

Grounding Semantics and Instructions with Vision

Xiujun Li

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2024

Reading Committee:

Yejin Choi, Chair

Luke Zettlemoyer

Ranjay Krishna

Program Authorized to Offer Degree:
Computer Science & Engineering

© Copyright 2024

Xiujun Li

University of Washington

Abstract

Grounding Semantics and Instructions with Vision

Xiujun Li

Chair of the Supervisory Committee:

Yejin Choi

Paul G. Allen School of Computer Science & Engineering

Learning cross-modal representations is fundamental to a wide range of vision language (V+L) tasks, such as visual question answering, image-text retrieval, image captioning. Previous studies ([LBPL19, TB19, CLY⁺20]) on vision-language pretraining (VLP) have shown that it can effectively learn generic representation from massive image-text pairs. However, there are several issues for the existing approaches. These methods simply concatenate image region features and text features as input and resort to the self-attention mechanism to learn semantic alignments between image regions and text. How to effectively fusion the two modalities, to learn the alignment between vision and text is important. The second issue is, these methods heavily rely on the pretrained object region features [AHB⁺18].

This thesis aims to ground semantics with vision from two perspectives: 1). OSCAR employs the object tags to help the alignment between the image objects and text sequence; 2). ViNVL pretrains a large, semantic object detection backbone with a unified object-attribute taxonomy from the public corpus, and shows that a rich, semantic visual representation matters in the vision-language pretraining.

Inspired by the instruction tuning in Large Language Models (LLMs), the latest Multimodal

Large Language Models (MLLMs) shows superb visual instruction following capability for multimodal tasks, however, the existing MLLMs heavily rely on the pretrained LLMs, inherit its language priors for instruction understanding.

The third part of this thesis aims to ground instructions with vision. It first proposes a new setting for probing the instruction following capabilities of MLLMs, and introduces a v-MLLM model for robust visual instruction following: VIM (VISUAL MODALITY INSTRUCTION) challenges the MLLMs by embedding the instructions into the visual pixel space, which demands for strong visual interpretative skills for instruction following. VIM probes the existing MLLMs from a vision perspective, and shows a significant disparity between the original Text Modality Instruction (TEM) and VIM settings for the existing open-source MLLMs on eight benchmarks; furthermore, v-MLLM provides a solution for the robust visual instruction following, fills this gap for the open-source MLLMs.

Acknowledgments

Time flies, marking the conclusion of my Ph.D. study. It is a long and rewarding journey, challenging not only my determination and perseverance, and even health. The last two year is quite difficult and atypical to me due to some chronic health issues, I am not able to maintain a good status. This unique experience, however, has been a significant learning opportunity for me on a personal level. I feel surreal to start to write the acknowledgement of this thesis, a tsunami of emotions and gratitude overwhelms me as I reflect on so many people that have help me in many ways.

First, I would like to give my tremendous thanks to Yejin Choi, my advisor. My Ph.D. would not have been possible without you. Not only the valuable research advice your provided, your insightful guidance, constant encouragement have significantly impacted me during this journey. You also provide many chances for us to develop our social skills, presentation and communication, especially as a non-native speaker student. Your open-mindedness and endless supports gave me ample space to explore and grow. Having you as my Ph.D. advisor is a great treasure for me, you are a role model. For me, I learned a lot from you on many aspects, not only knowledge. Ph.D. is just a milestone in my life, there are still a lot to learn from you as a life mentor beyond this Ph.D journey.

In a similar vein, I would like to give the massive thanks to Jianfeng Gao, and also Bill Dolan. I met Jianfeng when I was a research intern at MSR on 2015, working on a deep RL project for

CRM. This was my intern project, and also a first attempt on deep RL inside MSR, eventually, the success of this project landed me a full time position at the DLTC at MSR. I would give my deep and sincere thanks to these people in MSR, Li Deng, Jianfeng Gao, Xiaodong He, Lihong Li, Michel Galley, Yizhe Zhang etc. MSR (ex-)colleagues. You are mentors, and also friends, the five-year stint at MSR would be a good memory for me for ever.

I would also like to express my sincere thanks to my distinguished committee members, Luke Zettlemoyer, Ranjay Krishna and Fei Xia. I am fortunate to have you on my dissertation committee to give me feedback and support. I am ready to serve the community and society as a researcher.

I also want to extend my gratitude to many collaborators, friends, mentors at UW, MSR, ByteDance, Apple and other institutions. I am honoured to be a part of the Xlab. I especially thank Yonatan Bisk, Kay (Liyiming) Ke, Rowan Zellers, Ari Holtzman, Peter West, Antoine Bosselut, Maarten Sap, Saadia Gabriel, Ximing Lu, Liwei Jiang, James Park, Melanie Sclar; I also want to thank my colleagues over Microsoft, ByteDance, and Apple for their support over these years: Baolin Peng, Chunyuan Li, Pengchuan Zhang, Xi Yin, Xiaowei Hu, Lijuan Wang, Lei Zhang from Microsoft, Chong Wang, Ruofan Ding, Ding Zhou, Yunzhe Tao, Xiang Gao, Hongxia Yang, Liang Xiang from ByteDance, Yinfei Yang, Zhe Gan, Haotian Zhang from Apple. I also would like to thank Yujie Lu, especially during this COVID period, I am inspired by many of conversation with you, including research and life. It has been amazing working with all of you.

Most importantly, I owe my gratitude to my parents Weimin Kong and Fusen Li, without whom my achievement would not have been possible. I thank my parents for their unconditional love, support, and patience over these years. And I also would thanks my elder sisters Xiufen Li and Xiudong Li for their support and contribution to the family.

In closing, I am deeply grateful to everyone who has been part of my Ph.D. journey. The shared moments of success, failure, and learning have made the past six years truly memorable, would be a great treasure for rest of my life. The completion of Ph.D. is not an end of the journey,

but the beginning of next chapter, I believe our paths will cross again.

To my parents

Contents

1	Introduction	1
2	Oscar	5
2.1	Background	6
2.2	Motivation	8
2.3	OSCAR Pre-training	9
2.4	Adapting to V+L Tasks	14
2.5	Results & Analysis	18
2.5.1	Comparison with SoTA	18
2.5.2	Qualitative Studies	20
2.5.3	Ablation Analysis	22
2.6	More Results	23
2.7	Related Work	23
2.8	Conclusion	25
3	VinVL	29
3.1	Introduction	30
3.2	Improving Vision in Vision Language	32
3.2.1	Object Detection Pre-training	34
3.2.2	Injecting attribute information into the model	36

3.2.3	Efficient region feature extractor for VL tasks	36
3.3	OSCAR+ Pre-training	37
3.3.1	Pre-training corpus	37
3.3.2	Pre-training Objectives	38
3.3.3	Pre-trained models	40
3.3.4	Adapting to VL Tasks	40
3.4	Experiments and Analysis	43
3.4.1	Main Results	43
3.4.2	Ablation Analysis	45
3.5	Qualitative study of three pre-trained vision models	51
3.6	Related Work And Discussion	55
3.7	Conclusion	56
3.8	OSCAR+ pre-training	57
3.8.1	Pre-training Corpus	57
3.8.2	OSCAR+ pre-training objectives	57
3.8.3	Ablation of the two new techniques	58
3.9	Downstream Tasks Fine-tuning	60
3.9.1	VQA	60
3.9.2	GQA	61
3.9.3	Image Captioning	61
3.9.4	NoCaps	62
3.9.5	Image-Text Retrieval	63
3.9.6	NLVR2	64
3.10	Effect of the Object-Attribute Vocabulary Size	65
3.10.1	Disentangling the effects of region proposals and model weights on R50-C4	66
3.10.2	Disentangling the effects of region proposals and model weights on the SoTA model	68

3.11	More on FPN and Comparison of C4 and FPN	69
3.11.1	Two reasons why FPN performs worse than C4 on VL tasks.	69
3.11.2	Effect of pooling methods in FPN on VQA performance.	70
3.11.3	Large-scale object-detection pre-training of C4 and FPN models	71
3.12	Grid feature	72
3.13	End-to-end inference efficiency	74
4	VIM-MLLM	75
4.1	Introduction	76
4.2	Method	80
4.2.1	VIM	80
4.2.2	v-MLLM	83
4.3	Experiments	85
4.3.1	VIM-Bench	86
4.3.2	Evaluation Setup	86
4.3.3	Baselines	87
4.3.4	Main Results	87
4.4	Ablation	88
4.4.1	Instruction Recognition	88
4.4.2	Instruction Following	90
4.4.3	Qualitative Observations	91
4.5	Discussion	91
4.6	Related Work	93
4.7	Conclusion	94
4.8	Benchmark Details	95
4.8.1	Source Datasets	95
4.8.2	Dataset Processing Pipeline	95

4.8.3	Evaluation Details	95
4.9	Full Analysis	97
4.9.1	Robustness Analysis	97
4.9.2	Details of Instruction Recognition	99
4.9.3	Qualitative Observations	99
4.10	Limitations	101
5	Conclusion	103
	Bibliography	105

Chapter 1

Introduction

For human, we perceive the world through many signals, such as images, text and voices *etc.*, we can naturally align information from multiple channels for a better understanding of the whole world. For artificial intelligence (AI), one core goal is to build algorithms that endow computers with an ability to effectively learn from multiple modality of data. For example, computers could retrieve the relevant images given a text query, describe the content of an image using natural language. Vision-and-Language (VL), an important research area that resides at the intersection of Computer Vision (CV) and Natural Language Processing (NLP).

Inspired by the success of NLP pre-training (e.g., BERT [DCLT19], RoBERTa [LOG⁺19], T5 [RSR⁺20], and GPT series [BMR⁺20, Ope23a], PaLM [ADF⁺23], Vision-Language Pre-training (VLP) has also experienced a rapid growth over the past five years, with the promise to learn universal, transferable multimodal representations.

Generally, there are three stages for VLP. The first stage is object-level VLP. Inspired by the success of BERT [DCLT19], the VL field has gradually shifted to using transformer-based models that are pre-trained on the public available corpus size, e.g., using the public text-image pairs (up to 4M images, 10M image-text pairs), with model sizes ranging from 110M to 340M. There




Figure 1.1: Timeline to visualize the main progress of VLP development.

are two main components in the object-level VLP, one is image encoding, the mainstream object features is from Bottom-up Attention model [AHB⁺18], which is pre-processed offline, the other component is image-text fusion module. Typical examples in this stage include UNITER [CLY⁺19], LXMERT [TB19] ViLBERT [LBPL19] *etc.*

The second stage pushes the VLP to a very large scale with billion-size scale corpus, starting from CLIP [RKH⁺21] and ALIGN [JYX⁺21], which train dual encoders from noisy image-text pairs with contrastive objective, showing promising zero-shot capability on vision and language tasks, and is becoming the foundation models for many VL tasks. A boom of claimed VL foundation models, *e.g.*, SimVLM [WYY⁺21], Florence [YCC⁺21], CoCa [YWV⁺22], emerges. The pretrained models can be transferred to a wide range of downstream tasks, and show robust generalization capabilities. The pretrain corpus of image-text pairs can be scaled up to 10B with a very large batch size (*e.g.* 32k, 65k).

From April 2023, after the release of GPT-4 [Ope23a], inspired by the success of (Large Language Models) LLMs, the third stage of VLP extends the open-source LLMs (*i.e.* Vicuna [CLL⁺23],

Llama [TLI⁺23, TMS⁺23]) with *visual* instruction tuning recipe to connect vision and LLMs for visual understanding, reasoning and generation. OpenAI releases its GPT-4V(ision) [Ope23c] via the ChatGPT interface (on September 2023) and API (on November 2023), and Google publishes its Gemini (on December 2023), these proprietary MLLMs attract broad attention over the internet. In the research community, there are many open-source MLLMs, including LLaVA [LLWL23], InstructBLIP [DLL⁺23], OpenFlamingo [AGG⁺23] Qwen-VL [BBY⁺23] *etc.*. Researchers have ventured into the multimodal domain with a similar recipe for Multimodal Large Language Models (MLLMs).

An overview for this dissertation My research falls in the first stage and third stage of VLP. OSCAR focuses on the multimodal fusion part, it introduces an object tag sequence as *anchor points* to ease the learning of semantic alignment between images and texts. ViNVL focuses on the vision encoding part, it unifies the four existing public object detection datasets (MSCOCO, Object365, Visual Genome, OpenImages) into one large corpus with more diverse objects and attributes, and pretrains a large object detection backbone first, the pretrained object detection backbone can deliver more rich and semantic object region representation; with the pretrained ViNVL representation, we continue to pretrain an OSCAR+ model on the public image-text pairs, and demonstrate significant improvement over previous state-of-the-art models with large margins. ViNVL pushes the object-level VLP to a very high level.

ViM falls on the third category, it challenges the existing Multimodal Large Language Models(MLLMs) by embedding the instructions into the visual pixel space, which demands for strong visual interpretative skills for instruction following. We adapt ViM to eight benchmarks, and probe diverse MLLMs in both the text-modality instruction (TEM) setting and ViM setting. We observe a significant performance disparity between the original TEM and ViM settings for open-source MLLMs. Furthermore, we train v-MLLM, a generalizable model that is capable to conduct robust

instruction following in both text-modality and visual-modality instructions.

Overall, OSCAR and VINVL are two representative work in the object-level VLP, VIM and v-MLLM will be testified by time.

Chapter 2

Oscar

This section is adapted from Oscar [LYL⁺ 20].

Inspired by BERT [DCLT19], starting from summer 2019, large-scale pre-training methods of learning cross-modal representations on image-text pairs was becoming popular for vision-language tasks. While existing methods simply concatenate image region features and text features as input to the model to be pre-trained and use self-attention to learn image-text semantic alignments in a brute force manner. One typical observation is that the salient objects in an image can be accurately detected, and are often mentioned in the paired text. In this work, we propose a new learning method OSCAR¹, which uses object tags detected in images as *anchor points* to significantly ease the learning of alignments. We pre-train an OSCAR model on the public corpus of 6.5 million text-image pairs, and fine-tune it on downstream tasks, creating new state-of-the-arts on six well-established vision-language understanding and generation tasks by that time.²

¹Object-Semantics Aligned Pre-training

²The code and pre-trained models are released: <https://github.com/microsoft/Oscar>

2.1 Background

Learning cross-modal representations is fundamental to a wide range of vision-language (V+L) tasks, such as visual question answering, image-text retrieval, image captioning. Recent studies [LBPL19, TB19, CLY⁺19, SZC⁺19, LYY⁺19, LDF⁺19, ZPZ⁺20] on vision-language pre-training (VLP) have shown that it can effectively learn generic representations from massive image-text pairs, and that fine-tuning VLP models on task-specific data achieves state-of-the-art (SoTA) results on well-established V+L tasks.

These VLP models are based on multi-layer Transformers [VSP⁺17]. To pre-train such models, existing methods simply concatenate image region features and text features as input and resort to the self-attention mechanism to learn semantic alignments between image regions and text in a brute force manner. However, the lack of explicit alignment information between the image regions and text poses alignment modeling a weakly-supervised learning task. In addition, visual regions are often over-sampled [AHB⁺18], noisy and ambiguous, which makes the task even more challenging.

In this study, we show that the learning of cross-modal representations can be significantly improved by introducing object tags detected in images as *anchor points* to ease the learning of semantic alignments between images and texts. We propose a new VLP method OSCAR, where we define the training samples as triples, each consisting of a word sequence, a set of object tags, and a set of image region features. Our method is motivated by the observation that the salient objects in an image can be accurately detected by modern object detectors [RHGS15], and that these objects are often mentioned in the paired text. For example, on the MS COCO dataset [LMB⁺14], the percentages that an image and its paired text share at least 1, 2, 3 objects are 49.7%, 22.2%, 12.9%, respectively. Our OSCAR model is pre-trained on a large-scale V+L dataset composed of 6.5 million pairs, and is fine-tuned and evaluated on seven V+L understanding and generation tasks.




Figure 2.1: OSCAR pipeline. The model takes a triplet as input, is pre-trained with two losses (a masked token loss over words & tags, and a contrastive loss between tags and others), and fine-tuned for 5 understanding and 2 generation tasks.

The overall setting is illustrated in Fig 2.1.

Although the use of anchor points for alignment modeling has been explored in natural language processing *e.g.*, [BLM91], to the best of our knowledge, this work is the first that explores the idea for VLP. There have been previous works that use object or image tags in V+L tasks for the sake of enhancing the feature representation of image regions, rather than for learning image-text alignments. For example, Zhou *et al.* [ZPZ⁺20] uses the object prediction probability as a soft label and concatenate it with its corresponding region features. Wu *et al.* [WSL⁺16] and You *et al.* [YJW⁺16] introduce image-level labels or attributes to improve image-level visual representations.

The main contributions of this work can be summarized as follows: (i) We introduce OSCAR, a powerful VLP method to learn generic image-text representations for V+L understanding and generation tasks. (ii) We have developed an OSCAR model that achieves new SoTA on multiple V+L benchmarks, outperforming existing approaches by a significant margin; (iii) We present extensive experiments and analysis to provide insights on the effectiveness of using object tags as anchor points for cross-modal representation learning and downstream tasks.

2.2 Motivation

The training data for many V+L tasks consists of image-text pairs, as shown in Fig. 2.2(a). We denote a dataset of size N by $\mathcal{D} = \{(\mathbf{I}_i, \mathbf{w}_i)\}_{i=1}^N$, with image $\mathbf{I}$ and text sequence $\mathbf{w}$. The goal of pre-training is to learn cross-modal representations of image-text pairs in a self-supervised manner, which can be adapted to serve various down-stream tasks via fine-tuning.

VLP typically employs multi-layer self-attention Transformers [VSP⁺17] to learn cross-modal *contextualized* representations, based on the *singular* embedding of each modality. Hence, the success of VLP fundamentally relies on the quality of the input singular embeddings. Existing VLP methods take visual region features $\mathbf{v} = \{v_1, \dots, v_K\}$ of an image and word embeddings $\mathbf{w} = \{w_1, \dots, w_T\}$ of its paired text as input, and relies on the self-attention mechanism to learn image-text alignments and produce cross-modal contextual representations.

Though intuitive and effective, existing VLP methods suffer from two issues: (i) *Ambiguity*. The visual region features are usually extracted from over-sampled regions [AHB⁺18] via Faster R-CNN object detectors [RHGS15], which inevitably results in overlaps among image regions at different positions. This renders ambiguities for the extracted visual embeddings. For example, in Fig. 2.2(a) the region features for dog and couch are not easily distinguishable, as their regions heavily overlap. (ii) *Lack of grounding*. VLP is naturally a weakly-supervised learning problem because there is no explicitly labeled alignments between regions or objects in an image and words or phrases in text. However, we can see that salient objects such as dog and couch are presented in both image and its paired text as in Fig. 2.2(a), and can be used as anchor points for learning semantic alignments between image regions and textual units as in Fig. 2.2(b). In this paper we propose a new VLP method that utilizes these anchor points to address the aforementioned issues.




Figure 2.2: Illustration on the process that OSCAR represents an image-text pair into semantic space via dictionary look up. (a) An example of input image-text pair (b) The object tags are used as anchor points to align image regions with word embeddings of pre-trained language models. (c) The word semantic space is more representative than image region features. In this example, dog and couch are similar in the visual feature space due to the overlap regions, but distinctive in the word embedding space.

2.3 OSCAR Pre-training

Humans perceive the world through many channels. Even though any individual channel might be incomplete or noisy, important factors are still perceivable since they tend to be shared among multiple channels (*e.g.*, dog can be described visually and verbally, as in Fig. 2.2). With this motivation, we propose a new VLP method OSCAR to learn representations that capture channel-invariant (or modality-invariant) factors at the semantic level. Oscar differs from existing VLP in the way that the input image-text pairs are represented and the pre-training objective, as outlined in Fig. 2.3.

Input

OSCAR represents each input image-text pair as a Word-Tag-Image triple $(\mathbf{w}, \mathbf{q}, \mathbf{v})$, where $\mathbf{w}$ is the sequence of word embeddings of the text, $\mathbf{q}$ is the word embedding sequence of the object tags (in text) detected from the image, and $\mathbf{v}$ is the set of region vectors of the image.

Existing VLP methods represent each input pair as $(\mathbf{w}, \mathbf{v})$. OSCAR introduces $\mathbf{q}$ as anchor points to ease the learning of image-text alignment. This is motivated by the observation that in training




Figure 2.3: Illustration of OSCAR. We represent the image-text pair as a triple [word tokens , object tags , region features], where the object tags (e.g., “dog” or “couch”) are proposed to align the cross-domain semantics; when removed, OSCAR reduces to previous VLP methods. The input triple can be understood from two perspectives: a *modality* view and a *dictionary* view.

data, important objects in an image are often also *presented* in the image-paired text, using either the same words as object tags or different but semantically similar or related words. Since the alignments between $\mathbf{q}$ and $\mathbf{w}$, both in text, are relatively easy to identified by using pre-trained BERT models [DCLT19], which are used as initialization for VLP in OSCAR, the image regions from which the object tags are detected are likely to have higher attention weights than other regions, when queried by the semantically related words in the text. This alignment learning process is conceptually illustrated in Fig. 2.2(b). The process can also be interpreted as learning to ground the image objects, which might be ambiguously represented in the vision space such as dog and couch in Fig. 2.2(a), in distinctive entities represented in the language space, as illustrated in Fig. 2.2(c).

Specifically, $\mathbf{v}$ and $\mathbf{q}$ are generated as follows. Given an image with K regions of objects (normally over-sampled and noisy), Faster R-CNN [RHGS15] is used to extract the visual semantics of each region as (v', z) , where region feature $v' \in \mathbb{R}^P$ is a P -dimensional vector (i.e., $P = 2048$),

and region position z a R -dimensional vector (*i.e.*, $R = 4$ or 6)³. We concatenate v' and z to form a position-sensitive region feature vector, which is further transformed into v using a linear projection to ensure that it has the same vector dimension as that of word embeddings. Meanwhile, the same Faster R-CNN is used to detect a set of high precision object tags. $\mathbf{q}$ is the sequence of word embeddings of the object tags.

Pre-Training Objective

The OSCAR input can be viewed from two different perspectives as

$$\mathbf{x} \triangleq \left[\underbrace{\mathbf{w}}_{\text{language}}, \underbrace{\mathbf{q}, \mathbf{v}}_{\text{image}} \right] = \left[\underbrace{\mathbf{w}, \mathbf{q}}_{\text{language}}, \underbrace{\mathbf{v}}_{\text{image}} \right] \triangleq \mathbf{x}' \quad (2.1)$$

where $\mathbf{x}$ is the *modality* view to distinguish the representations between a text and an image; while $\mathbf{x}'$ is the *dictionary* view⁴ to distinguish the two different semantic spaces, in which the input is represented. The two-view perspective allows us to design a novel pre-training objective.

