Coarse-to-Fine Vision-Language Pre-training with Fusion in the Backbone

Zi-Yi Dou[*‡], **Aishwarya Kamath**[*♮], **Zhe Gan**[*†♠], **Pengchuan Zhang**[§], **Jianfeng Wang**[†]
Linjie Li[†], **Zicheng Liu**[†], **Ce Liu**[†], **Yann LeCun**[♮], **Nanyun Peng**[‡], **Jianfeng Gao**[†], **Lijuan Wang**[†]
[†]Microsoft [‡]University of California, Los Angeles [♮]New York University
{zdou,violetpeng}@cs.ucla.edu, {aish,yann.lecun}@nyu.edu, pengchuanzhang@fb.com
{zhgan,jianfw,linjli,zliu,liuce,jfgao,lijuanw}@microsoft.com

Abstract

Vision-language (VL) pre-training has recently received considerable attention. However, most existing end-to-end pre-training approaches either only aim to tackle VL tasks such as image-text retrieval, visual question answering (VQA) and image captioning that test high-level understanding of images, or only target region-level understanding for tasks such as phrase grounding and object detection. We present FIBER (**F**usion-**I**n-the-**B**ackbone-based transform**ER**), a new VL model architecture that can seamlessly handle both these types of tasks. Instead of having dedicated transformer layers for fusion after the uni-modal backbones, FIBER pushes multimodal fusion deep into the model by inserting cross-attention into the image and text backbones, bringing gains in terms of memory and performance. In addition, unlike previous work that is either only pre-trained on image-text data or on fine-grained data with box-level annotations, we present a two-stage pre-training strategy that uses both these kinds of data efficiently: (i) *coarse*-grained pre-training based on image-text data; followed by (ii) *fine*-grained pre-training based on image-text-box data. We conduct comprehensive experiments on a wide range of VL tasks, ranging from VQA, image captioning, and retrieval, to phrase grounding, referring expression comprehension, and object detection. Using deep multimodal fusion coupled with the two-stage pre-training, FIBER provides consistent performance improvements over strong baselines across all tasks, often outperforming methods using magnitudes more data. Code is available at https://github.com/microsoft/FIBER.

1 Introduction

Inspired by the success of language model pre-training [14, 56, 46], coupled with the unification of architectures used in the NLP and computer vision communities [15, 6], vision-language pre-training (VLP) [68, 49, 37, 8] has been receiving an increasing amount of attention. It has been proven that VLP can establish state-of-the-art performance on visual question answering [4], visual reasoning [66], image captioning, and image-text retrieval [45]. The pre-training objectives commonly used for these tasks, such as image-text matching, image conditioned masked language modeling and image-text constrastive learning, require multimodal understanding at the image level. Typically, this means the pre-training is done using images at lower resolution (*e.g.*, 384×384), making it possible to scale up training by using large batch sizes.

Recently, it has also been shown that tasks such as image classification and object detection (OD), which have been traditionally viewed as vision-only tasks, can benefit from being cast as VL tasks [55, 28, 38, 29]. Inspired by MDETR [29], GLIP [38] reformulates standard classification-based OD as phrase grounding. This opens up the possibility to leverage VLP for OD, and vice versa,

[*]Equal Technical Contribution [♠]Project Lead [§]Work done while at Microsoft

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

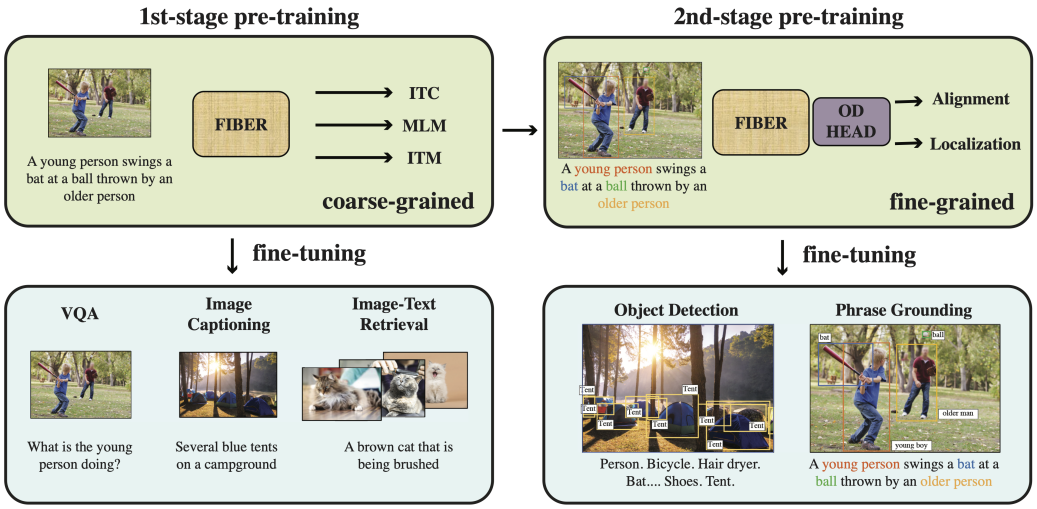

Figure 1: The proposed coarse-to-fine pre-training framework for vision-language tasks. We first perform *coarse*-grained pre-training with image-text data for VQA, image captioning and retrieval tasks, and then perform *fine*-grained pre-training with image-text-box data for phrase grounding and object detection tasks. The same FIBER architecture is used for both stages. OD: object detection. MLM: masked language modeling. ITM: image-text matching. ITC: image-text contrastive loss.

and this unification has led to impressive performance on several established OD as well as phrase grounding benchmarks [54]. Since these tasks involve fine-grained image understanding between regions in the image and phrases in the text, and also require prediction of precise bounding boxes at the output, the pre-training typically involves using high resolution input images (*e.g.*, 800×1,333).

Existing multimodal architectures typically do not support both kinds of tasks. Specifically, the fully end-to-end VLP models such as ALBEF [36], METER [16], and SimVLM [75] can achieve the state of the art (SoTA) on image-level understanding tasks, but it is non-trivial to extend them for region-level VL tasks because predicting bounding boxes is typically hard in end-to-end settings. On the other hand, MDETR [29] and GLIP [38] are designed to predict bounding boxes, but have not been shown to support tasks such as image captioning and retrieval. Further, fine-grained pre-training not only requires data with bounding box annotations that are cumbersome to acquire, but the requirement of high input image resolution makes pre-training very costly, especially when using standard Transformer architectures [69] that have quadratic complexity in the size of the image. A natural but challenging question arises: *can we have a unified framework for efficient VL pre-training that benefits both image-level and region-level VL tasks (e.g., both VQA and OD)?*

We answer this question by proposing two ideas: (i) a novel model architecture that can handle various types of tasks and pre-training strategies (high and low resolution inputs, image and region level outputs) more efficiently than previous work (see Section 3.1 and 4), and (ii) a two-stage pre-training pipeline.

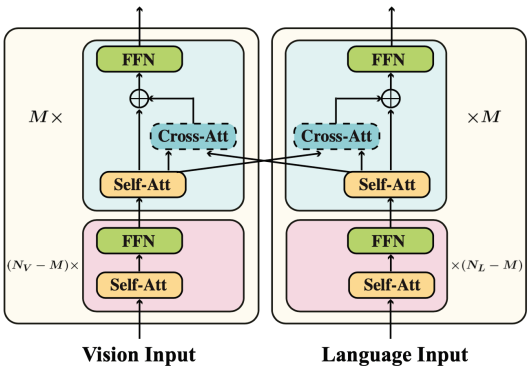

Figure 2: Model architecture for FIBER. Swin transformer is used as the image backbone, simplified here for illustration purposes.

In terms of *architecture*, we present FIBER, shown in Figure 2, which performs deep multimodal fusion in the backbone. Specifically, instead of having a few dedicated transformer layers on top of the image and text encoders for fusion (*e.g.*, as is commonly done in previous work [40, 8, 16, 29, 38]), we propose to directly insert cross-attention modules into the image and text backbones. Additionally, we support the ability to switch between a dual encoder (for fast image retrieval) and a fusion encoder (for VQA and captioning) readily, by switching on or off the cross-attention modules. With the same model architecture, by simply adding an object detection

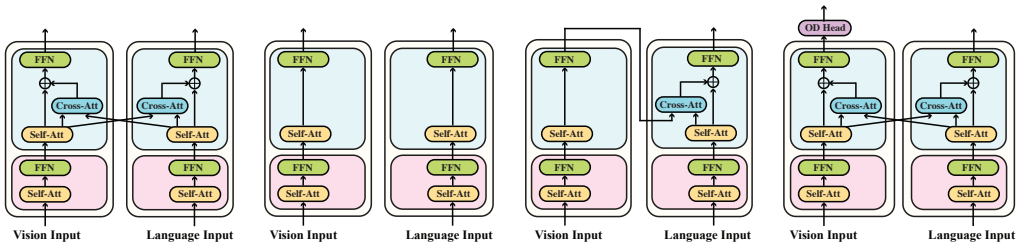

(a) VQA/Visual Reasoning **(b)** Retrieval **(c)** Captioning **(d)** Grounding/OD

Figure 3: FIBER can be readily adapted to various downstream VL tasks, ranging from VQA, image captioning and retrieval, to phrase grounding and object detection (OD).

head (*e.g.*, Dynamic Head [12]) on top, FIBER can be readily extended to visual grounding, referring expression comprehension and (open-vocabulary) OD tasks as well.

