

EVE: Efficient Vision-Language Pre-training with Masked Prediction and Modality-Aware MoE

Junyi Chen^{1*}, Longteng Guo², Jia Sun³, Shuai Shao³, Zehuan Yuan³, Liang Lin¹, Dongyu Zhang^{1†}

¹Sun Yat-sen University

²Institute of Automation, Chinese Academy of Sciences (CASIA)

³Bytedance Inc

chenjy765@mail2.sysu.edu.cn, longteng.guo@nlpr.ia.ac.cn, {sunjia.ly, shaoshuai.0516, yuanzehuan}@bytedance.com
linliang@ieee.org, zhangdy27@mail.sysu.edu.cn

Abstract

Building scalable vision-language models to learn from diverse, multimodal data remains an open challenge. In this paper, we introduce an Efficient Vision-languageE foundation model, namely EVE, which is one unified multimodal Transformer pre-trained solely by one unified pre-training task. Specifically, EVE encodes both vision and language within a shared Transformer network integrated with modality-aware sparse Mixture-of-Experts (MoE) modules, which capture modality-specific information by selectively switching to different experts. To unify pre-training tasks of vision and language, EVE performs masked signal modeling on image-text pairs to reconstruct masked signals, i.e., image pixels and text tokens, given visible signals. This simple yet effective pre-training objective accelerates training by 3.5x compared to the model pre-trained with Image-Text Contrastive and Image-Text Matching losses. Owing to the combination of the unified architecture and pre-training task, EVE is easy to scale up, enabling better downstream performance with fewer resources and faster training speed. Despite its simplicity, EVE achieves state-of-the-art performance on various vision-language downstream tasks, including visual question answering, visual reasoning, and image-text retrieval.

Introduction

Vision-Language Pre-training aims to learn a general multimodal representation that can be transferred to various vision-language downstream tasks, such as vision-language understanding and image-text retrieval. A vision-language foundation model should have excellent performance while being easy to train and scale up, which can be achieved through the model architecture and the pre-training tasks.

The model architectures of recent methods can be roughly divided into two categories: dual-encoder architecture and unified architecture. Dual-encoder methods (Radford et al. 2021; Zeng, Zhang, and Li 2022) employ modality-specific models (e.g. BERT (Devlin et al. 2019), ViT (Dosovitskiy et al. 2021)) to encode different modalities separately and a fusion module to integrate them. As for the fusion module, some methods (Radford et al. 2021) employ shallow




Figure 1: Performance of different models on VQA test-dev under different training hours. Training hours of all models are reproduced by us on A100 GPUs.

fusion (e.g., dot product) for the interaction of vision and language. Some alternative methods (Zeng, Zhang, and Li 2022) use deep neural networks, such as Transformer Encoders, to perform deep fusion on modality interaction, but lead to difficulties in scaling up and low efficiency. Unified methods (Kim, Son, and Kim 2021; Wang et al. 2022b) use a modality-shared Transformer to encode different modalities jointly. This approach simplifies the framework and improves the speed, helping with model scaling up. However, they overlook the inherent gap between modalities, leading to lower overall performance. Image is continuous, redundant, and low-level on the raw signals, while text is discrete, refined, and high-level. Directly using a shared Transformer to encode different modalities with semantic gap poses problems. Therefore, it is necessary to consider the differences between different modalities carefully.

Previous methods also have explored numerous pre-training tasks for vision-language pre-training, including Image-Text Contrastive Learning (Radford et al. 2021), Image-Text Matching (Li et al. 2021), Word-Patch Alignment (Chen et al. 2020), Masked Language Modeling (Su et al. 2020), Masked Image Modeling (Bao et al. 2022b), and so on. They have been widely used to improve vision-language pre-training. While incorporating more pre-

*This work was done while Junyi was an intern at ByteDance.

†Corresponding author.

training tasks can enhance performance, adding too many tasks can also lead to some problems. Foremost, it significantly prolongs the pre-training time and increases the computational resources required. Additionally, it necessitates manual weight adjustments for different objectives. Furthermore, excessive pre-training objectives can result in a reduction in the model’s scalability, which is crucial in designing pre-training models, as the recent success has shown in large language models (Ouyang et al. 2022; Wei et al. 2022b; Zhao et al. 2023b). Therefore, it is necessary to use effective and scalable pre-training tasks.

In this paper, we propose an Efficient Vision-language foundation model (EVE) with a unified modality-aware Transformer pre-trained with a single unified pretraining task, i.e., masked signal modeling.

In terms of model architecture, we use a unified modality-aware Transformer and revisit the integration of Mixture-of-Experts in vision-language pre-training. We employ a shared Multi-Head Self-Attention module and a Modality-Aware MoE module for the modality-aware Transformer to encode and fuse various modalities jointly. Using a unified shared Transformer is more concise and flexible, which simplifies the extension to additional modalities and facilitates cross-modal alignment. By incorporating MoE, we can take into account the differences between modalities and capture more modality-specific information. We also introduce a modality routing technique in MoE that enables the router select more appropriate experts for processing.

In terms of pre-training tasks, we propose a unified masked signal modeling technique combining pixel and language modeling, which significantly improves training speed and reduces scaling difficulty. Some methods (Wang et al. 2023; Kwon et al. 2023; Zhao et al. 2023a) have applied generative pre-training paradigm to vision-language pre-training. While they either add the generative objective with other complex objectives like ITC and ITM (Kwon et al. 2023) or employ more complicated targets such as visual tokens (Wang et al. 2023) or momentum features (Zhao et al. 2023a), which require a nontrivial visual tokenizer or momentum model. All of these increase the complexity of pre-training. In contrast to them, we just utilize the *raw signals* from the image-text pairs themselves to minimize the complexity of pre-training and achieve better scalability. Pre-training speed is 3.5x faster than incorporating ITC and ITM.

EVE can greatly enhance pre-training speed, as shown in Figure 1. It decreases the demand for extensive computational resources while being easy to scale up. We demonstrate the effectiveness of EVE on various vision-language downstream tasks, including visual question answering, visual reasoning, and image-text retrieval. EVE achieves state-of-the-art performance on Image-Text Retrieval and Vision-Language Understanding (VQA and NLVR2) tasks.