A Dictionary View: Masked Token Loss. The use of different dictionaries determines the semantic spaces utilized to represent different sub-sequences. Specifically, the object tags and word tokens share the same linguistic semantic space, while the image region features lie in the visual semantic space. We define the *discrete token sequence* as $\mathbf{h} \triangleq [\mathbf{w}, \mathbf{q}]$, and apply the Masked Token Loss (MTL) for pre-training. At each iteration, we randomly mask each input token in $\mathbf{h}$ with probability 15%, and replace the masked one h_i with a special token [MASK]. The goal of

³It includes coordinates of top-left & bottom-right corners, and/or height & width.

⁴A semantic space can be viewed a vector space defined by a dictionary, which maps an input to a vector representation in the semantic space. For example, BERT can be viewed as a dictionary that defines a linguistic semantic space. BERT maps an input word or word sequence into a feature vector in the semantic space.

training is to predict these masked tokens based on their surrounding tokens $\mathbf{h}_{\setminus i}$ and all image features $\mathbf{v}$ by minimizing the negative log-likelihood:

$$\mathcal{L}_{\text{MTL}} = -\mathbb{E}_{(\mathbf{v}, \mathbf{h}) \sim \mathcal{D}} \log p(h_i | \mathbf{h}_{\setminus i}, \mathbf{v}) \quad (2.2)$$

This is similar to masked language model used by BERT. The masked word or tag needs to be recovered from its surroundings, with additional image information attended to help ground the learned word embeddings in the vision context.

A Modality View: Contrastive Loss. For each input triple, we group $\mathbf{h}' \triangleq [\mathbf{q}, \mathbf{v}]$ to represent the image modality, and consider $\mathbf{w}$ as the language modality. We then sample a set of “polluted” image representations by replacing $\mathbf{q}$ with probability 50% with a different tag sequence randomly sampled from the dataset $\mathcal{D}$. Since the encoder output on the special token [CLS] is the fused vision-language representation of $(\mathbf{h}', \mathbf{w})$, we apply a fully-connected (FC) layer on the top of it as a binary classifier $f(\cdot)$ to predict whether the pair contains the original image representation ($y = 1$) or any polluted ones ($y = 0$). The contrastive loss is defined as

$$\mathcal{L}_{\text{C}} = -\mathbb{E}_{(\mathbf{h}', \mathbf{w}) \sim \mathcal{D}} \log p(y | f(\mathbf{h}', \mathbf{w})). \quad (2.3)$$

During the cross-modal pre-training, we utilize object tags as the proxy of images to adjust the word embedding space of BERT, where a text is similar to its paired image (or more specifically, the object tags detected from the image), and dissimilar to the polluted ones.

The full pre-training objective of OSCAR is:

$$\mathcal{L}_{\text{Pre-training}} = \mathcal{L}_{\text{MTL}} + \mathcal{L}_{\text{C}}. \quad (2.4)$$

Discussion. Although other loss function designs can be considered as pre-training objectives, we perform experiments with these two losses for two reasons: (i) Each loss provides a representative learning signal from its own perspective. We deliberately keep a clear and simple form for the joint loss to study the effectiveness of the proposed dictionary and modality views, respectively. (ii) Though the overall loss is much simpler than those of existing VLP methods, it yields superior performance in our experiments.

Pre-training Corpus

We have built the pre-training corpus based on the existing V+L datasets, including COCO [LMB⁺14], Conceptual Captions (CC) [SDGS18], SBU captions [OKB11], flicker30k [YLHH14], GQA [HM19b] *etc.*. In total, the unique image set is 4.1 million, and the corpus consists of 6.5 million text-tag-image triples. Table 2.1 shows the statistics of image and text of the corpus.

Table 2.1: Statistics of the pre-training corpus.

Source	COCO (train)	CC (all)	SBU (all)	Flicker30k (train)	VQA (train)	GQA (bal-train)	VG-QA (train)	Total
Image/Text	112k/560k	3.0M/3.0M	840k/840k	29k/145k	83k/444k	79k/1026k	48k/484k	4.1M/6.5M

Implementation Details

We pre-train two model variants, denoted as OSCAR_B and OSCAR_L, initialized with parameters θ_{BERT} of BERT base ($H = 768$) and large ($H = 1024$), respectively, where H is the hidden size. To ensure that the image region features have the same input embedding size as BERT, we transform the

position-sensitive region features using a linear projection via matrix $\mathbf{W}$. The trainable parameters are $\theta = \{\theta_{\text{BERT}}, \mathbf{W}\}$. The AdamW Optimizer is used. OSCAR_B is trained for at least 1.0M steps, with learning rate $5e^{-5}$ and batch size 768. OSCAR_L is trained for at least 900k steps, with learning rate $1e^{-5}$ and batch size 512. The sequence length of discrete tokens $\mathbf{h}$ and region features $\mathbf{v}$ are 35 and 50, respectively.

2.4 Adapting to V+L Tasks

We adapt the pre-trained models to seven downstream V+L tasks, including five understanding tasks and two generation tasks. Each task poses different challenges for adaptation. We introduce the tasks and our fine-tuning strategy in this section, and leave the detailed description of datasets and evaluation metrics to Appendix.

Image-Text Retrieval heavily relies on the joint representations. There are two sub-tasks: *image retrieval* and *text retrieval*, depending on which modality is used as the retrieved target. During training, we formulate it as a binary classification problem. Given an aligned image-text pair, we randomly select a different image or a different caption to form an unaligned pair. The final representation of [CLS] is used as the input to the classifier to predict whether the given pair is aligned or not. We did not use ranking losses [KFF15, LCH⁺18], as we found that the binary classification loss works better, similarly as reported in [QSS⁺20]. In the testing stage, the probability score is used to rank the given image-text pairs of a query. Following [LDF⁺19], we report the top- K retrieval results on both the 1K and 5K COCO test sets.

We adopt the widely used Karpathy split [KFF15] on the COCO caption dataset [LMB⁺14] to conduct our experiments. Specifically, the dataset consists of 113,287 images for training, 5,000 images for validation, and 5,000 images for testing. Each image is associated with 5 human-

generated captions. For the OSCAR_B model, we fine-tune with a batch size of 256 for 40 epochs. The initial learning rate is set to $2e^{-5}$ and linearly decreases. For the OSCAR_L model, we fine-tune with a batch size of 128 for 40 epochs. The initial learning rate is set to $1e^{-5}$ and linearly decreases. We use the validation set for parameter tuning. We compare with several existing methods, including DVSA [KFF15], VSE++ [FFKF17], DPC [ZZG⁺17], CAMP [WLL⁺19], SCAN [LCH⁺18], SCG [SJL⁺19], PFAN [WYQ⁺19], Unicoder-VL [LDF⁺19], 12-in-1 [LGR⁺19], UNITER [CLY⁺19].

Image Captioning requires the model to generate a natural language description of the content of an image. To enable sentence generation, we fine-tune OSCAR using the seq2seq objective. The input samples are processed to triples consisting of image region features, captions, and object tags, in the same way as that during the pre-training. We randomly mask out 15% of the caption tokens and use the corresponding output representations to perform classification to predict the token ids. Similar to VLP [ZPZ⁺20], the self-attention mask is constrained such that a caption token can only attend to the tokens before its position to simulate a uni-directional generation process. Note that all caption tokens will have full attentions to image regions and object tags but not the other way around.

During inference, we first encode the image regions, object tags, and a special token [CLS] as input. Then the model starts the generation by feeding in a [MASK] token and sampling a token from the vocabulary based on the likelihood output. Next, the [MASK] token in the previous input sequence is replaced with the sampled token and a new [MASK] is appended for the next word prediction. The generation process terminates when the model outputs the [STOP] token. We use beam search (*i.e.*, beam size = 5) [AHB⁺18] in our experiments and report our results on the COCO image captioning dataset.

Though the training objective (*i.e.*, seq2seq) for image captioning is different from that used in pre-training (*i.e.*, bidirectional attention-based mask token loss), we directly fine-tune OSCAR for

image captioning on COCO without additional pre-training on Conceptual Captions [SDGS18]. This is to validate the generalization ability of the OSCAR models for generation tasks. We use the same Karpathy split [KFF15]. During training, we randomly select 15% of caption tokens with a maximum of 3 tokens per caption to be masked out. For the OSCAR_B model, we fine-tune with cross-entropy loss for 40 epochs with a batch size of 256 and an initial learning rate of $3e^{-5}$ and then with CIDEr optimization [RMM⁺17] for 5 epochs with a batch size of 64 and initial learning rate of $1e^{-6}$. For the OSCAR_L model, we fine-tune for 30 epochs with a batch size of 128 and an initial learning rate of $1e^{-5}$ and then with CIDEr optimization for another 3 epochs with a batch size of 48 and learning rate of $\{1e^{-6}, 5e^{-7}\}$. We compare with several existing methods, including BUTD [AHB⁺18], VLP [ZPZ⁺20], AoANet [HWCW19].

Novel Object Captioning (NoCaps) [ADW⁺19] extends the image captioning task, and provides a benchmark with images from the Open Images dataset [KRA⁺18] to test models' capability of describing novel objects which are not seen in the training corpus. Following the restriction guideline of NoCaps, we use the predicted Visual Genome and Open Images labels to form tag sequences, and train OSCAR on COCO without the initialization of pre-training.

Since NoCaps images are collected from Open Images. We train an object detector using the Open Images training set and applied it to generate the tags. We conduct experiments from BERT model directly without pre-training as required by the task guidelines. For the OSCAR_B model, we train 40 epoch with a batch size of 256 and learning rate $3e^{-5}$; further we perform CIDEr optimization with learning rate $1e^{-6}$ and batch size 64 for 5 epochs. During inference, we use constrained beam search for decoding. We compare OSCAR with UpDown [ADW⁺19] on this task.

VQA [GKSS⁺17] requires the model to answer natural language questions based on an image. Given an image and a question, the task is to select the correct answer from a multi-choice list.

Here we conduct experiments on the widely-used VQA v2.0 dataset [GKSS⁺17], which is built based on the MSCOCO [LMB⁺14] image corpus. The dataset is split into training (83k images and 444k questions), validation (41k images and 214k questions), and test (81k images and 448k questions) sets. Following [AHB⁺18], for each question, the model picks the corresponding answer from a shared set consisting of 3,129 answers.

When fine-tuning on the VQA task, we construct one input sequence, which contains the concatenation of a given question, object tags and region features, and then the [CLS] output from OSCAR is fed to a task-specific linear classifier for answer prediction. We treat VQA as a multi-label classification problem [AHB⁺18] – assigning a soft target score to each answer based on its relevancy to the human answer responses, and then we fine-tune the model by minimizing the cross-entropy loss computed using the predicted scores and the soft target scores. At inference, we simply use a Softmax function for prediction.

For VQA training, we random sample a set of 2k images from the MS COCO validation set as our validation set, the rest of images in the training and validation are used in the VQA finetuning. For the OSCAR_B model, we fine-tune for 25 epochs with a learning rate of $5e^{-5}$ and a batch size of 128. For the OSCAR_L model, we fine-tune for 25 epochs with with a learning rate of $3e^{-5}$ and a batch size of 96.

GQA [HM19b] is similar to VQA, except that GQA tests the reasoning capability of the model to answer a question. We conduct experiments on the public GQA dataset [HM19b]. For each question, the model chooses an answer from a shared set of 1,852 candidate answers. We develop two fine-tuned models using OSCAR_B. One is similar to that of VQA. The other, denoted as OSCAR_B^{*} in Table 2.3(d), is first fine-tuned on unbalanced “all-split” for 5 epochs, and then fine-tuned on the “balanced-split” for 2 epochs, as suggested in [CGL⁺19].

The fine-tuning procedure of GQA is similar to that of VQA. For the OSCAR_B model, we fine-

tune for 5 epochs with a learning rate of $5e^{-5}$ and a batch size of 128. We compare with four existing methods, including LXMERT [TB19], MMN [CGL⁺19], 12-in-1 [LGR⁺19], NSM [HM19a].

Natural Language Visual Reasoning for Real (NLVR2) [SZZ⁺18] takes a pair of images and a natural language, and the goal is to determine whether the natural language statement is true about the image pair. When fine-tuning on the NLVR2 task, we first construct two input sequences, each containing the concatenation of the given sentence (the natural language description) and one image, and then two [CLS] outputs from OSCAR are concatenated as the joint input for a binary classifier, implemented by an MLP⁵.

For the OSCAR_B model, we fine-tune for 20 epochs with learning rate $\{2e^{-5}, 3e^{-5}, 5e^{-5}\}$ and a batch size of 72. For the OSCAR_L model, we fine-tune for 20 epochs with learning rate of $\{2e^{-5}, 3e^{-5}\}$ and a batch size of 48.

2.5 Results & Analysis

2.5.1 Comparison with SoTA

To account for parameter efficiency, we compare OSCAR against three types of SoTA’s: (i) SoTA_S indicates the best performance achieved by small models prior to the Transformer-based VLP models. (ii) SoTA_B indicates the best performance achieved by VLP models of similar size to BERT base. (iii) SoTA_L indicates the best performance yielded by models that have a similar size to BERT large. To the best of our knowledge, UNITER [CLY⁺19] is the only model of BERT large size.

Table 2.2 summarizes the overall results on all tasks⁶. For all the tables in this paper, **Blue**

⁵This is not necessarily the best fine-tuning choice for NLVR2, please refer to the *Pair-biattn* finetuning in UNITER [CLY⁺19] for a better choice, which introduces a multi-head attention layer to look back the concatenated text-image sequences.

⁶All the (single-model) SoTAs are from the published results.

Table 2.2: Overall results on six tasks. Δ indicates the improvement over SoTA. SoTA with subscript S, B, L indicates performance achieved by small models, VLP of similar size to BERT base and large model, respectively. Most results are from [CLY⁺19], except that image captioning results are from [HWCW19, ZPZ⁺20], NoCaps results are from [ADW⁺19], VQA results are from [TB19].

Task	Image Retrieval			Text Retrieval			Image Captioning				NoCaps		VQA	NLVR2
	R@1	R@5	R@10	R@1	R@5	R@10	B@4	M	C	S	C	S	test-std	test-P
SoTA _S	39.2	68.0	81.3	56.6	84.5	92.0	38.9	29.2	129.8	22.4	61.5	9.2	70.90	53.50
SoTA _B	48.4	76.7	85.9	63.3	87.0	93.1	39.5	29.3	129.3	23.2	73.1	11.2	72.54	78.87
SoTA _L	51.7	78.4	86.9	66.6	89.4	94.3	—	—	—	—	—	—	73.40	79.50
OSCAR _B	54.0	80.8	88.5	70.0	91.1	95.5	40.5	29.7	137.6	22.8	78.8	11.7	73.44	78.36
OSCAR _L	57.5	82.8	89.8	73.5	92.2	96.0	41.7	30.6	140.0	24.5	80.9	11.3	73.82	80.37
Δ	5.8 $\uparrow$	4.4 $\uparrow$	2.9 $\uparrow$	6.9 $\uparrow$	2.8 $\uparrow$	1.7 $\uparrow$	2.2 $\uparrow$	1.3 $\uparrow$	10.7 $\uparrow$	1.3 $\uparrow$	7.8 $\uparrow$	0.5 $\uparrow$	0.42 $\uparrow$	0.87 $\uparrow$

indicates the best result for a task, and gray background indicates results produced by OSCAR. As shown in the table, our base model outperforms previous large models on most tasks, often by a significantly large margin. It demonstrates that the proposed OSCAR is highly parameter-efficient, partially because the use of object tags as anchor points significantly eases the learning of semantic alignments between images and texts. Note that OSCAR is pre-trained on 6.5 million pairs, which is less than 9.6 million pairs used for UNITER pre-training and 9.18 million pairs for LXMERT.

We report the detailed comparison on each task in Table 2.3. (i) VLP methods dominate empirical performance across many V+L tasks, compared with small models. OSCAR outperforms all existing VLP methods on all seven tasks, and achieves new SoTA on six of them. On GQA, neural state machine (NSM) [HM19a] relies on a strong structural prior, which can also be incorporated into OSCAR for improvement in the future. (ii) 12-in-1 is a recently proposed multi-task learning model [LGR⁺19] for V+L, implemented on BERT base. We see that OSCAR_B outperforms 12-in-1 on almost all the tasks, except on Test-P of NLVR2. Given that our method is based on single task fine-tuning, the result demonstrates the effectiveness of our proposed pre-training scheme. (iii) overall, OSCAR is the best performer on both understanding and generation tasks. On the captioning task, we further fine-tune OSCAR with self-critical sequence training (SCST) [RMM⁺17] to improve sequence-level learning. The only comparable VLP method for captioning is [ZPZ⁺20]. The results




Figure 2.4: 2D visualization using t -SNE. The points from the same object class share the same color. Please refer Appendix for full visualization.

in Table 2.3 (e) show that OSCAR yields a much better performance, *e.g.*, improving BLEU@4 and CIDEr by more than 2 and 10 points, respectively. (iv) The NoCaps guideline requires to only use the COCO captioning training set. Hence, we initialize with BERT, and train OSCAR on the COCO training set. Constrained beam search (CBS) is used. The results in Table 2.3 (f) show that the variants of OSCAR consistently outperform the previous SoTA method UpDown [ADW⁺19]. The gap is much larger on the near-domain or out-of-domain cases, demonstrating the strong generalization ability of OSCAR.

2.5.2 Qualitative Studies

We visualize the learned semantic feature space of image-text pairs of the COCO test set on a 2D map using t -SNE [MH08]. For each image region and word token, we pass it through the model, and use its last-layer output as features. Pre-trained models with and without object tags are compared. The results in Fig 2.4 reveal some interesting findings. (i) *Intra-class*. With the aid of object tags, the distance of the same object between two modalities is substantially reduced. For example, the visual and textual representations for person (or zebra) in OSCAR is much closer




Figure 2.5: Examples of image captioning. Objects are colored, based on their appearance against the ground-truth (GT): all, OSCAR & tags, tags only.

than that in the baseline method. (ii) *Inter-class*. Object classes of related semantics are getting closer (but still distinguishable) after adding tags, while there are some mixtures in the baseline, such as animal (person, zebra, sheep, bird), furniture (chair, couch, bench), and transportation (bus, train, truck, motorcycle, car). This verifies the importance of object tags in alignment learning: it plays the role of anchor points in linking and regularizing the cross-modal feature learning.

We compare generated captions of different models in Fig. 2.5. The baseline method is VLP without object tags. We see that OSCAR generates more detailed descriptions of images than the baseline, due to the use of the accurate and diverse object tags detected by Faster R-CNN. They are the anchor points in the word embedding space, guiding the text generation process.




Figure 2.6: The learning curves of fine-tuning downstream tasks with different object tags. Each curve is with 3 runs.

2.5.3 Ablation Analysis

We perform ablation experiments over a number of design choices of OSCAR in both pre-training and fine-tuning to better understand their relative importance to four representative downstream tasks. All the ablation experiments are conducted on the base model.

The Effect of Object Tags To study the effect of object tags, we experiment three different settings: (i) *Baseline (No Tags)*: this reduces the models to their previous VLP counterparts, where no tag information is exploited. (ii) *Predicted Tags*: we use an off-the-shelf object detector (trained on COCO dataset) to predict object tags. (iii) *Ground-truth Tags*: The ground-truth tags from COCO dataset are utilized to serve as a performance “upper bound” for our method. The experiments are conducted with the same BERT base model on three representative tasks, including VQA, image retrieval, and image captioning. As shown in Fig. 2.6, the learning curves for fine-tuning with object tags converges significantly faster and better than the VLP method without tags on all tasks. On the VQA and retrieval tasks, training using tags only takes half of the training time to achieve the final performance of the baseline, showing that OSCAR is a more practical and efficient scheme for VLP. With more accurate object detectors developed in the future, OSCAR can achieve even better performance, closing the gap demonstrated by using

the ground-truth tags.

2.6 More Results

The enlarged t -SNE visualization results of OSCAR and baseline (no tags) are shown in Fig. 2.7 and Fig. 2.8, respectively.

2.7 Related Work

Vision-Language Pre-training There is a growing interest in pre-training generic models to solve a variety of V+L problems, such as visual question-answering (VQA), image-text retrieval and image captioning *etc.* The existing methods [SMV⁺19, TB19, LBPL19, CLY⁺19, ZPZ⁺20, SZC⁺19, LDF⁺19, HLL⁺20] employ BERT-like objectives [DCLT19] to learn cross-modal representations from a concatenated-sequence of visual region features and language token embeddings. They heavily rely on the self-attention mechanism of Transformers to learn joint representations that are appropriately contextualized in both modalities. For example, early efforts such as [LBPL19, TB19] propose a two-stream and three-stream Transformer-based framework with co-attention to fuse the two modalities, respectively. Chen *et al.* [CLY⁺19] conduct comprehensive studies on the effects of different pre-training objectives for the learned generic representations. Zhou *et al.* [ZPZ⁺20] propose the first unified model to deal with both understanding and generation tasks, using only VQA and image captioning as the downstream tasks. In this paper, the VIM models have been applied to a wider range of downstream tasks, including both understanding and generation tasks, and have achieved new SoTA in most of them. Compared to existing VLP methods, the most salient difference of the proposed VIM is the use of object tags for aligning elements in two modalities. It alleviates the challenge of VLP models having to figure out the cross-modal semantic

alignment from scratch, and thus improves the learning efficiency. In fact, our base model already outperforms the existing large VLP models on most V+L tasks.

Object Tags Anderson *et al.* [AHB⁺18] introduce the bottom-up mechanism to represent an image as a set of visual regions via Faster R-CNN [RHGS15], each with an associated feature vector. It enables attention to be computed at the object level, and has quickly become the de facto standard for fine-grained image understanding tasks. In this paper, we propose to use object tags to align the object-region features in [AHB⁺18] in the pre-trained linguistic semantic space. The idea of utilizing object tags has been explored for image understanding [WSL⁺16, YJW⁺16, ZPZ⁺20]. Based on grid-wise region features of CNNs, Wu *et al.* [WSL⁺16] employ the predicted object tags only as the input to LSTM for image captioning, while You *et al.* [YJW⁺16] consider both tags and region features. Based on salient regions proposed by object detectors, Zhou *et al.* [ZPZ⁺20] concatenate the object prediction probability vector with region features as the visual input for VLP. Unfortunately, the tags in these works are not simultaneously associated with both object regions and word embeddings of text, resulting in a lack of grounding. Our construction of object tags with their corresponding region features & word embeddings yields more complete and informative representations for objects, particularly when the linguistic entity embeddings are pre-trained, as described next.

