention are used to enable cross-modal interaction. For region-based VLP models, as shown in [3], the *merged attention* and *co-attention* models can achieve comparable performance. Yet, the *merged attention* module is more parameter-efficient, as the same set of parameters are used for both modalities. Since end-to-end VLP models are becoming increasingly popular, in this work, we re-examine the impact of both types of fusion modules in our new context.

Encoder-Only vs. Encoder-Decoder. Many VLP models such as VisualBERT [24] adopt the encoder-only architecture, where the cross-modal representations are directly fed into an output layer to generate the final outputs. Recently, VL-T5 [6] and SimVLM [51], on the other hand, advocate the use of a transformer encoder-decoder architecture, where the cross-modal representations are first fed into a decoder and then to an output layer. In their models,

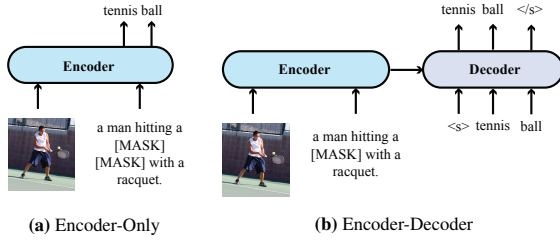


Figure 3. Illustration of the *encoder-only* and *encoder-decoder* model architectures for VLP.

the decoder attends to both the encoder representations and the previously generated tokens, producing the outputs autoregressively. Figure 3 shows the difference between them when performing the masked language modeling task. For the encoder-decoder model, when performing classification tasks such as VQA, we feed the text inputs into its encoder and feed a classification token into the decoder, and the decoder then generates the output class accordingly.

3.2. Pre-training Objectives

Now, we introduce how we pre-train our models. Specifically, we will first briefly review masked language modeling and image-text matching, which have been used extensively in almost every VLP model. Then, we will shift our focus to how we can design and explore interesting masked image modeling tasks.

Masked Language Modeling. The masked language modeling (MLM) objective is first introduced in pure language pre-training [9, 29]. In VLP, MLM with images has also proven to be useful. Specifically, given an image-caption pair, we randomly mask some of the input tokens, and the model is trained to reconstruct the original tokens given the masked tokens \mathbf{I}^{mask} and its corresponding visual input \mathbf{v} .

Image-Text Matching. In image-text matching, the model is given a batch of matched or mismatched image-caption pairs, and the model needs to identify which images and captions correspond to each other. Most VLP models treat image-text matching as a binary classification problem. Specifically, a special token (*e.g.*, [CLS]) is inserted at the beginning of the input sentence, and it tries to learn a global cross-modal representation. We then feed the model with either a matched or mismatched image-caption pair $\langle \mathbf{v}, \mathbf{I} \rangle$ with equal probability, and a classifier is added on top of the [CLS] token to predict a binary label y , indicating if the sampled image-caption pair is a match.

Masked Image Modeling. Similar to the MLM objective, researchers have tried masked image modeling (MIM) on the vision side. For example, many previous work, such as LXMERT [44] and UNITER [5], mask some of the input regions, and the model is trained to regress the original

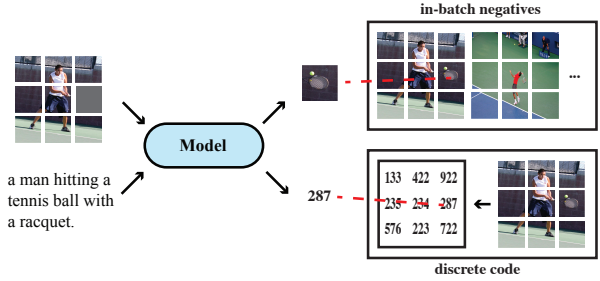


Figure 4. Illustration of masked patch classification with *in-batch negatives* and with *discrete code*.

region features. Formally, given a sequence of visual features $\mathbf{v} = \langle v_1, \dots, v_M \rangle$, where v_i is typically a region feature, we randomly mask some of the visual features, and the model outputs the reconstructed visual features \mathbf{o}_v given the rest of the visual features and the unmasked tokens \mathbf{t} , and regression aims to minimize the mean squared error loss. Researchers [5, 32, 44] have also tried to first generate object label for each region using a pre-trained object detector, which can contain high-level semantic information, and the model is trained to predict the object labels for the masked regions instead of the original region features.