By considering the nature of different VL tasks, FIBER is pre-trained with a coarse-to-fine two-stage pipeline, as detailed in Figure 1. Specifically,

- During *coarse*-grained pre-training, FIBER takes low-resolution (384×384) images as input, and is pre-trained with image-text matching, masked language modeling, and image-text contrastive losses, as used in previous work [16, 74, 72]. The pre-trained model can then be directly finetuned for VQA and image captioning tasks (Figure 3a and 3c). By switching off the cross-attention modules, FIBER also automatically functions as a dual encoder for fast image-text retrieval (Figure 3b).
- During *fine*-grained pre-training, FIBER uses the coarse pre-trained model as initialization, in addition to randomly initialized parameters for the OD head. At this stage, the model takes high-resolution (800×1,333) images as input, and is pre-trained with bounding box localization loss and word-region alignment loss, as used in GLIP [38]. We use image-text-box data with ground-truth box annotations for pre-training, and the model can be directly fine-tuned for grounding and detection tasks (Figure 3d).

Compared to fine-grained pre-training, coarse-grained pre-training is easier to scale up, as it only requires paired image-text data which can be easily harvested from the web. Crucially, we show that re-using all the parameters from our coarse-grained pre-trained model for fine-grained pre-training alleviates the requirement for large amounts of box-level annotated data. In our experiments, we show that on fine-grained tasks such as Flickr30k Entities, FIBER using coarse-grained pre-training achieves gains even over previous SoTA (GLIP [38]) that uses 25× more box-level annotated images during the fine-grained pre-training stage. We also show that our architecture is much more efficient in terms of training time on OD tasks, as compared to GLIP .

FIBER is the first end-to-end VLP model that can support VL tasks encompassing image-level and region-level outputs. We conduct experiments on VQAv2 [4], NLVR^2 [66], COCO captioning [45], NoCaps [1], COCO and Flickr30k image-text retrieval [54], as well as on phrase grounding [54], referring expression comprehension [83], COCO and LVIS detection [20], and a suite of 13 object detection in the wild datasets [38]. We show that our model can provide consistent performance improvement over strong baselines (*e.g.*, METER [16] and GLIP [38]) across tasks.

2 Related Work

VLP for Classical VL Tasks. ViLBERT [49] and LXMERT [68] were the first two methods to introduce using transformers for VLP. Since then, we have witnessed a boom of VLP methods [37, 34, 65, 81, 25, 79, 89, 42, 9, 39]. Early methods mainly focus on the use of pre-trained object detectors to extract image region features offline, such as UNITER [8], OSCAR [40], VILLA [18] and VinVL [87]. More recently, end-to-end VLP methods that use the image directly as input have become popular. In these approaches, convolution networks or vision transformers [15] are used as the image backbone, with additional transformer layers for modeling multimodal fusion [27, 26, 32, 76, 36, 72]. Prominent examples along this line include ViLT [32], ALBEF [36], SimVLM [75], METER [16], X-VLM [85] and BLIP [35]. These models have achieved the current SoTA on major VL benchmarks such as VQA and image captioning. However, they cannot be directly used for tasks such as object detection.

Model	VQA†	$O(n+m)$ Retrieval‡	Captioning	Grounding	OD	End2End
ViLBERT [49], LXMERT [68], UNITER [8]	✓	×	×	✓	×	×
OSCAR [40], VinVL [87]	✓	×	✓	✓	×	×
PixelBERT [27], CLIP-ViL [63], ViLT [32]	✓	×	×	×	×	✓
CLIP [55]*, ALIGN [28]	×	✓	×	×	×	✓
VL-T5 [9]	✓	×	✓	×	×	×
METER [16], SimVLM [75]	✓	×	✓	×	×	✓
ALBEF [36], FLAVA [64], VLMo [74]	✓	✓	×	×	×	✓
BLIP [35], CoCa [82], Flamingo [2]	✓	✓	✓	×	×	✓
MDETR [29], GLIP [38]	✓	×	×	✓	✓	✓
UNICORN [78], X-VLM [85], OFA [73]	✓	×	✓	✓	×	✓
FIBER	✓	✓	✓	✓	✓	✓

Table 1: Comparison among different VLP models. FIBER is the only VLP model that can support all tasks considered. (†) VQA is used as a representative VL classification task. (‡) $O(n+m)$ retrieval denotes model backbones process inputs $O(n+m)$ times given n images and m text sentences during image-text retrieval. ($*$) Here, we mainly focus on what tasks CLIP can be directly used for.

VLP for Vision Tasks. Recently, it has been shown that image-text data can be used to learn image encoders from scratch [13, 59]. By performing large-scale contrastive pre-training, CLIP [55] and ALIGN [28] display strong zero-shot image classification capabilities. While these models mainly tackle image-level understanding tasks, MDETR [29] extends the end-to-end OD model DETR [6], and uses contrastive learning along with an alignment loss to learn correspondences between image regions and text phrases, opening up the possibility to tackle tasks such as phrase grounding and long-tailed OD using VL models. This has inspired many follow-up works to further enhance the pre-training [41, 80, 50, 77], among which GLIP [38] shows that OD can also be cast as a VL task (*i.e.*, phrase grounding). However, it has not been shown how traditional VL tasks such as VQA, captioning and retrieval can be well supported in GLIP [38] and MDETR [29].

Unified VL Modeling. There have been a few recent attempts that try to develop unified VL models. VL-T5 [9] unifies VL tasks as text generation; however, pre-trained object detectors are used for image feature extraction, so the model cannot be end-to-end pre-trained. UniT [23] proposes a multimodal multi-task framework with a unified transformer; however, it can only support VQA and object detection tasks, but not captioning and grounding. GPV [21] proposes a general-purpose vision system, and FLAVA [64] presents a VL system similar to METER [16]; however, they did not evaluate on grounding and detection tasks, and their performance on other VL tasks is still far from SoTA. UNICORN [78] and OFA [73] reformulate grounding as a sequence generation task, by borrowing ideas from Pix2Seq [7]. However, these approaches have not been demonstrated to work on standard OD benchmarks, and also cannot be used as dual encoders for fast image retrieval. Our model is the first work that can support not only VQA, image captioning and $O(n+m)$ retrieval, but also visual grounding and object detection, with impressive performance across all tasks. A detailed comparison is provided in Table 1.

3 Method

In this section, we first describe the proposed model architecture in Section 3.1. We then illustrate our two-stage pre-training paradigm in Section 3.2, followed by fine-tuning strategies for all the tasks supported by FIBER in Section 3.3.

3.1 Fusion in the Backbone

The architecture of FIBER is shown in Figure 2. Different from models that stack a modality fusion module on top of the vision or language backbones [8, 16], we insert multimodal fusion inside the backbones, and include a gating mechanism for the cross-modal layers (shown in Figure 4). Specifically, at each encoding layer, we have:

$$\begin{aligned} \tilde{\boldsymbol{x}} &= \text{SELF-ATT}(\boldsymbol{x}), \\ \boldsymbol{x} &= \boldsymbol{x} + \tilde{\boldsymbol{x}} + \alpha * \text{CROSS-ATT}(\tilde{\boldsymbol{x}}, \boldsymbol{y}), \\ \boldsymbol{x} &= \boldsymbol{x} + \text{FFN}(\boldsymbol{x}), \end{aligned} \tag{1}$$

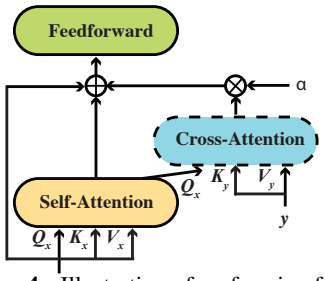

Figure 4: Illustration of performing fusion in the backbone. $(\boldsymbol{x}, \boldsymbol{y})$ are the (image, text) or (text, image) representations, and α is a learnable scalar.

where α is a learnable parameter initialized to 0. For simplicity, we insert the same number of cross-attention layers into the vision and language backbones.

By inserting cross-attention layers with the gating mechanism, we enable cross-modal interactions without affecting the original computational flow of the backbones at the beginning of model training. Also, we can easily switch off the interactions by setting α to 0, and the backbones can be used in the dual-encoder setting. In addition, compared to stacking a large number of transformer layers on top of the backbones, our approach of inserting cross-attention layers is relatively light-weight and thus more memory-efficient. To illustrate, both GLIP [38] and METER [16] use an additional 110M modality fusion parameters for a base-size model, while FIBER only adds about 26M parameters. During training, the fusion module of FIBER only consumes half of the FLOPs needed by METER (12.35 vs. 24.04 GFLOPs for one instance). We experimented with two other model variants for fusion in the backbone, the details of which are provided in Appendix.

3.2 Coarse-to-Fine Pre-training

We divide VL tasks into two categories based on whether or not we need to generate region-level outputs on the image side. While these two kinds of tasks are characteristically different, they both require fusion between the vision and language modalities, and we hypothesize that sharing as many parameters as possible between the model used for these two sets of tasks will be beneficial. Based on this motivation, we propose a two-stage pre-training paradigm, where we first pre-train models with image-level objectives on images at low resolution, and then perform further pre-training with region-level objectives where the input images are at a higher resolution. In this way, the coarse-grained supervision from the first stage can provide good initialization for the second stage for all the shared parameters. FIBER with the same architecture (Swin Transformer [47] and RoBERTa [46]) is used as the backbone for both stages of pre-training.

Coarse-grained Pre-training. For tasks like VQA and captioning, it has been demonstrated [36, 16, 74] that masked language modeling (MLM), image-text matching (ITM), and image-text contrastive (ITC) objectives are helpful for ViT-based VLP models. Following previous work, we use all the three objectives during pre-training. Specifically,

- **For ITC**, the inserted cross-attention modules are switched off, so FIBER functions as a dual encoder. Given a batch of N image-caption pairs, we first compute their representations with our vision and language encoders independently without modality fusion, and then maximize the similarities between N positive image-text pairs while minimizing the similarities between the rest $N^2 - N$ negative pair, via a contrastive loss.
- **For MLM and ITM**, the inserted cross-attention modules are switched on, so FIBER now functions as a fusion encoder. For MLM, we randomly mask 15% of the input tokens and the model is trained to reconstruct the original tokens. For image-text matching, the model is given an image-text pair and predicts whether they are matched. Following previous work [36, 74], we sample hard negatives based on the similarities computed from the above ITC loss.