Our contributions are summarized as follows:

- We introduce EVE, an efficient vision-language foundation model that achieves state-of-the-art performance while improving training speed, with one unified multimodal Transformer and one unified pre-training task.

- We integrate Modality-Aware MoE with a shared multimodal Transformer to achieve a more profound fusion of different modalities and capture more modality-specific information simultaneously, resulting in better performance and faster inference speed within a unified architecture.
- We propose a unified masked signal modeling technique, simplifying vision-language pre-training into a single unified objective, resulting in significantly improved pre-training speed and competitive performance.

Related Work

Model architecture and pre-training tasks are crucial factors in the representation learning of vision-language.

Model Architecture Dual-encoder with a fusion module (Li et al. 2021; Liu et al. 2021; Dou et al. 2022b; Zhao et al. 2023a) performs well on vision-language tasks but with higher time and architecture complexity. Unified architecture methods (Kim, Son, and Kim 2021; Wang et al. 2022b; Bao et al. 2022a,b) can flexibly encode different modalities as a fusion encoder or process a single modality as a unimodal encoder, demonstrating faster inference speed and promising performance. Some of them (Kim, Son, and Kim 2021; Wang et al. 2022b) use a shared standard Transformer (Vaswani et al. 2017) to jointly encode different modalities, while they ignore the modality gap and lead to worse performance. Others (Bao et al. 2022a,b) use MoME Transformer instead and prove that shared attention is better for multimodal learning. However, MoME Transformer uses modality-shared FFN in the deep layers may neglect some modality-specific information.

Considering the simplicity, effectiveness, and flexibility of the unified architecture, we adopt a unified architecture with Modality-Aware MoE to better capture modality specifics during fusion for multimodal representation learning. We achieve state-of-the-art performance with approximately the same inference cost.

Masked Signal Modeling Recently, several methods (Bao et al. 2022b; Zhao et al. 2023a; He et al. 2022b; Diao et al. 2023; Geng et al. 2022) explore the "mask then predict" paradigm in the vision for vision-language pre-training. While VLBEiT (Bao et al. 2022b) introduces training on the visual modality through masked image modeling, their reconstruction target is the visual token, which may significantly influence performance depending on the visual tokenizer employed. DAVINCI (Diao et al. 2023) extends prefix language modeling further to vision, but it also uses the discrete visual token as the target. MAMO (Zhao et al. 2023a) enriches multimodal representation by using momentum features in masked representation modeling, which relies heavily on a momentum teacher model to avoid divergence. Some methods (Kwon et al. 2023; He et al. 2022b; Gui et al. 2023) use masked pixel modeling, but they all require additional costly pre-training tasks such as ITC (Radford et al. 2021) and ITM (Li et al. 2019). Among these methods, VLMAE (He et al. 2022b) only applies masked pixel modeling to the image encoder. M3AE (Geng et al. 2022) leverages a

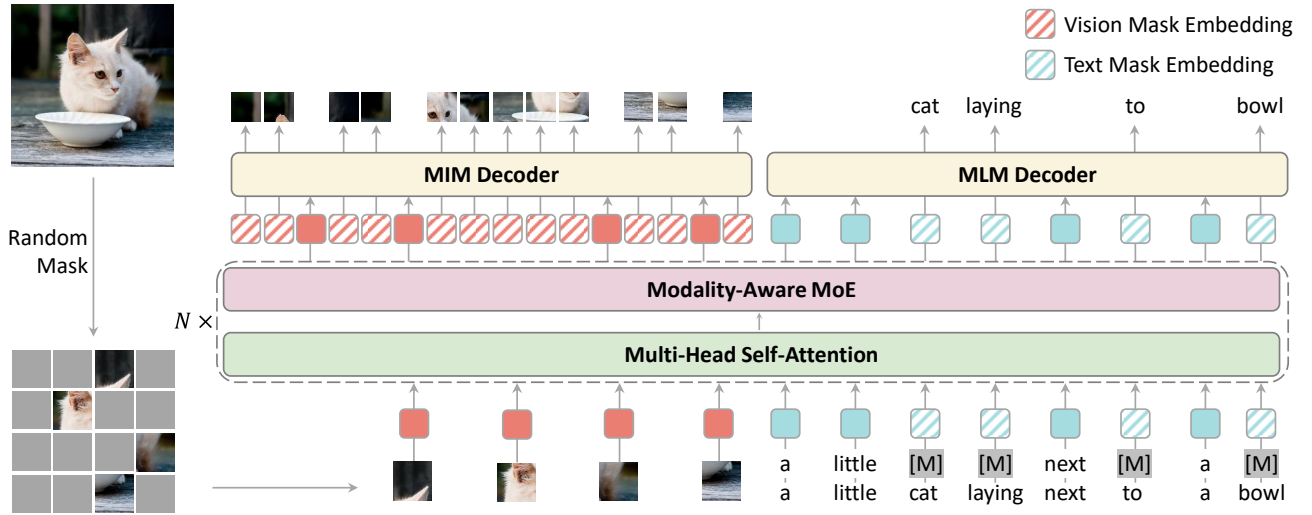


Figure 2: Overview of EVE and Masked Signal Modeling. We use a unified architecture with shared attention and Modality-Aware MoE for EVE and a single unified masked signal modeling for pre-training. We employ random masking on both image and text. Masked image and complete text are used in masked image modeling, vice versa.

unified Image-Language masking approach to mask and reconstruct both images and text simultaneously, but it is not used in multimodal downstream tasks.

We unify masked pixel and language modeling into masked signal modeling, reconstructing masked raw signals from visible signals. This simplifies and accelerates training, achieving better performance and scalability.

Mixture-of-Experts (MoE) Mixture-of-Experts has been extensively explored in computer vision (Riquelme et al. 2021) and natural language processing (Shazeer et al. 2017; Lepikhin et al. 2021). These methods generally aim to improve performance by learning a better routing using auxiliary losses (Lepikhin et al. 2021; Zoph et al. 2022), converting it into a linear assignment problem (Lewis et al. 2021), or making it differentiable (Hazimeh et al. 2021). MoE seems well-suited for multimodal learning, but the differences between modalities present some challenges. LIMoE (Mustafa et al. 2022) involves more auxiliary losses to balance different modalities, uni-perceiver-moe (Zhu et al. 2022) employs conditional MoE, VLMO (Bao et al. 2022a) and VL-MoE (Shen et al. 2023) use shared expert in the deep layers.