Notably, recent state-of-the-art models (*e.g.*, ALBEF [23], VinVL [57]) do not apply MIM during VLP. In addition, in ViLT [20], the authors also demonstrate that masked patch regression is not helpful in their setting. These results make it questionable whether MIM is truly effective for VLP models. To further investigate this, we treat masked image modeling as a masked patch classification task, and propose two ways of implementing the idea.

1) Masked Patch Classification with In-batch Negatives.

By imitating MLM which uses a text vocabulary, we first propose to let the model reconstruct input patches by using a dynamically constructed vocabulary constructed with in-batch negatives. Concretely, at each training step, we sample a batch of image-caption pairs $\{(\mathbf{v}^k, \mathbf{I}^k)\}_{k=1}^B$, where B is the batch size. We treat all the patches in $\{\mathbf{v}^k\}_{k=1}^B$ as candidate patches, and for each masked patch, we mask 15% of the input patches, and the model needs to select the original patch within this candidate set. Denoting the original patch representations and our model’s output representations of $\{\mathbf{v}^k\}_{k=1}^B$ as $\{c(\mathbf{v}^k)\}_{k=1}^B$ and $\{h(\mathbf{v}^k)\}_{k=1}^B$, respectively, we can represent the output probability at position i for the k -th instance as:

$$p(\mathbf{v}_i^k | [\mathbf{v}^{k,mask}; \mathbf{I}^k]) = \frac{e^{h(\mathbf{v}_i^k)^T c(\mathbf{v}_i^k)}}{\sum_{j,k'} e^{h(\mathbf{v}_i^k)^T c(\mathbf{v}_j^{k'})}}. \quad (1)$$

The model is trained to maximize its probability similar to noise contrastive estimation [12, 18].

2) Masked Patch Classification with Discrete Code. Inspired by BEiT [2], we also propose to obtain discrete repre-

representations of the input patches, and the model is then trained to reconstruct the discrete tokens. Specifically, we first use the VQ-VAE [47] model in DALL-E [37] to tokenize each image into a sequence of discrete tokens. We resize each image so that the number of patches is equal to the number of tokens, and thus each patch corresponds to a discrete token. Then, we randomly mask 15% of the patches and feed the masked image patches to the model as before, but now the model is trained to predict the discrete tokens instead of the masked patches.

3.3. Our Default Settings for METER

There are many different model designs under METER, and we detail our default settings in this part.

Model Architecture. The default setting of model architecture is shown in Figure 2a. In the bottom part, there are one pre-trained visual encoder and one pre-trained text encoder. On top of each encoder, we stack $M = 6$ transformer encoding layers, with each layer consisting of one self-attention block, one cross-attention block, and one feed-forward network block. Unless otherwise stated, the hidden size is set to 768, and the number of heads is set to 12 for the top layers. Note that there is no decoder and no parameter sharing between the vision and language branches.

Pre-training Objectives. Unless otherwise stated, we pre-train the models with masked language modeling (MLM) and image-text matching (ITM) only. For MLM, we mask 15% of the input text tokens, and the model tries to reconstruct the original tokens. For ITM, we feed the model with either matched or mismatched image-caption pairs with equal probability, and the model needs to identify whether the input is a match.

Pre-training Datasets. Following previous work [5, 20, 23], we pre-train models on four commonly used datasets, including COCO [27], Conceptual Captions [40], SBU Captions [33], and Visual Genome [21]. The statistics of these datasets is shown in Appendix. The combined training data consists of about 4M images in total.