Multimodal Embeddings It has been shown that V+L tasks can benefit from a shared embedding space to align the inter-modal correspondences between images and text. Early attempts from Socher *et al.* [SFF10] project words and image regions into a common space using kernelized canonical correlation analysis, and achieve good results for annotation and segmentation. Similar ideas are employed for image captioning [KFF15] and text-based image retrieval [RJL⁺16]. In particular, the seminal work DeVISE [FCS⁺13] proposes to identify visual objects using semantic

information gleaned from un-annotated text. This semantic information is exploited to make predictions of image labels that are not observed during training, and improves zero-shot predictions dramatically across thousands of novel labels that have never been seen by the vision model. The idea has been extended in [SGMN13, KSZ14, NMB⁺13], showing that leveraging pre-trained linguistic knowledge is highly effective for aligning semantics and improving sample efficiency in cross-modal transfer learning. Inspired by this line of research, we revisit the idea and propose to leverage the rich semantics from the learned word embeddings in the era of neural language model pre-training. Indeed, our results on novel object captioning demonstrate that VIM helps improve the generalizability of the pre-trained models.

2.8 Conclusion

In this work, we have presented a new pre-training method VIM, which uses object tags as anchor points to align the image and language modalities in a shared semantic space. We validate the schema by pre-training VIM models on a public corpus with 6.5 million text-image pairs. The pre-trained models archive new state-of-the-arts on six established V+L understanding and generation tasks.

Table 2.3: Detailed results on V+L tasks.

Method	Size	Text Retrieval			Image Retrieval			Text Retrieval			Image Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
		1K Test Set						5K Test Set					
DVSA [14]	-	38.4	69.9	80.5	27.4	60.2	74.8	-	-	-	-	-	-
VSE++ [7]	-	64.7	-	95.9	52.0	-	92.0	41.3	-	81.2	30.3	-	72.4
DPC [46]	-	65.6	89.8	95.5	47.1	79.9	90.0	41.2	70.5	81.1	25.3	53.4	66.4
CAMP [42]	-	72.3	94.8	98.3	58.5	87.9	95.0	50.1	82.1	89.7	39.0	68.9	80.2
SCAN [18]	-	72.7	94.8	98.4	58.8	88.4	94.8	50.4	82.2	90.0	38.6	69.3	80.4
SCG [33]	-	76.6	96.3	99.2	61.4	88.9	95.1	56.6	84.5	92.0	39.2	68.0	81.3
PFAN [41]	-	76.5	96.3	99.0	61.6	89.6	95.2	-	-	-	-	-	-
Unicoder-VL [19]	B	84.3	97.3	99.3	69.7	93.5	97.2	62.3	87.1	92.8	46.7	76.0	85.3
12-in-1 [24]	B	-	-	-	65.2	91.0	96.2	-	-	-	-	-	-
UNITER [5]	B	-	-	-	-	-	-	63.3	87.0	93.1	48.4	76.7	85.9
UNITER [5]	L	-	-	-	-	-	-	66.6	89.4	94.3	51.7	78.4	86.9
OSCAR	B	88.4	99.1	99.8	75.7	95.2	98.3	70.0	91.1	95.5	54.0	80.8	88.5
	L	89.8	98.8	99.7	78.2	95.8	98.3	73.5	92.2	96.0	57.5	82.8	89.8

(a) Image-text retrieval

Method	ViLBERT	VL-BERT	VisualBERT	LXMERT	12-in-1	UNITER _B	UNITER _L	OSCAR _B	OSCAR _L
Test-dev	70.63	70.50	70.80	72.42	73.15	72.27	73.24	73.16	73.61
Test-std	70.92	70.83	71.00	72.54	-	72.46	73.40	73.44	73.82

(b) VQA

Method	MAC	VisualBERT	LXMERT	12-in-1	UNITER _B	UNITER _L	OSCAR _B	OSCAR _L
Dev	50.8	67.40	74.90	-	77.14	78.40	78.07	79.12
Test-P	51.4	67.00	74.50	78.87	77.87	79.50	78.36	80.37

(c) NLVR2

Method	Test-dev	Test-std	Method	cross-entropy optimization				CIDEr optimization			
				B@4	M	C	S	B@4	M	C	S
LXMERT [39]	60.00	60.33	BUTD [2]	36.2	27.0	113.5	20.3	36.3	27.7	120.1	21.4
MMN [4]	-	60.83	VLP [47]	36.5	28.4	117.7	21.3	39.5	29.3	129.3	23.2
12-in-1 [24]	-	60.65	AoANet [11]	37.2	28.4	119.8	21.3	38.9	29.2	129.8	22.4
NSM [12]	-	63.17	OSCAR _B	36.5	30.3	123.7	23.1	40.5	29.7	137.6	22.8
OSCAR _B	61.19	61.23	OSCAR _L	37.4	30.7	127.8	23.5	41.7	30.6	140.0	24.5
OSCAR _B *	61.58	61.62									

(d) GQA

(e) Image captioning on COCO

Method	in-domain		near-domain		out-of-domain		overall	
	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE
UpDown [1]	78.1	11.6	57.7	10.3	31.3	8.3	55.3	10.1
UpDown + CBS [1]	80.0	12.0	73.6	11.3	66.4	9.7	73.1	11.1
UpDown + ELMo + CBS [1]	79.3	12.4	73.8	11.4	71.7	9.9	74.3	11.2
OSCAR _B	79.6	12.3	66.1	11.5	45.3	9.7	63.8	11.2
OSCAR _B + CBS	80.0	12.1	80.4	12.2	75.3	10.6	79.3	11.9
OSCAR _B + SCST + CBS	83.4	12.0	81.6	12.0	77.6	10.6	81.1	11.7
OSCAR _L	79.9	12.4	68.2	11.8	45.1	9.4	65.2	11.4
OSCAR _L + CBS	78.8	12.2	78.9	12.1	77.4	10.5	78.6	11.8
OSCAR _L + SCST + CBS	85.4	11.9	84.0	11.7	80.3	10.0	83.4	11.4

(f) Evaluation on NoCaps Val. Models are trained on COCO only without pre-training.




Figure 2.7: Feature visualization of OSCAR. We observe small distances between text and image features of the same object; some of them are perfectly aligned, as demonstrated by the overlap-




Figure 2.8: Feature visualization of baseline (no tags). For several object classes, their text and image features are largely separated (e.g., person, umbrella, zebra). The distance of image features between some objects is too small (e.g., bench, chair, couch).

Chapter 3

VinVL

This section is adapted from VinVL [ZLH⁺ 21].

This chapter presents a detailed study of improving visual representations for vision language (VL) tasks and develops an improved object detection model to provide object-centric representations of images. Compared to the most widely used *bottom-up and top-down* model [AHB⁺18], the new model is bigger, better-designed for VL tasks, and pre-trained on much larger training corpora that combine multiple public annotated object detection datasets. Therefore, it can generate representations of a richer collection of visual objects and concepts. While previous VL research focuses mainly on improving the vision-language fusion model and leaves the object detection model improvement untouched, we show that visual features matter significantly in VL models. In our experiments we feed the visual features generated by the new object detection model into a Transformer-based VL fusion model OSCAR [LYL⁺20], and utilize an improved approach OSCAR+ to pre-train the VL model and fine-tune it on a wide range of downstream VL tasks. Our results show that the new visual features significantly improve the performance across all VL tasks, creating new state-of-the-art results on seven public benchmarks. Code, models and pre-extracted features are released at <https://github.com/pzzhang/VinVL>.

3.1 Introduction

Vision language pre-training (VLP) has proved effective for a wide range of vision-language (VL) tasks [LBPL19, TB19, CLY⁺19, SZC⁺19, LYY⁺19, LDF⁺19, ZPZ⁺20, LYL⁺20]. VLP typically consists of two stages: (1) an object detection model is pre-trained to encode an image and the visual objects in the image to feature vectors, and (2) a cross-modal fusion model is pre-trained to blend text and visual features. While existing VLP research focuses mainly on improving the cross-modal fusion model, this paper focuses on improving the object-centric visual representations and presents a comprehensive empirical study to demonstrate that visual features matter in VL models.

Visual feature	VQA		GQA		Image Captioning				NoCaps		Image Retrieval			Text Retrieval			NLVR2	
	test-dev	test-std	test-dev	test-std	B@4	M	C	S	C	S	R@1	R@5	R@10	R@1	R@5	R@10	dev	test-P
Anderson <i>et al.</i> [AHB ⁺ 18]	73.16	73.44	61.58	61.62	40.5	29.7	137.6	22.8	86.58	12.38	54.0	80.8	88.5	70.0	91.1	95.5	78.07	78.36
Ours	75.95	76.12	65.05	64.65	40.9	30.9	140.6	25.1	92.46	13.07	58.1	83.2	90.1	74.6	92.6	96.3	82.05	83.08
Δ	2.79 $\uparrow$	2.68 $\uparrow$	3.47 $\uparrow$	3.03 $\uparrow$	0.4 $\uparrow$	1.2 $\uparrow$	3.0 $\uparrow$	2.3 $\uparrow$	5.9 $\uparrow$	0.7 $\uparrow$	4.1 $\uparrow$	2.4 $\uparrow$	1.6 $\uparrow$	4.6 $\uparrow$	1.5 $\uparrow$	0.8 $\uparrow$	3.98 $\uparrow$	4.71 $\uparrow$

Table 3.1: Uniform improvements on seven VL tasks by replacing visual features from Anderson *et al.* [AHB⁺18] with ours. The NoCaps baseline is from VIVO [HYL⁺20], and our results are obtained by directly replacing the visual features. The baselines for rest tasks are from OSCAR [LYL⁺20], and our results are obtained by replacing the visual features and performing OSCAR+ pre-training. All models are Bert-Base size. As analyzed in Section 3.4.2, the new visual features contributes 95% of the gain.

Among the aforementioned work, a widely-used object detection (OD) model [AHB⁺18] is trained on the Visual Genome dataset [KZG⁺17]. The OD model provides an object-centric representation of images, and has been used in many VL models as a black box. In this work, we pre-train a large-scale object-attribute detection model based on the ResNeXt-152 C4 architecture (short as X152-C4). Compared to the OD model of [AHB⁺18], the new model is better-designed for VL tasks, and is bigger and trained on much larger amounts of data, combining multiple public object detection datasets, including COCO [LMB⁺14], OpenImages (OI) [KRA⁺18], Objects365 [SLZ⁺19] and Visual Genome (VG) [KZG⁺17]. As a result, our OD model achieves much better results on a wide range of VL tasks, as shown in Table 3.1. Compared to other typical OD models, such as



X152-FPN trained on OpenImages, our new model can encode a more diverse collection of visual objects and concepts (e.g., producing visual representations for 1848 object categories and 524 attribute categories), as illustrated by an example in Figure 3.1.

To validate the effectiveness of the new OD model, we pre-train a Transformer-based cross-modal fusion model OSCAR [LYL⁺20] on a public dataset consisting of 8.85 million text-image pairs, where the visual representations of these images are produced by the new OD model and are fixed during OSCAR+ pre-training. We then fine-tune the pre-trained OSCAR+ for a wide range of downstream tasks, including VL understanding tasks such as VQA [GKSS⁺17], GQA [HM19b],

NLVR2 [SZZ⁺18], and COCO text-image retrieval [LMB⁺14], and VL generation tasks such as COCO image captioning [LMB⁺14] and NoCaps [ADW⁺19]. Our results show that the object-centric representations produced by the new OD model significantly improve the performance across all the VL tasks, often by a large margin over strong baselines using the classical OD model [AHB⁺18], creating new state of the arts on all these tasks, including GQA on which none of the published pre-trained models has surpassed the deliberately designed neural state machine (NSM) [HM19a].

The main contributions of this work can be summarized as follows: (i) We present a comprehensive empirical study to demonstrate that visual features matter in VL models. (ii) We have developed a new object detection model that can produce better visual features of images than the classical OD model [AHB⁺18] and substantially uplifts the state-of-the-art results on all major VL tasks across multiple public benchmarks. (iii) We provide a detailed ablation study of our pre-trained object detection model to investigate the relative contribution to the performance improvement due to different design choices regarding diversity of object categories, visual attribute training, training data scale, model size, and model architecture.

3.2 Improving Vision in Vision Language

Deep learning-based VL models typically consist of two modules: an image understanding module **Vision** and a cross-modal understanding module **VL**:

$$(\mathbf{q}, \mathbf{v}) = \mathbf{Vision}(Img), \quad \mathbf{y} = \mathbf{VL}(\mathbf{w}, \mathbf{q}, \mathbf{v}), \quad (3.1)$$

where Img and $\mathbf{w}$ are the inputs of the vision and language modalities, respectively. The output of the **Vision** module consists of $\mathbf{q}$ and $\mathbf{v}$. $\mathbf{q}$ is the semantic representation of the image, such as tags

or detected objects, and $\mathbf{v}$ the distributional representation of the image in a high-dimensional latent space represented using e.g., the box or region¹ features produced by a VG-pre-trained Faster-RCNN model [AHB⁺18]. Most VL models use only the visual features $\mathbf{v}$, while the recently proposed OSCAR [LYL⁺20] model shows that $\mathbf{q}$ can serve as anchors for learning better vision-language joint representations and thus can improve the performance on various VL tasks. $\mathbf{w}$ and y of the VL module of Equation equation 3.1 vary among different VL tasks. In VQA, $\mathbf{w}$ is a question and y is an answer to be predicted. In text-image retrieval, $\mathbf{w}$ is a sentence and y is the matching score of a sentence-image pair. In image captioning, $\mathbf{w}$ is not given and y is a caption to be generated.

Inspired by the great success of pre-trained language models to various natural language processing tasks, vision-language pre-training (VLP) has achieved remarkable success in improving the performance of the cross-modal understanding module VL by (1) unifying vision and language modeling VL with Transformer and (2) pre-training the unified VL with large-scale text-image corpora. However, most recent works on VLP treat the image understanding module **Vision** as a black box and leave the visual feature improvement untouched since the development of the classical OD model [AHB⁺18] three years ago, despite that there has been much research progress on improving object detection by 1) developing much more diverse, richer, and larger training datasets (e.g. OpenImages and Objects 365), 2) gaining new insights in object detection algorithms such as feature pyramid network [LDG⁺17], one-stage dense prediction [LGG⁺17], and anchor-free detectors [TSCH20], and 3) leveraging more powerful GPUs for training bigger models.

In this work, we focus on improving **Vision** for better visual representations. We developed a new OD model by enriching the visual object and attribute categories, enlarging the model size and training on a much larger OD dataset, and thus advanced the state of the arts on a wide

¹We use the terms region and box interchangeably.

range of VL tasks. We detail how the new OD model is developed in the rest of this section and then describe the use of OSCAR+ for VL pre-training in Section 3.3.

3.2.1 Object Detection Pre-training

To improve the OD model for VL tasks, we utilize four public object detection datasets. As most datasets do not have attribute annotations, we adopt a *pre-training and fine-tuning* strategy to build our OD model. We first pre-train an OD model on a large-scale corpus consisting of four public datasets, and then fine-tune the model with an additional attribute branch on Visual Genome, making it capable of detecting both objects and attributes.

Data. Table 3.2 summarizes the statistics of the four public datasets used in our object detection pre-training, including COCO, OpenImagesV5 (OI), Objects365V1, and Visual Genome (VG). These datasets have complementary characters, and are extremely unbalanced in terms of data size, object vocabulary, and the number of annotations in each class. For example, the VG dataset has a rich and diverse set of annotations for both objects and their attributes with an open vocabulary. But its annotations are noisy and suffer from the missing-annotation problem. The COCO dataset, on the other hand, is very well annotated. But the coverage of visual objects and attributes is much lower than that in VG although we use both its 80 object classes and 91 stuff classes to include as diverse visual concepts as possible. We take the following steps to build a unified corpus by combining the four datasets.

1. First of all, to enhance visual concepts of tail classes, we perform class-aware sampling for OpenImages and Objects365 to get at least 2000 instances per class, resulting in 2.2M and 0.8M images, respectively.
2. To balance the contribution of each dataset, we merge the four datasets with 8 copies of COCO ($8 \times 0.11\text{M}$), 8 copies of VG ($8 \times 0.1\text{M}$), 2 copies of class-aware sampled Objects365

(2×0.8M) and one copy of the class-aware sampled OpenImages (2.2M).

3. To unify their object vocabularies, we use the VG vocabulary and its object aliases as the base vocabulary, merge a class from the other three datasets into a VG class if their class names or aliases match, and add a new class if no match is found.
4. Finally, we keep all VG classes that contain at least 30 instances, resulting in 1594 VG classes and 254 classes from the other three datasets that cannot be mapped to the VG vocabulary, resulting in a merged object detection dataset that contains 1848 classes.

Source	VG	COCO w/ stuff	Objects365	OpenImagesV5	Total
Image	97k	111k	609k	1.67M	2.49M
classes	1594	171	365	500	1848
Sampling	×8	×8	CA-2k, ×2	CA-2k	5.43M

Table 3.2: The Vision pre-training datasets. In sampling, $\times k$ means k copies in one epoch and “CA-2k” means class-aware sampling with at least 2K instances per class.

Model Architecture (FPN vs C4). Although [LDG⁺17] shows that the FPN model outperforms the C4 model for object detection, recent studies [JMR⁺20] demonstrate that FPN does not provide more effective region features for VL tasks than C4, which is also confirmed by our experimental results ². We thus conduct a set of carefully designed experiments, as to be detailed in Appendix 3.11, and find two main reasons for this. The first is that all layers in the C4 model used for region feature extraction are pre-trained using the ImageNet dataset while the multi-layer-perceptron (MLP) head of the FPN model are not. It turns out that the VG dataset is still too small to train a good enough visual features for VL tasks and using ImageNet-pre-trained

²We find in our experiments that using the same training process, the X152-C4 model even produces better object detection result than the X152-FPN model. See Appendix 3.11 for details.

weights is beneficial. The second is due to the different network architectures (CNN vs. MLP). The convolutional head used in C4 has a better inductive bias for encoding visual information than the MLP head of FPN. Therefore, in this study we use C4 architecture for VLP.

Model Pre-Training. Following the common practice in object detection training, we freeze the first convolution layer, the first residual block, and all the batch-norm layers. We also use several data augmentation methods, including horizontal flipping and multi-scale training. To train a detection model with the X152-C4 architecture, we initialize the model backbone from an ImageNet-5K checkpoint [WKM⁺19] and train for 1.8M iterations with a batch size of 16 images.

3.2.2 Injecting attribute information into the model

Following [AHB⁺18], we add an attribute branch to the pre-trained OD model, and then fine-tune the OD model on VG to inject attribute information (524 classes). Since the object representations are pre-trained in the object detection pre-training stage, we can focus the VG fine-tuning on learning attributes by picking a much larger attribute loss weight 1.25, compared to 0.5 used in [AHB⁺18, JMR⁺20]. Thus, our fine-tuned model significantly outperforms previous models [AHB⁺18, JMR⁺20] in detecting objects and attributes on VG.

3.2.3 Efficient region feature extractor for VL tasks

With a richer set of visual objects and attributes, the classical class-aware non-maximal suppression (NMS) post-processing takes a significantly larger amount of time to remove overlapped bounding boxes, making the feature extraction process extremely slow. To improve the efficiency, we replace the class-aware NMS with the class-agnostic NMS that only conducts the NMS operation once³. We also replace the time-consuming conv layers with dilation=2 used in [AHB⁺18] with conv

³Counting the NMS in the RPN module, there are in total 2 NMS operations in our efficient region feature extractor.

layers without dilation. These two replacements make the region feature extraction process much faster than that in [AHB⁺18] without any accuracy drop on VL downstream tasks. We report the end-to-end inference time of VL models with different vision models on a Titan-X GPU and a CPU with a single thread in Table 3.22 in Appendix 3.12.

In summary, the pre-trained OD model serves as the image understanding module, as in Equation equation 3.1, to produce vision presentations $(\mathbf{q}, \mathbf{v})$ for downstream VL tasks. Here, $\mathbf{q}$ is the set of detected object names (in text) and $\mathbf{v}$ is the set of region features. Each region feature is denoted as $(\hat{v}, z)$, where $\hat{v}$ is a P -dimensional representation from the input of the last linear classification layer of the detection head (*i.e.*, $P = 2048$) and z is a R -dimensional position encoding of the region (*i.e.*, $R = 6$)⁴.

3.3 OSCAR+ Pre-training

The success of VLP lies in the use of a unifying model architecture for a wide range of VL tasks and the large-scale pre-training of the unified model using objectives that correlate with the performance metrics of these downstream VL tasks. In this study we pre-train an improved version of OSCAR [LYL⁺20], known as OSCAR+ models, to learn the joint image-text representations using image tags as anchors for image-text alignment.

3.3.1 Pre-training corpus

We build our pre-training corpus based on three types of existing vision and VL datasets: (1) image captioning datasets with human-annotated captions as $\mathbf{w}$ and machine-generated ⁵ image tags as $\mathbf{q}$, including COCO [LMB⁺14], Conceptual Captions (CC) [SDGS18], SBU captions [OKB11]

⁴It includes coordinates of the bounding boxes, and height & width.

⁵We use the same model to extract visual features.

and flicker30k [YLHH14]; (2) visual QA datasets with questions as $\mathbf{w}$ and human-annotated answers as $\mathbf{q}$, including GQA [HM19b], VQA [GKSS⁺17] and VG-QAs; (3) image tagging datasets with machine-generated ⁶ captions as $\mathbf{w}$ and human-annotated tags as $\mathbf{q}$, including a subset of OpenImages (1.67M images). In total, the corpus contains 5.65 million unique images, 8.85 million text-tag-image triples. The detailed statistics are presented in Table 2.1 in the Appendix. The size of the pre-training corpus could have been significantly increased by combining large-scale image tagging datasets, such as the full set of OpenImages (9M images) and YFCC (92M images). We leave it to future work to leverage much larger corpora for model pre-training.

Loss $\mathbf{w}'/\mathbf{q}'$	$(\mathbf{w}, \mathbf{q}/\mathbf{q}', \mathbf{v})$		$(\mathbf{w}/\mathbf{w}', \mathbf{q}, \mathbf{v})$	3-way contrastive	
	All $\mathbf{q}'$'s (OSCAR)	$\mathbf{q}'$'s from QA	All $\mathbf{w}'$'s	All (OSCAR+)	$\mathbf{q}'$'s from QA
VQA (dev)	69.8 $\pm$ 0.08	70.1 $\pm$ 0.08	69.5 $\pm$ 0.05	69.8 $\pm$ 0.06	69.7 $\pm$ 0.06
COCO-IR	73.9 $\pm$ 0.2	75.0 $\pm$ 0.2	75.0 $\pm$ 0.7	78.3 $\pm$ 0.3	77.7 $\pm$ 0.7

Table 3.3: Effects of different pre-training contrastive losses on downstream tasks (R50-C4 as **Vision** module and 4-layer Transformer as **VL** module in equation 3.1). COCO-IR metric is Image-to-Text retrieval R@1 at COCO 1K test set. **Blue** indicates the best result for a task and **Black** indicates the runner-up.