However, existing methods increase complexity or limit performance due to manual routing and ignoring modality information. Therefore, we propose Modality-Aware MoE as a simple way to apply MoE to multimodal learning. We simplify the auxiliary loss and capture more modality specifics by expert switching.

Methods

Backbone Network

As shown in Figure 2, we adopt a unified multimodal Transformer with shared attention and Modality-Aware Mixture-of-Experts as the backbone network, which is capable of encoding different modalities. After pre-training, the model

can be utilized as either a fusion encoder or a unimodal encoder for various downstream tasks through fine-tuning.

For Image I , following ViT (Dosovitskiy et al. 2021), we first split the Image I into N patches with a patch size of P . The resulting $N = HW/P^2$ patches are projected into a shared embedding space using a linear projector. A special token I_{cls} is added at the beginning of all visual tokens. We employ learnable visual position embeddings I_{pos} and visual type embeddings I_{type} on visual tokens. Image embedding can be summarized as follows.

$$I_{emb} = [I_{cls}, I_1, \dots, I_N] + I_{pos} + I_{type} \quad (1)$$

For Text T , following BERT (Devlin et al. 2019), we tokenize the text into discrete tokens with the maximum length of n and project them into the joint embedding space. We add a special token T_{cls} at the beginning of all text tokens and use learnable text position embeddings T_{pos} and text type embeddings T_{type} for text encoding. Text embedding can be summarized as follows.

$$T_{emb} = [T_{cls}, T_1, \dots, T_n] + T_{pos} + T_{type} \quad (2)$$

We concatenate I_{emb} and T_{emb} as the input to the model:

$$P_{emb} = [I_{emb}, T_{emb}] \quad (3)$$

Modality-Aware Mixture-of-Experts

Multimodal learning differs significantly from unimodal learning, as the differences between modalities cannot be ignored. Using the same Feed-Forward Network for all modalities can lead to inappropriate fusion of modalities, resulting in degraded performance. Conversely, using modality-specific MoE in all layers may not benefit the alignment of different modalities. Therefore, we propose the Modality-Aware Mixture-of-Experts (MoE) as shown in Figure 3, which incorporates the modality routing technique on top

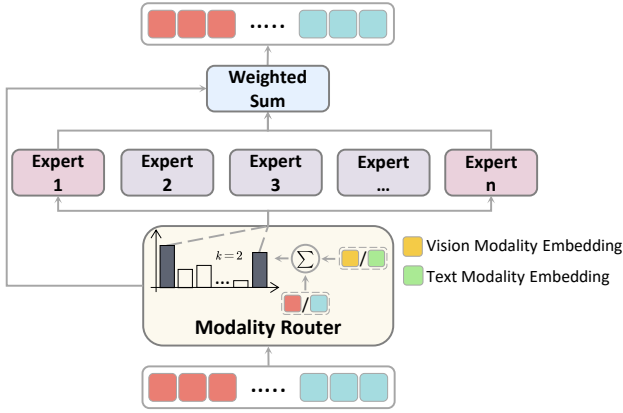


Figure 3: Architecture of Modality-Aware MoE.

of the general MoE to capture modality-specific information while fusing by selectively switching to different experts.

In the general MoE, each MoE block typically consists of N experts, and each input token is processed by k experts selected from the N experts. A lightweight router g is used to select the k experts for each token, which employs a simple linear-softmax predictor to calculate the routing weight. This can be formulated as:

$$g(\mathbf{x}) = \text{softmax}(\mathbf{W} \cdot \mathbf{x}) \quad (4)$$

$\mathbf{W} \in \mathbb{R}^{D \times N}$ is a learnable projector for input $\mathbf{x} \in \mathbb{R}^D$.

The final output of the MoE block is the weighted average of the k selected experts, which can be formulated as:

$$\text{MoE}(\mathbf{x}) = \sum_{i=1}^k g(\mathbf{x})_i \cdot \text{FFN}_i(i) \quad (5)$$

Modality Routing General MoE does not impose any restrictions on the router, which can easily lead to unbalanced routing. LIMoE (Mustafa et al. 2022) points out that this phenomenon can be exacerbated in multimodal learning due to the difference in token count across different modalities.

To address this issue, we propose a modality-aware routing approach to enhance the router. We adopt a best-effort strategy for routing to preserve all tokens while explicitly providing modality information to the router by adding modality-specific embeddings. The new routing function can be formulated as follows:

$$g(\mathbf{x}) = \text{softmax}(\mathbf{W} \cdot (\mathbf{x} + \mathbf{b}_m)) \quad (6)$$

Here, we use modality-specific embeddings $\mathbf{b}_m \in \mathbb{R}^D$ for different modalities, i.e., \mathbf{b}_I for images and \mathbf{b}_T for text.

Auxiliary Loss In addition to modality routing, we use a single simple auxiliary loss to balance routing and avoid carefully tuning the weight. Following Shazeer et al. (2017), we add Load-Balancing Loss as the auxiliary loss to train the router. It can be formulated as follows:

$$\mathcal{L}_{aux} = \alpha \cdot N \sum_i^N f_i \times p_i \quad (7)$$

MLM	Pre-training Tasks			MIM Pixel	Batch size	Time
	ITC	ITM	MIM Token			
✓				✓	224	2.14h
✓			✓	✓	152	3.09h
✓	✓			✓	132	3.26h
✓	✓	✓		✓	80	6.88h
✓	✓	✓		✓	64	7.73h

Table 1: Maximum batch size per GPU and pre-training time per epoch of different pre-training tasks on 8 A100 GPUs with the same architecture as EVE-Base. We add vision mask tokens in the encoder during masked token modeling.

This objective encourages uniform routing of tokens, where N denotes the number of experts, f_i denotes the fraction of tokens dispatched to the i^{th} expert, and p_i denotes the average routing weight for the i^{th} expert. The weight α is a hyperparameter that we set at 0.001 by default to avoid overwhelming other objectives.