Downstream Tasks. For ablation and analysis, we mainly focus on VQAv2 [1], arguably the most popular dataset for VLP evaluation. We also test on Flickr30k zero-shot image-text retrieval to remove any confounders that may be introduced during finetuning [14]. For VQAv2, we follow the standard practice [5] to train the models with both training and validation data, and test the models on the test-dev set. For Flickr30k, we follow the standard splits.

For comparison with state of the arts, we also evaluate models on visual reasoning (NLVR² [43]), visual entailment (SNLI-VE [52]), and image-text retrieval (COCO [27] and Flickr30k [34]) tasks. For retrieval tasks, we evaluate models in both zero-shot and finetuning settings.

Text Enc.	VQAv2	VE	IR	TR	SQuAD	MNLI
	Acc.	Acc.	R@1	R@1	EM	Acc.
Emb-only	67.13	74.85	49.06	68.20	-	-
ELECTRA	69.22	76.57	41.80	58.30	86.8	88.8
CLIP	69.31	75.37	54.96	73.80	-	-
DeBERTa	69.40	76.74	51.50	67.70	87.2	88.8
BERT	69.56	76.27	49.60	66.60	76.3	84.3
RoBERTa	69.69	76.53	49.86	68.90	84.6	87.6
ALBERT	69.94	76.20	52.20	68.70	86.4	87.9

Table 2. Comparisons of different text encoders without VLP. CLIP-ViT-224/32 is used as the vision encoder. All the text encoders are in base model size, except ALBERT, which is xlarge. Emb-only: only using word embeddings as text encoder. IR/TR: Flickr30k image/text retrieval. EM: exact match. The results of SQuAD and MNLI are copied from their corresponding papers. All the results on VL tasks are from their test-dev/val sets.

Vision Encoder	VQAv2	VE	IR	TR	ImageNet
Dis. DeiT B-384/16	67.84	76.17	34.84	52.10	85.2
BEiT B-224/16	68.45	75.28	32.24	59.80	85.2
DeiT B-384/16	68.92	75.97	33.38	50.90	82.9
ViT B-384/16	69.09	76.35	40.30	59.80	83.97
CLIP B-224/32	69.69	76.53	49.86	68.90	-
VOLO 4-448/32	71.44	76.42	40.90	61.40	86.8
CaiT M-384/32	71.52	76.62	38.96	61.30	86.1
CLIP B-224/16	71.75	77.54	57.64	76.90	-
Swin B-384/32	72.38	77.65	52.30	69.50	86.4

Table 3. Comparisons of different vision encoders without VLP. RoBERTa is used as the default text encoder. IR/TR: Flickr30k image/text retrieval; B: Base. The results of ImageNet classification are copied from their corresponding papers. All the results on VL tasks are from their test-dev/val sets. N and M in ViT-N/M denote the image resolution and patch size, respectively.

Implementation Details. We pre-train our models using AdamW [31] for 100k steps. The learning rates for the bottom and top layers are set to 1e-5 and 5e-5 respectively during pre-training. The warm-up ratio is set to 10%, and the learning rate is linearly decayed to 0 after 10% of the total training steps. We use center-crop to resize each image into the size of 224×224 or 384×384 depending on the adopted vision transformers.

4. Experiments

In this section, we provide comprehensive analysis of each individual module design. Specifically, (i) we study the impact of vision and language encoders in Section 4.1, (ii) we perform analysis on multimodal fusion designs in Section 4.2, (iii) we compare encoder-only and encoder-decoder architectures in Section 4.3, and (iv) we ablate pre-training objectives in Section 4.4. Finally, we compare with state of the arts in Section 4.5.