3.3.2 Pre-training Objectives

There are two terms in the OSCAR+ pre-training loss as in Equation equation 3.2.

$$\mathcal{L}_{\text{Pre-training}} = \mathcal{L}_{\text{MTL}} + \mathcal{L}_{\text{CL3}}. \quad (3.2)$$

$\mathcal{L}_{\text{MTL}}$ is the Masked Token Loss defined on the text modality ($\mathbf{w}$ and $\mathbf{q}$), following closely [LYL⁺20]. (See Appendix 3.8.2 for details.) $\mathcal{L}_{\text{CL3}}$ is a novel *3-way Contrastive Loss*. Different from the binary contrastive loss used in OSCAR [LYL⁺20], the proposed *3-way Contrastive Loss* to effectively

⁶We use the captioning model released by OSCAR [LYL⁺20].

optimize the training objectives used for VQA [YHG⁺16] and text-image matching [FGI⁺15]⁷. As shown in Equation 3.3, $\mathcal{L}_{\text{CL3}}$ takes into account two types of training samples $\mathbf{x}$: the {caption, image-tags, image-features} triplets of the image captioning and image tagging data, and the {question, answer, image-features} triplets of the VQA data.

$$\mathbf{x} \triangleq \left(\underbrace{\mathbf{w}}_{\text{caption}}, \underbrace{\mathbf{q}, \mathbf{v}}_{\text{tags\&image}} \right) \quad \text{or} \quad \left(\underbrace{\mathbf{w}, \mathbf{q}}_{\text{Q\&A}}, \underbrace{\mathbf{v}}_{\text{image}} \right) \quad (3.3)$$

To compute contrastive losses, negative examples need to be constructed. We construct two types of negative (unmatched) triplets for the two types of training samples, respectively. One is the polluted “captions” $(\mathbf{w}', \mathbf{q}, \mathbf{v})$ and the other the polluted “answers” $(\mathbf{w}, \mathbf{q}', \mathbf{v})$. To classify whether a caption-tags-image triplet contains a polluted caption is a text-image matching task. To classify whether a question-answer-image triplet contains a polluted answer is an answer selection task for VQA. Since the encoding of [CLS] can be viewed as a representation of the triplet $(\mathbf{w}, \mathbf{q}, \mathbf{v})$, we apply a fully-connected (FC) layer on top of it as a 3-way classifier $f(\cdot)$ to predict whether the triplet is matched ($c = 0$), contains a polluted $\mathbf{w}$ ($c = 1$), or contains a polluted $\mathbf{q}$ ($c = 2$). The 3-way contrastive loss is defined as

$$\mathcal{L}_{\text{CL3}} = -\mathbb{E}_{(\mathbf{w}, \mathbf{q}, \mathbf{v}; c) \sim \tilde{\mathcal{D}}} \log p(c | f(\mathbf{w}, \mathbf{q}, \mathbf{v})), \quad (3.4)$$

where the dataset $(\mathbf{w}, \mathbf{q}, \mathbf{v}; c) \in \tilde{\mathcal{D}}$ contains 50% matched triples, 25% $\mathbf{w}$ -polluted triples, and 25% $\mathbf{q}$ -polluted triples. For efficient implementation, the polluted $\mathbf{w}'$ is uniformly sampled from all $\mathbf{w}$'s (captions and questions) and $\mathbf{q}'$ is uniformly sampled from all $\mathbf{q}$'s (tags and answers) in the corpus.

⁷[FGI⁺15] uses a deep-learning-based text-image matching model to select the best caption candidate for a given image.

As demonstrated in Table 3.3, when only the answer-polluted triplets are used, i.e., $(\mathbf{w}, \mathbf{q}', \mathbf{v})$ with $\mathbf{q}'$ sampled from $\mathbf{q}$'s from QA corpus, the contrastive loss simulates closely the objective for the VQA task but not the text-image retrieval task. As a result, the pre-trained model can be effectively adapted to VQA, but not so to text-image retrieval. By contrast, the proposed 3-way contrastive loss transfers well to both tasks.

3.3.3 Pre-trained models

We pre-train two model variants, denoted as OSCAR+B and OSCAR+L, which are initialized with parameters θ_{BERT} of BERT base ($L = 12, H = 768, A = 12$) and large ($L = 24, H = 1024, A = 16$), respectively, where L is the number of layers, H the hidden size, and A the number of self-attention heads. To ensure that the image region features have the same input embedding size as BERT, we transform the position-augmented region features using a linear projection via matrix $\mathbf{W}$. The trainable parameters are $\theta = \{\theta_{\text{BERT}}, \mathbf{W}\}$. OSCAR+B is trained for at least 1M steps, with learning rate $1e^{-4}$ and batch size 1024. OSCAR+L is trained for at least 1M steps, with learning rate $3e^{-5}$ and batch size 1024. The sequence length of language tokens $[\mathbf{w}, \mathbf{q}]$ and region features $\mathbf{v}$ are 35 and 50, respectively.

3.3.4 Adapting to VL Tasks

We adapt the pre-trained models to seven downstream VL tasks, including five understanding tasks and two generation tasks. Each task poses different challenges for adaptation. We refer to Appendix 3.9 for details about the seven tasks and our fine-tuning strategies.

VQA & GQA These two are the most widely used understanding task for evaluating VL models in the research community. The tasks require the model to answer natural language ques-

tions based on an image. In this study, we perform experiments on the widely-used VQA v2.0 dataset [GKSS⁺17] and GQA dataset [HM19b], Following the setting of [AHB⁺18], for each question, the model picks an answer from a shared answer set (i.e., 3, 129 candidates for VQA, 1, 852 candidates for GQA). When adapting a VLP model to the VQA task, we construct the input by concatenating a given question, object tags and object region features, and then feed the [CLS] output from ViM to a task-specific linear classifier with a softmax layer for answer prediction.

Image Captioning & NoCaps The captioning task is to generate a natural language caption for an image. This is the most widely used VL generation task in the research community – the Image Captioning Leaderboard ⁸ hosts more than 260 models as of December 10, 2020. To enable caption generation, we fine-tune ViM using the seq2seq objective. Each training sample is converted to a triplet consisting of a caption, a set of image region features, and a set of object tags. We randomly mask out 15% of the caption tokens, and use the encoding of the remaining context (the triplet) to predict the masked tokens. Similar to VLP [LYL⁺20, ZPZ⁺20], the self-attention mask is constrained such that a caption token can only attend to the tokens before its position to simulate a uni-directional generation process. All caption tokens have full attentions to image regions and object tags but not the other way around. During inference, we first encode the image regions, object tags, and a special token [CLS] as input. Then the model starts to generate a caption by feeding in a [MASK] token and sampling a token from a vocabulary based on the token probability output. Next, the [MASK] token in the previous input sequence is replaced with the sampled token and a new [MASK] is appended for the next word prediction. The generation process terminates when the model outputs the [STOP] token or the generated sentence exceeds a pre-defined max length. We perform image captioning experiments on the COCO image captioning dataset [LMB⁺14]. **Novel Object Captioning at Scale** [ADW⁺19] extends the image captioning task

⁸Image Captioning Leaderboard: <https://competitions.codalab.org/competitions/3221>

to test a model’s capability of describing novel objects from the Open Images dataset [KRA⁺18] which are unseen in the training corpus. Following the restriction guideline of NoCaps, we use the predicted Visual Genome and Open Images labels to form the input tag sequences, and directly train VIM on COCO without the initialization from pre-training. VIVO [HYL⁺20] proposed a VLP technique by only using image tagging data, and achieved SOTA results on NoCaps by fine-tuning on COCO captions. We reproduced VIVO with only one change, i.e., replacing its original vision model with our new vision model, and improved the VIVO performance significantly (short as VinVL+VIVO), as reported in Table 3.9.

Image(-to-Text) Retrieval & Text(-to-Image) Retrieval Both tasks require the model to calculate a similarity score between an image and a sentence. Thus, the task is widely used to directly measure the quality of the cross-modal VL representation. Following [LYL⁺20], we formulate the task as a binary classification problem, where given a matched image-text pair, we randomly select a different image or a different sentence to form an unmatched pair. The representation of [CLS] is used as the input to a classifier to predict a score indicating how likely the given pair is matched. In testing, the predicted score is used to rank a given image-text pairs of a query. Following [LDF⁺19], we report the top- K retrieval results on both the 1K and 5K COCO test sets.

NLVR2 The dataset is developed for joint reasoning about natural language and images [SZZ⁺18]. The task is to determine whether a text description is true about a pair of images. For fine-tuning, we first construct two input sequences, each containing the concatenation of the given text description and one of the images, and then two [CLS] outputs from VIM are concatenated to form the input to a binary classifier for prediction.

3.4 Experiments and Analysis

Task	VQA		GQA		Image Captioning				NoCaps		Image Retrieval			Text Retrieval			NLVR2	
	test-dev	test-std	test-dev	test-std	B@4	M	C	S	C	S	R@1	R@5	R@10	R@1	R@5	R@10	dev	test-P
SoTAS	70.55	70.92	—	63.17	38.9	29.2	129.8	22.4	61.5	9.2	39.2	68.0	81.3	56.6	84.5	92.0	54.10	54.80
SoTAS _B	73.59	73.67	61.58	61.62	40.5	29.7	137.6	22.8	86.58	12.38	54.0	80.8	88.5	70.0	91.1	95.5	78.39	79.30
SoTAS _L	74.75	74.93	—	—	41.7	30.6	140.0	24.5	—	—	57.5	82.8	89.8	73.5	92.3	96.0	79.76	81.47
VINVL _B	75.95	76.12	65.05	64.65	40.9	30.9	140.6	25.1	92.46	13.07	58.1	83.2	90.1	74.6	92.6	96.3	82.05	83.08
VINVL _L	76.52	76.60	—	—	41.0	31.1	140.9	25.2	—	—	58.8	83.5	90.3	75.4	92.9	96.2	82.67	83.98
Δ	1.77 ↑	1.67 ↑	3.47 ↑	1.48 ↑	0.7 ↓	0.5 ↑	0.9 ↑	0.7 ↑	5.9 ↑	0.7 ↑	1.3 ↑	0.7 ↑	0.5 ↑	1.9 ↑	0.6 ↑	0.3 ↑	2.91 ↑	2.51 ↑

Table 3.4: An overall comparison with SoTAs on seven tasks. Δ indicates the improvement over SoTA. SoTA with subscript *S*, *B*, *L* indicates performance achieved by small models, and models with the model size similar to BERT base and large, respectively. SoTAs: VQA is from ERNIE-VIL [YTY⁺20], GQA is from NSM [HM19a], NoCaps is from VIVO [HYL⁺20], NLVR2 is from VILLA [GCL⁺20], the rest tasks are from OSCAR [LYL⁺20].

3.4.1 Main Results

To account for model parameter efficiency, we group the SoTA models in three categories: (i) SoTAS indicates the best performance achieved by small models prior to the Transformer-based VLP models. (ii) SoTAS_B indicates the best performance produced by VLP models of BERT base size. (iii) SoTAS_L indicates the best performance yielded by VLP models of BERT large size.

Table 3.4 gives an overview of the results of OSCAR+ with VINVL(short for VINVL) on seven VL tasks, compared to previous SoTAs⁹. VINVL outperforms previous SoTA models on all tasks¹⁰, often by a significantly large margin. The result demonstrates the effectiveness of the region features produced by the new OD model.

In Tables 3.5 to 3.11, we report the detailed results for each downstream task, respectively. (i) The **VQA** results are shown in Table 3.5, where our single OSCAR+_B model outperforms the best ensemble model (InterBERT large [LYZ⁺20]) on the VQA leaderboard as of Dec. 12, 2020

⁹All the (single-model) SoTAs are from the published results. For all the tables in this paper, **Blue** indicates the best result for a task, and gray background indicates results produced by VINVL.

¹⁰The only exception is B@4 on image captioning.

Method	ViLBERT Base	VL-BERT Base	VisualBERT Base	LXMERT Base	12-in-1 Base	UNITER Base Large	OSCAR Base Large	VILLA Base Large	ERNIE-ViL Base Large	InterBERT Ensemble*	OSCAR+w/ VINVL Base Large
Test-dev	70.63	70.50	70.80	72.42	73.15	72.27 73.24	73.16 73.61	73.59 73.69	72.62 74.75	-	75.95 76.52
Test-std	70.92	70.83	71.00	72.54	-	72.46 73.40	73.44 73.82	73.67 74.87	72.85 74.93	76.10	76.12 76.60

Table 3.5: Evaluation results on VQA. * denotes the No.1 ensemble model of InterBERT Large on the VQA leaderboard.

Method	LXMERT MMN [CGL ⁺ 19]	12-in-1	OSCAR _B	NSM [HM19a]	OSCAR+B w/ VINVL
Test-dev	60.00	-	-	61.58	-
Test-std	60.33	60.83	60.65	61.62	63.17

Table 3.6: Evaluation results on GQA.

¹¹. (ii) The **GQA** results are shown in Table 3.6, where OSCAR+w/VINVL is the first VLP model that outperforms the neural state machine (NSM) [HM19a] which contains some sophisticated reasoning components deliberately designed for the task. (iii) The **Image Captioning** results on the public “Karpathy” 5k test split are shown in Table 3.7. Table 3.8 shows on a concise version of the COCO image captioning online leaderboard¹². The online testing setting reports the results on 40K images, with 5 reference captions (c5) and 40 reference captions (c40) per image. At the time of submitting this paper, our single model achieves No.1 on the entire leaderboard, outperforming all 263 models, including many ensemble (and anonymous) models. (iv) The Novel Object Captioning (**NoCaps**) results are shown in Table 3.9. Without any VLP, i.e. by directly training a BERT-based captioning model on COCO, the model with our new visual features (denoted as VinVL) already surpasses the human performance in CIDEr¹³. By adding VIVO [HYL⁺20] pre-training, our VinVL improves the original VIVO result by 6 CIDEr points and creates a new SoTA. (v) Overall, on all these tasks (VQA in Table 3.5, Image Captioning in Table 3.7, NoCaps in Table 3.9, Image-Text

¹¹VQA leaderboard: <https://eval.ai/web/challenges/challenge-page/514/leaderboard/1386>

¹²Image Captioning Leaderboard: <https://competitions.codalab.org/competitions/3221#results>

¹³NoCaps leaderboard: <https://eval.ai/web/challenges/challenge-page/355/leaderboard/1011>

Method	cross-entropy optimization				CIDEr optimization			
	B@4	M	C	S	B@4	M	C	S
BUTD [AHB ⁺ 18]	36.2	27.0	113.5	20.3	36.3	27.7	120.1	21.4
VLP [ZPZ ⁺ 20]	36.5	28.4	117.7	21.3	39.5	29.3	129.3	23.2
AoANet [HWCW19]	37.2	28.4	119.8	21.3	38.9	29.2	129.8	22.4
OSCAR _B [LYL ⁺ 20]	36.5	30.3	123.7	23.1	40.5	29.7	137.6	22.8
OSCAR _L [LYL ⁺ 20]	37.4	30.7	127.8	23.5	41.7	30.6	140.0	24.5
OSCAR _B w/ VINVL	38.2	30.3	129.3	23.6	40.9	30.9	140.4	25.1
OSCAR _L w/ VINVL	38.5	30.4	130.8	23.4	41.0	31.1	140.9	25.2

Table 3.7: Image captioning evaluation results (single model) on COCO “Karpathy” test split. (Note: B@4: BLEU@4, M: METEOR, C: CIDEr, S: SPICE.)

Retrieval in Table 3.10, NLVR2 in Table 3.11), we show that OSCAR_B can match or outperform previous SoTA large models, and OSCAR_L substantially uplifts the SoTA.

3.4.2 Ablation Analysis

We select the VQA task for the ablation study because its evaluation metric is well-defined and the task has been used as a testbed for all VLP models. To assist our analysis, we create a local validation set, vqa-dev, out of the standard validation set to select the best model during training for evaluation. vqa-dev contains randomly sampled 2K images and their corresponding questions, amounting to 10.4K image-QA pairs in total. Except for Table 3.4 and 3.5, all our VQA results are reported on this vqa-dev set. Unless otherwise specified, the reported STD is half of the difference of two runs of the VQA training with different random seeds.

In VQA, the VL model $y = \text{VL}(\mathbf{w}, \mathbf{q}, \mathbf{v})$ has $\mathbf{w}$ as the question and y as the answer. We focus on studying the effect of visual features $\mathbf{v}$ produced by different Vision models $\text{Vision}(\text{Img})$ to better understand their relative contribution in the VQA performance. To eliminate the impact of using

Method	BLEU@4		METEOR		ROUGE-L		CIDEr-D	
	c5	c40	c5	c40	c5	c40	c5	c40
BUTD [AHB ⁺ 18]	36.9	68.5	27.6	36.7	57.1	72.4	117.9	120.5
AoANet [HWCW19]	39.4	71.2	29.1	38.5	58.9	74.5	126.9	129.6
X-Transformer [PYLM20]	40.3	72.4	29.6	39.2	59.5	75.0	131.1	133.5
OSCAR _L w/ VINVL	40.4	74.9	30.6	40.8	60.4	76.8	134.7	138.7

Table 3.8: Leaderboard of the state-of-the-art image captioning models on the COCO online testing.

Method	CIDEr SPICE		CIDEr SPICE	
	Validation Set		Test Set	
UpDown ⁺	74.3	11.2	73.1	11.2
OSCAR _B [*]	81.1	11.7	78.8	11.7
OSCAR _L [*]	83.4	11.4	80.9	11.3
Human [ADW ⁺ 19]	87.1	14.2	85.3	14.6
VIVO [*] [HYL ⁺ 20]	88.3	12.4	86.6	12.4
VinVL [*]	90.9	12.8	85.5	12.5
VinVL+VIVO	98.0	13.6	92.5	13.1

Table 3.9: NoCaps evaluation “overall” results. All the models are trained on COCO without additional image-caption pairs following the restriction of NoCaps. (UpDown⁺ is UpDown+ELMo+CBS, the models with ^{*} is +SCST+CBS, VinVL+VIVO is with SCST only.) We refer to Table 3.18 in Appendix 3.9 for results on different subsets.

different tags $\mathbf{q}$, we use the same tags in the VQA models of OSCAR [LYL⁺20]. All the ablation experiments are conducted using models of the BERT-base size.

How much do the V and VL matter to the SoTA? Table 3.12 shows the VQA results with different vision models, i.e., R101-C4 model from [AHB⁺18] and our X152-C4 model pre-trained with 4 datasets (VinVL), and with different VLP methods, i.e., no VLP, OSCAR [LYL⁺20] and our OSCAR+. Taking the OSCAR_B model with R101-C4 features as the baseline, the OSCAR_B model with

Method	BERT	1K Test Set					
		Text Retrieval			Image Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10
Unicoder-VL [LDF ⁺ 19]	B	84.3	97.3	99.3	69.7	93.5	97.2
OSCAR	B	88.4	99.1	99.8	75.7	95.2	98.3
	L	89.8	98.8	99.7	78.2	95.8	98.3
OSCAR+ w/ ViNVL	B	89.8	98.8	99.7	78.2	95.6	98.0
	L	90.8	99.0	99.8	78.8	96.1	98.5

Method	BERT	5K Test Set					
		Text Retrieval			Image Retrieval		
		R@1	R@5	R@10	R@1	R@5	R@10
Unicoder-VL [LDF ⁺ 19]	B	62.3	87.1	92.8	46.7	76.0	85.3
UNITER [CLY ⁺ 19]	B	63.3	87.0	93.1	48.4	76.7	85.9
	L	66.6	89.4	94.3	51.7	78.4	86.9
OSCAR	B	70.0	91.1	95.5	54.0	80.8	88.5
	L	73.5	92.2	96.0	57.5	82.8	89.8
OSCAR+ w/ ViNVL	B	74.6	92.6	96.3	58.1	83.2	90.1
	L	75.4	92.9	96.2	58.8	83.5	90.3

Table 3.10: Text and Image retrieval evaluation on the COCO 1K and 5K test sets. (B for Base, L for Large)

Method	MAC	VisualBERT base	LXMERT base	12-in-1 base	UNITER		OSCAR		VILLA		OSCAR+w/ ViNVL	
					base	large	base	large	base	large	base	large
Dev	50.8	67.40	74.90	—	77.14	78.40	78.07	79.12	78.39	79.76	82.05	82.67
Test-P	51.4	67.00	74.50	78.87	77.87	79.50	78.36	80.37	79.47	81.47	83.08	83.98

Table 3.11: Evaluation results on NLVR2.

Vision	no VLP	OSCAR _B [LYL ⁺ 20]	OSCAR _B (ours)
R101-C4 [AHB ⁺ 18]	68.52 ±0.11	72.38	72.46±0.05
VinVL (ours)	71.34 ±0.17	—	74.90±0.05

Table 3.12: Effects of vision (V) and vision-language (VL) pre-training on VQA.

our X152-C4 features improves the absolute accuracy from 72.38 to 74.90, in which the OSCAR+ pre-training contributes 5% of the gain (i.e., 72.38 → 72.46) and the vision pre-training (improved

visual features) 95% (i.e., $72.46 \rightarrow 74.90$). This demonstrates that vision representations matter significantly in VLP and downstream tasks.

Taking the “no VLP” model with R101-C4 features as the baseline, Table 3.12 shows that the gains of VinVL ($71.34 - 68.52 = 2.82$) and VLP ($72.46 - 68.52 = 3.94$) are additive ($74.90 - 68.52 \approx 2.82 + 3.94$). This is intuitive because vision pre-training and VLP improve the Vision model $\text{Vision}(Img)$ and VL model $\text{VL}(\mathbf{w}, \mathbf{q}, \mathbf{v})$ separately. This also indicates that our pre-trained vision model can be utilized in any VL models by directly replacing their vision models, such as R101-C4 [AHB⁺18], with ours.

How much do data and model sizes matter to the new vision model? The improvement of VQA from R101-C4 [AHB⁺18] to VinVL (ours) in Table 3.12 is a compound effect of increasing model size (from R101-C4 to X152-C4) and data size (from VG to our merged four OD datasets). Table 3.13 shows the ablation of the two factors without VLP. Although VG’s large object and attribute vocabulary allows to learn rich semantic concepts, VG does *not* contain large amounts of annotations for effective training of deep models. Vision models trained using the merged four OD datasets perform much better than VG-only-trained models, and the improvement is larger with the increase of the model size.¹⁴

How much does OD model architecture matter? The choice of model architecture affects the VQA performance. Table 3.13 shows that R50-FPN under-performs R50-C5 when they are trained only on VG; but the performance gap diminishes when both are trained on the merged dataset (4Sets). A detailed comparison between FPN and C4 architectures is presented in Appendix 3.11.