Considering efficiency, we use a soft router with top- $k = 2$ in the deep layers and a hard router in the shallow layers. An MoE module equipped with a hard router has the same number of experts as the number of modalities. The hard router directly selects the corresponding expert based on the modality of each token.

Pre-training Task: Masked Signal Modeling

Previous multimodal models (Li et al. 2021; Radford et al. 2021; Bao et al. 2022a; Li et al. 2019; Zhao et al. 2023a) typically involve complex pre-training tasks like Image-Text Contrastive Learning (ITC) (Radford et al. 2021), Image-Text Matching (ITM) (Li et al. 2019), and Masked Representation Modeling (MRM) (Zhao et al. 2023a). These methods have shown good performance, but pre-training still requires significant computational resources, and is challenging to scale up.

Table 1 shows the efficiency comparison between different pre-training tasks, which indicates a significant difference in time consumption and batch size. Compared to pre-training without ITC and ITM, including them requires four times more computational resources to achieve a similar speed. Moreover, ITC and ITM tasks are similar to other contrastive learning-based methods that typically require a larger batch size to achieve better performance. Incorporating additional pre-training tasks can significantly decrease training speed, increase training difficulty, and have an impact on the scalability of the model.

Thus, we pre-train our model with only one unified masked signal modeling objective on image-text pairs to reconstruct masked signals by visible signals as shown in Figure 2. Specifically, masked signal modeling combines masked image modeling and masked language modeling, and only utilizes the raw signals from image-text pairs themselves without relying on any additional techniques. We use masked image and complete text in masked image modeling, while complete image and masked text in masked language modeling. Despite its simplicity, our ap-

proach achieves competitive performance compared to previous methods and can be easily scaled up.

In this section, we use $h(\cdot)$ and $\theta(\cdot)$ to denote the encoder and the decoder. \hat{I} and \hat{T} are represented for masked image and masked text. D indicates the dataset.

Masked Language Modeling (MLM) Following BERT (Devlin et al. 2019), we randomly mask some of the text tokens and predict them based on the information provided by the image and corrupted text. The Masked Language Modeling (MLM) objective can be formulated as follows:

$$\mathcal{L}_{mlm} = \mathbb{E}_{(I,T) \sim D} \ell_{mlm} \left(\theta_t \left(h(I, \hat{T}) \right), T \right) \quad (8)$$

ℓ_{mlm} computes the cross-entropy loss between the prediction probability P_{mlm} , obtained from the text decoder g_t , and the ground truth on each masked token. We use a two-layer MLP with a softmax layer as the text decoder.

Masked Image Modeling (MIM) Previous methods (Zhao et al. 2023a; Zhang et al. 2023; Wang et al. 2023) have typically employed semantically rich visual features obtained from the model itself or discrete visual tokens obtained from visual tokenizers as the targets for MIM. However, both approaches have their drawbacks. Training visual tokenizers (Ramesh et al. 2021; Peng et al. 2022) is a challenging task as different tokenizers can have varying impacts on performance and may lead to error propagation. Meanwhile, using visual features (Zhao et al. 2023a; Zhang et al. 2023) requires either applying momentum distillation techniques or employing other loss functions and techniques to prevent the model from diverging during training. These MIM targets make the overall framework more complex.

In visual self-supervised learning, some works use other information as the MIM targets, such as RGB pixels (He et al. 2022a), scene depth (Bachmann et al. 2022), HOG (Wei et al. 2022a), etc. However, using targets such as scene depth and HOG requires additional techniques, which increases the complexity of the training process. In order to maintain simplicity and effectiveness, we choose to utilize the image pixels themselves as the reconstruction target.

Following MAE (He et al. 2022a), we adopt an asymmetric design for MIM, where only observed image patches and all text tokens are fed into the encoder. A lightweight decoder is used to reconstruct raw pixels on masked positions from partial image representation and masked tokens, as shown in Figure 2. We use multiple Transformer blocks with narrower hidden widths as the decoder. The MIM objective can be formulated as:

$$\mathcal{L}_{mim} = \mathbb{E}_{(I,T) \sim D} \ell_{mim} \left(\theta_i \left(h(\hat{I}, T) \right), I \right) \quad (9)$$

ℓ_{mim} calculates the mean square error between the raw pixels and the reconstructed result generated by the image decoder. We compute the loss on masked image patches.

The overall objective of masked signal modeling is:

$$\mathcal{L} = \mathcal{L}_{mlm} + \mathcal{L}_{mim} \quad (10)$$

Experiments

Pre-training Datasets

Following Previous methods, we pre-train EVE on four widely used public datasets: MSCOCO Captions (Lin et al. 2014), Visual Genome (Krishna et al. 2017), SBU Captions (Ordonez, Kulkarni, and Berg 2011) and Conceptual Captions (Sharma et al. 2018). There are about 4M images and 10M image-text pairs in all datasets. Since some downstream tasks are based on COCO, we exclude all images in the test sets of downstream tasks from the pre-training data. We also pre-train EVE-Large on a larger dataset with 21M image-text pairs by adding CC12M (Changpinyo et al. 2021).

Implementation Details

EVE-Base has 12 Transformer blocks and EVE-Large has 24 Transformer blocks. We employ a soft router with 32 experts in EVE-Base on top-2 blocks, EVE-Large on top-3 blocks, and a hard router on the other blocks. We pre-train EVE-Base for 480k steps with a batch size of 2048 and EVE-Large with the same batch size for 280k steps. We use AdamW (Loshchilov and Hutter 2019) optimizer. The peak learning rate is 5e-4 for EVE-Base and 2e-4 for EVE-Large. During pre-training, the image resolution is 224×224 . We use random resized cropping and horizontal flipping for data augmentation. We mask 75% of image in MIM and 50% of text in MLM. EVE is initialized with BEiTv2. More details are provided in Appendix ¹.

Vision-Language Downstream Tasks

We evaluate our pre-trained model on three common Vision-Language Tasks. More implementation details and comparison on inference speed are provided in Appendix.