4.1. Impact of Vision and Language Encoders

4.1.1 Explorations without VLP

Since pre-training is time-consuming, we first perform an exploration study by comparing different text and visual en-

Text Enc.	Vision Enc.	VQAv2	Flickr-ZS	
			IR	TR
Emb-only	CLIP-32	73.99	60.32	74.10
BERT	CLIP-32	74.98	66.08	78.10
	CLIP-16	76.70	74.52	87.20
RoBERTa	CLIP-32	74.67	65.50	76.60
	CLIP-16	77.19	76.64	89.60
	Swin	76.43	71.68	85.30

Table 4. Comparisons of different vision and text encoders with VLP. Results on VQAv2 are on test-dev set. ZS: zero-shot.

coders without VLP for efficiency. Concretely, we initialize the bottom layers with specific pre-trained vision and text encoders, and randomly initialize the top layers. Then, we directly finetune the models on three tasks, including VQAv2, SNLI-VE, and Flickr30k retrieval. The learning rates for the bottom and top layers are set to $1e-5$ and $1e-4$, and the number of training epochs is set to 10 for all the tasks. We choose CLIP-ViT-224/32 [35] and RoBERTa [29] as the default encoders. Here, N and M in ViT-N/M denote image resolution and patch size, respectively.

Impact of Text Encoders. As shown in Table 2, there are no significant differences between the model performance of different text encoders. RoBERTa seems to achieve the most robust performance in this setting. Also, as can be seen from the Emb-only results, it is necessary to have a pre-trained encoder because otherwise the downstream task performance will be degraded.

Impact of Vision Encoders. As summarized in Table 3, both CLIP-ViT-224/16 and Swin Transformer can achieve decent performance in this setting. Notably, Swin Transformer can achieve an VQA score of **72.38** on the test-dev set *without any VLP*, which is already comparable to some VLP models after pre-training.

Conclusion. If we directly finetune the models on downstream tasks without any VLP, RoBERTa and Swin Transformer or CLIP-ViT perform the best. While models such as DeBERTa and BEiT can achieve better performance than the two models on pure language or vision tasks such as MNLI [49] or ImageNet classification [8], that does not necessarily indicate that they are more suitable for VL tasks.

4.1.2 Results with VLP

Now, we follow the default setting in Section 3.3, and compare different vision/text encoders with VLP. Based on the previous results, we compare Embed-only, BERT, and RoBERTa on the text side, and CLIP-ViT-224/32, CLIP-ViT-224/16, and Swin Transformer on the image side.

Results. As shown in Table 4, after VLP, the difference between BERT and RoBERTa seems to be diminished, but it is still important to have a pre-trained text encoder on the bottom (Embed-only vs. RoBERTa). For vision encoder,

Bottom LR	Top LR	VQAv2	Flickr-ZS	
			IR	TR
$1e-5$	$1e-5$	73.16	48.80	63.70
$2e-5$	$2e-5$	73.66	53.14	67.20
$3e-5$	$3e-5$	73.77	56.48	70.90
$5e-5$	$5e-5$	73.54	52.48	65.90
$1e-5$	$5e-5$	74.98	66.08	78.10

Table 5. Using different learning rates for the randomly-initialized and pre-trained parameters is better than using the same learning rate. Results on VQAv2 are on test-dev set. ZS: zero-shot.

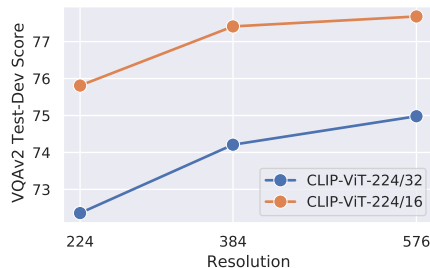


Figure 5. Increasing the image resolution during finetuning greatly improves the performance on the VQAv2 test-dev set.

both CLIP-ViT-224/16 and Swin Transformer can achieve pretty good performance. Especially, CLIP-ViT-224/16 can achieve a VQA score of **77.19/77.20** on the test-dev/test-std sets, respectively, outperforming the previous state-of-the-art region-based VinVL [57] models.

