
OFA: Unifying Architectures, Tasks, and Modalities Through a Simple Sequence-to-Sequence Learning Framework

Peng Wang An Yang Rui Men Junyang Lin Shuai Bai
Zhikang Li Jianxin Ma Chang Zhou Jingren Zhou Hongxia Yang

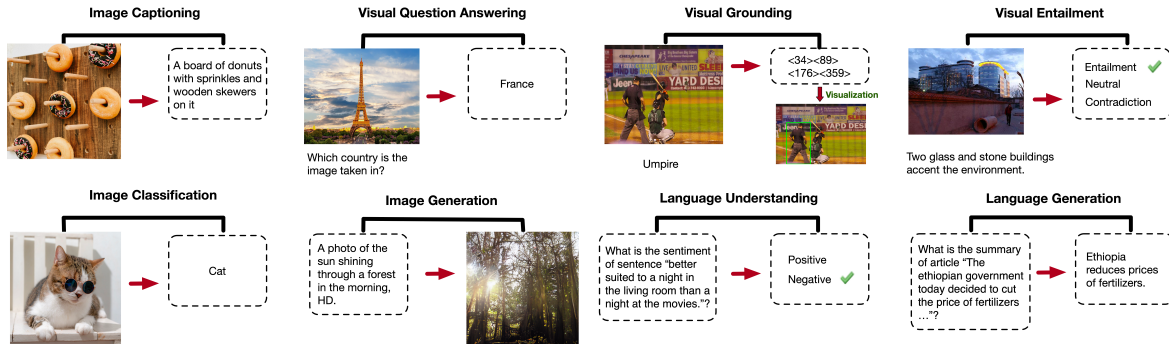


Figure 1. Examples of various tasks supported by OFA.

Abstract

In this work, we pursue a unified paradigm for multimodal pretraining to break the shackles of complex task/modality-specific customization. We propose OFA, a Task-Agnostic and Modality-Agnostic framework that supports Task Comprehensiveness. OFA unifies a diverse set of cross-modal and unimodal tasks, including image generation, visual grounding, image captioning, image classification, language modeling, etc., in a simple sequence-to-sequence learning framework. OFA follows the instruction-based learning in both pre-training and finetuning stages, requiring no extra task-specific layers for downstream tasks. In comparison with the recent state-of-the-art vision & language models that rely on extremely large cross-modal datasets, OFA is pretrained on only 20M publicly available image-text pairs. Despite its simplicity and relatively small-scale training data, OFA achieves new SOTAs in a series of cross-modal tasks while attaining highly competitive performances on uni-modal tasks. Our further analysis indicates that OFA can also effec-

tively transfer to unseen tasks and unseen domains. Our code and models are publicly available at <https://github.com/OFA-Sys/OFA>.

1. Introduction

Building an omnipotent model that handles many tasks and modalities is an attractive goal in the AI community. The possibilities of achieving this goal may largely depend on whether massive varieties of modalities, tasks and training regimes can be represented with only a few forms that can be unified and managed by a single model or system.

Recent developments of the Transformer (Vaswani et al., 2017) architecture have shown its potential for being a universal computation engine (Devlin et al., 2019; Brown et al., 2020; Cobbe et al., 2021; Schneider et al., 2019; Dosovitskiy et al., 2020; Jaegle et al., 2021b; Su et al., 2019). In the settings of supervised learning, the “pretrain-finetune” paradigm achieves excellent success in many domains. In the regimes of few-/zero-shot learning, language models with prompt / instruction tuning prove powerful zero-/few-shot learners (Brown et al., 2020; Wei et al., 2021; Sanh et al., 2021). These advances have provided more significant than ever opportunities for the emergence of an omni-model.

To support better generalization for open-ended problems while maintaining high multitask performance and ease of use, we advocate that an omnipotent model should have the

DAMO Academy, Alibaba Group, China. Correspondence to: Chang Zhou <ericzhou.zc@alibaba-inc.com>.

following three properties: 1. Task-Agnostic (TA): unified task representation to support different types of tasks, including classification, generation, self-supervised pretext tasks, etc., and to be agnostic to either pretraining or finetuning. 2. Modality-Agnostic (MA): unified input and output representation shared among all tasks to handle different modalities. 3. Task Comprehensiveness (TC): enough task variety to accumulate generalization ability robustly.