Fine-grained Pre-training. Most existing VL architectures [8, 68, 29, 35, 73, 9] use vanilla transformers both for encoding the vision as well as language inputs. However, in contrast to tokens in text, the entities of interest in images do not all occur at the same scale. Being able to accurately model the image at different scales is especially important for tasks such as object detection and phrase grounding. To handle this, it is typical in object detection literature to use input images at higher resolutions (800×1333), which becomes problematic when using vanilla transformers that scale quadratically in the length of the input sequence. As mentioned earlier, we use a Swin Transformer [47] as our image encoder, which provides hierarchical representations of the image while having linear complexity in the size of the image. We combine these multi-scale representations using an FPN [43] for object detection training. For fine-grained pre-training, we switch on the cross-attention modules, using FIBER as a fusion encoder. This ensures that the image representations that are passed to the FPN are already text-aware, and is a crucial difference compared to GLIP [38], where the image-text fusion takes place in the object detection head. Once the text-aware image features are extracted by the Swin backbone and image-aware text features are extracted using RoBERTa [46], the image features after the FPN are fed to a DynamicHead [12] which predicts a set of regions. Just as in [38], we compute the dot product between the image region features $\boldsymbol{R}_{\text{TA}}$ and the contextualized

token representations $\boldsymbol{T}_{\text{IA}}$ to compute the grounding score:

$$\boldsymbol{I}_{\text{TA}}, \boldsymbol{T}_{\text{IA}} = \text{FIBER}(I, T), \ \boldsymbol{R}_{\text{TA}} = \text{OD-HEAD}(\boldsymbol{I}_{\text{TA}}), \ S_{\text{GROUNDING}} = \boldsymbol{R}_{\text{TA}} \boldsymbol{T}_{\text{IA}}^{\top}, \tag{2}$$

where $\boldsymbol{R}_{\text{TA}}$ represents regions that are text aware, produced using the OD-Head that takes as input $\boldsymbol{I}_{\text{TA}}$, which are image representations that are already text-aware and $\boldsymbol{T}_{\text{IA}}$ are the text features that have already attended to the image features. The typical object detection model has a classification head that predicts the label of the object, and a localization head that predicts the bounding box. We follow GLIP [38] by substituting the classification head with the grounding score $S_{\text{GROUNDING}}$. The localization loss is composed of two parts: a centerness loss and GIoU loss, which are used to supervise the box prediction. Taken together, FIBER learns the correspondence between regions in the image and phrases in the text, making it possible to tackle tasks such as phrase grounding and object detection using the same framework. We use ATSS framework [88] in our paper, but our method can be combined easily with other object detectors such as Faster-RCNN [57] and RetinaNet [44] as well.

3.3 Adaptation to Downstream Tasks

We now describe how we adapt FIBER to different downstream tasks as depicted in Figure 3.

- **For VL classification tasks such as VQA**, we use FIBER as a fusion encoder. Specifically, the top M layers of the vision and language backbones interact with each other and produce multimodal representations. The final layer representations of the two modalities are concatenated together to generate the final outputs for tasks such as VQA and visual reasoning.
- **For retrieval tasks**, we switch off the inserted cross-attention modules to use FIBER as a dual encoder for fast image-text retrieval.
- **For captioning**, we adapt FIBER by only keeping the image-to-text cross-attentions and using causal masks in the decoding side. The representations of the final image encoding layer are fed into the cross-attention modules. In this way, the model is turned into a seq2seq model [67, 10] and performs captioning in an auto-regressive way.
- **For phrase grounding, object detection and referring expression comprehension**, we use FIBER as a fusion encoder, and the OD-Head introduced during fine-grained pre-training receives image features that are already language aware due to the multimodal representations extracted by FIBER. The pre-trained model is directly used without any modifications for these tasks.

4 Experiments

Pre-training Datasets. Following previous work [8, 32, 36, 16, 72, 74], we perform coarse-grained pre-training on COCO [45], Conceptual Captions [62], SBU Captions [51], and Visual Genome [33]. The four datasets consist of about 4M images in total. For fine-grained pre-training, we use two data sources: data curated by MDETR [29] after removing the COCO images, and the Objects365 [61] detection dataset, together consisting of about 0.8M images. We ensure that we exclude any data that exists in the validation or test splits of downstream tasks.

Type of Fusion	COCO Val2017	GPU-hours V100 (32GB)	Sec/Iter
No Fuse	53.9	511	1.31
GLIP-B [38]	54.6	840	2.14
FIBER-B	54.5	540	1.38

Table 2: Object detection on COCO [45], without vision-language pre-training. We initialize the text encoder and image backbones using a pre-trained RoBERTa and a Swin transformer pre-trained on ImageNet22k. Our proposed FIBER model achieves the same performance as GLIP [38] while taking much less time to train.

Architecture. We adopt Swin-Base [47] and RoBERTa-Base [46] as our vision and text backbones, which are initialized with weights from uni-modal pre-training. We insert cross-attention blocks into the top 6 blocks of the vision and text encoders. The input resolution is 384×384 for coarse-grained pre-training and $800 \times 1,333$ for fine-grained pre-training. Using a hierarchical vision transformer enables us to efficiently tackle these high resolution tasks, which would be expensive in models such as BLIP [35] that rely on the vanilla transformer architecture. In METER [16], which does explore using a Swin transformer as the image encoder, the multi-modal fusion occurs in layers specifically designed to align the modalities, only after the image and text features are extracted from the uni-modal backbones. This is in contrast to our approach where the hierarchical image features that are used in the FPN for fine-grained training

Model	#Pretrain Images	VQAv2		NLVR^2		Flickr30k		COCO	
		test-dev	test-std	dev	test-P	IR@1	TR@1	IR@1	TR@1
Base-size models pre-trained on COCO, VG, SBU, and CC datasets									
UNITER-B [8]	4M	72.70	72.91	77.18	77.85	72.5	85.9	50.3	64.4
VILLA-B [18]	4M	73.59	73.67	78.39	79.30	74.7	86.6	-	-
UNIMO-B [39]	4M	73.79	74.02	-	-	-	-	-	-
ViLT-B [32]	4M	71.26	-	75.70	76.13	64.4	83.5	42.7	61.5
ALBEF-B [36]	4M	74.54	74.70	80.24	80.50	82.8†	94.3†	56.8†	73.1†
VLMo-B [74]	4M	76.64	76.89	82.77	83.34	79.3	92.3	57.2	74.8
METER-Swin-B [16]	4M	76.43	76.42	82.23	83.47	79.02	92.4	54.85	72.96
X-VLM [85]	4M	78.22	78.37	84.41	84.76	86.9†	97.0†	63.4†	81.2†
Models pre-trained on more data and/or with larger size									
VLMo-L [74]	4M	79.94	79.98	85.64	86.86	84.5	95.3	60.6	78.2
BLIP$_{\text{CapFilt-L}}$ [35]	129M	78.25	78.32	82.15	82.24	87.5†	97.2†	64.1†	81.2†
SimVLM-B [75]	1.8B	77.87	78.14	81.72	81.77	-	-	-	-
SimVLM-H [75]	1.8B	80.03	80.34	84.53	85.15	-	-	-	-
FIBER-B	4M	**78.55**	**78.46**	**84.59**	**85.52**	**81.44**	**92.90**	**58.01**	**75.38**

Table 3: Results on VL classification and retrieval. Complete results on image-text retrieval are provided in Appendix. We also include models pre-trained on more data and/or with larger size. FIBER and VLMo use dual encoders for retrieval. (†) ALBEF, X-VLM, and BLIP first use its dual encoder to obtain top-k candidates, and then use its fusion encoder to re-rank the candidates. Our retrieval results with re-ranking are provided in Table 4. All the other models use fusion encoders.

are already language aware, due to the multi-modal fusion being in the backbone. This also lets us avoid adding additional "language-aware deep fusion layers" [38] as part of the OD head as in GLIP, resulting in 1.5x faster training while maintaining performance as shown in Table 2. While in principle it would be possible to use the image features extracted by METER's backbone for object detection, it would be necessary as in GLIP to add additional layers to make the visual features "language-aware" for good detection performance, especially on datasets with limited training data and with rare and infrequent objects.

Implementation Details. We perform coarse-grained pre-training for 100k steps with 4,096 batch size on 64 A100 GPUs. We use AdamW [48] with the peak learning rates of 1e-4 for the backbones and 5e-4 for the cross-modal parameters. We use linear warmup over the first 1k steps and linear decay. For fine-grained pre-training, we train for 800k steps on 64 V100 GPUs, with a batch size of 64. We use a learning rate of 1e-5 for the language backbone, and 1e-4 for the rest of the model with a weight decay of 0.01. We use a linear warmup over the first 2k steps and then a constant learning rate, with two learning rate drops by a factor of 10 at 67% and 89% of the total number of steps.

4.1 Results on Downstream Tasks

Vision-Language Classification. We first experiment on two representative VL classification tasks, including VQAv2 [4] and NLVR^2 [66]. As reported in Table 3, we achieve the best performance compared to other models in the same setting. It is worth noting that FIBER pre-trained with 4M images can achieve better performance than BLIP trained with 129M images and SimVLM trained with 1.8B images. The results indicate that introducing fusion modules into the backbone is an effective alternative to appending them on the top of uni-modal backbones.