Useful Tricks. In experiments, we found two tricks for ViT-based VLP models that can greatly boost the performance. First, it is better to use a *larger* learning rate for the randomly initialized parameters than parameters initialized with pre-trained models, which is also found useful in some other NLP tasks [28]. As shown in Table 5, using the same learning rate for all parts of the model will lead to degraded performance, possibly because the pre-trained parameters already contain certain amounts of knowledge about vision and language, and finetuning them aggressively can result in the loss of these valuable information.

Second, similar to several previous work [20,56], we find that increasing the image resolution during finetuning can improve the model performance by a large margin, especially when the ratio of image resolution to patch size is low. Figure 5 shows that increasing the image resolution from 224 to 576 can improve the VQA score by about 3 and 1 points for the CLIP-ViT-224/32 and CLIP-ViT-224/16 model, respectively.

4.2. Analysis of the Multimodal Fusion Module

Now, following the default setting in Section 3.3, we perform investigations on multimodal fusion. First, we design both *merged attention* and *co-attention* models and investigate their performance. For the merged attention model (Figure 2b), the top transformer consists of M_{merged} encod-

Fusion	Decoder	VQAv2	Flickr-ZS	
			IR	TR
Merged attention	✗	74.00	57.46	73.10
Co-attention	✓	74.73	48.96	71.60

Table 6. Co-attention performs better than merged attention in our setting, and adding a decoder is not helpful for our discriminative VL tasks. Results on VQAv2 are on test-dev set. ZS: zero-shot.

ing layer, with each layer consisting of one self-attention block and one feed-forward network block. To help the model distinguish between the two modalities, we add a modality embedding to the input features before feeding them to the top transformer. For the co-attention model (Figure 2a), we feed the text and visual features to two M_{co} -layer transformers separately, and each top transformer encoding layer consists of one self-attention block, one cross-attention block, and one feed-forward network block. Compared with merged attention, co-attention allows separate transformation functions for the vision and language modalities. We set $M_{merged} = 12$ and $M_{co} = 6$ so that the numbers of parameters of the two models are roughly comparable to each other.

Results. Table 6 reports the downstream performance of the two models. The co-attention model performs better than the merged attention model in our setting, indicating that it is important to have different sets of parameters for the two modalities. Note that this contradicts with the findings in region-based VLP models [3], possibly because (i) findings of region-based VLP models cannot directly apply to ViT-based VLP models; (ii) most region-based VLP models only use pre-trained visual encoders, and also do not have a pre-trained text encoder included, thus the inconsistency between the two modalities will not favor symmetrical architecture like the co-attention model.

4.3. Encoder-Only vs. Encoder-Decoder

We then compare the encoder-only and encoder-decoder architecture. For the encoder-only model, we use the same co-attention model as in Section 4.2. For the encoder-decoder model, we set the number of layers to 3 for both the encoder and decoder, and each decoding layer has two separate cross-attention blocks that attend to the vision and text representations, respectively. According to [6], we adopt T5-style [36] language modeling objective as it works well for their model. Specifically, we mask 15% of input text tokens and replace contiguous text span with sentinel tokens, and the decoder is trained to reconstruct the masked tokens. For image-text matching, we feed the decoder with a special class token and it generates a binary output.

Results. As shown in Table 6, the encoder-only model can outperform the encoder-decoder model on our two discriminative tasks, which is consistent with the findings in [6].

Pre-training Objectives	VQAv2	Flickr-ZS	
		IR	TR
MLM	74.19	-	-
ITM	72.63	53.74	71.00
MLM+ITM	74.98	66.08	78.10
MLM+ITM + MIM (In-batch Negatives)	74.01	62.12	76.90
MLM+ITM + MIM (Discrete Code)	74.21	59.80	76.30

Table 7. Masked language modeling (MLM) and image-text matching (ITM) can both improve the model performance, but both of our designed masked image modeling (MIM) objectives lead to degraded performance on downstream tasks. Results on VQAv2 are on test-dev set. ZS: zero-shot.