However, it is challenging to satisfy these properties while maintaining superior performance in downstream tasks. Current language and multimodal pretrained models readily fail at parts of these properties, due to their following design choices. 1. Extra learnable components for finetuning, e.g., task-specific heads (Devlin et al., 2019), adapters (Houlsby et al., 2019), soft prompts (Lester et al., 2021). This makes the model structure task-specific and poses discrepancy between pretraining and finetuning. Such designs are also not friendly to supporting unseen tasks in a zero-shot manner. 2. Task-specific formulation. For most current methods, pretraining, finetuning and zero-shot tasks usually differ in task formulation and training objectives. This violates TA and it is burdensome to scale up the task population to achieve TC. 3. Entangling modality representation with downstream tasks. It is a common practice for Vision & Language models to take the detected objects as part of the image input features (Su et al., 2019; Lu et al., 2019; Chen et al., 2020c; Li et al., 2020; Gan et al., 2020; Zhang et al., 2021a). Though it demonstrates better downstream task performance on some closed-domain datasets, it depends on an extra object detector which usually fails at open-domain data.

Therefore, we explore an omni-model for multimodal pretraining and propose **OFA**, hopefully “One For All”, which achieves the objectives of unifying architectures, tasks, and modalities, and supports the three properties above. We formulate both pretraining and finetuning tasks in a unified sequence-to-sequence abstraction via handcrafted instructions (Wei et al., 2021; Sanh et al., 2021) to achieve Task-Agnostic. A Transformer is adopted as the Modality-Agnostic compute engine, with a constraint that no learnable task- or modality-specific components will be added to downstream tasks. It is available to represent information from different modalities within a globally shared multimodal vocabulary across all tasks. We then support Task Comprehensiveness by pretraining on varieties of uni-modal and cross-modal tasks.

To summarize:

- We propose OFA, a Task-Agnostic and Modality-Agnostic framework that supports Task Comprehensiveness. OFA is the first attempt to unify the following vision & language, vision-only and language-only tasks, including understanding and generation,

e.g., text-to-image generation, visual grounding, visual question answering (VQA), image captioning, image classification, language modeling, etc., via a simple sequence-to-sequence learning framework with a unified instruction-based task representation.

- OFA is pretrained on the publicly available datasets of 20M image-text pairs, in comparison with recent models that rely on paired data of a much larger scale (Wang et al., 2021b; Yuan et al., 2021). OFA achieves state-of-the-art performances in a series of vision & language downstream tasks, including image captioning, visual question answering, visual entailment, referring expression comprehension, etc.
- OFA, as a multimodal pretrained model, achieves comparable performances on unimodal tasks with SOTA pretrained models in language or vision, e.g., RoBERTa, ELECTRA and DeBERTa for natural language understanding, UniLM, Pegasus and ProphetNet for natural language generation, and MoCo-v3, BEiT and MAE for image classification.
- We verify that OFA achieves competitive performance in zero-shot learning. Also, it can transfer to unseen tasks with new task instructions and adapt to out-of-domain information without finetuning.

2. Related Work

Language Pretraining & Vision Pretraining Natural language pretraining has revolutionized the whole NLP research community. A representation of this track is the birth of BERT (Devlin et al., 2019) and GPT (Radford et al., 2018). A number of studies have been progressively advancing pretraining by improving pretraining tasks and designing more sophisticated model architectures (Yang et al., 2019; Sun et al., 2019a;b; Liu et al., 2019; Dong et al., 2019; Raffel et al., 2020; Lewis et al., 2020). Having witnessed the success of natural language pretraining, researchers have promoted self-supervised learning (SSL) in computer vision (Chen et al., 2020a;b; Grill et al., 2020). Recently, mirroring masked language modeling (MLM) in language pretraining, generative pretraining (Bao et al., 2021; He et al., 2021a) with ViT architecture (Dosovitskiy et al., 2020) boosts downstream performance.