Image-Text Retrieval. In Table 3 we report image retrieval performance in the dual encoder setting, achieving competitive performance on both Flickr30k [54] and COCO [45] retrieval tasks. However, previous work has shown that fusion encoders obtain superior performance, albeit at the cost of efficiency as it involves feeding every image-text pair into the model. To illustrate, on the COCO test data, ranking the similarities between 5K images and 25K captions requires the model to process each image-caption pair 75M times, whereas the dual encoder model only needs 30K forward passes. As shown in Table 4, the fusion encoder can indeed surpass the dual encoder on retrieval tasks by a large margin. In addition, directly ensembling the two models by summing their similarity scores together for each image-caption pair can bring us huge improvements.

Model	Flickr30k						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
FIBER-ITC	81.44	96.72	98.48	92.90	99.50	99.90	58.01	83.45	90.11	75.38	94.04	97.36
FIBER-ITM	84.10	97.54	98.88	95.10	99.60	99.90	59.03	84.04	91.03	75.14	93.88	97.36
FIBER-ITC+ITM Ensemble	90.96	98.44	99.14	96.00	99.70	**100.00**	**69.73**	**90.66**	**94.59**	**80.10**	**95.60**	**97.98**
ALBEF [36]	82.8	96.7	98.4	94.3	99.4	99.8	56.8	81.5	89.2	73.1	91.4	96.0
X-VLM [85]	86.1	97.4	98.7	**96.8**	**99.8**	100.0	63.1	85.7	91.6	**80.4**	95.5	**98.2**
FIBER-Rerank-10	90.94	98.16	98.48	95.80	99.60	99.90	68.71	87.69	90.09	79.66	95.34	97.36
FIBER-Rerank-20	90.10	98.38	99.14	95.90	**99.80**	100.00	69.32	89.52	93.33	79.78	95.20	97.66
FIBER-Rerank-50	**91.08**	98.50	**99.37**	96.10	99.70	100.00	69.58	90.41	94.35	79.98	95.40	97.76
FIBER-Rerank-100	91.02	**98.54**	99.34	96.00	99.70	100.00	**69.63**	**90.54**	**94.47**	80.06	**95.60**	97.96

Table 4: Additional results on image-text retrieval, where (i) the fusion encoder is used for retrieval, or (ii) the dual encoder is first used to obtain top-k candidates, and then the fusion encoder is used to re-rank the candidates. We also provide a full set of results on all evaluation metrics.

Model	#Pretrain Images	COCO				NoCaps Val		NoCaps Test	
		B@4	M	C	S	C	S	C	S
Models trained without CIDEr optimization									
UFO-B [72]	4M	36.0	28.9	122.8	22.2	80.7	12.5	78.8	12.5
ViTCAP [17]	4M	36.3	29.3	125.2	22.6	-	-	-	-
METER-CLIP-B [16]	4M	38.8	30.0	128.2	23.0	-	-	-	-
X-VLM [85]	4M	39.8	-	133.1	-	-	-	-	-
VinVL-B [87]	5.7M	38.2	30.3	129.3	**23.6**	-	-	-	-
BLIP$_{\text{CapFilt-L}}$ [35]	129M	39.7	-	133.3	-	109.6	14.7	-	-
LEMON-B [24]	200M	40.3	30.2	133.3	23.3	106.8	14.1	-	-
SimVLM-B [75]	1.8B	39.0	32.9	134.8	24.0	-	-	94.8	13.1
FIBER-B	4M	39.1	30.4	128.4	23.1	88.6	13.0	86.0	12.9
FIBER-GOLD-B	4M	**40.3**	**30.7**	**133.6**	**23.6**	**92.8**	**13.4**	**90.6**	**13.4**
Models trained with CIDEr optimization									
ViTCAP [17]	4M	41.2	30.1	138.1	24.1	89.2	12.7	-	-
X-VLM [85]	4M	41.3	-	140.8	-	-	-	-	-
VinVL-B [87]	5.7M	40.9	30.9	140.4	**25.1**	94.3*	13.1*	92.5*	13.1*
LEMON-B [24]	200M	41.6	31.0	142.7	25.1	-	-	-	-
FIBER-B	4M	42.8	31.0	142.8	24.3	96.7	13.4	94.1	13.4
FIBER-GOLD-B	4M	**43.4**	**31.3**	**144.4**	24.6	**99.2**	**13.7**	**97.1**	**13.8**

Table 5: Results of base-size models on image captioning. We grey models pre-trained on larger magnitudes of data. Numbers with '*' are obtained with constrained beam search during inference and without VLP. The complete results on all metrics are provided in Appendix. B@4: BLEU@4, M: METEOR, C: CIDEr, S: SPICE.

Further, we explore combining the strengths of both strategies by performing re-ranking as in [19, 35, 36]. Specifically, we first retrieve the top-k most similar instances using the dual encoder setup, and then add the similarity scores between the given instance and the top-k candidates provided by the fusion encoder to the original scores to perform retrieval. From Table 4, we can see that this strategy provides a balance between efficiency and performance, and that just re-ranking the top-10 instances can achieve comparable performance with ensembling.

Image Captioning. We also evaluate our models on COCO [45] and NoCaps [1] captioning to test whether FIBER can be adapted to generation tasks. As in Table 5, FIBER can achieve better performance than models trained on the same data with and without CIDEr optimization [58]. We find that integrating GOLD [52] into FIBER can bring significant improvements, outperforming models trained with hundreds of millions of images. Notably, we establish the absolute state-of-the-art CIDEr scores on COCO for base-size models. Considering that FIBER is not pre-trained to perform captioning, the results demonstrate the strong generalization ability of FIBER.

Phrase Grounding. Our fine-grained pre-training stage incorporates Flickr30k entities grounding data, and we achieve 87.4 on the Recall@1 metric on the test set without any subsequent fine-tuning. This not only surpasses the current SoTA [38] using a smaller sized model (Swin-B compared to their Swin-L), but also uses 25x less fine-grained data. Our FIBER model is able to leverage the image-text coarse-grained pre-training stage better, instead of relying on expensive pseudo-labelling of large

Model	Image Backbone	#Pretrain Images (fine-grained)	Flickr30k Val			Flickr30k Test		
			R@1	R@5	R@10	R@1	R@5	R@10
Visual-BERT [37]	ResNet-101	120k	70.4	84.5	86.3	71.3	85.0	86.5
MDETR [29]	EN-B5	200k	83.6	93.4	95.1	84.3	93.9	95.8
GLIP [38]	Swin-B	860k	85.7	95.0	96.2	86.1	95.5	96.4
Models pre-trained on more data and/or with larger size								
GLIP [38]	Swin-L	27M	86.7	96.4	97.9	87.1	96.9	98.1
FIBER-B	Swin-B	860k	**87.1**	**96.1**	97.4	**87.4**	**96.4**	97.6
w/o C.G. VLP	Swin-B	860k	86.2	96.0	**97.6**	86.5	**96.4**	**97.7**

Table 6: Phrase grounding performance on Flickr30k entities dataset. We reproduce GLIP-Base sized results, and GLIP-Large sized results are taken from [38]. FIBER with Base size outperforms a GLIP-L which is trained with 25x more fine-grained data on the R@1 metric. Further, FIBER without coarse-grained VL pretraining outperforms GLIP-B when trained on the same fine-grained data.

Model	Pre-training data		RefCOCO			RefCOCO+			RefCOCOg	
	Im-Txt	Im-Txt-Box	val	testA	testB	val	testA	testB	val	test
MDETR-B [29]		✓	87.51	90.40	82.67	81.13	85.52	72.96	83.35	83.31
UNICORN-B [78]		✓	88.29	90.42	83.06	80.30	85.05	71.88	83.44	83.93
Models pre-trained on more data and/or with larger size										
UNITER-L [8]	✓		81.41	87.04	74.17	75.90	81.45	66.70	74.86	75.77
VILLA-L [18]	✓		82.39	87.48	74.84	76.17	81.54	66.84	76.18	76.71
OFA-L [73]	✓	✓	90.05	92.93	85.26	84.49	90.10	77.77	84.54	85.20
FIBER-B	✓	✓	**90.68**	**92.59**	**87.26**	**85.74**	**90.13**	**79.38**	**87.11**	**87.32**

Table 7: Results on referring expression comprehension datasets.

web-scale corpus and subsequent high-resolution training on this generated fine-grained data as in [38]. We also compare our approach without using any coarse-grained VL training (image encoder initialized to Swin-B weights from IN22k, and text encoder initialized to pre-trained RoBERTa), and even in this setting, we are able to outperform a similarly sized GLIP model (GLIP-B), proving that our fusion in the backbone is better at capturing fine-grained image-text understanding.

Referring Expression Comprehension (REC). In contrast to many previous works [8, 18, 49] that tackle the REC task by re-ranking object proposals provided by an off-the-shelf detector, we follow [29] to directly predict the bounding box for the given referring expression. Using our proposed two stage pre-training, FIBER achieves better performance than current SoTA [73] that uses a Large sized model. Notably, on RefCOCOg [84], which contains much longer referring expressions than in RefCOCO/RefCOCO+ [31], we observe more than 2 points boost over OFA-L. On the challenging testB split of both RefCOCO and RefCOCO+, FIBER outperforms current SoTA, OFA-L.

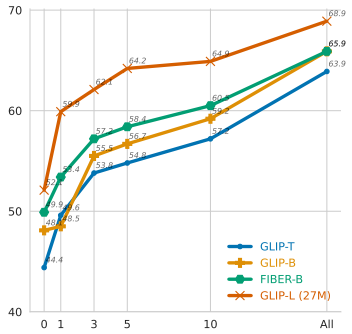

Figure 5: Few-shot results on the aggregated 13 ODinW datasets.