How much does OD pre-training matter for object detection tasks? Table 3.14 presents the object detection results on COCO and the object-attribute detection results on VG (1594 object classes, 524 attribute classes). The results show that OD pre-training benefits the object detection

¹⁴The R101-C4 model in Table 3.13 is exactly the VG-pre-trained model from [AHB⁺18]. We do not train this model on our merged OD dataset because this model architecture is old-fashioned and is slow to train.

Data	R50-FPN	R50-C4	R101-C4 [AHB ⁺ 18]	X152-C4
VG	67.35±0.26	67.86±0.31	68.52 ±0.11	69.10±0.06
4Sets→VG	68.3±0.11	68.39±0.16	–	71.34 ±0.17

Table 3.13: Ablation of model size and data size on training vision models.

Model Pre-training dataset	R50-FPN		R50-C4		X152-C4	
	ImageNet	4Sets	ImageNet	4Sets	ImageNet5k	4Sets
COCO <i>mAP</i>	40.2 [WKM ⁺ 19]	44.78	38.4 [WKM ⁺ 19]	42.4	42.17	50.51
VG obj <i>mAP</i> ⁵⁰	9.6	11.3	9.6	12.1	11.2	13.8
attr <i>mAP</i> with gt boxes	5.4	5.5	6.3	6.1	6.6	7.1

Table 3.14: Effect of vision pre-training on object detection tasks. Since our 4 pre-training datasets contain Objects365, it is not surprising that we obtain better results than *mAP*⁵⁰ 42.3 reported in [SLZ⁺19].

Dataset name #obj & #attr	ImageNet 1000 & 0	VG-obj 317 & 0	VG w/o attr 1594 & 0	VG [AHB ⁺ 18] 1600 & 400	VG 1594 & 524	4Sets→VG 1848 & 524
R50-C4 + BERT _B	66.13±0.04	64.25±0.16	66.51±0.11	67.63±0.25	67.86±0.31	68.39±0.16

Table 3.15: Effect of object-attribute vocabulary. We use all grid features (maximal 273) for the ImageNet classification model (first column), and maximal 50 region features for OD models (other columns).

tasks. Note that the mAP on VG is much lower than that on typical OD datasets (such as COCO) due to two reasons: (1) VG contains a large number of object classes with limited and extremely unbalanced annotations, (2) there are many missing annotations in the VG evaluation data.¹⁵ Although the mAP numbers are low, the detection result using X152-C4 is reasonably good; see Appendix 3.5 for visualizations. We also see that FPN models perform consistently worse in attribute detection than C4 models, neither do FPN models show any advantage in object detection

¹⁵As a reference, the R101-C4 model from [AHB⁺18] on VG with 1600 objects and 400 attributes has mAP of 8.7/7.8 evaluated in our code, whereas it was reported as 10.2/7.8 due to differences in OD evaluation pipeline.

on VG. This contributes to the inferior performance of FPN, compared to C4, on downstream VL tasks, as discussed in Section 3.2.1.

How much does the diversity of visual concepts, i.e., object and attribute vocabularies, matter? We directly train vision models on different datasets, including (1) standard ImageNet with 1K classes (ImageNet), (2) Visual Genome with 317 object classes (VG-obj) that are shared with COCO 80 classes and OpenImagesV5 500 classes, (3) VG with all 1594 object classes (VG w/o attr), (4) VG with 1594 object classes and 524 attribute classes (VG), and (5) the merged OD dataset (4Sets) for pre-training and VG for fine-tuning. For all the OD models (the last four columns in Table 3.15), we initialize the OD training with an ImageNet-pre-trained classification model, and use maximal 50 *region* features per image as input to the VL fusion module. For the ImageNet pre-trained classification model (the second column in Table 3.15), we use all the *grid* features (maximal 273) for each image. The results show that

- In general, vocabularies with richer objects lead to better VQA results: VG-obj < ImageNet < VG w/o attr. The VG-obj vocabulary contains 79 of 80 COCO classes (only missing potted plant) and 313 of 500 OpenImagesV5 classes, and is a good approximation of common object classes of typical OD tasks. However, our results show that this vocabulary is not rich enough for VL tasks because it misses many important visual concepts (e.g., sky, water, mountain, etc.) which are crucial for VL tasks, as also illustrated by the comparison of detected regions in Figure 3.1.¹⁶
- Attribute information is crucial to VL tasks: models trained with attributes (VG and 4Sets→VG) are significantly better than those without attributes.

¹⁶Using the same training procedure on VG, we trained an R50-C4 model on the OpenImagesV5 dataset (500 classes). Using the region features produced by this model, the VQA performance is 63.55 ± 0.14 . The result is slightly worse than that of VG-obj because both VG and VQA images are from the COCO dataset but OpenImages images are not.

- Even for the small vision model R50-C4, vision pre-training improves visual features for VQA, i.e., 4Sets→VG is the best performer.

In Table 3.16, we use different kinds of region proposals to extract image features. COCO groundtruth object regions (GT-Obj, 80 classes) and object-stuff regions (GT-Obj&Stuff, 171 classes) are perfect in terms of localization, but their vocabulary sizes are limited. Regions proposed by VG-trained models ([2] and VinVL) are imperfect in localization but using a larger vocabulary. For the VQA task, COCO GT boxes are much worse than the proposals generated by VG-trained models. The result demonstrates the difference between the typical OD tasks and the OD tasks in VL: OD in VL requires much richer visual semantics to align with the rich semantics in the language modality. This further echoes our claim that an image understanding module trained using richer vocabularies performs better for VL tasks.

Model	GT-Obj	GT-Obj&Stuff	Anderson et al. [2]	VinVL (ours)
Anderson et al. [2]	63.81 $\pm$ 0.94	66.68 $\pm$ 0.16	68.52 $\pm$ 0.11	69.05 $\pm$ 0.06
VinVL (ours)	65.60 $\pm$ 0.21	68.13 $\pm$ 0.26	70.25 $\pm$ 0.05	71.34 $\pm$ 0.17

Table 3.16: Effect of different region proposals on VQA.

3.5 Qualitative study of three pre-trained vision models

We apply three (pre-trained) object detection models on the image in Figure 3.1 and list their detection results for a more detailed comparison.

Detections from X152-FPN trained on Open Images V5. See Figure 3.2:

Surfboard; Surfboard; Surfboard; Surfboard; Man; Human leg; Human leg;

Swimwear; Swimwear; Shorts; Shorts; Boy; Human arm.





Figure 3.2: Predictions from X152-FPN trained on OpenImages. Test image: COCO_test2015_000000028839

Detections from R101-C4 trained on VG by Anderson et.al. [AHB⁺18]. There are obviously wrong detections, marked in **red**. See Figure 3.3 (top):

black shorts; young, shirtless, standing, barefoot, surfing, little, playing boy; shirtless, standing, barefoot, walking, wet, surfing, young man; tan, bare, shirtless back; blue, clear, cloudy, hazy, light blue sky; young, shirtless, standing, surfing, barefoot, little boy; brown, short, wet, blond hair; brown, short, wet, blond hair; small, crashing wave; white, wet surfboard; white, crashing, big, rolling wave; wet, tan surfboard; green, blue fin; blue, calm, choppy, wavy, ocean,



splashing, foamy, water, rough, sandy, wet ocean; wet, calm, sandy, splashing, wavy water; white, wet surfboard; bare, wet foot; blue, colorful, multi colored, floral shorts; calm, choppy, water, rough, foamy, wavy water; distant, rocky, hazy mountains; standing, shirtless, young, barefoot, wet, surfing, walking, smiling boy; calm ocean; distant, rocky mountain; white, bare, wet surfboard; wet, sandy, calm, tan beach; gray, big rock; blue, calm background; wet, brown, tan, sandy sand; wet shadow; blue, colorful, floral, multi colored swim trunks; yellow, plastic hand.

Detections from our pre-trained X152-C4 model pre-trained on four datasets and fine-tuned on VG. There are some repetitive detections, but no obvious wrong detections. See Figure 3.3 (bottom): blue, green fin; young, barefoot, shirtless, standing, surfing, smiling, little, playing, looking, blond boy; young, barefoot, standing, shirtless, smiling, surfing, blond, playing, looking, little, walking, riding boy; shirtless, barefoot, standing, young, smiling, surfing, walking, wet, playing man; bare, wet foot; black, white surfboard; small, large, white, crashing, big, water, rolling, splashing, rough, foamy wave; bare, wet foot; dark, black, wet, cast shadow; blue, clear, hazy, cloudy, cloudless sky; black, gray, white, raised surfboard; black, wet, short short; brown, short, blond, wet, curly, wavy hair; distant, brown, large, rocky, hazy, big mountain; brown, short, dark, blond, wet hair; blue, white, calm, wavy, choppy, ocean, splashing, water, rough, clear, shallow water; bare, tan, light, beige back; black, blue, wet surfboard; small, dark, water, crashing, rolling, splashing, big wave; wet, white, sandy, tan surfboard; blue, colorful, floral, multi colored, patterned trunk; wet, brown, sandy, tan sand; white, blue, calm, foamy, choppy, splashing, wavy, ocean, rough, wa-

ter, clear, shallow water; wet, brown, sandy, calm, tan, shallow, smooth, muddy, rough beach; black, white, young board; shirtless, young, standing, barefoot, smiling, surfing, looking, walking, playing boy; blue, calm, choppy, wavy, ocean, clear, rough, splashing, water, foamy, shallow, rippled ocean; yellow, gold bracelet; white, silver, black logo; wet, bare, bent, tan, crossed, hairy, short, skinny, back, muscular, extended, outstretched leg; black, gray, white board; brown, distant, large, rocky, big hill; brown, short, blond, wet, curly head; red, black logo; bare, raised, extended, holding, open, up, bent, outstretched hand; black, wet swim trunks; bare, wet, bent, tan, crossed, skinny, short, back, muscular leg; wet, brown, muddy, sandy, tan, shallow reflection.

3.6 Related Work And Discussion

Recently there is a resurgence of interests in grid-feature based VL models [JMR⁺20, HZL⁺20, PHA⁺20] because its simplicity and efficiency enables end-to-end vision-language training, instead of fixing the vision module. [HZL⁺20, PHA⁺20] try to learn good grid features from weakly-supervised image-caption corpus. However, the end-to-end vision-language pre-training with grid features is still computationally intensive and its performance is sub-optimal compared with region-feature based models. Given that vision pre-training with object detection data is difficult to scale up beyond our setting, we are working on using large-scale weakly annotated image tagging data to further enrich the visual concepts and improve visual representations. Recent large scale vision pre-training work [SSSG17, KBZ⁺19, DBK⁺20] are pioneering works in this direction.

After the Faster-RCNN [RHGS15] was introduced, tremendous progresses on both accuracy and efficiency have been achieved in object detection, e.g., [LDG⁺17, HGDG17, LGG⁺17, RDGF16, RF17, RF18, LPY⁺17, TPL20, GLL19]. As analyzed in our work, vision in VL is closely related but

differs from standard object detection tasks. It is an important but widely open to leverage these works improve the visual understanding and representations in VL tasks.

3.7 Conclusion

In this paper we have presented a new recipe to pre-train an OD model for VL tasks. Compared to the most widely used *bottom-up and top-down* model [AHB⁺18], the new model is bigger, better-designed for VL tasks, and pre-trained on much larger text-image corpora, and thus can generate visual features for a richer collection of visual objects and concepts that are crucial for VL tasks. We validate the new model via a comprehensive empirical study where we feed the visual features to a VL fusion model which is pre-trained on a large-scale paired text-image corpus and then fine-tuned on seven VL tasks. Our results show that the new OD model can substantially uplift the SoTA results on all seven VL tasks across multiple public benchmarks. Our ablation study shows that the improvement is mainly attributed to our design choices regarding diversity of object categories, visual attribute training, training data scale, model size, and model architecture.

Appendix to VinVL

3.8 OSCAR+ pre-training

3.8.1 Pre-training Corpus

Small	0.22M Images, 2.5M QAs, 0.7M captions							
Medium	1.89M Images, 2.5M QAs, 0.7M captions, 1.67M pseudo-captions							
Large	5.65M Images, 2.5M QAs, 4.68M captions, 1.67M pseudo-captions							
Source	VQA (train)	GQA (bal-train)	VG-QA (train)	COCO (train)	Flicker30k (train)	OpenImages (od train)	CC (train)	SBU (all)
Image/Text	83k/545k	79k/1026k	87k/931k	112k/559k	29k/145k	1.67M/1.67M	3.1M/3.1M	875k/875k
$\mathbf{w}, \mathbf{q}, \mathbf{v}$	Question, Answer, ImageFeatures			(Generated) Caption, (Generated) ImageTags, ImageFeatures				

Table 3.17: Statistics of the pre-training corpus.

Table 3.17 shows the statistics of image and text of the pre-training corpora. In our ablation study, we use corpora of three different sizes: ‘Small’, ‘Medium’, ‘Large’. Different from OSCAR [LYL⁺20], we make use of image tagging datasets OpenImages, by generating captions using OSCAR’s image captioning model to form triplets of (generated caption, image tags, image features) for OSCAR+ pre-training. By self-training technique, our pre-training corpora can be scaled to a much larger amount by making use of large-scale image tagging datasets, e.g., OpenImages (9M) and YFCC (92M).

3.8.2 OSCAR+ pre-training objectives

Masked Token Loss: A Loss Mimics Image Captioning. The word tokens of image captions (questions) $\mathbf{w}$ and word tokens of object tags (answers) $\mathbf{q}$ share the same linguistic semantic space, and the Masked Token Loss (MTL) is applied on tokens of both $\mathbf{w}$ and $\mathbf{q}$. We define the *discrete*

token sequence as $\mathbf{h} \triangleq [\mathbf{w}, \mathbf{q}]$, and apply the Masked Token Loss (MTL) for pre-training. At each iteration, we randomly mask each input token in $\mathbf{h}$ with probability 15%, and replace the masked one h_i with a special token [MASK]. The goal of training is to predict these masked tokens based on their surrounding tokens $\mathbf{h}_{\setminus i}$ and image features $\mathbf{v}$ by minimizing the negative log-likelihood:

$$\mathcal{L}_{\text{MTL}} = -\mathbb{E}_{(\mathbf{v}, \mathbf{h}) \sim \mathcal{D}} \log p(h_i | \mathbf{h}_{\setminus i}, \mathbf{v}) \quad (3.5)$$

This is the same MTL as in OSCAR [LYL⁺20] and similar to the masked language model used by BERT. The masked word or tag needs to be recovered from its surrounding context, with additional image information to help ground the learned word embeddings in the vision context.

3-way Contrastive Loss: A Loss Mimics Text-Image Retrieval and Visual Question Answering Simultaneously. We present our 3-way contrastive loss in Section 3.3.2 in the main paper.

3.8.3 Ablation of the two new techniques

Effect of self-training: Leveraging Image Tagging data. In Figure 3.4, we show the effect of self-training by making use of tagging data in OSCAR+, by fine-tuning OSCAR+ pre-training checkpoints on VQA. Compared with “OSCAR+, Small; VinVL” (green), “OSCAR+, Medium; VinVL” (yellow) adds the 1.7M OpenImages Tagging data into pre-training and its performance gets improved significantly, demonstrating the effect of self-training by making use of tagging data. As baselines, we also provide performance of OSCAR and OSCAR+ with image features from [AHB⁺18], which clearly demonstrates that the new image features pre-trained by VinVL matter significantly in the VL pre-training and VL downstream tasks.




Figure 3.4: Effect of OSCAR+ pre-training corpus size and effect of self-training by making use of tagging data in OSCAR+. Each curve, with legend “VLP, Corpus; VisionFeature”, denotes a VLP experiment where the VLP method is either OSCAR or OSCAR+, the VLP pre-training Corpus is Small/Medium/Large (defined in Table 2.1), and VisionFeature is either our new vision features (VinVL for short) or those from [AHB⁺18] ([AHB⁺18] for short). X-axis denotes the pre-training iterations of OSCAR+ checkpoints. Y-axis is the vqa-dev accuracy of a VQA model initialized from the corresponding pre-training checkpoint and fine-tuned with a fixed scheme. Compared with “OSCAR+, Small; VinVL” (green), “OSCAR+, Medium; VinVL” (yellow) adds the 1.7M OpenImages Tagging data into the pre-training and its performance gets improved significantly, demonstrating the effect of self-training by making use of tagging data. The “OSCAR+, Large; VinVL” (blue) further scales up the pre-training corpus by adding Google Conceptual Captions and SBU datasets with generated tags and its performance gets further improved, demonstrating the effect of OSCAR+ pre-training corpus size. As baselines, we also provide performance of OSCAR and OSCAR+ with image features from [AHB⁺18], which clearly demonstrates that our new image features (VinVL) matter significantly in the VL pre-training and VL downstream tasks.

Effect of the new 3-way contrastive loss. As illustrated in Table 3.3, with the new 3-way contrastive loss, the VQA performance is the same as the OSCAR pre-training, while the Text-Image Retrieval performance improves significantly compared with the OSCAR pre-training.

Overall improvement from OSCAR to OSCAR+. We point out that the improvement from OSCAR to OSCAR+ with image features from [AHB⁺18] is minor, because (1) we only add 1.7M OpenImages’ tagging data to enlarge the pre-training corpus, which is a small portion compared with OSCAR’s original pre-training corpus (i.e., Large\OI, 3.98M images and 7.18M image-caption pairs), and (2) the new 3-way contrastive loss has more significant improvements in Text-Image Retrieval tasks than that in the VQA task, as illustrated in Table 3.3. We would expect much more significant improvements when we scale up the OSCAR+’s pre-training corpus to a much larger scale by adding large scale image tagging datasets, e.g., OpenImages (9M) and YFCC (92M).

3.9 Downstream Tasks Fine-tuning

We follow the downstream task fine-tuning recipes in OSCAR [LYL⁺20].

3.9.1 VQA

Given an image and a question, the task is to select the correct answer from a multi-choice list, it requires the model to answer natural language questions based on an image. Here we conduct experiments on the widely-used VQA v2.0 dataset [GKSS⁺17], which is built on the MSCOCO [LMB⁺14] images. Following [AHB⁺18], for each question, the model picks the corresponding answer from a shared set of 3,129 candidates.

When fine-tuning on the VQA task, the input sequence contains the concatenation of a given question, object tags and object region features, and then the [CLS] output from OSCAR+ is fed to

a task-specific linear classifier for answer prediction. Similarly as the literature [AHB⁺18], we treat VQA as a multi-label classification problem – assigning a soft target score to each answer based on its relevancy to the human answer responses, and then we fine-tune the model by minimizing the cross-entropy loss computed using the predicted scores and the soft target scores. During inference, we simply use Softmax for answer prediction.

For VQA training, we random sample a set of 2k images from the MS COCO validation set as our validation set, the rest of images in the training and validation are used in the VQA fine-tuning. For the OSCAR_B model, we fine-tune for 25 epochs with a learning rate of $5e^{-5}$ and a batch size of 128. For the OSCAR_L model, we fine-tune for 25 epochs with a learning rate of $3e^{-5}$ and a batch size of 96.

3.9.2 GQA

Similarly as VQA, GQA tests the reasoning capability of the model to answer a question. We conduct experiments on the public GQA dataset [HM19b]. For each question, the model chooses an answer from a shared set of 1,852 candidates. Our fine-tuning procedure is following Oscar [LYL⁺20, CGL⁺19], which first fine-tunes the model on unbalanced “all-split” for 5 epochs with a learning rate of $5e^{-5}$ and a batch size of 128, and then fine-tuned on the “balanced-split” for 2 epochs.

3.9.3 Image Captioning

An image captioning model generates a natural language description for a given image. To enable sentence generation, we fine-tune OSCAR⁺ using the seq2seq objective. The input samples are processed to triples consisting of image region features, captions, and object tags, in the same way as that during the pre-training. We randomly mask out 15% of the caption tokens and use the corresponding output representations to perform classification to predict the token ids. Similar

to previous works [LYL⁺20, ZPZ⁺20], the self-attention mask is constrained such that a caption token can only attend to the tokens before its position to simulate a uni-directional generation process. Note that all caption tokens will have full attentions to image regions and object tags but not the other way around.

During inference, we first encode the image regions, object tags, and a special token [CLS] as input. Then the model starts the generation by feeding in a [MASK] token and selecting a token from the vocabulary based on the likelihood output. Next, the [MASK] token in the previous input sequence is replaced with the selected token and a new [MASK] is appended for the next word prediction. The generation process terminates when the model outputs the [SEP] token. We use beam search (*i.e.*, beam size = 5) [AHB⁺18] in our experiments and report our results on the COCO image captioning dataset.

Though the training objective (*i.e.*, seq2seq) for image captioning is different from that used in pre-training (*i.e.*, bidirectional attention-based masked token loss), we directly fine-tune OSCAR+ for image captioning on COCO without additional pre-training on Conceptual Captions [SDGS18]. This is to validate the generalization ability of the OSCAR+ models for generation tasks. We use the same Karpathy split [KFF15]. For the OSCAR+B model, we fine-tune with cross-entropy loss for 30 epochs with a batch size of 256 and an initial learning rate of $1e^{-5}$ and then with CIDEr optimization [RMM⁺17] for 10 epochs with a batch size of 128 and initial learning rate of $2e^{-6}$. We compare with several existing methods, including BUTD [AHB⁺18], VLP [ZPZ⁺20], AoANet [HWCW19], OSCAR [LYL⁺20].

3.9.4 NoCaps

Novel Object Captioning [ADW⁺19] extends the image captioning task, is to test models' capability of describing novel objects from the Open Images dataset [KRA⁺18] which are not seen in the

Method	in-domain		near-domain		out-of-domain		overall		in-domain		near-domain		out-of-domain		overall	
	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE
Validation Set								Test Set								
UpDown ⁺	79.3	12.4	73.8	11.4	71.7	9.9	74.3	11.2	76.0	11.8	74.2	11.5	66.7	9.7	73.1	11.2
Oscar _B [*]	83.4	12.0	81.6	12.0	77.6	10.6	81.1	11.7	81.3	11.9	79.6	11.9	73.6	10.6	78.8	11.7
Oscar _L [*]	85.4	11.9	84.0	11.7	80.3	10.0	83.4	11.4	84.8	12.1	82.1	11.5	73.8	9.7	80.9	11.3
Human [ADW ⁺ 19]	84.4	14.3	85.0	14.3	95.7	14.0	87.1	14.2	80.6	15.0	84.6	14.7	91.6	14.2	85.3	14.6
VIVO ⁺ [HYL ⁺ 20]	92.2	12.9	87.8	12.6	87.5	11.5	88.3	12.4	89.0	12.9	87.8	12.6	80.1	11.1	86.6	12.4
VinVL [*]	96.8	13.5	90.7	13.1	87.4	11.6	90.9	12.8	93.8	13.3	89.0	12.8	66.1	10.9	85.5	12.5
VinVL+VIVO	103.7	13.7	95.6	13.4	83.8	11.9	94.3	13.1	98.0	13.6	95.2	13.4	78.0	11.5	92.5	13.1

Table 3.18: NoCaps evaluation results. All the models are trained on COCO without additional image-caption pairs following the restriction of NoCaps. (UpDown⁺ is UpDown+ELMo+CBS, the models with ^{*} is +SCST+CBS, VinVL+VIVO is with SCST only.)

training corpus. Following the restriction guideline of NoCaps, we train OSCAR+ on COCO without the initialization from pre-training, so no additional image-text pairs are used for training except COCO.