However, it should be noted that the encoder-decoder architecture is more flexible, as it can perform tasks such as image captioning which may not be that straightforward for an encoder-only model to be applied to.

4.4. Ablations on Pre-training Objectives

In all the previous experiments, we pre-train our models with different objectives, following the default setting in Section 3.3. Now, we alter the pre-training objectives.

Results. As summarized in Table 7, both masked language modeling and image-text matching can bring performance improvements on downstream tasks. However, both of our masked image modeling objectives can lead to degraded performance on both VQAv2 and Flickr30k retrieval tasks. This further indicates that conclusions in region-based VLP models may not necessarily hold in vision transformer-based models. We hypothesize that the performance drop is due to the conflicts between different objectives, and some techniques in multi-task optimization [50, 54] may be borrowed to resolve the conflicts, which we list as one of the future directions. Another possible reason is that image patches can be noisy, thus the supervisions on reconstructing these noisy patches can be uninformative.

4.5. Comparison with Prior Arts

In this section, we evaluate our best-performing models (*i.e.*, RoBERTa-base+Swin Transformer and RoBERTa-base+CLIP-ViT-224/16 with co-attention fusion module, and with image resolutions set to 384 and 288, respectively), and compare them with previous work. We evaluate the models on visual question answering (VQAv2), visual reasoning (NLVR²), visual entailment (SNLI-VE), Flickr30k retrieval tasks in zero-shot and finetuning settings, and COCO retrieval tasks in the finetuning setting.

Main Results. As in Table 8 and 9, compared with models pre-trained with fewer than 10M images, our CLIP-based model (METER-CLIP-ViT_{BASE}) can achieve either the best or the second best scores on all the downstream tasks. Notably, our model can achieve a VQA score of 77.64% on the VQAv2 test-std set using only 4M images for pre-training,

Model	VQAv2		NLVR ²		SNLI-VE		Flickr-ZS					
	test-dev	test-std	dev	test	dev	test	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10
<i>Pre-trained with >10M images</i>												
ALBEF (14M) [23]	75.84	76.04	82.55	83.14	80.80	80.91	82.8	96.3	98.1	94.1	99.5	99.7
SimVLM _{BASE} (1.8B) [51]	77.87	78.14	81.72	81.77	84.20	84.15	-	-	-	-	-	-
SimVLM _{HUGE} (1.8B) [51]	80.03	80.34	84.53	85.15	86.21	86.32	-	-	-	-	-	-
<i>Pre-trained with <10M images</i>												
UNITER _{LARGE} [5]	73.82	74.02	79.12	79.98	79.39	79.38	68.74	89.20	93.86	83.60	95.70	97.70
VILLA _{LARGE} [11]	74.69	74.87	79.76	81.47	80.18	80.02	-	-	-	-	-	-
UNIMO _{LARGE} [25]	75.06	75.27	-	-	81.11	<u>80.63</u>	-	-	-	-	-	-
VinVL _{LARGE} [57]	<u>76.52</u>	76.60	82.67	83.98	-	-	-	-	-	-	-	-
PixelBERT [16]	74.45	74.55	76.5	77.2	-	-	-	-	-	-	-	-
CLIP-ViL (ResNet50x4) [41]	76.48	<u>76.70</u>	-	-	80.61	80.20	-	-	-	-	-	-
ViLT [57]	71.26	-	75.70	76.13	-	-	55.0	82.5	89.8	73.2	93.6	96.5
Visual Parsing [53]	74.00	74.17	77.61	78.05	-	-	-	-	-	-	-	-
ALBEF (4M) [23]	74.54	74.70	80.24	80.50	80.14	80.30	<u>76.8</u>	<u>93.7</u>	<u>96.7</u>	<u>90.5</u>	98.8	99.7
METER-Swin _{BASE}	76.43	76.42	82.23	82.47	80.61	80.45	71.68	91.80	95.30	85.30	97.70	99.20
METER-CLIP-ViT _{BASE}	77.68	77.64	<u>82.33</u>	<u>83.05</u>	<u>80.86</u>	81.19	79.60	94.96	97.28	90.90	<u>98.30</u>	<u>99.50</u>

Table 8. Comparisons with models pre-trained with <10M images on visual question answering, visual reasoning, visual entailment, and zero-shot image retrieval (IR) and text retrieval (TR) tasks. The best scores are in **bold**, and the second best scores are underlined.