Visual Question Answering (VQA) VQA requires the model to predict an answer based on the given image and question. We use VQA2.0 dataset (Goyal et al. 2017) to evaluate our model. Following previous work (Bao et al. 2022a), we view the task as a classification task.

Natural Language for Visual Reasoning (NLVR2) Given a sentence and two images, NLVR2 asks the model to judge whether the sentence accurately describes the relationship between the two images. We evaluate our model on NLVR2 dataset (Suhr et al. 2019). Following Chen et al. (2020), we convert the triplet input into two image-text pairs with the same text description and different images.

Image-Text Retrieval Retrieval task contains two sub-tasks: Image-to-Text Retrieval (TR) and Text-to-Image Retrieval (IR). We evaluate the model on widely used Flickr30K (Plummer et al. 2015) and MSCOCO (Lin et al. 2014) benchmarks following Karpathy split (Karpathy and Fei-Fei 2015). Following Li et al. (2021), we apply ITC and ITM losses in the fine-tuning stage and we use rerank strategy during inference.

¹Full manuscript with appendix is in arXiv:2308.11971.

Model	#Images	VQA		NLVR2		COCO		Flickr30K	
		test-dev	test-std	dev	test-P	TR@1	IR@1	TR@1	IR@1
ALBEF (Li et al. 2021)	4M	74.54	74.70	80.24	80.50	73.1	56.8	94.3	82.8
Triple (Yang et al. 2022)	4M	74.90	74.92	80.54	81.33	75.6	59.0	94.9	84.0
Codebook (Duan et al. 2022)	4M	74.86	74.97	80.50	80.84	75.3	58.7	95.1	83.3
METER (Dou et al. 2022a)	4M	77.68	77.64	82.33	83.05	76.2	57.1	94.3	82.2
MAMO (Zhao et al. 2023a)	4M	76.12	76.20	81.86	81.53	77.1	60.3	95.6	85.4
VLMO (Bao et al. 2022a)	4M	76.64	76.89	82.77	83.34	74.8	57.2	92.3	79.3
VL-BEiT (Bao et al. 2022b)	4M	77.53	77.75	81.93	82.66	79.5	61.5	95.8	83.9
VLMAE (He et al. 2022b)	4M	75.30	75.40	80.50	81.20	77.3	59.6	95.2	83.6
MaskVLM (Kwon et al. 2023)	4M	75.45	75.40	81.58	81.98	76.3	60.1	95.6	84.5
VLC-Base (Li et al. 2021)	5.6M	74.02	74.00	77.70	79.04	72.4	50.7	89.2	71.3
DAVINCI (Diao et al. 2023)	631.8M	76.32	76.44	80.03	80.25	-	-	-	-
SimVLM-Base (Wang et al. 2022b)	1.8B	77.87	78.14	81.72	81.77	-	-	-	-
BEiT3-Base (Wang et al. 2023)	3.1B	77.65	-	83.60	84.40	79.1	61.4	96.3	86.2
EVE-Base (Ours)	4M	78.00	78.02	83.34	83.93	79.6	62.0	95.6	84.1

Table 2: Comparison with state-of-the-art base-size models on VQA, NLVR2, MSCOCO and Flickr30K. Gray lines indicate the model pre-trained with much more data (more than 400M).

MIM Target	NLVR2		Flickr30K		VQA
	dev	test-P	TR	IR	
BEiTv2 Token	78.0	78.5	92.6	78.3	76.6
DALL-E Token	×	×	92.4	77.4	75.8
Pixel (Ours)	79.7	80.1	93.9	80.7	77.3

Table 3: Ablation study on MIM target. × denotes divergence during fine-tuning.

Results on Downstream Tasks

We present the results of VQA, NLVR2, COCO, and Flickr30K with state-of-the-art base models in Table 2 and large models in Table 4. We report the accuracy for VQA and NLVR2, top-1 recall for TR and IR.

Results on Vision-Language Understanding EVE-Base outperforms all previous methods on Understanding tasks and even marginally outperforms BEiT3-Base (Wang et al. 2023) pre-trained with 3.1B data on VQA. EVE-Base outperforms VLMO (Bao et al. 2022a), which also employs a unified architecture with more pre-training objectives by 1.77% on VQA test-dev and 0.70% on NLVR2 test-P. EVE-Large4M shows similar performance to SimVLM-Large (Wang et al. 2022b), whereas EVE-Large16M surpasses SimVLM-Huge which is larger and pre-trained on much more data.

Results on Image-Text Retrieval EVE-Base achieves competitive results on Flickr and outperforms the previous state-of-the-art methods on COCO. Compared to VLMO, EVE-Base achieves improvements of 6.42% on COCO text retrieval R@1 and 8.39% on COCO image retrieval R@1. In addition, EVE-Large demonstrates better performance on both COCO and Flickr30K to other Large or even Huge models with very limited data. Notably, Image-Text Contrastive Learning and Image-Text Matching are not involved in the pre-training of EVE.




Figure 4: Ablation study on masking ratio. Left and right y-axis denote VQA accuracy and Flickr mean recall.

Ablation Studies

For all ablation studies, we pre-train the model for 25 epochs with a similar architecture to EVE-Base and report accuracy on NLVR2, VQA dev set, and top-1 recall on Flickr30K. We use the soft router with top- $k = 2$ by default. We present some more ablation studies in Appendix.

MIM Target We compare different MIM targets in Table 3, including image token and pixel. We use the tokenizer from BEiT v2 (Peng et al. 2022) and DALL-E (Ramesh et al. 2021). It is observed that reconstructing pixels is better than reconstructing image tokens in all tasks. Using a more complex MIM target does not achieve the expected effect.

Masking Ratio In Figure 4, we investigate the impact of different masking ratios on both vision and language. Results indicate that a higher vision masking ratio leads to improved performance. We hypothesize that the raw signals are highly redundant for image, and a higher masking ratio is needed to facilitate representation learning. The noteworthy difference from previous work (Zhao et al. 2023a) is that we achieve better performance at a higher text masking ratio. Our interpretation is that with a more profound integration of vision and language, the model can more easily predict masked text tokens with the aid of vision.