Since NoCaps images are collected from Open Images, we train an object detector using the Open Images training set and apply it to generate the tags. We conduct experiments from BERT model directly without pre-training as required by the task guidelines. For the OSCAR_B model, we train 30 epochs with a batch size of 256 and learning rate $1e^{-4}$; further we perform CIDEr optimization with learning rate $5e^{-6}$ and batch size 112 for 10 epochs. During inference, we use constrained beam search for decoding. We compare OSCAR+ with OSCAR [LYL⁺20] on this task.

3.9.5 Image-Text Retrieval

There are two sub-tasks: *image retrieval* and *text retrieval*, depending on which modality is used as the retrieved target. Both tasks calculate a similarity score between an image and a sentence, which heavily relies on the cross-modal representations.

Following Oscar [LYL⁺20], we formulate the retrieval as a binary classification problem, where given an aligned image-text pair, we randomly select a different image or a different sentence to form an unaligned pair. The final representation of [CLS] is used as the input to the classifier to

predict whether the given pair is aligned or not. In the testing stage, the probability score is used to rank the given image-text pairs of a query.

Following [LDF⁺19], we report the top- K retrieval results on both the 1K and 5K COCO test sets. We adopt the widely used Karpathy split [KFF15] on the COCO caption dataset [LMB⁺14] to conduct our experiments. Specifically, the dataset consists of 113,287 images for training, 5,000 images for validation, and 5,000 images for testing. Each image is associated with 5 human-generated captions. For the OSCAR_B model, we fine-tune with a batch size of 256 for 40 epochs. The initial learning rate is set to $2e^{-5}$ and linearly decreases. For the OSCAR_L model, we fine-tune with a batch size of 128 for 40 epochs. The initial learning rate is set to $1e^{-5}$ and linearly decreases. We use the validation set for parameter tuning. We compare with several existing methods, including DVSA [KFF15], VSE++ [FFKF17], DPC [ZZG⁺17], CAMP [WLL⁺19], SCAN [LCH⁺18], SCG [SJL⁺19], PFAN [WYQ⁺19], Unicoder-VL [LDF⁺19], 12-in-1 [LGR⁺19], UNITER [CLY⁺19].

3.9.6 NLVR2

Given a pair of images and a natural language, the goal of NLVR2 [SZZ⁺18] is to determine whether the natural language statement is true about the image pair. For NLVR2 fine-tuning, we first construct two input sequences, each containing the concatenation of the given sentence (the natural language description) and one image, and then two [CLS] outputs from OSCAR+ are concatenated as the joint input for a binary classifier, implemented by an MLP.

For the OSCAR_B model, we fine-tune for 20 epochs with learning rate $\{2e^{-5}, 3e^{-5}, 5e^{-5}\}$ and a batch size of 72. For the OSCAR_L model, we fine-tune for 20 epochs with learning rate of $\{2e^{-5}, 3e^{-5}\}$ and a batch size of 48.

Mean	63.85	64.33	64.27	65.13	66.15	66.25	66.44	65.18
4Sets->VG	64.59	65.05	64.59	66.22	66.44	67.26	68.39	66.00
VG	64.91	65.95	65.24	66.68	67.40	68.17	68.13	66.53
VGw/oAttr	65.04	65.64	64.95	66.61	67.57	67.28	67.46	66.29
VG1600/400	64.32	65.25	64.42	65.85	67.63	67.19	67.62	65.93
Grid-273	65.11	63.44	66.17	66.09	67.29	66.85	66.84	65.87
VG-obj	63.55	64.62	64.30	65.17	66.20	66.44	66.76	65.15
GT-ObjStuff	63.95	64.54	64.34	64.89	66.17	66.22	66.59	65.11
OI	63.41	64.14	63.31	64.15	65.38	65.38	65.38	64.35
Grid-50	62.55	63.54	64.02	63.92	64.68	64.72	64.76	63.91
GT-Obj	62.53	62.52	62.43	62.46	63.69	63.74	63.63	62.88
	OI (O:500)	VG-obj (O:317)	ImageNet (O:1000)	VG w/o Attr (O:1594)	VG (O:1600, A:400)	VG (O:1594, A:524)	4Sets->VG (O:1594, A:524)	Mean

Figure 3.5: Overall comparison of vocabulary effect on VQA. X-axis: how the R50-C4 model is trained; Y-axis: how the feature is extracted (grid or region features, different kinds of boxes to extract region features). All region features have maximal 50 regions. The top row “Mean” is the average over all rows, showing the overall quality of different vision models. The far-right column “Mean” is the average over all columns, showing the overall quality of different feature extraction methods.

3.10 Effect of the Object-Attribute Vocabulary Size

In Section 3.4.2, we demonstrate that the more diverse the visual concepts (object and attribute vocabularies) are, the better the visual region features for VL tasks. The better performance may come from the more diverse proposed regions where the region features are extracted (see the comparison in Figure 3.1, “region” for short), or from the better model weights that can produce better high-dimensional region representation even for the same region (“model” for short). In this section, we disentangle effects of region proposals and model weights, by performing synthetic experiments in which we use region proposals from one vision model and model weights from another vision model. Our results show that both the region proposals and model weights matter




Figure 3.6: Left: comparison of object vocab and attribute vocab, average over all types of bounding boxes. Right: comparison of feature extraction methods, average over all types of pre-trained vision models. X-axis is the number of iterations when we take the checkpoint for evaluation. Y-axis is the VQA accuracy on our vqa-dev.

for VL tasks.

3.10.1 Disentangling the effects of region proposals and model weights on R50-C4

As in Section 3.4.2, We train vision models $v = \text{Vision}(Img)$ on different datasets, i.e., OpenImages with 500 object classes (OI:O500), standard ImageNet with 1K classes (ImageNet:O1000), Visual Genome with 317 object classes (VG-obj), Visual Genome with 1594 object classes (VG:O1594), VG with 1594 object classes and 524 attribute classes (VG:O1594A524), pretrain on the merged 4 datasets and finetune on VG:O1594A524 (4Sets $\rightarrow$ VG:O1594A524). For *each* model, we also try different ways to extract features: (1) *region* features from different models' proposed regions (same notations with models) where each image has maximal 50 region features, and (2) *grid*

features where we use all grid features (Grid-273) or randomly sampled 50 grid features (Grid-50) for each image. We present the results of these model-region cross-combination experiments in Figure 3.5. We also present the mean accuracy over all box types to obtain a robust ranking of different checkpoints and the mean accuracy over all checkpoints to obtain a robust ranking of different box types. We have the following observations:

- The richer the object vocabulary is, the better for VQA: $\text{OI:500} \approx \text{VG-obj:O317} < \text{ImageNet:O1000} < \text{VG:O1594}$.
- Attribute information is crucial to VL tasks: all features trained with attributes (Columns with VG:O1594A524) are significantly better than those without attributes.
- Even for small vision backbone R50, vision pre-training makes vision features better: Column “4Sets→VG:O1594A524” are better than all other columns. Notice that the vision pre-training improves both the region features and the grid features.
- It is crucial to extract features from semantically diverse regions: regions from OI and VG-obj are significantly worse than all other regions, and is even worse than grid features.
- Grid features perform worse than region features with regions proposed by VG models. By comparing Row “Grid-273” with rows with VG regions, it seems hopeful to close this gap while paying more hardware memory and computational cost in cross-modal models VL. It is three times slower to **train** the “Grid-273” models than training models with region features.

In Figure 3.6, instead of just showing one final number, we provide the *mean evaluation curves along training trajectories* to demonstrate the ranking, as an even more robust evidence. These results further confirm the conclusions we draw in Section 3.4.2.

3.10.2 Disentangling the effects of region proposals and model weights on the SoTA model

In Table 3.19, we alternate the combination of region proposals and model weights, and evaluate them on VQA. As we can see, the improvement of using boxes from the R101-C4 model [AHB⁺18] to extract features from our X152-C4 model is much bigger than that of using boxes from our X152-C4 model to extract features from the R101-C4 model [AHB⁺18], indicating pre-trained model weights are more important than regions. Inspired by this analysis, we propose the class-agnostic NMS for region selection in the box head of the OD model, which does not sacrifice any VQA performance but greatly improves the model’s inference speed. This analysis also suggests that large-scale OD pre-training should improve performance for grid-feature based VL models, as supported by more results in Appendix 3.12.

In Table 3.19, We also report VQA results with COCO groundtruth object regions (GT-Obj, 80 classes) and object-stuff regions (GT-Obj&Stuff, 171 classes). For VQA task, COCO GT boxes are much worse than proposals from VG trained models. This shows the difference between typical OD tasks and OD in VL: OD in VL requires much richer visual semantics to align with the rich semantics in the language modality. This further echoes with our claim that an image understanding module trained with rich semantics is crucial for VL tasks.

Model	GT-Obj	GT-Obj&Stuff	Anderson et al. [2]	VinVL (ours)
Anderson et al. [2]	63.81 $\pm$ 0.94	66.68 $\pm$ 0.16	68.52 $\pm$ 0.11	69.05 $\pm$ 0.06
VinVL (ours)	65.60 $\pm$ 0.21	68.13 $\pm$ 0.26	70.25 $\pm$ 0.05	71.34 $\pm$ 0.17

Table 3.19: Ablation of region and model on VQA.

	no image feature w	R50-C4 w/ box head randomly initialized	R50-FPN	R50-C4	4Sets→ R50-FPN	4Sets→R50-C4
VG-trained	–	67.6 $\pm$ 0.13	67.6 $\pm$ 0.30	68.0 $\pm$ 0.16	68.3 $\pm$ 0.11	68.2 $\pm$ 0.05
Initial	55.5 $\pm$ 0.50	61.8 $\pm$ 0.47	57.6 $\pm$ 0.16	64.8 $\pm$ 0.44	66.1 $\pm$ 0.23	66.8 $\pm$ 0.21

Table 3.20: C4 vs FPN architecture on VQA. Boxes used to extract features v and tags q used in VL model are the same with those used in OSCAR [LYL⁺20]. Row “Initial” means using the initialization model without VG training for feature extraction.

3.11 More on FPN and Comparison of C4 and FPN

3.11.1 Two reasons why FPN performs worse than C4 on VL tasks.

Our experimental results confirm the conclusion of [JMR⁺20] that the FPN model does not provide better region features for VL tasks than the C4 model (Columns “R50C4” vs. “R50FPN” in Table 3.20). Our analysis reveals two reasons. First of all, all layers involved in feature extraction in the C4 model have been pre-trained using ImageNet while the MLP head of FPN does not. It turns out that the VG dataset is still small to train a good visual features for VL tasks and using ImageNet-pre-trained weights is beneficial. This can be verified by two experiments: (1) When the R50-C4 model is trained on VG with its box head randomly initialized (VG-trained - R50C4 w/ box head randomly initialized), the C4 model’s performance is the same as FPN; and (2) C4 and FPN achieve the same performance after vision pre-training on 4 datasets (68.3 vs. 68.2). The second reason is due the network architecture (CNN vs. MLP) of the box head in the OD model. The convolutional head in C4 has a better inductive bias in encoding visual information than the MLP head in FPN. This can be verified by the fact that when vision features from randomly initialized models are used (Row “Initial” in Table 3.20), R50-C4 performs much better than R50-FPN, indicating that the initial C4 features encode much more useful visual information than the initial FPN features. The “random” C4 features nearly match the feature from ImageNet pre-trained model (Row “Initial” Column “R50C4”), while “random” FPN features are close to the performance without visual

features as input (Row “Initial” Column “no image feature $\mathbf{w}$ ”).

3.11.2 Effect of pooling methods in FPN on VQA performance.

Different from C4 models that extract region features from a single scale (the end of C4 block), FPN models extract region features from multiple scales adaptively based on the area of the region. Therefore, there is some in-homogeneity in FPN’s region features since they may come from different scales. In Figure 3.7, we show that this is not the cause of FPN’s worse performance than C4 on the VQA task. More specifically, we experiment with 4 pooling methods for FPN architecture. (1) adapt: the original FPN’s pooling method that extract features adaptively from different scales; (2) max: extract features from all scales and then do a max-pool; (3) avg: extract features from all scales and then do an average-pool; (4) concat: extract features from all scales and then concatenate them together. We also train multiple FPN models on VG with these pooling methods, with or without pre-training on the Objects365 dataset. We experiment on all possible combinations (in total 8×4) of 8 vision models and 4 pooling methods on the VQA task. When there is a parameter dimension mis-match, e.g., non-concat FPN models but use concat pooling methods in VQA and vice versa, we specify those parameter randomly with PyTorch’s default initialization method. The results in Figure 3.7 shows that (1) there is no obvious difference in different pooling methods, with the default “adapt” and the “concat” methods perform slightly better than “max” and “avg”; (2) (without surprise) the performance is significantly worse when there is a parameter dimension mismatch between vision models and VL task feature extraction methods, i.e., non-concat FPN models but use concat pooling methods in VQA and vice versa. These results show that the pooling method (no matter in vision model training or in VL task feature extraction) is not the root cause of FPN’s worse performance than C4 on the VQA task.

o365->vg-max	65.96	67.00	65.20	62.68	o365->vg-max	67.12	66.27	65.71	63.63
o365->vg-avg	66.58	66.96	66.88	61.03	o365->vg-avg	66.55	66.62	66.15	63.62
o365->vg-adapt	67.10	66.63	67.43	63.96	o365->vg-adapt	66.57	67.01	67.60	63.52
o365->vg-concat	61.14	62.17	59.92	66.90	o365->vg-concat	60.98	62.70	60.25	66.87
vg-max	66.64	66.76	64.55	64.22	vg-max	66.94	65.76	64.64	62.89
vg-avg	66.14	66.88	67.06	62.60	vg-avg	67.00	67.06	66.74	63.77
vg-adapt	66.10	66.74	66.81	63.71	vg-adapt	66.76	66.38	66.29	63.77
vg-concat	61.03	64.34	60.64	67.19	vg-concat	61.71	62.64	59.91	67.18
	max	avg	adapt	concat		max	avg	adapt	concat

Figure 3.7: Pooling methods in FPN feature extraction are not the root cause of FPN’s worse performance than C4. X-axis: the pooling method when extracting features for VL tasks; Y-axis: the pooling method (vision model) when pre-training the visual feature extraction model. All experiments are using regions from the Bottom-up Top-down model [AHB⁺18]. Each combination is experimented twice with two random seeds, i.e. seed=42 on the left and seed=88 on the right. The results from two random seeds are consistent.

	ImageNet-5k [WKM ⁺ 19]	4Sets	VG with Attr	4Sets→VG
grid feature (273)	68.3±0.29	65.2±2.47	67.5±0.20	69.4 [*]
region feature (50)	67.7±0.16	68.5±0.13	69.8±0.23	70.6±0.13

^{*} The other run failed and thus there is no std for this experiment.

Table 3.21: Ablation study of X152 models on VQA. Vision models in the last three columns are trained with initialization from the ImageNet-5k checkpoint in the first column. All the region features are extracted with boxes proposed by our best X152-C4 model (pre-trained on 4Sets and fine-tuned on VG). By comparing the first column and the last column, we see that our proposed vision pre-training (first on 4 sets and then on VG with attributes) improves performance for both the grid-feature based model and the region-feature based model. Since the X152 backbone is much larger than the R50 backbone in Figure 3.5, the larger model can make better use of the large pre-training datasets and thus have more significant improvements.

3.11.3 Large-scale object-detection pre-training of C4 and FPN models

In this paper, we have trained R50-C4, R50-FPN, R152-C4 and R152-FPN models on the merged object detection datasets described in Table 3.2. In Figure 3.8, we report the mAP^{50} of checkpoints from these 4 experiments on 4 validation sets: COCO with stuff (top left), Objects365 (top right),




Figure 3.8: Checkpoints’ mAP^{50} on 4 validation sets: COCO with stuff (top left), Objects365 (top right), OpenImages (bottom left) and Visual Genome (1594 object classes, bottom right). For R50 models, the R50-FPN model is slightly better than C4 on COCO and Objects365 but slightly worse than C4 on Visual Genome. For R152 models, the R152-FPN model is consistently worse than the R152-C4 model on all 4 different datasets.

OpenImages (bottom left) and Visual Genome (1594 object classes, bottom right). For R50 models, the R50-FPN model is slightly better than C4 on COCO and Objects365 but slightly worse than C4 on Visual Genome. For R152 models, the R152-FPN model is consistently worse than the R152-C4 model on all 4 different datasets. Therefore, we finally use the R152-C4 model for downstream vision-language tasks.

3.12 Grid feature

In Table 3.21, we train grid-feature based and region-feature based X152 models for VQA, with the vision models pre-trained on different vision datasets, i.e., “ImageNet-5k” from [WKM⁺19], our

4-dataset merged OD dataset 3.2 (4Sets), our VG dataset with 1594 object classes and 524 attribute classes (VG with Attr), and first 4Sets and then VG (4Sets→VG). Vision models in the last three cases are trained with initialization from the same ImageNet-5k checkpoint from [WKM⁺19]. All the region features are extracted with boxes proposed by our best X152-C4 model (pre-trained on 4Sets and fine-tuned on VG). By comparing “ImageNet-5k” and “4Sets→VG”, we see that our proposed vision pre-training improves performance for both the grid-feature based model and the region-feature based model. Since the X152 backbone is much larger than the R50 backbone in Figure 3.5, the larger model makes better use of the large pre-training datasets and thus has more significant improvements. It is interesting to see that for grid-feature based models, the “ImageNet-5k” model performs better than the “4Sets” model and the “VG with Attr”, while it is not the case for region-feature based models. This may indicate that how the vision model is trained (grid-feature wise or region-feature wise) may have big impact on the downstream VL tasks.

Model	R50-C4		R101-C4 [AHB ⁺ 18]		X152-C4	
	Vision	VL	Vision	VL	Vision	VL
Grid-50	0.059±0.018	0.029±0.002	0.083±0.025	0.030±0.003	0.355±0.022	0.031±0.003
Grid-273	0.056±0.005	0.027±0.002	0.082±0.022	0.034±0.001	0.344±0.036	0.037±0.004
Object	0.373±0.040	0.031±0.005	0.663±0.042	0.034±0.003	0.687±0.064	0.036±0.005
Object-eff	0.165±0.029	0.029±0.002	0.442±0.119	0.036±0.003	0.475±0.049	0.037±0.005
Grid-50 (cpu)	1.943±0.244	0.480±0.042	4.050±0.398	0.469±0.046	17.765±1.693	0.501±0.047
Grid-273 (cpu)	2.032±0.230	1.368±0.056	4.052±0.372	1.283±0.067	17.664±1.713	1.326±0.053
Object (cpu)	11.808±1.322	0.500±0.045	31.863±7.932	0.585±0.044	29.641±3.097	0.565±0.044
Object-eff (cpu)	11.729±1.280	0.510±0.044	31.791±8.027	0.587±0.043	29.687±3.011	0.574±0.036

Table 3.22: Time cost of end-to-end inference on VQA. All cross-modal models are BERT-Base. On the SOTA number obtained with X152-C4 region features, the performance *keeps the same* when changing to the efficient way to extract the feature while the efficiency greatly improves on GPU. The efficient version does not lead to time saving on CPU, because nearly all inference time is taken by the backbone and C4 head and the time from NMS operations is nearly ignorable on CPU.

3.13 End-to-end inference efficiency

We report the end-to-end inference time of different VQA models on a Titan-X GPU and a Xeon E5 CPU in Table 3.22. For CPU evaluation, we force that the inference use only one CPU thread. The input image size is 800×1333 , and we run the inference with batch size 1 (one image-question pair per batch). We can see that (1) vision models dominate the inference time, especially for large models; (2) models based on grid-feature are faster than those based on region feature; (3) with our proposed fast inference trick, region-feature models are greatly sped up and their inference time can be brought to within 3 times of that of grid-feature models on GPU. We find that on CPU with a single thread, our class-agnostic trick does not lead to time saving, because nearly all inference time is taken by the backbone and C4 head and the time from NMS operations is nearly ignorable on CPU.

Chapter 4

VIM-MLLM

Recent multimodal large language models (MLLMs) have shown promising instruction following capabilities on multimodal tasks. However, the existing MLLMs are built on top of large language models (LLMs), inherit its language priors for instruction understanding. In this chapter, we introduce VISUAL MODALITY INSTRUCTION (VIM)¹, challenging the MLLMs by embedding the instructions into the visual pixel space, which demands for strong visual interpretative skills for instruction following. We adapt VIM to eight benchmarks, including OKVQA, MM-Vet, MathVista, MMMU, *etc.*, and probe diverse MLLMs in both the text-modality instruction (TEM) setting and VIM setting. We observe a significant performance disparity between the original TEM and VIM settings for open-source MLLMs, implying that their proficiency in visual instruction comprehension is not up to par. Furthermore, we train v-MLLM², a generalizable model that is capable to conduct robust instruction following in both text-modality and visual-modality instructions.³

¹VIM is short for **V**isual **m**odality instruction.

²v-MLLM is short for VIM-MLLM.

³Here, **robust** is to instruction modality, i.e., text-modality instruction and visual-modality instruction in the paper.


4.1 Introduction

In recent years, the triumph of Large Language Models (LLMs) such as GPT series [BMR⁺20, Ope23a], PaLM [ADF⁺23] and other open-source LLMs [TLI⁺23, TMS⁺23] showcased the power of large-scale pretraining in understanding, reasoning and generating intricate textual patterns. Their instruction-tuned counterparts, InstructGPT [OWJ⁺22], Flan-PaLM [CHL⁺22], Alpaca [TGZ⁺23], Vicuna [CLL⁺23], GPT4-LLM [PLH⁺23], further demonstrated the power of aligned LLMs in following human instructions.