Multimodal Pretraining Multimodal pretraining has been developing rapidly with the transfer of BERT to cross-modal representation learning (Li et al., 2019b; Lu et al., 2019; Zhou et al., 2020; Tan & Bansal, 2019; Chen et al., 2020c; Li et al., 2019a; Lin et al., 2020; Lu et al., 2020; Xu et al., 2021; Li et al., 2020; Gan et al., 2020; Zhang et al., 2021a; Yu et al., 2021; Li et al., 2021; Huang et al., 2020). To adapt models to generation tasks, researchers

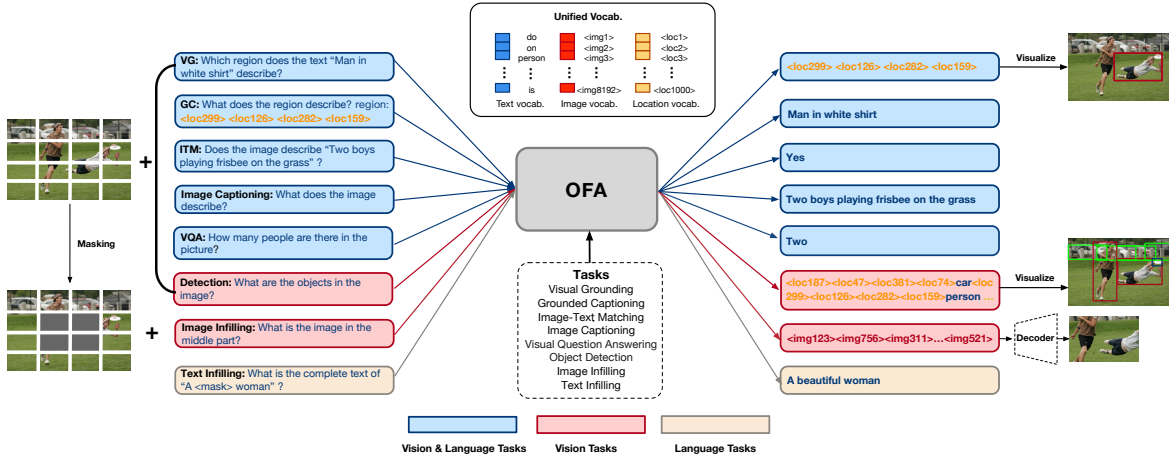


Figure 2. A demonstration of the pretraining tasks, including visual grounding, grounded captioning, image-text matching, image captioning, VQA, object detection, image infilling as well as text infilling.

have applied specific masking strategies or the encoder-decoder architecture (Li et al., 2020; Zhang et al., 2021a; Lin et al., 2021; Wang et al., 2021b). Besides, to simplify preprocessing, patch projection has received attention and helped Transformer achieve SOTA performance in downstream tasks (Wang et al., 2021b;a). To make full use of large-scale weakly supervised data, Radford et al. (2021) train a bi-encoder on 400 million pairs and demonstrates excellent performance in retrieval tasks. Another line of work is text-to-image synthesis. A bunch of works (Ramesh et al., 2021; Ding et al., 2021; Lin et al., 2021; Wu et al., 2021) incorporate Transformer with VQVAE (van den Oord et al., 2017) or VQGAN (Esser et al., 2021) to generate high-quality images with high resolution. However, the previously mentioned methods are limited in processing a single type of data, such as cross-modal data only or limited in their capabilities. Also, the discrepancy between pre-training and finetuning behaviors limits the transferability to open-ended data.

Unified Frameworks To pursue the unified models, Kaiser et al. (2017) demonstrate a uniform format to represent tasks. In NLP, recent studies unify diverse tasks, covering natural language understanding and generation, to text-to-text transfer (Raffel et al., 2020) or language modeling (Brown et al., 2020). Following this idea, Cho et al. (2021) and Yang et al. (2021) demonstrate text-generation-based multimodal pretrained models. Jaegle et al. (2021b) and Jaegle et al. (2021a) propose a simple framework that can process information from multiple modalities with a uniform byte-sequence representation. Hu & Singh (2021) and Singh et al. (2021) unify tasks of different modalities by designing various task-specific layers. Zhu et al. (2021) explore a retrieval-based unified paradigm. However,

the multimodal pretrained models suffer from performance degradation in downstream tasks, and they have no image generation capability.

3. OFA

In this work, we propose OFA, a unified Seq2Seq framework for the unification of I/O & architectures, tasks, and modalities. The overall framework is illustrated in Figure 2.