Object Detection. We report FIBER results on two standard object detection benchmarks, COCO [45] and LVIS [20], in zero-shot transfer[1] as well as fine-tuned settings in Table 8. The LVIS dataset consists of a long-tail of object classes, and is a popular test-bed for evaluating models on their generalization capabilities and robustness to class imbalance. On the APr metric, which is the Average Precision on

[1]Following [55, 86], we consider zero-shot transfer to mean that during pre-training we may have seen relevant data but it is not used for training for the task of interest. For instance, our coarse-grained pre-training includes some images from COCO (without any box information), but we do not have any COCO images in our fine-grained training that we use to train the object detection head.

Model	COCO Val 2017	LVIS MiniVal				ODinW
	AP	APr	APc	APf	AP	
	Zero-shot/Fine-tune	Zero-shot/Fine-tune				Zero-shot/Fine-tune
Mask R-CNN [22]	-	- /26.3	- /34.0	- /33.9	- /33.3	-
MDETR [29]	-	- /20.9	- /24.9	- /24.3	- /24.2	-
GLIP-T [38]	46.7/55.1	17.7/ -	19.5/ -	31.0/ -	24.9/ -	44.4/63.9
GLIP-B [38]	48.1/57.0	17.0/31.3	23.9/48.3	35.9/56.9	29.1/51.0	44.8/65.8
Models pre-trained on more data and/or with larger size						
GLIP-L [38]	49.8/60.8	28.2/ -	34.3/ -	41.5/ -	37.3/ -	52.1/68.9
FIBER-B	**49.3/58.4**	**29.5/50.0**	**32.2/56.9**	**40.1/58.1**	**35.8/56.9**	**47.0/65.9**

Table 8: Zero-shot transfer and fine-tuning results for object detection on COCO, LVIS and the average over 13 datasets for object detection in the wild. Detailed scores on the 13 datasets are presented in the Appendix. FIBER achieves better AP across the board compared to similarly sized GLIP-B, trained on the same amount of fine-grained data. On rare objects in LVIS, FIBER outperforms GLIP-L trained on 25x more fine-grained data. Results without coarse-grained pre-training are provided in the Appendix.

rare objects, FIBER outperforms GLIP-L which is a bigger model and also trained with 25$\times$ more fine-grained data.

We also report zero-shot and fine-tuned results on a suite of 13 ODinW (object detection in the wild) datasets, spanning various domains and show consistent performance improvements over previous SoTA. Additionally, in Figure 5, we report few-shot results aggregated across these 13 datasets and show better data efficiency over GLIP-B trained with the same fine-grained data.

Ablation Study. In Appendix A.2 and A.3, we have provided detailed ablations that guided our architecture design, including ablations on fusion strategies, pre-training objectives, architecture for captioning, and additional results on open-ended VQA, and detailed few-shot ODinW results. Due to the space limit, these ablations and additional results are only provided in the Appendix. Some important observations are summarized below. (i) Co-attention works similarly to merged attention for fusion in the backbone. (ii) Adding a gating parameter in co-attention allows the addition of fusion in more layers, and also gives better performance than merged attention. (iii) Adding co-attention in the last 6 layers provides a balance between performance and efficiency. (iv) MLM, ITM with hard negative mining, and ITC are all important pre-training objectives for training FIBER-style models.

5 Conclusion

We propose (i) FIBER, a novel architecture and (ii) a coarse-to-fine pre-training pipeline. We perform extensive experiments and show consistent improvements over strong baselines across a diverse set of tasks. The results demonstrate the effectiveness of FIBER coupled with our pre-training strategy, by setting new SoTA scores while at the same time reducing the requirement of expensive box-level annotations. Future directions include scaling our models and extending our framework to other modalities.

The approach introduced in our work can potentially inherit undesirable societal biases that exist in our pre-training data. Careful debiasing and filtering of data should be undertaken before real-life deployment of our work. Additionally, pre-training can induce environmental costs, and minimizing these costs is an avenue that we plan to explore further.

Acknowledgement

We would like to thank Nguyen Bach, Jiayuan Huang, and Luis Vargas for their support. We also thank Wenhui Wang, Li Dong, Furu Wei, Bin Xiao, and Lu Yuan for their helpful discussions. We also thank Liunian Harold Li and Te-Lin Wu for their feedback on the manuscript. Aishwarya is supported in part by the National Science Foundation under NSF Award 1922658. Zi-Yi is supported in part by the DARPA Machine Common Sense (MCS) program under Cooperative Agreement N66001-19-2-4032 and NIH R01HL152270.

References

[1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. nocaps: Novel object captioning at scale. In *ICCV*, 2019. 3, 8

[2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint*, 2022. 4

[3] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: Semantic propositional image caption evaluation. In *ECCV*, 2016. 16

[4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *ICCV*, 2015. 1, 3, 7

[5] Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL Workshops*, 2005. 16

[6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1, 4

[7] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. In *ICLR*, 2022. 4

[8] 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 *ECCV*, 2020. 1, 2, 3, 4, 5, 6, 7, 9

[9] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *ICML*, 2021. 3, 4, 5, 19

[10] Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *EMNLP*, 2014. 6

[11] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR Workshops*, 2020. 16

[12] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *CVPR*, 2021. 3, 5

[13] Karan Desai and Justin Johnson. VirTex: Learning visual representations from textual annotations. In *CVPR*, 2021. 4

[14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 1

[15] 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 *ICLR*, 2021. 1, 3

[16] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, Zicheng Liu, and Michael Zeng. An empirical study of training end-to-end vision-and-language transformers. In *CVPR*, 2022. 2, 3, 4, 5, 6, 7, 8, 16, 17, 18, 20

[17] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lin Liang, Zhe Gan, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Injecting semantic concepts into end-to-end image captioning. In *CVPR*, 2022. 8

[18] 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 *NeurIPS*, 2020. 3, 7, 9

[19] Gregor Geigle, Jonas Pfeiffer, Nils Reimers, Ivan Vulić, and Iryna Gurevych. Retrieve fast, rerank smart: Cooperative and joint approaches for improved cross-modal retrieval. *Transactions of the Association for Computational Linguistics*, 10:503–521, 2022. 8

[20] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 3, 9

[21] Tanmay Gupta, Amita Kamath, Aniruddha Kembhavi, and Derek Hoiem. Towards general purpose vision systems. In *CVPR*, 2022. 4

[22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 10

[23] Ronghang Hu and Amanpreet Singh. UniT: Multimodal multitask learning with a unified transformer. In *ICCV*, 2021. 4

[24] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *CVPR*, 2022. 8

[25] Xiaowei Hu, Xi Yin, Kevin Lin, Lijuan Wang, Lei Zhang, Jianfeng Gao, and Zicheng Liu. VIVO: Surpassing human performance in novel object captioning with visual vocabulary pre-training. In *AAAI*, 2021. 3

[26] 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 *CVPR*, 2021. 3

[27] 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. 3, 4

[28] 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. In *ICML*, 2021. 1, 4

[29] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. MDETR-modulated detection for end-to-end multi-modal understanding. In *ICCV*, 2021. 1, 2, 4, 5, 6, 9, 10

[30] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 20

[31] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014. 9

[32] Wonjae Kim, Bokyung Son, and Ildoo Kim. ViLT: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 3, 4, 6, 7

[33] 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. *IJCV*, 2017. 6

[34] Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, 2020. 3

[35] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint*, 2022. 3, 4, 5, 6, 7, 8

[36] 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 *NeurIPS*, 2021. 2, 3, 4, 5, 6, 7, 8, 17

[37] 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. 1, 3, 9

[38] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10

[39] 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 *ACL*, 2021. 3, 7

[40] 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 *ECCV*, 2020. 2, 3, 4

[41] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *ICLR*, 2022. 4

[42] Yehao Li, Yingwei Pan, Ting Yao, Jingwen Chen, and Tao Mei. Scheduled sampling in vision-language pretraining with decoupled encoder-decoder network. In *AAAI*, 2021. 3

[43] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 5

[44] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *TPAMI*, 2020. 6

[45] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1, 3, 6, 7, 8, 9

[46] 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. 1, 5, 6, 20

[47] 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 *ICCV*, 2021. 5, 6

[48] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 7

[49] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 1, 3, 4, 9

[50] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. SLIP: Self-supervision meets language-image pre-training. In *CVPR*, 2022. 4

[51] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2Text: Describing images using 1 million captioned photographs. In *NeurIPS*, 2011. 6

[52] Richard Yuanzhe Pang and He He. Text generation by learning from demonstrations. In *ICLR*, 2021. 8, 16

[53] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. In *ACL*, 2002. 16

[54] 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 *ICCV*, 2015. 2, 3, 7

[55] 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 *ICML*, 2021. 1, 4, 9

[56] 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. *JMLR*, 2020. 1

[57] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *TPAMI*, 2016. 6

[58] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. 8

[59] Mert Bulent Sariyildiz, Julien Perez, and Diane Larlus. Learning visual representations with caption annotations. In *ECCV*, 2020. 4

[60] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017. 21

[61] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 2019. 6

[62] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 6

[63] 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? In *ICLR*, 2022. 4

[64] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. FLAVA: A foundational language and vision alignment model. In *CVPR*, 2022. 4

[65] 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 *ICLR*, 2019. 3

[66] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *ACL*, 2019. 1, 3, 7

[67] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NeurIPS*, 2014. 6

[68] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. 1, 3, 4, 5

[69] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2

[70] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. In *CVPR*, 2015. 16

[71] 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 *ICLR*, 2019. 20

[72] Jianfeng Wang, Xiaowei Hu, Zhe Gan, Zhengyuan Yang, Xiyang Dai, Zicheng Liu, Yumao Lu, and Lijuan Wang. UFO: A unified transformer for vision-language representation learning. *arXiv preprint*, 2021. 3, 6, 8

[73] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *arXiv preprint*, 2022. 4, 5, 9