Model	#Images	VQA		NLVR2		COCO		Flickr30K	
		test-dev	test-std	dev	test-P	TR@1	IR@1	TR@1	IR@1
VinVL-Large (Zhang et al. 2021)	8.9M	76.52	76.60	82.67	83.98	75.4	58.8	-	-
BLIP-CapFiltL (Li et al. 2022)	129M	78.25	78.32	82.15	82.24	81.2	64.1	97.2	87.5
BLIP-Large (Li et al. 2022)	129M	-	-	-	-	82.4	65.1	97.4	87.6
Uni-PerceiverMoE-L (Zhu et al. 2022)	44.1M	-	-	-	-	74.7	58.3	94.1	83.7
FILIP-Large (Yao et al. 2022)	340M	-	-	-	-	78.9	61.2	96.6	87.1
Prismer-Large (Liu et al. 2023)	12.7M	78.4	78.5	-	-	-	-	-	-
GIT (Wang et al. 2022a)	800M	75.5	-	-	-	-	-	-	-
ALIGN-Large (Jia et al. 2021)	1.8B	-	-	-	-	77.0	59.9	95.3	84.9
SimVLM-Large (Wang et al. 2022b)	1.8B	79.32	79.56	84.13	84.84	-	-	-	-
SimVLM-Huge (Wang et al. 2022b)	1.8B	80.03	80.34	84.53	85.15	-	-	-	-
Florence-Huge (Yuan et al. 2021)	900M	80.16	80.36	-	-	81.8	63.2	97.2	87.9
EVE-Large (Ours)	4M	79.25	79.20	84.03	84.69	82.5	65.2	96.3	86.3
EVE-Large (Ours)	16M	80.17	80.18	85.63	86.22	83.5	66.7	98.0	87.9

Table 4: Comparison with state-of-the-art large-size models on VQA, NLVR2, MSCOCO and Flickr30K. Gray lines indicate the model pre-trained with much more data (more than 400M).




Figure 5: Ablation study on the number of experts and top- k design. We use soft router in [8, 10, 12] Transformer blocks.

Tasks		NLVR2		Flickr30K		VQA
MIM	MLM	dev	test-P	TR	IR	
✓		57.2	57.4	30.4	22.9	60.9
	✓	78.8	79.3	92.2	79.2	77.0
✓†	✓†	75.4	75.7	88.6	74.2	74.6
✓	✓	79.7	80.1	93.9	80.7	77.3

Table 5: Ablation study on MIM and MLM. † denotes the model is pre-trained by MIM and MLM simultaneously with masked image and text inputs. Masking ratio is set to 50% for both image and text in †, but 75% for image in others.

Number of Experts and Top-K The number of experts and the selection of top- k are crucial aspects of MoE design, as they determine the model’s parameters, computational complexity, and performance. Figure 5 clearly demonstrates that performance deteriorates as the number of selected experts decreases from 2 to 1. When $k = 1$, increasing the number of experts can actually lead to a decrease in performance, which is more evident in retrieval tasks. When $k = 2$, increasing the number of experts leads to corresponding improvements in the performance of both VQA and retrieval tasks, with a more significant improvement observed in the retrieval task.

Pre-training Tasks				Flickr30K		VQA
MIM	MLM	ITC	ITM	TR	IR	
✓	✓	✓		94.0	80.0	76.8
✓	✓		✓	94.0	80.7	77.0
✓	✓	✓	✓	94.2	80.8	77.1
✓	✓			94.4	81.2	77.4

Table 6: Ablation study on more pre-training tasks. All models are pre-trained with the same pre-training GPU hours.

Pre-training Tasks We explore the use of different pre-training tasks for masked signal modeling in Table 5. Experiments reveal that MLM with a high masking ratio is sufficient for learning the interaction between vision and language. The addition of MIM further improves the results by reducing bias, as observed in (Kwon et al. 2023). Pre-training with MIM alone results in a minimal fusion between vision and language. We hypothesize that text descriptions are typically coarse-grained and may not offer significant assistance in fine-grained vision reconstruction. Simultaneously masking both modalities and performing MIM and MLM is not recommended. This task reduces the amount of vision and language information available, which in turn increases the difficulty of MLM and MIM, resulting in performance decline.

We further explore more pre-training tasks under the same pre-training GPU hours in Table 6. Pre-training only on MIM and MLM achieves better results in both retrieval tasks and understanding tasks, thereby demonstrating the efficiency of Masked Signal Modeling. Performance on NLVR task is provided in Appendix.

Deep FFN We compare different designs of FFN in the deep layers in Table 7. Modality-shared FFN performs better than modality-specific MoE in the deep layers, as deep features require more alignment between modalities. Using a soft router can align different modalities while obtaining more modality-specific information, thereby further enhancing performance compared to deeper architecture. When we set $\mathcal{L}_{aux} = 0$, there is a noticeable decline in the model’s

Deep FFN	NLVR2		Flickr30K		VQA
	dev	test-P	TR	IR	
Shared FFN	79.6	80.1	93.5	80.1	77.0
Shared FFN [†]	80.1	80.2	93.9	80.6	77.1
Hard Router	79.8	80.1	93.2	79.3	77.0
Soft Router	80.3	80.7	94.4	81.2	77.4
Soft Router [‡]	79.2	80.0	93.6	79.7	77.1

Table 7: Ablation study on deep (top-2 layers) FFN design. Shared FFN indicates different modalities use the same FFN. We additionally add one more Transformer block to investigate the impact of parameters per token for [†]. We set $\mathcal{L}_{aux} = 0$ for [‡].




Figure 6: Frequency distribution of different modal tokens routed to specific experts in the last layer of EVE-Base.

performance across various tasks.

Figure 6 illustrates the frequency distribution of tokens from various modalities routed to specific experts in the last layer of EVE-Base during inference in retrieval tasks. Most experts evenly process both modalities, yet some are specialized in only images or text (e.g., experts 1 and 24). This reveals a modality gap within the model’s deeper layers. Shared FFN fails to address this issue, whereas the application of MoE can help mitigate it by routing to experts that specialize in processing vision or language tokens, thereby improving overall performance.