Inspired by the success of LLMs, especially after the release of GPT-4, researchers have ventured into the multimodal domain with a similar recipe for Multimodal Large Language Models (MLLMs). A thread of research [LLWL23, DLL⁺23, AGG⁺23, YXX⁺23, ZHZ⁺23, SLL⁺23] extends the open-source LLMs with *visual* instruction tuning to connect vision and LLMs for visual understanding, reasoning and generation. Concurrently, a number of benchmarks [FCS⁺23, YYL⁺23, LDZ⁺23, LWW⁺23, BBH⁺23, WZZ⁺23] have been proposed to assess the capability of the MLLMs. And recently, OpenAI releases its GPT-4V(ision) [Ope23c] via the ChatGPT interface (on September 2023) and API (on November 2023), and Google launches its Gemini (on December 2023), these proprietary MLLMs attract broad attention over the internet. A comprehensive report [YLL⁺23] showcases its exceptional perception, reasoning and generation capability across a variety of real-world scenarios.⁴

Despite the impressive performance of MLLMs on a wide variety of benchmarks [FCS⁺23, YYL⁺23, LDZ⁺23, LWW⁺23, BBH⁺23, WZZ⁺23], they still remain poorly understood, and also these MLLMs are not as impressive as their LLM counterparts, let alone their landing business applications. The current MLLMs are built on top of the pretrained LLMs, and *visual* instruction tuning follows the recipe from its LLM counterparts, specifically, the instruction data is synthesized

⁴There are two terms widely used on this topic, MLLMs and LMMs. In this work, we use MLLMs for consistency.



Text-based Probing Paradigm



Our VIM Paradigm




Figure 4.1: Model performance on the original TEM vs. VIM settings. Each data point corresponds to one model’s number on one benchmark, we probe eight MLLMs on eight benchmarks, in total, there are 64 data points. This plot reveals: 1). There is a significant drop for open-source MLLMs from original TEM to VIM setting; 2). GPT-4V, Gemini Pro and our v-MLLM exhibit robust instruction following capability, as their data points are consistently close to the diagonal line.




Figure 4.2: Evaluation paradigm comparison for MLLMs. (a) Left is TEM setting, where Image + Text instruction as two separate modalities are fed into MLLMs for inference; an LLM model (for example, Vicuna) can also make correct prediction, even without accessing to the image. (b) Right: VIM **only** takes the image modality with the text instruction embedded in the image, no additional text prompt is required, LLMs are not applicable. The above example is from OKVQA (question #209725). Note: Image modality input, Text modality input.

by the LLMs (mostly from GPT-4 or GPT-4V) in the text format. As illustrated in the left part of Figure 4.2, the instruction and image are expressed in two modalities, we denote this kind of *visual* instruction data as Text-Modality Instruction (TEM). Under this setting, a pure LLM, for example, Llama 2 or Vicuna in Figure 4.2 can still make plausible or correct prediction, even without accessing to the image input. All the current benchmarks [FCS⁺23, YYL⁺23, LDZ⁺23, LWW⁺23, BBH⁺23] follow the same format.

This raises a question - *how proficiently these MLLMs can follow instructions if we embed the text instruction into visual format?* As shown in the right part of Figure 4.2, we name it as VISUAL MODALITY INSTRUCTION, where the image and instruction are in the visual modality. Under the VIM setting, LLMs are not applicable, and LLaVA-1.5 simply repeats the question for the image, may not understand the visual-modality instruction.




Figure 4.3: Performance comparison of six MLLMs for visual instruction following between text-modality instruction (TEM —) and our introduced VISUAL MODALITY INSTRUCTION (VIM —) settings on eight benchmarks. There exists a disparity between TEM and VIM settings for all open-source MLLMs (the first row); GPT-4V, Gemini Pro and our v-MLLM are robust to instruction modality.

Motivated by this, we introduce a new setting, called VISUAL MODALITY INSTRUCTION (short for VIM), evaluating the capability of MLLMs for visual-modality instruction following. We adapt VIM to various benchmarks [MRFM19, FCS⁺23, YYL⁺23, LBX⁺23, YNZ⁺23], and compose a new benchmark - VIM-Bench. As highlighted in Figure 4.1 and 4.3, there exists a performance disparity between the TEM and VIM settings for all open-source MLLMs, all of them are not robust enough at visual-modality instruction following. To summarize, our main contributions are:

- We propose VISUAL MODALITY INSTRUCTION, a challenging setting to probe the capability of Multimodal Large Language Models for visual-modality instruction following.

- We adapt the VIM to various benchmarks, and reveal a significant disparity between their text-modality instruction setting and VIM setting.
- We train a v-MLLM, which demonstrates robust visual instruction following capabilities.

4.2 Method

Instruction following, is viewed as one key capability of high-performing MLLMs. In this section, we first introduce VIM, to examine the instruction following capability of MLLMs, specifically the visual-modality instruction following. Then, we introduce v-MLLM, enhancing the MLLMs with visual-modality instruction following.

4.2.1 VIM

Visual-Modality Instruction

As illustrated in the left part of Figure 4.2, the current evaluation norm of MLLMs takes two modalities as input: image and text (as instruction). The existing MLLMs are built on top of the LLMs, benefiting from its strong text understanding capability. For the current MLLM evaluation paradigm, instruction is presented in the text modality, which can utilize the strong language priors from the LLMs for understanding. As shown in Table 4.3, even a pure LLM model (GPT-4, Llama 2 or Vicuna) can get some success without accessing to the images. Interestingly, on most of eight tasks, Llama 2 shows better numbers over GPT-4 (gpt-4-1106-preview). We manually check some response, and find that the output from GPT-4 are more reasonable than Llama 2. This might rise several interesting issues, we leave a discussion in Section 4.5.

VIM challenges the MLLMs by embedding the textual instruction into the visual pixel space (image), this enhancement demands not just textual but also strong visual comprehension for

instruction understanding. It asks for the strong visual interpretation capability to recognize and follow the embedded instruction in the image.

Design Choices

How to lay out the image and embedded instruction in one visual space? For zero-shot setting, there are many combinations to position the instruction and image in the same visual space. Here, we enumerate two important elements we investigated.

Instruction Location Potentially, there are many options to place the instruction into the image. To narrow down the search choices of instruction’s placement, we focus on three primary positions: {top, right and bottom}. Additionally, we add a random selection from these three positions to ensure robustness. Preliminary experiments⁵ on these locations reveal that both GPT-4V and LLaVA-1.5 are robust to the locations of the embedded instruction, as shown in Figure 4.4. For the sake of simplicity, we take the bottom position as the default placement for the embedded instruction.

Image Resolution The resolution of the image is also the key for the model to understand the visual-modality instruction. For the image-text pairs, we aim to keep the resolution of the raw images, we add the text with the same font size for all images and add extra white space fillings to render the text regarding the length of the instruction. The resolution of image is minimally changed, and the origin image quality is maintained.

⁵Here the small subset for preliminary experiments is 21 examples from MM-Vet.




Figure 4.4: Exploration setup for instruction location on zero shot evaluation for MM-Vet.

Text Prompting

To probe the visual instruction following of MLLMs, we prompt the models under three settings: 1) **Text-Modality Instruction (TEM)**: given the image, the instruction is fed to the models as prompt in text format; 2) **Mix Instruction**: given the image with embedded instruction, an extra instruction as text prompt can be taken as input to guide models to follow the embedded instruction in the image, e.g., “*Answer the question in the image*”; 3) **Visual-Modality Instruction**: given the image with the embedded instruction, no extra text instruction is provided as text prompt, the model needs to recognize, and understand the embedded instruction in the image automatically and follow the instruction to deliver the answer.

The **Text-Modality Instruction** setting is the standard setup in MLLMs, that the model will take the image and text question or instruction separately through vision and language encoders respectively. In the **Mix Instruction** and **Visual-Modality Instruction** settings, the models

Models	LLM	Tem	Mix	Vim
<i>Small set</i>				
GPT-4V	-	66.7	54.4	63.5
<i>Full Set</i>				
LLaVA-1.5	Vicuna-7b	30.5*	10.3 (−20.2)	11.0 (−19.5)
LLaVA-1.5	Vicuna-13b	35.4*	14.8 (−20.6)	14.6 (−20.8)
InstructBLIP	FlanT5 _{XXL}	22.9	12.5 (−10.4)	4.4 (−18.5)
Qwen-VL-Chat	-	41.1	26.2 (−14.9)	13.5 (−27.6)

Table 4.1: Exploration setup for text prompt on zero-shot evaluation for MM-Vet. * denotes from the paper reported [LLLL23].

are required to understand the embedded instruction in the image, while, Mix Instruction is a relaxed setting between Text-Modality Instruction and Visual-Modality Instruction. Preliminary results in Table 4.1 demonstrates that GPT-4V is robust to all three prompt settings, while, VIM is more challenging for the existing open-source MLLMs, for example, LLaVA-1.5 and Qwen-VL-Chat drop more (>20) when transferring from Text-Modality Instruction to the VIM setting, similarly for InstructBLIP. In this paper, we will use the VIM setting for the rest of experiments.

4.2.2 v-MLLM

VIM Corpus

The key ingredient of high-performing MLLMs is high-quality instruction tuning data. There are two categories of visual instruction tuning data, one is the synthetic data by GPT-4, like LLaVA [LLWL23], the other one is the synthetic data generated by GPT-4V, like LVIS-Instruct4V [WMW⁺23]

Human:  Question: What is the animal to the right of the sheep the collar is on? Answer the question using a single word or phrase. Assistant: Goat.

Table 4.2: An VIM training data example from GQA.

and ShareGPT-4V [CLD⁺23]. Here, we use the LVIS-Instruct4V-LLaVA-Instruct-mix880k [WMW⁺23]⁶ as our origin instruction tuning corpus D_R , and convert it into VIM format (Table 4.2 shows an VIM training example from GQA), we only consider the first turn for the multiple turn conversation data. In total, we get 876k VIM training data D_V after filtering the unavailable image links.

VIM Training

v-MLLM adopts a similar model architecture with LLaVA-1.5 [LLLL23] and LVIS-Instruct4V [WMW⁺23]. The model follows an autoregressive training approach, focusing on optimizing the sequential prediction of the answer series $y_1, y_2, \dots, y_n$ by minimizing the loss function $L = \sum_{i=1}^n \text{loss}(\text{LM}(y_{<i}, T, V), y_i)$, where $y_{<i}$ signifies all tokens preceding the i -th token, T and V represent the textual (e.g., text-modality instruction or prompt) and visual (e.g., visual-modality instruction and image context) tokens. T is optional in the VIM training.

⁶We may consider to scale up the data in the future, for example, including ShareGPT4V.

Models	LLM	Res.	MM-Vet	MME	OKVQA	VizWiz	TextVQA	MathVista*	ChartQA	MMMU*
<i>TEM Setting</i>										
GPT-4	-	-	9.8	74.6	8.37	2.76	3.36	18.7	4.12	28.8
Llama 2	Llama2-7b	-	11.1	1609.5	16.21	5.67	7.18	23.2	0.0	6.2
Vicuna	Vicuna-7b	-	11.7	1120.6	4.5	1.88	1.88	18.1	0.0	2.0
GPT-4V	-	-	67.7	1926.6	22.28	17.59	43.14	46.08	28.0	42.9
Gemini Pro	-	-	56.3	1864.2	30.46	4.17	44.83	30.2	13.2	16.2
Qwen-VL-Chat	-	-	41.1	1848.3	56.6	52.74	61.5	36.1	21.28	35.1
LLaVA-1.5	Vicuna-7b	336	30.5	1851.5	58.41	32.08	45.36	25.1	18.08	36.1
LLaVA-1.5	Vicuna-13b	336	35.4	1808.4	61.27	33.92	48.04	26.9	18.56	38.0
InstructBLIP	FlanT5 _{XXL}	224	22.9	1497.5	47.46	25.75	30.91	1.4	11.8	4.4
v-MLLM	Vicuna-7b	336	27.5	1726.8	56.23	49.62	43.81	25.7	16.80	34.1
v-MLLM	Vicuna-13b	336	36.2	1704.7	58.65	51.95	46.27	27.5	16.60	34.1
<i>VIM Setting</i>										
GPT-4V	-	-	63.5	1713.1	28.32	22.18	42.5	12.8	27.44	37.3
Gemini Pro	-	-	50.6	1434.6	26.43	4.93	33.24	11.7	15.44	21.9
Qwen-VL-Chat	-	-	13.5	21.2	0.01	0.15	0.27	6.1	0.0	8.8
LLaVA-1.5	Vicuna-7b	336	11.0	2.9	0	0	0	0.9	0.0	1.4
LLaVA-1.5	Vicuna-13b	336	14.6	24.4	0.38	0	1.51	1.8	0.0	4.6
InstructBLIP	FlanT5 _{XXL}	224	4.4	0	0.07	0	0.04	0.6	0.0	0.0
v-MLLM	Vicuna-7b	336	24	1488.7	52.9	66.5	39.67	6.7	12.88	23.7
v-MLLM	Vicuna-13b	336	26.5	1546.2	54.21	67.78	43.28	6.9	14.12	27.1

Table 4.3: Main quantitative results over each benchmark under TEM and VIM settings. : LLM models, : proprietary models, : the proposed models. *We use a more strict evaluation protocol to remove randomness when mapping from open-ended responses to multiple choices.

To train a unified model that can robustly follow the text-modality and visual modality instructions, we first train the v-MLLM on the origin corpus D_R first, then continue to train the model on the VIM corpus D_V . We train two versions of the model: v-MLLM-7B and v-MLLM-13B.

4.3 Experiments

We first build our VIM -Bench based on eight existing representative benchmarks, then compare the v-MLLM with six representative MLLMs under two settings (TEM and VIM) across all the tasks.

4.3.1 VIM-Bench

Benchmarks To assess the generalization capability of MLLMs, we adapt VIM to eight representative benchmarks, including MME [FCS⁺23], MM-Vet [YYL⁺23], OKVQA [MRFM19], VizWiz [BJJ⁺10], TextVQA [SNS⁺19], MathVista [LBX⁺23], ChartQA [MLT⁺22], and MMMU [YNZ⁺23]. The details of source datasets, data processing pipeline, and evaluations can be found in Appendix 4.8.

Data Reformatting Given the above mentioned benchmarks, we try to do minimal changes (i.e., keeping the image resolution) for evaluation. This process involves re-formatting text instruction into visual embedded instruction by moving the text instruction into the image modality. In reformatting, we retain the original task’s goal while maintaining the original images and text instructions, see the example in Figure 4.2. These repurposed benchmarks are integrated into our VIM-Bench. Theoretically, VIM can be applied to any existing benchmarks, even for pure NLP tasks. We choose eight representative MLLM benchmarks, although our selections are not exhaustive, they provide a broad basis for MLLM evaluation.

4.3.2 Evaluation Setup

Evaluation Settings For each benchmark, we have two main evaluation settings, one is the standard evaluation setting, which is Text-Modality Instruction, `image + text` prompt as the input, denoted as *TEM* setting; the other one is the *VIM* setting with only `image` as input. For the *VIM* setting, it also can accept any text prompts, we call *Mix Instruction* in the ablation experiments.

Evaluation Metrics For the evaluation metrics, we follow the evaluation protocols and metrics of each benchmark. For MathVista and MMMU, the open-ended responses will be parsed into options before calculating the accuracy. We take a more strict protocol that requires the models

to follow explicit instructions (e.g., “Answer with the option letter.”), instead of mapping to the closest options no matter whether the response is related or not. Details of metrics are illustrated in Appendix 4.8.3.

4.3.3 Baselines

We evaluate six representative open-source and proprietary MLLMs, including InstructBLIP [DLL⁺23], LLaVA-1.5 [LLLL23], Qwen-VL-Chat [BBY⁺23], GPT-4V [Ope23a], Gemini Pro [TAB⁺23], and GPT-4V. For LLaVA-1.5, we use its latest versions with Vicuna-v1.5 as the LLM backbone. We use Flan-T5 XXL [CHL⁺22] as LLM backbone for InstructBLIP. All experiments are conducted on NVIDIA A100 and H100 80G GPUs, we use as large models as possible, model training was integrated with Deepspeed Zero-3. For 7b size models, we use batch size 80 with 3 epochs; for 13b models, batch size is 48 with 2 epochs on a single node with 8 GPUs.

4.3.4 Main Results

Table 4.3 summarizes the overall results for two settings. 1). In the original TEM setting, the backbone LLM models can get decent performance on these benchmarks, even without access to the image modality. Interestingly, on most of eight tasks, Llama 2 is much better than GPT-4, we will briefly discuss this in Section 4.5. 2). For all open-source MLLMs, there is a significant disparity between the TEM setting and VIM setting. 3). GPT-4V and Gemini Pro are robust to the instruction modality, while, open-source MLLMs struggle in the VIM setting, achieve significantly low scores. 4). Our proposed v-MLLM shows robust instruction following capabilities in two settings across all the tasks, especially in the VIM setting, significant gain over open-source MLLMs.

4.4 Ablation

We have demonstrated that VIM is a challenging setting for the current open-source MLLMs. To better understand the gap between the standard Text-Modality Instruction and Visual Embedded Instruction, we decompose the VIM into two steps for ablation.

4.4.1 Instruction Recognition

One hypothesis is the vision encoder of MLLMs cannot discern the instruction in the image, only generates the long description for the image, as the example shown in Table 4.4. *Can the MLLMs recognize the embedded instruction in the image?* Here, we conduct an experiment to verify the visual embedded instruction recognition ability of the models, we choose an external dataset VQA_{v2} for test (beyond our eight test benchmarks)⁷, we explicitly prompt the models with “*What is the text in the image?*”, then manually check the outputs with the ground-truth text instructions in the zero-shot setting for VQA_{v2}. We select 30 images as our subsets for both models. Here, the images we choose only have the embedded instructions since we want to verify the instruction recognition capability of these models in an ideal setup.




Figure 4.5: Instruction recognition results on VQA_{v2}. We report the correct number of matches out of total 30 samples for LLaVA-1.5-7B & 13B, and GPT-4V.

We do word match and semantic match for the results. For example, the origin instruction

⁷Due to the large size of VQA_{v2}, it is expensive to run the full *test-dev* and *test* splits for GPT-4V and Gemini Pro for fair comparison, we did not include it in the VIM-Bench.




Table 4.4: LLaVA’s response to the question #3575865 in OKVQA.

of question #393225000 in VQA_{v2} is “*What website copyrighted the picture*”, the output “*What website copied the picture?*” is counted as a correct word match, since most of the words are recognized. However, it will be considered as a wrong semantic match. As shown in Figure 4.5, GPT-4V can recognize the embedded instructions nearly perfectly for both settings on word match and semantic match, the only “*failure*” case is from an image containing a text logo inside as shown in Table 4.10 (in the Appendix). While, LLaVA-1.5 can detect words of the instructions, but semantically different. A detailed analysis of instruction recognition capability can be found in Appendix 4.9.2.

Models	LLM	OKVQA	MM-Vet	VizWiz
LLaVA-1.5	Vicuna-7b	0.00	11.0	0
		14.28	10.3	22.47
InstructBLIP	FlanT5 _{XXL}	0.07	4.4	0
		25.44	12.5	18.79
Qwen-VL-Chat	-	0.01	13.5	0.15
		30.75	26.2	34.03
GPT-4V	-	28.32	63.5	22.18
		27.70	54.4	23.37
v-MLLM	Vicuna-7b	52.9	24	66.5
		53.21	22.5	47.09

Table 4.5: Ablation on OKVQA, MM-Vet, VizWiz tasks.  highlights the numbers for mix instruction setting.

4.4.2 Instruction Following

We further relax the VIM setting to the Mix Instruction setting with an extra text prompt “Answer the question in the image in one word or phrase.”, these open-source MLLMs exhibit some degree of success (the highlighted gray rows in Table 4.5), which proves that the existing MLLMs rely more on their LLM backbones for instruction following. For example, on OKVQA and VizWiz, LLaVA-1.5, InstructBLIP and Qwen-VL-Chat can achieve moderate success in the Mix Instruction setting, large improvement over their VIM setting.

4.4.3 Qualitative Observations

We highlight several intriguing observations during the experiments, and a comprehensive qualitative analysis is provided in the Appendix 4.9.3.

4.5 Discussion

GPT-4V and Gemini Pro demonstrate robust instruction following under the original TEM and VIM settings, especially in the VIM setting, they show robust response, while, most of MLLMs struggle under the VIM setting. Our proposed v-MLLM narrows down this gap for the open-source MLLMs, and bring robust instruction following capability to different instruction modalities. There could be several directions to be explored along this routine. One direction is integrating VIM with grounding abilities of MLLMs discussed by several works [YZL⁺23, CZZ⁺23, YZG⁺23]. Another direction is scaling up, in terms of data scale and model scale.

MLLM Evaluation In the LLM exploration setting, Table 4.3 shows that Llama 2 is much better than GPT-4 on six of eight multimodal tasks, without image input. We manually check some results, and find that the responses from GPT-4 are more reasonable than Vicuna and Llama 2. For example, in Table 4.6, we only provide the text Prompt to the LLMs, the output from Vicuna and Llama 2 is more like text continuation from the training corpus, while, GPT-4 makes more reasonable response to the Prompt question.

Another potential issue exposed from this exploration experiment is about MLLM evaluation, though GPT-4 makes more reasonable response, its score is low on all six tasks, which may be contrary to the objective of these benchmarks. In-depth analysis and discussion of evaluation protocol and metrics are beyond the scope of this work, which may leave for future work.

MM-Vet #1

Image

Solve the following equations:

1) $8x + 11 = 4x + 14$

2) $7d - 4 = 11d - 9$

Prompt: What is d in the last equation? Answer the question in one word or phrase.

GPT-4: I'm sorry, but I need more context or information about the "last equation" you are referring to in order to provide an answer. Could you please provide the equation or more details surrounding it?

Vicuna: d is the derivative of the function with respect to x.

Llama 2: D is the constant of integration.

OKVQA #2971475

Image

Prompt: What sport can you use this for? Answer the question in one word or phrase.

GPT-4: Please provide the context or the object you're referring to in order to answer the question about which sport it can be used for.

Vicuna: Basketball.

Llama 2: Basketball

Table 4.6: LLMs' responses to the questions in MM-Vet and OKVQA. Note: *Image* is not feed into the LLMs.

4.6 Related Work

Multimodal Large Language Models With the success of Large Language Models (LLMs) [Ope23a, Goo23, TLI⁺23, TMS⁺23], there is growing interest in studying Multimodal Large Language Models [LLWL23, DLL⁺23, ZCS⁺23, AGG⁺23, YXX⁺23, ZHZ⁺23, SLL⁺23, SSC⁺23, BBY⁺23] to improve multimodal understanding, reasoning and generation by leveraging the strong capability of LLMs. Following the recipe of the instruction tuning [TGZ⁺23, CLL⁺23, PLH⁺23] in LLMs, LLaVA [LLWL23] and MiniGPT-4 [ZCS⁺23] propose to use GPT-4 [Ope23a] to synthesize the instruction, and employ the open-source LLMs for instruction tuning to connect the pretrained vision encoder and open-source LLMs. Recently, GPT-4V [Ope23c] and Gemini [TAB⁺23] were released, followed by work [YLL⁺23, ZLW⁺23, TLZ⁺24] that explore their impressive multimodal capabilities. LVIS-Instruct4V [WMW⁺23] and ShareGPT-4V [CLD⁺23] are created with GPT-4V for high-quality instruction tuning data. Our work implies the limitations of existing open-source MLLMs, and provides a solution to enhance robust visual instruction following.