[74] Wenhui Wang, Hangbo Bao, Li Dong, and Furu Wei. VLMo: Unified vision-language pre-training with mixture-of-modality-experts. *arXiv preprint*, 2021. 3, 4, 5, 6, 7, 17

[75] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. SimVLM: Simple visual language model pretraining with weak supervision. In *ICLR*, 2022. 2, 3, 4, 7, 8, 19, 20

[76] Hongwei Xue, Yupan Huang, Bei Liu, Houwen Peng, Jianlong Fu, Houqiang Li, and Jiebo Luo. Probing inter-modality: Visual parsing with self-attention for vision-language pre-training. In *NeurIPS*, 2021. 3

[77] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *CVPR*, 2022. 4

[78] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Faisal Ahmed, Zicheng Liu, Yumao Lu, and Lijuan Wang. Crossing the format boundary of text and boxes: Towards unified vision-language modeling. *arXiv preprint*, 2021. 4, 9

[79] Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florencio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. TAP: Text-aware pre-training for text-vqa and text-caption. In *CVPR*, 2021. 3

[80] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. FILIP: Fine-grained interactive language-image pre-training. In *ICLR*, 2022. 4

[81] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. ERNIE-ViL: Knowledge enhanced vision-language representations through scene graph. In *AAAI*, 2021. 3

[82] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. CoCa: Contrastive captioners are image-text foundation models. *arXiv preprint*, 2022. 4

[83] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MAttNet: Modular attention network for referring expression comprehension. In *CVPR*, 2018. 3

[84] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. Modeling context in referring expressions. In *ECCV*, 2016. 9

[85] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. In *ICML*, 2022. 3, 4, 7, 8

[86] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. LiT: Zero-shot transfer with locked-image text tuning. In *CVPR*, 2022. 9

[87] 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 *CVPR*, 2021. 3, 4, 8

[88] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z. Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 2020. 6

[89] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *AAAI*, 2020. 3

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 5
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Appendix

A.1 Implementation Details

Vision-Language Classification. For the VL classification tasks, we follow METER [16] to set the hyper-parameters. Specifically, we fine-tune our models with the peak learning rates of 2e-5 for the backbones, 1e-4 for the cross-modal parameters, and 1e-3 for the head layer for 10 epochs. The batch size is set to 512. The image resolutions are set to 576 for VQAv2 and 384 for NLVR^2. We evaluate models with the VQA scores for VQAv2 and accuracy for NLVR^2. RandAugment [11] is used during the downstream fine-tuning stage.

Image-Text Retrieval. For image-text retrieval, we remove the cross-attention layers in the backbones and use the dual encoder architecture. We set the peak learning rates to 2e-5 for the backbones and 1e-4 for the head layer. The batch size is set to 1024. The image resolutions are set to 576 for both COCO and Flickr30k. We evaluate on the Recall@1,5,10 metrics for both text and image retrieval.

Image Captioning. For image captioning, we only keep the image-to-text attentions and feed the image representations in the last layer of the image encoder to the cross-attention modules. In this way, the model is turned into a standard seq2seq model, and we use the causal mask in the decoding side and predict outputs auto-regressively. We first train our models with the cross-entropy loss for 5 epochs with the peak learning rates of 5e-5 for the backbones, and 2.5e-4 for the rest of the parameters. Then, we fine-tune it with GOLD [52] for 5 epochs as it is efficient and has proven to be effective when the model input can correspond to different outputs. We set the peak learning rate to 1e-5 for the backbones during GOLD training. For CIDEr optimization, the learning rate is further reduced to 1e-6 and we train the models for 3 epochs. The batch size is set to 512. We use a beam size of 5 during inference and do not use constrained beam search. We use the same model when testing on COCO and NoCaps, and we evaluate on BLEU [53], METEOR [5], CIDEr [70], and SPICE [3] metrics.

Phrase Grounding. For phrase grounding on Flickr30k, we do not further fine-tune the model after fine-grained pre-training, and just directly evaluate on the Recall@ 1,5,10 metrics.

Referring Expression Comprehension (REC). For the REC datasets, we use batch size 16 and fine-tune on the respective dataset for 20 epochs. We use a warmup of 2000 steps, with a peak learning rate of 1e-5 for both the OD head as well as the rest of the model's paramaters, with two learning rate drops at 67% and 89% of the total number of steps. We switch off the horizontal flip augmentation during REC training, as we find that it adversely affects the performance, especially on the RefCOCO dataset, which includes many examples having degenerate language such as just "left" or "right" rather than using descriptive words for the referring expressions.

Object Detection. For both COCO and LVIS detection, we train for 24 epochs, with batch size 32, with a learning rate of 1e-5 for the whole model, with two learning rate drops at 67% and 89% of the total number of steps. For the ODinW datasets, we fine-tune for 12 epochs, with early stopping based on the validation accuracy.

The object detection data is constructed as follows - The object category names are directly used in their text form separated by full stops as input to the text encoder. We follow the same protocol as in GLIP [3] to be comparable to their experiments. More specifically, the input text will look like this: "person. bicycle. car. .... toothbrush", and the model will learn how to ground image regions into these object names. An example of input and output predicted by the model can be seen in Fig. 9.

A.2 Ablation Study

Ablation Study on the Fusion Strategies. We perform ablation studies on our fusion module. We investigate three different fusion strategies as shown in Figure 6. Merged attention concatenates representations from the two input modalities and feeds them into the self-attention layer for fusion. Note that here the key and value matrices for the two modalities are different. On the other hand, co-attention inserts a cross-attention layer into each of the encoding layer. The insertion of the

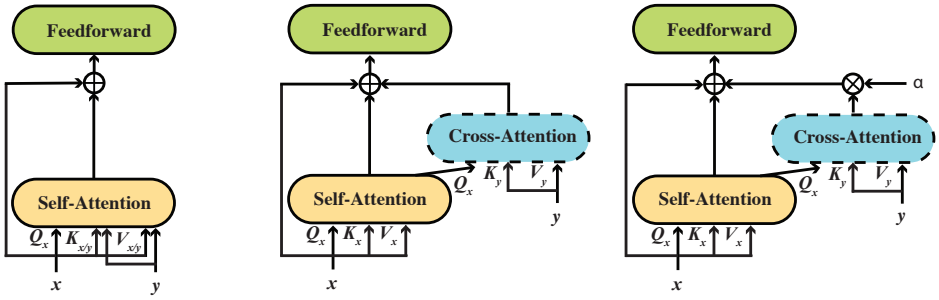

(a) merged attention **(b)** co-attention w/o α **(c)** co-attention w/ α

Figure 6: Different strategies for fusion in the backbone. ($\boldsymbol{x}$, $\boldsymbol{y}$) are the (image, text) or (text, image) representations, and α is a learnable scalar.

Model	#Fusion Layers	#Fusion Params.	VQAv2
No Fusion	0	0M	65.65
Merged Attention	3	7.9M	71.24
	6	12.6M	70.67
Co-attention w/o α	3	16.1M	70.84
	6	26.0M	68.13
Co-attention w/ α	3	16.1M	71.20
	6	26.0M	71.97
	9	35.8M	72.10
	12	45.6M	72.08

Table 9: Ablation study on the fusion strategies. Results are obtained by directly fine-tuning models initialized with uni-modally pre-trained parameters and without VLP. Results on VQAv2 are on test-dev set.

Pre-training Objectives				VQAv2	Flickr30k	
MLM	ITM	ITM-hard	ITC	test-dev	IR@1	TR@1
✓	✓	×	×	72.47	65.50	79.30
✓	×	×	✓	74.16	73.74	87.70
×	×	✓	✓	67.45	75.20	87.00
✓	✓	×	✓	74.49	73.58	87.80
✓	×	✓	✓	75.98	75.26	87.50

Table 10: Ablation study on the pre-training objectives and whether the hard negative mining strategy is necessary in the coarse-grained pre-training stage.

cross-attention layer offers the flexibility of controlling to what extent we want the two modalities to fuse together as we can easily introduce an α term into the module as in Figure 4.

As shown in Table 9, we compare the three fusion strategies by directly fine-tuning our models without performing VLP for efficiency. We use Swin Transformer and RoBERTa as our vision and text backbones and load their pre-trained parameters for initialization. We set the image resolution to 224×224. We can see that merged attention and co-attention achieve comparable performance without α. For both strategies, increasing the number of fusion layers can lead to performance drop. However, after introducing α, we can see significant improvements of co-attention, indicating the importance of having an explicit controlling/gating mechanism for fusion in the backbone.

After the α term is introduced, we can increase the number of fusion layers and achieve robust performance. Based on the ablation results, we choose to fuse the top 6 layers of the backbones as it can achieve a good accuracy-efficiency trade-off.

Ablation Study on Pre-training Objectives. Following previous work [36, 16, 74], we pre-train our models with image conditioned masked language modeling, image-text matching with hard negative mining, and image-text contrastive losses during the coarse-grained pre-training stage. In

Model	OD on COCO	OD on LVIS	ODinW	RefCOCO+
	Zero-shot/Fine-tune	Zero-shot/Fine-tune	Zero-shot	Val/TestA/TestB
OFA-L	-	-	-	84.49/90.10/ 77.77
GLIP-B	48.1/57.0	29.1/51.0	44.8	-
FIBER-B w/o C.G. VLP	48.9/57.8	31.6/55.8	45.1	85.04/88.82/78.59
FIBER-B	49.3/58.4	35.8/56.9	47.0	85.74/90.13/79.38

Table 11: Ablation study on our proposed two-stage pre-training strategy.