Modality Routing We compare the performance of the model whether use modality routing in the soft router or not in Table 8, and the results show that our proposed modality routing can help the router to distinguish the inputs of different modalities and thus achieve better performance.

Visualization

We use Grad-CAM (Selvaraju et al. 2017) heatmap to visualize the self-attention maps of EVE in masked signal modeling and VQA Task. Results are provided in Appendix.

Conclusion

In this paper, we present a new multimodal foundation model EVE only pre-trained by Maked Signal Modeling with Modality-Aware MoE which is flexible and capable of encoding different modalities in a unified manner. We accelerate pre-training speed 3.5x faster than pre-training with ITC and ITM. Additionally, it is easy to scale up with a larger model or more pre-training data. Extensive experiments demonstrate that EVE outperforms existing methods in various Vision Language downstream tasks.

Modality Routing	NLVR2		Flickr30K		VQA
	dev	test-P	TR	IR	
EVE-Base	80.3	80.7	94.4	81.2	77.4
w/o MR	79.7	80.0	93.7	80.8	77.3

Table 8: Ablation study on modality routing technique.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant No. 61876224.

References

- Bachmann, R.; Mizrahi, D.; Atanov, A.; and Zamir, A. 2022. MultiMAE: Multi-modal Multi-task Masked Autoencoders. arXiv:2204.01678.
- Bao, H.; Wang, W.; Dong, L.; Liu, Q.; Mohammed, O. K.; Aggarwal, K.; Som, S.; Piao, S.; and Wei, F. 2022a. VLMO: Unified Vision-Language Pre-Training with Mixture-of-Modality-Experts. In *Advances in Neural Information Processing Systems*.
- Bao, H.; Wang, W.; Dong, L.; and Wei, F. 2022b. VL-BEiT: Generative Vision-Language Pretraining. arXiv:2206.01127.
- Changpinyo, S.; Sharma, P.; Ding, N.; and Soricut, R. 2021. Conceptual 12M: Pushing Web-Scale Image-Text Pre-Training To Recognize Long-Tail Visual Concepts. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3558–3568.
- Chen, Y.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020. UNITER: UNiversal Image-Text Representation Learning. In *European Conference on Computer Vision*, 104–120.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 4171–4186.
- Diao, S.; Zhou, W.; Zhang, X.; and Wang, J. 2023. Write and Paint: Generative Vision-Language Models are Unified Modal Learners. In *International Conference on Learning Representations*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Dou, Z.; Xu, Y.; Gan, Z.; Wang, J.; Wang, S.; Wang, L.; Zhu, C.; Zhang, P.; Yuan, L.; Peng, N.; Liu, Z.; and Zeng, M. 2022a. An Empirical Study of Training End-to-End Vision-and-Language Transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18145–18155.
- Dou, Z.-Y.; Kamath, A.; Gan, Z.; Zhang, P.; Wang, J.; Li, L.; Liu, Z.; Liu, C.; LeCun, Y.; Peng, N.; Gao, J.; and Wang, L. 2022b. Coarse-to-Fine Vision-Language Pre-training with