Model	Flickr						COCO					
	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10
<i>Pre-trained with >10M images</i>												
ALBEF (14M) [23]	85.6	97.5	98.9	95.9	99.8	100.0	60.7	84.3	90.5	77.6	94.3	97.2
<i>Pre-trained with <10M images</i>												
UNITER _{LARGE} [5]	75.56	94.08	96.76	87.30	98.00	99.20	52.93	79.93	87.95	65.68	88.56	93.76
VILLA _{LARGE} [11]	76.26	94.24	96.84	87.90	97.50	98.80	-	-	-	-	-	-
UNIMO _{LARGE} [25]	78.04	94.24	97.12	89.40	98.90	<u>99.80</u>	-	-	-	-	-	-
VinVL _{LARGE} [57]	-	-	-	-	-	-	58.8	83.5	90.3	<u>75.4</u>	<u>92.9</u>	<u>96.2</u>
PixelBERT [16]	71.5	92.1	95.8	87.0	98.9	99.5	50.1	77.6	86.2	63.6	87.5	93.6
ViLT [57]	64.4	88.7	93.8	83.5	96.7	98.6	42.7	72.9	83.1	61.5	86.3	92.7
Visual Parsing [53]	73.5	93.1	96.4	87.0	98.4	99.5	-	-	-	-	-	-
ALBEF (4M) [23]	82.8	96.7	98.4	94.3	<u>99.4</u>	<u>99.8</u>	56.8	81.5	89.2	73.1	91.4	96.0
METER-Swin _{BASE}	79.02	95.58	98.04	92.40	99.00	99.50	54.85	81.41	89.31	72.96	92.02	96.26
METER-CLIP-ViT _{BASE}	<u>82.22</u>	<u>96.34</u>	98.36	94.30	99.60	99.90	<u>57.08</u>	<u>82.66</u>	<u>90.07</u>	76.16	93.16	96.82

Table 9. Comparisons with models pre-trained with <10M images on Flickr30k and COCO image retrieval (IR) and text retrieval (TR) tasks in the finetuning setting. The best scores are in **bold**, and the second best scores are underlined.

Model (#Pre-training Images)	test-dev	test-std
SimVLM _{BASE} (1.8B)	77.87	78.14
SimVLM _{HUGE} (1.8B)	80.03	80.34
METER-CoSwin _{HUGE} (14M)	80.33	80.54

Table 10. Pre-training a huge model under the METER framework with 14M images can lead to state-of-the-art performance on VQAv2, surpassing previous models trained with 1.8B images.

surpassing the state-of-the-art region-feature-based VinVL model by 1.04%, and outperforming the previous best fully transformer-based model (*i.e.*, ALBEF) by 1.6%. In addition, while ALBEF has specially-designed objectives for retrieval, our model can still outperform ALBEF on text and image retrieval tasks, further demonstrating the effectiveness of METER. Also, as shown in Appendix, we can maintain the fast inference speed of ViT-based models.

Scaling the Model. We also investigate if the METER framework is scalable. To this end, we pre-train our model with more images and larger vision backbone. Specifically, we pre-train the model with COCO, CC, CC12M [4], SBU, and VG datasets, consisting of about 14M images and 20M

image-caption pairs in total. We use CoSwin-Huge [55] as our vision backbone and RoBERTa-base as our text backbone. The hidden size of the fusion module remains unchanged. As shown in Table 10, our model can achieve state-of-the-art performance on VQAv2, surpassing previous models trained with 1.8B images. The results indicate that our METER framework is scalable.