Vision Encoder	Text Encoder	VQAv2
Swin	RoBERTa	71.97
Swin	BERT	71.86
CLIP-ViT	RoBERTa	71.37

Table 12: Results of different vision and text backbones for FIBER without VLP.

this part, we ablate each of the pre-training objectives and evaluate our models on both VQAv2 and Flickr30k retrieval tasks. Specifically, we use Swin Transformer and RoBERTa as our vision and text backbones and load their pre-trained parameters for initialization. The image resolution is set to 224×224 and we pre-train models for 100k steps with 1,024 batch size. We use AdamW with the peak learning rates of 1e-4 for the backbones and 5e-4 for the cross-modal parameters. We use linear warmup over the first 1k steps and linear decay.

As shown in Table 10, we can see that removing any of the pre-training objectives can lead to performance drop, and hard negative mining can bring improvements on both VQA and retrieval tasks. Masked language modeling is most effective for VQA, while removing it will not hurt the retrieval performance. This set of experiments demonstrates that all of the objectives are necessary for our models to obtain good performance.

Ablation Study on the Two-Stage Pre-training. In this paper, we propose a coarse-to-fine pre-training strategy for handling VL tasks of different kinds. In this paragraph, we remove the coarse-grained pre-training stage and only pre-train the models with image-text-box data and see how it performs. As shown in Table 11, we see gains across both tasks when utilizing the coarse-grained pre-training. Similar to the case of Flickr30k, on RefCOCO+ the coarse-grained pre-training helps FIBER to get better performance than large-sized model trained with more data. In addition, note that without the coarse-grained pre-training, the only difference between FIBER and GLIP is the architectural difference, and the fact that FIBER can still outperform GLIP in this setting demonstrates the effectiveness of our proposed architecture.

Ablation Study on Different Backbones. While previous work [16] has compared different vision and text backbones for VLP models, we investigate if their conclusions stil apply in our settings. Specifically, we try BERT and RoBERTa for our text encoder and CLIP-ViT and Swin Transformer for our image encoder. As shown in Table 12, we can see that RoBERTa and Swin Transformer perform slightly better than BERT and CLIP-ViT before VLP, which is consistent with previous findings in METER [16]. Note that while CLIP-ViT has the potential to perform better than Swin Transformer after VLP, it is hard to be adapted for region-level tasks such as object detection. Therefore, pairing Swin Transformer with RoBERTa is the optimal configuration in our settings.

A.3 Additional Results

Additional Results on Image Captioning. For image captioning, we evaluate the models with BLEU-4, METEOR, ROUGE-L, CIDEr, and SPICE metrics on COCO and NoCaps. On NoCaps, we have fine-grained evaluation results on different domains, including in-domain, near-domain, out-domain, and entire domain settings. In this part, we provide the complete evaluation results in Table 13, 14 and 15. We can see that both GOLD and CIDEr optimization can improve the model performance across metrics. We also see a noticeable performance drop when evaluating our models on out-of-domain data, but complementary methods such as constrained beam search can be used to

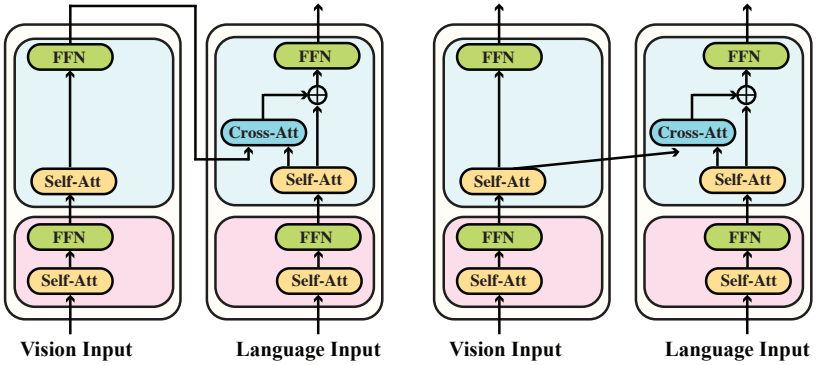

(a) seq2seq **(b)** ladder

Figure 7: When adapting FIBER to image captioning, we can either use the seq2seq structure or the ladder architecture as in pre-training.

Model	COCO				
	BLEU@4	METEOR	ROUGE-L	CIDEr	SPICE
Models trained without CIDEr optimization					
FIBER-Ladder-B	38.6	30.1	58.8	127.5	22.8
FIBER-B	39.1	30.4	59.3	128.4	23.1
FIBER-GOLD-B	40.3	30.7	60.0	133.6	23.6
Models trained with CIDEr optimization					
FIBER-B	42.8	31.0	61.5	142.8	24.3
FIBER-GOLD-B	43.4	31.3	61.8	144.4	24.6

Table 13: The complete set of results on COCO image captioning, with another model variant FIBER-Ladder. See Figure 7 for details.

alleviate the issue. Also, training our models with more captioning data should also be helpful in these settings.

Also, in the main paper, we adapt FIBER for image captioning by turning it into a standard seq2seq model as in Figure 7a, where the output of the final encoding layer will be fed into the image-to-text cross-attention modules. Another possible design is to keep the ladder structure as we used in pre-training (Figure 7b), so there can be less mismatching between pre-training and fine-tuning. As shown in Table 13, the two architectures can achieve comparable performance. Considering that the seq2seq architecture is more widely adopted in the current literature, we decide to use the seq2seq architecture for image captioning.

Open-ended VQA. In most existing literature, VQA is treated as a classification task, where a vocabulary of some most frequent answers are constructed and VL models predict which answer corresponds to the given question based the constructed vocabulary. However, question answering is inherently open-ended. Since we can turn our models into a generative model by fine-tuning on image captioning, we also investigate if our models can perform open-ended VQA in this part.

Following [9], we break down VQA questions into in-domain and out-of-domain questions,

Model	Open-ended VQA		
	In-D	Out-D	overall
VL-T5 [9]	71.4	13.1	67.9
VL-BART [9]	72.1	13.2	68.6
SimVLM-B [75]	78.3	25.8	75.2
FIBER-B	75.9	14.7	71.6

Table 16: Results on open-ended VQA. We follow [9] to split the data into in domain (In-D) and out of domain (out-D).

Model	in-domain					near-domain					out-domain					entire				
	B@4	M	R	C	S	B@4	M	R	C	S	B@4	M	R	C	S	B@4	M	R	C	S
Models trained without CIDEr optimization																				
FIBER-B	29.7	30.0	58.2	98.5	13.9	24.4	27.5	55.6	88.2	13.0	18.0	25.4	53.5	82.8	12.2	23.9	27.5	55.6	88.6	13.0
FIBER-GOLD-B	29.7	30.1	58.2	100.6	14.0	26.8	28.2	57.0	92.9	13.5	18.3	25.8	54.3	86.6	12.8	25.5	28.0	56.6	92.8	13.4
Models trained with CIDEr optimization																				
FIBER-B	34.2	30.9	60.0	108.9	14.0	28.8	28.4	58.2	96.0	13.5	19.8	26.0	55.6	90.1	12.7	27.7	28.3	57.9	96.7	13.4
FIBER-GOLD-B	35.4	31.2	60.6	110.3	14.3	30.5	29.0	58.9	99.5	13.8	20.4	26.0	55.6	90.2	12.8	29.1	28.7	58.5	99.2	13.7

Table 14: The complete set of results on the NoCaps validation set. B@4: BLEU@4, M: METEOR, R: ROUGE-L, C: CIDEr, S: SPICE.

Model	in-domain					near-domain					out-domain					entire				
	B@4	M	R	C	S	B@4	M	R	C	S	B@4	M	R	C	S	B@4	M	R	C	S
Models trained without CIDEr optimization																				
FIBER-B	28.6	29.5	57.7	92.8	13.6	25.6	28.0	56.1	87.3	13.0	16.2	24.4	52.1	76.4	11.6	24.3	27.6	55.6	86.0	12.9
FIBER-GOLD-B	29.9	30.1	58.4	95.9	14.1	28.0	28.7	57.4	92.0	13.5	18.4	25.3	53.4	81.0	12.3	26.5	28.3	56.8	90.6	13.4
Models trained with CIDEr optimization																				
FIBER-B	33.3	30.4	59.9	102.7	14.1	29.8	28.9	58.6	95.3	13.6	20.7	25.5	55.1	83.4	12.3	28.6	28.5	58.2	94.1	13.4
FIBER-GOLD-B	34.6	30.9	60.6	104.7	14.4	31.3	29.4	59.4	98.7	13.9	21.2	25.9	55.4	85.7	12.7	29.9	29.0	58.8	97.1	13.8

Table 15: The complete set of results on the NoCaps test set. B@4: BLEU@4, M: METEOR, R: ROUGE-L, C: CIDEr, S: SPICE.

Text Encoder	QQP	MNLI	QNLI	SST2	CoLA	MRPC	STSB	RTE
RoBERTa-B [46]	91.31	87.53	92.61	94.38	58.72	91.03	90.15	71.24
METER-RoBERTa-B [16]	91.34	87.38	92.67	93.92	57.88	90.57	89.93	70.28
SimVLM-B [75]	90.4	83.4	88.6	90.9	46.7	84.4	-	63.9
FIBER-RoBERTa-B	91.60	86.23	91.34	92.66	59.56	90.72	89.77	62.09

Table 17: Performance of text encoders on the GLUE dev sets.

where the answers to the out-of-domain questions do not appear in the top-k ($k = 3, 129$) candidates. We use the Karpathy split [30] in this setting.

As shown in Table 16, our generative model can perform better than VL-T5 and VL-BART, while lagging behind SimVLM especially in out-of-domain settings, possibly because SimVLM is trained with over a billion image-caption pairs and is more robust in this setting. The results indicate that our model can be turned into a general open-ended VQA model as well.

Uni-modal Performance. It can be interesting to see whether our backbones can still perform uni-modal tasks after VLP. Therefore, in this part, we also evaluate our language backbones on uni-modal tasks.