Benchmarks In parallel with the MLLMs development, a trend emerges in proposing a variety of benchmarks. MME [FCS⁺23] proposes 14 tasks with Yes/No questions based on the images. MMBench [LDZ⁺23] and SEED-Bench [LWW⁺23] introduce benchmarks that cover a variety of multiple-choice questions. MM-Vet [YYL⁺23] extends to evaluate the open-ended outputs from MLLMs. VisIT-Bench [BBH⁺23] accesses a range of tasks from recognition to complex reasoning, while Q-Bench [WZZ⁺23] accesses low-level visual perception and understanding. MathVista [LBX⁺23] focuses on systematically studying the mathematical reasoning capability in visual context. More recently, MMMU [YNZ⁺23] is proposed to evaluate multimodal models in a broad range of college-level subject knowledge. All these benchmarks focus on text instruction evaluation for MLLMs, and our proposed ViM integrates the text instruction into the visual modality

space, asking for strong visual interpretation capability for embedded instruction recognition and following. Additionally, our VIM is orthogonal with these existing benchmarks, can be seamlessly adapted to any of them.

4.7 Conclusion

In this work, we present VIM, a challenging setting to assess the visual instruction following ability of Multimodal Large Language Models. We adapt VIM to eight benchmarks, leading to VIM-Bench. Through in-depth probing under zero-shot setting, we observed that the existing open-source MLLMs fall short in the VIM setting, in most cases performing not as good as those in the original TEM setting. Furthermore, we train a v-MLLM, which demonstrates robust instruction following capabilities under text and visual modality instruction settings across all the tasks.

Appendix to VIM-MLLM

4.8 Benchmark Details

4.8.1 Source Datasets

We consider eight representative datasets, MM-Vet [YYL⁺23], MME [FCS⁺23], OKVQA, VizWiz, TextVQA, MathVista, ChartQA, and MMMU. We also provide probing analysis on VQA_{v2} *test-dev* split, RefCOCO *testA* split, RefCOCO+ *testA* split, and RefCOCOg *test* split to better illustrate how the data is formatted.

In Figure 4.6, we showcase an example from the source dataset VQA_{v2}, which is with instruction probing setting “Text”. We also consider other two probing settings, “MIX” that have an additional text prompt to guide the visual embedded instruction following, and our proposed “VIM” that only allow the image with embedded instruction as input. We also showcases dataset examples sampled from source datasets RefCOCO, MME, and MM-Vet. All the VIM test samples does not include any additional instruction input in text modality (noted as “NA”).

4.8.2 Dataset Processing Pipeline

For each source dataset, we start by building up zero shot by embedding instructions into the input image to concatenate as a new image which contains instructions in image modality. In this way, we obtain a new image with embedded instructions for each image-question pair.

4.8.3 Evaluation Details

Metrics We follow the evaluation pipeline for each benchmark. We use parsing and accuracy for MathVista and MMMU with a more strict protocol. For example, when the model is outputting

Setting	Source Dataset	Image Modality Input	Text Modality Input
Text	VQAv2		Question: What sport is the man participating in? Answer the question using a single word or phrase.
MIX	VQAv2	 Question: What sport is the man participating in? Answer the question using a single word or phrase.	Answer the question in the image.
VIM (Zero Shot)	VQAv2	 Question: What sport is the man participating in? Answer the question using a single word or phrase.	NA
	RefCOCO series	 Question: What is the normalized coordinate of "the catcher" in the format [x1, y1, x2, y2] rounded to four decimal places?	NA
	MME	 Question: Is this artwork titled the adoration of the shepherds? Answer the question using a single word or phrase.	NA
	MM-Vet	 Question: What will the girl on the right write on the board? Answer the question using a single word or phrase.	NA

Figure 4.6: Dataset example comparison of three instruction probing settings: Text, MIX and VIM.

an empty or random string for a multiple-choice question, in MathVista, the original evaluation protocol will use Levenshtein distance to map to a most similar prediction option, and in MMMU, a random choice from the candidate list will be applied. This will introduce noise and randomness for the evaluation, may not correctly reflect the model performance. In our strict evaluation protocol, we eliminate this random match strategy.

For OKVQA and TextVQA, we follow the leaderboard evaluation to use an evaluation metric that is robust to inter-human variability: $\text{Acc}(\text{ans}) = \min \left\{ \frac{\#\text{humans that said ans}}{3}, 1 \right\}$. For ChartQA, we use relaxed accuracy on human and augmented split.

For MME, the standard metric (*Score*) proposed in [FCS⁺23] is the summed up Accuracy (*Acc*) and Accuracy+ (*Acc₊*) as: $\text{Score} = \text{sum}(\text{Acc} \times 100\%, \text{Acc}_+ \times 100\%)$, where the former one count each correct answer as correct, while the latter one only considers correct when both “Yes” and “No” questions for each image are answered correctly.

For MM-Vet, we use GPT-4 (“gpt-4-0613” version) to automatically provide the score for each sample. The final Accuracy reported as: $\text{Acc} = \frac{\sum_{i=1}^N s_i}{N} \times 100\%$, where s_i is the score at scale 0 – 1 for sample i .

4.9 Full Analysis

4.9.1 Robustness Analysis

Prompt for MIX Probing Setting

In mix probing setting, the MLLMs can accept an extra text instruction input as guidance. The model performance may vary when given different prompts. We report the results using four relevant but diversified prompts (Prompt #1-#5) in Table 4.7. To be specific, the detailed prompts we use are: 1) Prompt #1: “*Answer the question in the image.*”, 2) Prompt #2: “*Please answer the*

question that is written in the image.”, 3) Prompt #3: “Follow the instruction embedded in the image.”, 4) Prompt #4: “Detect the question in the image and directly answer to it.”.

Task	LLaVA-1.5-7b				LLaVA-1.5-13b			
	Prompt #1	Prompt #2	Prompt #3	Prompt #4	Prompt #1	Prompt #2	Prompt #3	Prompt #4
artwork	5	5	0	0	0	0	0	0
celebrity	5	5	0	0	1	0	0	0
code reasoning	4	3	0	0	0	0	0	0
color	5	5	0	0	1	1	0	0
commonsense reasoning	4	4	0	0	0	0	0	0
count	5	5	0	0	1	0	0	0
existence	6	6	0	0	0	0	0	0
landmark	5	5	0	0	3	0	0	0
numerical calculation	5	5	0	0	0	0	0	0
OCR	5	5	0	0	0	0	0	0
position	4	5	0	0	0	0	0	0
posters	5	5	0	0	1	0	0	0
scene	5	5	0	0	0	0	0	0
text translation	5	5	0	0	0	0	0	0
<i>Correct</i>	68	68	0	0	7	1	0	0

Table 4.7: Zero shot evaluation results on MME subset under mix setting. We compare the performance of LLaVA-1.5-7b and LLaVA-1.5-13b across four different prompts.

Image Embedded with Instruction

To investigate whether the model performance is robust to the minimal changes introduced by the embedded instruction, we give both the original instruction in the text modality and the image with instruction embedded as the image modality to the model.

In Table 4.8, we present comparative results of LLaVA-1.5 using Vicuna-7b and Vicuna-13b language backbones. It’s observed that embedding images with instructions leads to a marginal decline in performance. This trend suggests that current MLLMs may not be entirely robust to variations in images. However, this performance degradation is minor and within acceptable limits. This implies that the disparity in performance between the Text and VIM probing settings is not solely attributable to changes in the images, but is largely due to the models’ capacity to follow visual embedded instructions.

Models	LLM	Embedded Instruction	<i>Zero shot</i>		
			VQA _{v2}	MME	MM-Vet
Sub set					
LLaVA-1.5	Vicuna-7b	w.o.	60.75	108	31.3
		w.	57.88	88	27.9
	Vicuna-13b	w.o.	61.00	106	35.2
		w.	58.00	87	32.7

Table 4.8: Zero shot evaluation results of Text probing setting on VQA_{v2}, MME, MM-vet. This is the popular setting for evaluating text instruction following capability of MLLMs, where the input image and text instruction are both provided.

4.9.2 Details of Instruction Recognition

Based on Section 4.4.1, we conduct an ablation to verify the instruction recognition capability of these MLLMs. Table 4.9 showcases some example results for Zero Shot instruction recognition. GPT-4V can recognize the text instruction in both settings, LLaVA can detect some words of the instructions, but may not perfectly recognize the instructions, especially in the one shot setting. Table 4.10 shows a failure example of GPT-4V for Instruction Recognition, it recognizes the logo texts on the bus as the text instruction.

4.9.3 Qualitative Observations

People related questions

In VQA and MM-Vet, there are some categories of questions about people or movies. For example, in Table 4.11, GPT-4V will response “*Sorry, I cannot help with that.*”, while, LLaVA-1.5 and InstructBLIP just simply repeat the question.

Image w. Embedded Instruction #42000 in VQA

Question: What color are the gym shoes?
Answer the question using a single word or phrase.

Text Prompt: What is the text in the image?

GPT-4V: Question: What color are they gym shoes? Answer the question using a single word or phrase.

LLaVA-1.5-13B: The text in the image is a question asking about the color of the gym shoes.

Image w. Embedded Instruction #132132000 in VQA

Question: How many men are shopping?
Answer the question using a single word or phrase.

Text Prompt: What is the text in the image?

GPT-4V: Question: How many men are shopping? Answer the question using a single word or phrase.

LLaVA-1.5-13B: The text in the image is a question asking how many men are shopping.

Image w. Embedded Instruction #757000 in VQA

Question: How many elephants are pictured in this photo? Answer the question using a single word or phrase.

Text Prompt: What is the text in the image?

GPT-4V: Question: How many elephants are pictured in this photo? Answer the question using a single word or phrase.

LLaVA-1.5-13B: The text in the image is a question asking how many elephants are pictured in the photo.

Image w. Embedded Instruction: #131612000 in VQA

Question: What is the man doing with the horse? Answer the question using a single word or phrase.

Text Prompt: What is the text in the image?

GPT-4V: Question: What is the man doing with the horse? Answer the question using a single word or phrase.

LLaVA-1.5-13B: The text in the image is a caption that reads "Question: What is the man doing with the horse? Answer the question using a single word or phrase."

Table 4.9: Zero Shot Instruct Recognition: MLLMs's recognition to the example questions in VQA.

Image w. Embedded Instruction

Question: How many girls are on the advertisement?
Answer the question using a single word or phrase.

Text Prompt:

What is the text in the image?

GPT-4V:

1. "anwb"
2. "Feel Good Bus"
3. "Verrassend leuk vrijwilligerswerk!"

Table 4.10: One "failure" examples to show the instruction recognition of GPT-4V.

4.10 Limitations

We discuss the limitations of our work as follows: 1). Though v-MLLM exhibits robust instruction following capability in both the TEM and VIM settings, it still has a gap with proprietary models, especially GPT-4V, there is still space to improve to be a generalist model. 2). In this work, the evaluation protocols and metrics of VIM setting follow these from the original TEM setting, it also inherits the shortcoming of these evaluations, we leave these for future work. 3). For proprietary


<i>Image w. Embedded Instruction</i>  Question: Can you give a short introduction to this person? GPT-4V: Sorry, I cannot help with that. LLaVA-1.5: Question: Can you give a short introduction to this person? InstructBLIP: albert einstein with the caption can you give a short introduction to this person? v-MLLM: albert einstein

Table 4.11: MLLMs’s response to the question #202 in MM-Vet.

models (GPT-4V and Gemini Pro), they are evolving with in-context learning as more queries are fed to the models, the results from these models may not be stable or potentially reproducible.

Chapter 5

Conclusion

In this dissertation, I presented three work. OSCAR shows that object tags can help the alignment between the text and image in the object-level vision-language pretraining; VINVL demonstrates that vision feature matters in the vision-language pretraining, by training a large object detection backbone with a large object-attribute taxonomy on a large corpus, it can deliver more rich, semantic vision representation, together with OSCAR+ pretraining, it achieves the state-of-the-art on many tasks by that time.

OSCAR and VINVL focus on the grounding semantics with vision, and push the object-level vision-language pretraining to a very high level. However, there are several issues for object-level vision-language pretraining, the first issue is scaling up, the object-level vision-language pretraining needs to extract the object region features offline first; this also brings another issue, which is end-to-end training, the object-level feature extraction process is non-differentiable, which may prevent the vision backbone for further finetuning. After the release of VINVL, OpenAI's CLIP [RKH⁺21] came out immediately, it fills the above gaps for the object-level pretraining with very large scale (billion size image-text pairs), and demonstrates superb performance on zero-shot learning for vision-language tasks.

CLIP [RKH⁺21] brings the vision-language pretraining to a new era, a new term like “vision foundation models”. On April 2023, the release of GPT-4 [Ope23b], brings a new wave of open-source MLLMs development. VIM challenges the MLLMs by embedding the instruction into the visual pixel space, we call VISUAL MODALITY INSTRUCTION, and shows that there is a significant disparity between text-modality instruction and VIM settings for the open-source MLLMs. Furthermore, v-MLLM fills this gap, and demonstrates strong robust visual instruction following capability on several benchmarks in both text-modality instruction and VIM settings.

In this dissertation, we witness the development for the vision-language pretraining, from object-level pretraining, to image-level contrastive learning pretraining, and to multimodal large language models. However, there still leave us a lot of questions unknown, for example, given the current evaluation benchmarks, how is the data contamination issue for the current MLLMs, there also calls for better evaluation benchmarks and protocols for MLLMs, and also the grounding capability for MLLMs.

Bibliography

- [ADF⁺23] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [ADW⁺19] 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.
- [AGG⁺23] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023.
- [AHB⁺18] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018.
- [BBH⁺23] Yonatan Bitton, Hritik Bansal, Jack Hessel, Rulin Shao, Wanrong Zhu, Anas Awadalla, Josh Gardner, Rohan Taori, and Ludwig Schimdt. Visit-bench: A benchmark for vision-language instruction following inspired by real-world use. *arXiv preprint arXiv:2308.06595*, 2023.
- [BBY⁺23] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023.
- [BJJ⁺10] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, et al. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 333–342, 2010.

- [BLM91] Peter F Brown, Jennifer C Lai, and Robert L Mercer. Aligning sentences in parallel corpora. In *Proceedings of the 29th annual meeting on Association for Computational Linguistics*, 1991.
- [BMR⁺20] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [CGL⁺19] Wenhui Chen, Zhe Gan, Linjie Li, Yu Cheng, William Wang, and Jingjing Liu. Meta module network for compositional visual reasoning. *arXiv preprint arXiv:1910.03230*, 2019.
- [CHL⁺22] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [CLD⁺23] Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. *arXiv preprint arXiv:2311.12793*, 2023.
- [CLL⁺23] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023.
- [CLY⁺19] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740*, 2019.
- [CLY⁺20] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020.
- [CZZ⁺23] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- [DBK⁺20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020.

- [DCLT19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2019.
- [DLL⁺23] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023.
- [FCS⁺13] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. DeViSE: A deep visual-semantic embedding model. In *NeurIPS*, 2013.
- [FCS⁺23] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023.
- [FFKF17] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improved visual-semantic embeddings. *arXiv preprint arXiv:1707.05612*, 2(7):8, 2017.
- [FGI⁺15] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482, 2015.
- [GCL⁺20] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *arXiv preprint arXiv:2006.06195*, 2020.
- [GKSS⁺17] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017.
- [GLL19] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V. Le. NAS-FPN: learning scalable feature pyramid architecture for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 7036–7045. Computer Vision Foundation / IEEE, 2019.
- [Goo23] Google. Bard - chat based ai tool from google, powered by palm 2. <https://bard.google.com/?hl=en>, 2023.
- [HGDG17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

- [HLL⁺20] Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. *CVPR*, 2020.
- [HM19a] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *NeurIPS*, 2019.
- [HM19b] Drew A Hudson and Christopher D Manning. GQA: A new dataset for real-world visual reasoning and compositional question answering. *arXiv preprint arXiv:1902.09506*, 2019.
- [HWCW19] Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on attention for image captioning. In *ICCV*, 2019.
- [HYL⁺20] 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. *arXiv preprint arXiv:2009.13682*, 2020.
- [HZL⁺20] 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 arXiv:2004.00849*, 2020.
- [JMR⁺20] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10267–10276, 2020.
- [JYX⁺21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021.
- [KBZ⁺19] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Large scale learning of general visual representations for transfer. *arXiv preprint arXiv:1912.11370*, 2019.
- [KFF15] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- [KRA⁺18] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018.

- [KSZ14] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [KZG⁺17] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017.
- [LBPL19] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlbnet: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019.
- [LBX⁺23] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023.
- [LCH⁺18] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *ECCV*, 2018.
- [LDF⁺19] 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. *arXiv preprint arXiv:1908.06066*, 2019.
- [LDG⁺17] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [LDZ⁺23] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023.
- [LGG⁺17] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [LGR⁺19] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-Task vision and language representation learning. *arXiv preprint arXiv:1912.02315*, 2019.
- [LLLL23] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023.

- [LLWL23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [LMB⁺14] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.
- [LOG⁺19] 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 arXiv:1907.11692*, 2019.
- [LPY⁺17] Zeming Li, Chao Peng, Gang Yu, Xiangyu Zhang, Yangdong Deng, and Jian Sun. Light-head R-CNN: in defense of two-stage object detector. *CoRR*, abs/1711.07264, 2017.
- [LWW⁺23] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023.
- [LYL⁺20] 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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 121–137. Springer, 2020.
- [LYY⁺19] 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 arXiv:1908.03557*, 2019.
- [LYZ⁺20] Junyang Lin, An Yang, Yichang Zhang, Jie Liu, Jingren Zhou, and Hongxia Yang. Interbert: Vision-and-language interaction for multi-modal pretraining. *arXiv preprint arXiv:2003.13198*, 2020.
- [MH08] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 2008.
- [MLT⁺22] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- [MRFM19] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019.

- [NMB⁺13] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650*, 2013.
- [OKB11] Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS*, 2011.
- [Ope23a] OpenAI. Gpt-4: Technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [Ope23b] OpenAI. Gpt-4 technical report, 2023.
- [Ope23c] OpenAI. Gpt-4v(ision) system card. <https://openai.com/research/gpt-4v-system-card>, 2023.
- [OWJ⁺22] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [PHA⁺20] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, João Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. *arXiv preprint arXiv:2010.02824*, 2020.
- [PLH⁺23] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023.
- [PYLM20] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10971–10980, 2020.
- [QSS⁺20] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*, 2020.
- [RDGF16] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [RF17] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [RF18] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.

- [RHGS15] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [RJL⁺16] Zhou Ren, Hailin Jin, Zhe Lin, Chen Fang, and Alan Yuille. Joint image-text representation by gaussian visual-semantic embedding. In *Multimedia*, 2016.
- [RKH⁺21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [RMM⁺17] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017.
- [RSR⁺20] 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. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [SDGS18] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Annual Meeting of the Association for Computational Linguistics*, 2018.
- [SFF10] Richard Socher and Li Fei-Fei. Connecting modalities: Semi-supervised segmentation and annotation of images using unaligned text corpora. In *CVPR*, 2010.
- [SGMN13] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *NeurIPS*, 2013.
- [SJL⁺19] Botian Shi, Lei Ji, Pan Lu, Zhendong Niu, and Nan Duan. Knowledge aware semantic concept expansion for image-text matching. In *IJCAI*, 2019.
- [SLL⁺23] Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.
- [SLZ⁺19] 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 *Proceedings of the IEEE international conference on computer vision*, pages 8430–8439, 2019.
- [SMV⁺19] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A joint model for video and language representation learning. *ICCV*, 2019.

- [SNS⁺19] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019.
- [SSC⁺23] Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*, 2023.
- [SSSG17] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017.
- [SZC⁺19] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.
- [SZZ⁺18] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018.
- [TAB⁺23] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [TB19] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019.
- [TGZ⁺23] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- [TLI⁺23] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [TLZ⁺24] Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms, 2024.
- [TMS⁺23] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al.

- Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [TPL20] Mingxing Tan, Ruoming Pang, and Quoc V. Le. Efficientdet: Scalable and efficient object detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10778–10787. IEEE, 2020.
- [TSCH20] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: A simple and strong anchor-free object detector. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [VSP⁺17] 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.
- [WKM⁺19] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [WLL⁺19] Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. CAMP: Cross-Modal adaptive message passing for text-image retrieval. In *ICCV*, 2019.
- [WMW⁺23] Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. To see is to believe: Prompting gpt-4v for better visual instruction tuning. *arXiv preprint arXiv:2311.07574*, 2023.
- [WSL⁺16] Qi Wu, Chunhua Shen, Lingqiao Liu, Anthony Dick, and Anton Van Den Hengel. What value do explicit high level concepts have in vision to language problems? In *CVPR*, 2016.
- [WYQ⁺19] Yaxiong Wang, Hao Yang, Xueming Qian, Lin Ma, Jing Lu, Biao Li, and Xin Fan. Position focused attention network for image-text matching. *arXiv preprint arXiv:1907.09748*, 2019.
- [WYY⁺21] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021.
- [WZZ⁺23] Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Chunyi Li, Wenxiu Sun, Qiong Yan, Guangtao Zhai, et al. Q-bench: A benchmark for general-purpose foundation models on low-level vision. *arXiv preprint arXiv:2309.14181*, 2023.

- [YCC⁺21] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [YHG⁺16] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016.
- [YJW⁺16] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016.
- [YLHH14] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.
- [YLL⁺23] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 2023.
- [YNZ⁺23] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. *arXiv preprint arXiv:2311.16502*, 2023.
- [YTY⁺20] 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. *arXiv preprint arXiv:2006.16934*, 2020.
- [YWV⁺22] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- [YXX⁺23] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- [YYL⁺23] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023.
- [YZG⁺23] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023.

- [YZL⁺23] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023.
- [ZCS⁺23] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- [ZHZ⁺23] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023.
- [ZLH⁺21] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588, 2021.
- [ZLW⁺23] Xinlu Zhang, Yujie Lu, Weizhi Wang, An Yan, Jun Yan, Lianke Qin, Heng Wang, Xifeng Yan, William Yang Wang, and Linda Ruth Petzold. Gpt-4v(ision) as a generalist evaluator for vision-language tasks, 2023.
- [ZPZ⁺20] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and VQA. *AAAI*, 2020.
- [ZZG⁺17] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, and Yi-Dong Shen. Dual-path convolutional image-text embedding with instance loss. *arXiv preprint arXiv:1711.05535*, 2017.