- Fusion in the Backbone. In *Advances in Neural Information Processing Systems*.
- Duan, J.; Chen, L.; Tran, S.; Yang, J.; Xu, Y.; Zeng, B.; and Chilimbi, T. 2022. Multi-modal Alignment using Representation Codebook. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15630–15639.
- Geng, X.; Liu, H.; Lee, L.; Schuurmans, D.; Levine, S.; and Abbeel, P. 2022. Multimodal Masked Autoencoders Learn Transferable Representations. arXiv:2205.14204.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6325–6334.
- Gui, L.; Chang, Y.; Huang, Q.; Som, S.; Hauptmann, A.; Gao, J.; and Bisk, Y. 2023. Training Vision-Language Transformers from Captions. arXiv:2205.09256.
- Hazimeh, H.; Zhao, Z.; Chowdhery, A.; Sathiamoorthy, M.; Chen, Y.; Mazumder, R.; Hong, L.; and Chi, E. H. 2021. DSelect-k: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning. In *Advances in Neural Information Processing Systems*, 29335–29347.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. B. 2022a. Masked Autoencoders Are Scalable Vision Learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15979–15988.
- He, S.; Guo, T.; Dai, T.; Qiao, R.; Wu, C.; Shu, X.; and Ren, B. 2022b. VLMAE: Vision-Language Masked Autoencoder. arXiv:2208.09374.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.; Parekh, Z.; Pham, H.; Le, Q. V.; Sung, Y.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 4904–4916.
- Karpathy, A.; and Fei-Fei, L. 2015. Deep visual-semantic alignments for generating image descriptions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3128–3137.
- Kim, W.; Son, B.; and Kim, I. 2021. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 5583–5594.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.; Shamma, D. A.; Bernstein, M. S.; and Fei-Fei, L. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision*, 123(1): 32–73.
- Kwon, G.; Cai, Z.; Ravichandran, A.; Bas, E.; Bhotika, R.; and Soatto, S. 2023. Masked Vision and Language Modeling for Multi-modal Representation Learning. In *International Conference on Learning Representations*.
- Lepikhin, D.; Lee, H.; Xu, Y.; Chen, D.; Firat, O.; Huang, Y.; Krikun, M.; Shazeer, N.; and Chen, Z. 2021. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. In *International Conference on Learning Representations*.
- Lewis, M.; Bhosale, S.; Dettmers, T.; Goyal, N.; and Zettlemoyer, L. 2021. BASE Layers: Simplifying Training of Large, Sparse Models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 6265–6274.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. arXiv:2201.12086.
- Li, J.; Selvaraju, R. R.; Gotmare, A.; Joty, S. R.; Xiong, C.; and Hoi, S. C. 2021. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *Advances in Neural Information Processing Systems*, 9694–9705.
- Li, L. H.; Yatskar, M.; Yin, D.; Hsieh, C.-J.; and Chang, K.-W. 2019. VisualBERT: A Simple and Performant Baseline for Vision and Language. arXiv:1908.03557.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *Proceedings of the European Conference on Computer Vision*, 740–755.
- Liu, J.; Zhu, X.; Liu, F.; Guo, L.; Zhao, Z.; Sun, M.; Wang, W.; Lu, H.; Zhou, S.; Zhang, J.; and Wang, J. 2021. OPT: Omni-Perception Pre-Trainer for Cross-Modal Understanding and Generation. arXiv:2107.00249.
- Liu, S.; Fan, L.; Johns, E.; Yu, Z.; Xiao, C.; and Anandkumar, A. 2023. Prism: A Vision-Language Model with An Ensemble of Experts. arXiv:2303.02506.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Mustafa, B.; Riquelme, C.; Puigcerver, J.; Jenatton, R.; and Houlsby, N. 2022. Multimodal Contrastive Learning with LIMoE: the Language-Image Mixture of Experts. In *Advances in Neural Information Processing Systems*.
- Ordonez, V.; Kulkarni, G.; and Berg, T. L. 2011. Im2Text: Describing Images Using 1 Million Captioned Photographs. In *Advances in Neural Information Processing Systems*, 1143–1151.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. arXiv:2203.02155.
- Peng, Z.; Dong, L.; Bao, H.; Ye, Q.; and Wei, F. 2022. BEiT v2: Masked Image Modeling with Vector-Quantized Visual Tokenizers. arXiv:2208.06366.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models. In *IEEE/CVF International Conference on Computer Vision*, 2641–2649.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 8748–8763.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-Shot Text-to-Image Generation. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 8821–8831.
- Riquelme, C.; Puigcerver, J.; Mustafa, B.; Neumann, M.; Jenatton, R.; Pinto, A. S.; Keyser, D.; and Hounsby, N. 2021. Scaling Vision with Sparse Mixture of Experts. In *Advances in Neural Information Processing Systems*, 8583–8595.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *IEEE/CVF International Conference on Computer Vision*, 618–626.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2556–2565.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q. V.; Hinton, G. E.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- Shen, S.; Yao, Z.; Li, C.; Darrell, T.; Keutzer, K.; and He, Y. 2023. Scaling Vision-Language Models with Sparse Mixture of Experts. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Su, W.; Zhu, X.; Cao, Y.; Li, B.; Lu, L.; Wei, F.; and Dai, J. 2020. VL-BERT: Pre-training of Generic Visual-Linguistic Representations. In *International Conference on Learning Representations*.
- Suhr, A.; Zhou, S.; Zhang, A.; Zhang, I.; Bai, H.; and Artzi, Y. 2019. A Corpus for Reasoning about Natural Language Grounded in Photographs. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 6418–6428.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wang, J.; Yang, Z.; Hu, X.; Li, L.; Lin, K.; Gan, Z.; Liu, Z.; Liu, C.; and Wang, L. 2022a. GIT: A Generative Image-to-text Transformer for Vision and Language. *Transactions on Machine Learning Research*.
- Wang, W.; Bao, H.; Dong, L.; Bjorck, J.; Peng, Z.; Liu, Q.; Aggarwal, K.; Mohammed, O. K.; Singhal, S.; Som, S.; and Wei, F. 2023. Image as a foreign language: BEiT pretraining for vision and vision-language tasks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2022b. SimVLM: Simple Visual Language Model Pre-training with Weak Supervision. In *International Conference on Learning Representations*.
- Wei, C.; Fan, H.; Xie, S.; Wu, C.; Yuille, A. L.; and Feichtenhofer, C. 2022a. Masked Feature Prediction for Self-Supervised Visual Pre-Training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14648–14658.
- Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; Chi, E. H.; Hashimoto, T.; Vinyals, O.; Liang, P.; Dean, J.; and Fedus, W. 2022b. Emergent Abilities of Large Language Models. arXiv:2206.07682.
- Yang, J.; Duan, J.; Tran, S.; Xu, Y.; Chanda, S.; Chen, L.; Zeng, B.; Chilimbi, T.; and Huang, J. 2022. Vision-Language Pre-Training with Triple Contrastive Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15650–15659.
- Yao, L.; Huang, R.; Hou, L.; Lu, G.; Niu, M.; Xu, H.; Liang, X.; Li, Z.; Jiang, X.; and Xu, C. 2022. FILIP: Fine-grained Interactive Language-Image Pre-Training. In *International Conference on Learning Representations*.
- Yuan, L.; Chen, D.; Chen, Y.-L.; Codella, N.; Dai, X.; Gao, J.; Hu, H.; Huang, X.; Li, B.; Li, C.; Liu, C.; Liu, M.; Liu, Z.; Lu, Y.; Shi, Y.; Wang, L.; Wang, J.; Xiao, B.; Xiao, Z.; Yang, J.; Zeng, M.; Zhou, L.; and Zhang, P. 2021. Florence: A New Foundation Model for Computer Vision. arXiv:2111.11432.
- Zeng, Y.; Zhang, X.; and Li, H. 2022. Multi-Grained Vision Language Pre-Training: Aligning Texts with Visual Concepts. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, 25994–26009.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. VinVL: Revisiting Visual Representations in Vision-Language Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5579–5588.
- Zhang, X.; Zeng, Y.; Zhang, J.; and Li, H. 2023. Toward Building General Foundation Models for Language, Vision, and Vision-Language Understanding Tasks. arXiv:2301.05065.
- Zhao, Z.; Guo, L.; He, X.; Shao, S.; Yuan, Z.; and Liu, J. 2023a. MAMO: Fine-Grained Vision-Language Representations Learning with Masked Multimodal Modeling. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Zhao, Z.; Guo, L.; Yue, T.; Chen, S.; Shao, S.; Zhu, X.; Yuan, Z.; and Liu, J. 2023b. ChatBridge: Bridging Modalities with Large Language Model as a Language Catalyst. arXiv:2305.16103.
- Zhu, J.; Zhu, X.; Wang, W.; Wang, X.; Li, H.; Wang, X.; and Dai, J. 2022. Uni-Perceiver-MoE: Learning Sparse Generalist Models with Conditional MoEs. In *Advances in Neural Information Processing Systems*.
- Zoph, B.; Bello, I.; Kumar, S.; Du, N.; Huang, Y.; Dean, J.; Shazeer, N.; and Fedus, W. 2022. ST-MoE: Designing Stable and Transferable Sparse Expert Models. arXiv:2202.08906.