Image Encoder	ImageNet	ADE20k
Swin-B	86.3	51.6
FIBER-Swin-B	86.0	52.0

Table 18: Performance of image encoders on image classification and semantic segmentation.

Specifically, we test our language backbone after the first coarse-grained pre-training stage on the GLUE [71] benchmark. As shown in Table 17, the uni-modal performance of our text encoder can drop marginally on some tasks, possibly because the model only encounters simple text captions during VLP. However, it is still better than SimVLM which is trained with 800GB of web-crawled documents from scratch.

On the other hand, as shown in Table 18, our image encoder can achieve comparable and sometimes even better performance on vision-only tasks including image classification and semantic segmenta-

Model	Shot	PascalVOC	AerialDrone	Aquarium	Rabbits	EgoHands	Mushrooms	Packages	Raccoon	Shellfish	Vehicles	Pistols	Pothole	Thermal	Avg
GLIP-B	1	51.7	25.0	34.0	69.2	67.4	63.4.	54.4	56.5	14.1	57.9	51.7	15.6	70.2	48.5
GLIP-B	3	57.4	25.2	44.9	65.1	69.3	88.3	67.3	52.8	28.9	60.7	62.1	31.3	67.9	55.5
GLIP-B	5	57.9	28.1	44.1	64.6	68.1	85.1	74.2	60.8	24.0	61.9	59.1	33.7	75.2	56.7
GLIP-B	10	60.6	27.8	49.7	67.8	65.2	87.4	67.5	54.5	42.3	65.1	63.9	39.2	78.7	59.2
GLIP-B	All	62.2	36.0	55.3	74.0	79.8	88.1	74.3	64.1	47.0	64.4	72.4	56.5	81.1	65.8
FIBER-B	1	55.7	25.0	37.9	69.8	67.2	83.0	73.2	54.7	29.7	58.0	44.8	27.8	67.0	53.4
FIBER-B	3	59.8	28.2	42.9	71.5	68.4	88.1	65.6	64.6	38.6	61.8	47.8	37.5	68.5	57.2
FIBER-B	5	61.6	30.9	49.5	72.3	69.2	87.5	73.2	57.4	38.2	62.7	55.3	40.3	61.8	58.4
FIBER-B	10	60.7	31.2	52.0	68.8	70.4	88.1	69.3	53.8	41.7	66.6	62.2	46.8	74.8	60.5
FIBER-B	All	68.7	35.4	58.3	75.9	79.4	88.1	72.2	52.9	45.2	66.1	72.2	60.9	80.6	65.9

Table 19: Few-shot and full fine-tuned results on the various ODinW datasets.

tion after coarse-grained pre-training. The results suggest that the image encoder can remain powerful even if the fusion modules are removed.

Using the Model Checkpoint After Fine-grained Pre-training for VQA. We also test what if we fine-tune the fine-grained pre-trained checkpoint on VQA in this part. We find that after the second-stage fine-grained pre-training, the model performance on the VQAv2 test-dev set can drop from 78.55 to 74.3, indicating that the two-stage pre-training paradigm is indeed necessary for different VL tasks, as tasks of different characteristics can require checkpoints from different pre-training stages.

Detailed Results on ODinW. Detailed results on the 13 ODinW datasets are provided in Table 19.

A.4 Visualization after coarse-grained pre-training

We also provide a qualitative analysis of our model. As shown in Figure 8, we use Grad-CAM [60] to visualize the cross-attention maps of our coarse-grained pre-trained checkpoint. We find that the model can correctly align concepts and image regions for some examples, suggesting that the model can learn visual grounding implicitly.

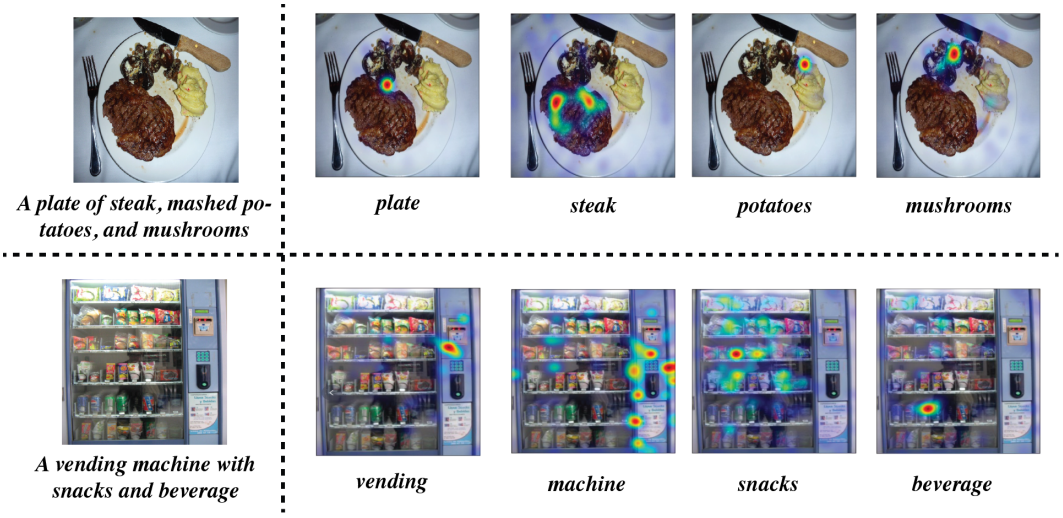

Figure 8: Visualizations of the cross-attention maps obtained by Grad-CAM [60]. Given each of the tokens in a caption, the model can attend to its corresponding regions. The figures are from the NoCaps validation set (ID: 253, 3766).

A.5 Visualization after fine-grained pre-training

We also provide visualization after fine-grained pre-training in Figure 9, 10, 11, and 12.

A.6 Text and Vision Backbones of Different Models

In this section, we list the backbones of different models in Table 20.

steak. mushrooms. plate. knife. spoon. fork. mashed potatoes.

vending machine. drinks. snacks.

Figure 9: The same images probed after fine-grained pre-training.

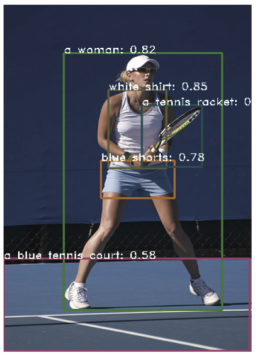

A woman in blue shorts and white shirt holds a tennis racket on a blue tennis court.

A woman wearing a long black sweater is standing near a pink bike on the sidewalk.

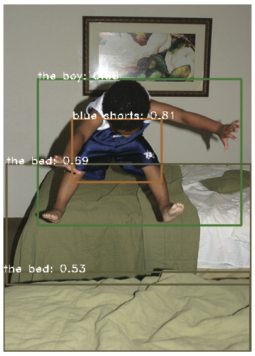

The boy in blue shorts is bouncing on the bed.

Figure 10: Some examples of phrase grounding from the validation set for Flickr30k entities.

Model	Vision Encoder	Text Encoder
UNITER	Frozen ResNet-based OD	XFM (init w/ BERT)
VILLA	Frozen ResNet-based OD	XFM (init w/ BERT)
UNIMO	Frozen ResNet-based OD	XFM (init w/ RoBERTa)
VinVL	Frozen ResNet-based OD	XFM (init w/ BERT)
ViLT	XFM (init 5w/ ImageNet-ViT)	XFM (init w/ ImageNet-ViT and BERT embeddings)
ALBEF	XFM (init w/ ImageNet-ViT)	XFM (init w/ BERT)
VLMo	XFM (init w/ ImageNet-ViT)	XFM (init w/ ImageNet-Language-ViT)
UFO	XFM (init w/ ImageNet-ViT)	XFM (init w/ ImageNet-ViT)
ViTCAP	XFM (init w/ ImageNet-ViT)	XFM (init w/ ImageNet-ViT and BERT embeddings)
METER-Swin	Swin (init w/ ImageNet-Swin)	XFM (init w/ RoBERTa)
METER-CLIP	XFM (init w/ CLIP-ViT)	XFM (init w/ RoBERTa)
MDETR	EfficientNet	XFM (init w/ RoBERTa)
GLIP	Swin (init w/ ImageNet-Swin)	XFM (init w/ BERT)
FIBER	Swin (init w/ ImageNet-Swin)	XFM (init w/ RoBERTa)

Table 20: Backbones of different models. XFM stands for Transformer

chair green

middle person

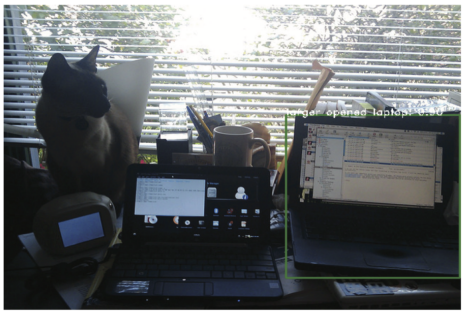

larger opened laptop

yellow and blue

Figure 11: Some examples of referring expression comprehension from the validation set of RefCOCO+.

a signboard that displays the speed limit.

signboard.

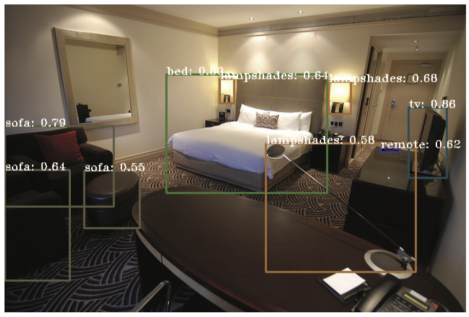

bed. lampshades. sofa. tv. fridge. remote.

a photo frame.

Figure 12: Some images with prompts for various items in the scene.