Further Analysis. We also conduct experiments on image captioning, investigate multi-scale feature fusion, study the model performance on unimodal tasks after VLP, and provide visualization of learned attention maps. All these results are provided in Appendix.

5. Conclusion

We present METER, through which we systematically investigate how to train a fully-transformer VLP model in an end-to-end manner. Experiments demonstrate that we can achieve competitive performance with state-of-the-art models with only 4M images for pre-training. When further scaled up, METER achieves new state of the art on VQA.

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiaseen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *International Conference on Computer Vision (ICCV)*, 2015.
- [2] Hangbo Bao, Li Dong, and Furu Wei. BEiT: Bert pre-training of image transformers. *arXiv preprint*, 2021.
- [3] Emanuele Bugliarello, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. Multimodal pretraining unmasked: Unifying the vision and language BERTs. *Transactions of the Association for Computational Linguistics (TACL)*, 2021.
- [4] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [5] 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 (ECCV)*, 2020.
- [6] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning (ICML)*, 2021.
- [7] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations (ICLR)*, 2020.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [10] 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. In *International Conference on Learning Representations (ICLR)*, 2021.
- [11] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [12] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- [13] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced bert with disentangled attention. *arXiv preprint*, 2020.
- [14] Lisa Anne Hendricks, John Mellor, Rosalia Schneider, Jean-Baptiste Alayrac, and Aida Nematzadeh. Decoupling the role of data, attention, and losses in multimodal transformers. *Transactions of the Association for Computational Linguistics (TACL)*, 2021.
- [15] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [16] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-BERT: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint*, 2020.
- [17] 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*, 2021.
- [18] Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. Exploring the limits of language modeling. *arXiv preprint*, 2016.
- [19] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *International Conference on Computer Vision (ICCV)*, 2021.
- [20] Wonjae Kim, Bokyung Son, and Ildoo Kim. ViLT: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning (ICML)*, 2021.
- [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 (IJCV)*, 2017.
- [22] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations (ICLR)*, 2020.
- [23] 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. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [24] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint*, 2019.
- [25] 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. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [26] 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 (ECCV)*, 2020.
- [27] 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 (ECCV)*, 2014.

- [28] Yang Liu and Mirella Lapata. Text summarization with pre-trained encoders. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019.
- [29] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint*, 2019.
- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *International Conference on Computer Vision (ICCV)*, 2021.
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2018.
- [32] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlb-ert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [33] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2011.
- [34] 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 *International Conference on Computer Vision (ICCV)*, 2015.
- [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. In *International Conference on Machine Learning (ICML)*, 2021.
- [36] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research (JMLR)*, 2020.
- [37] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning (ICML)*, 2021.
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2016.
- [39] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016.
- [40] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [41] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint*, 2021.
- [42] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations (ICLR)*, 2019.
- [43] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [44] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019.
- [45] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint*, 2020.
- [46] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *arXiv preprint*, 2021.
- [47] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [49] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations (ICLR)*, 2019.
- [50] Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models. In *International Conference on Learning Representations (ICLR)*, 2020.
- [51] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint*, 2021.
- [52] Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint*, 2019.
- [53] Hongwei Xue, Yupan Huang, Bei Liu, Houwen Peng, Jianlong Fu, Houqiang Li, and Jiebo Luo. Probing intermodality: Visual parsing with self-attention for vision-language pre-training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [54] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [55] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang,

Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision. *arXiv preprint*, 2021.

- [56] Li Yuan, Qibin Hou, Zihang Jiang, Jiashi Feng, and Shuicheng Yan. Volo: Vision outlooker for visual recognition. *arXiv preprint*, 2021.
- [57] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. VinVL: Revisiting visual representations in vision-language models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.