3.1. I/O & Architecture

I/O The most common practice of multimodal pretraining is the pretraining of Transformer models on image-text pair corpus at scale. This requires data preprocessing or modality-specific adaptors to enable the joint training of both visual and linguistic information with the Transformer architecture. Compared with the complex, resource&time-consuming object feature extraction, we aim for simplicity and directly use ResNet modules to convolve $x_v \in \mathbb{R}^{H \times W \times C}$ to P patch features of the hidden size, following Dai et al. (2021) and Wang et al. (2021b). As to processing the linguistic information, we follow the practice of GPT (Radford et al., 2018) and BART (Lewis et al., 2020) that apply byte-pair encoding (BPE) (Sennrich et al., 2016) to a given text sequence and then embed the subword tokens to features.

To process different modalities without task-specific output schema, it is essential to represent data of various modalities in a unified space. A possible solution is to discretize text, image, and object and represent them with tokens in a unified vocabulary. Recent advances in image quantization (van den Oord et al., 2017; Esser et al., 2021) have demonstrated effectiveness in text-to-image synthe-

sis (Ramesh et al., 2021; Lin et al., 2021; Ding et al., 2021; Zhang et al., 2021b), and thus we utilize this strategy for the target-side image representations. Sparse coding is effective in reducing the sequence length of image representation. For example, an image of the resolution of 256×256 is represented as a code sequence of the length of 16×16 . Each discrete code strongly correlates with the corresponding patch (Bao et al., 2021).

Apart from representing images, it is also essential to represent objects within images as there are a series of region-related tasks. Following Chen et al. (2021), we represent objects as a sequence of discrete tokens. To be more specific, for each object, we extract its label and its bounding box. The continuous corner coordinates (the top left and the bottom right) of the bounding box are uniformly discretized to integers as location tokens $\langle x_1, y_1, x_2, y_2 \rangle$. As to the object labels, they are intrinsically words and thus can be represented with BPE tokens.

Architecture Following the previous successful practices in multimodal pretraining (Chen et al., 2020c; Zhang et al., 2021a; Wang et al., 2021b), we choose Transformer as the backbone architecture, and we adopt the encoder-decoder framework as the unified architecture for all the pretraining, finetuning, and zero-shot tasks. Specifically, both the encoder and the decoder consist of stacks of Transformer layers. A Transformer encoder layer consists of a self attention and a feed-forward network (FFN), while a Transformer decoder layer consists of a self attention, an FFN and a cross attention for building the connection between the decoder and the encoder output representations. To stabilize training and accelerate convergence, we add head scaling to self attention, a post-attention layer normalization (LN) (Ba et al., 2016), and an LN following the first layer of FFN (Shleifer et al., 2021). For the positional information, we use two absolute position embeddings for text and images, respectively. Instead of simply adding the position embeddings, we decouple the position correlation between token embeddings and patch embeddings (Ke et al., 2020). In addition, we also use the 1D relative position bias for text (Raffel et al., 2020) and the 2D relative position bias for image (Wang et al., 2021b; Dai et al., 2021).

3.2. Tasks & Modalities

A unified framework is designed to provide architecture compatibility across different modalities and downstream tasks, so that opportunities can arise to generalize to unseen tasks within the same model. Then we have to represent the possible downstream tasks concerning different modalities in a unified paradigm. Therefore, an essential point for the design of pretraining tasks is the consideration of multitask and multimodality.

To unify tasks and modalities, we design a unified sequence-to-sequence learning paradigm for pretraining, finetuning, and inference on all tasks concerning different modalities. Both pretraining and downstream tasks of cross-modal and uni-modal understanding and generation are all formed as Seq2Seq generation. It is available to perform multitask pretraining on multimodal and uni-modal data, which can endow the model with comprehensive capabilities. Specifically, we share the identical schema across all tasks, while we specify handcrafted instructions for discrimination (Wei et al., 2021).

For cross-modal representation learning, we design 5 tasks, including visual grounding (VG), grounded captioning (GC), image-text matching (ITM), image captioning (IC), and visual question answering (VQA). For VG, the model learns to generate location tokens specifying the region position $\langle x_1, y_1, x_2, y_2 \rangle$ based on the input of the image x^i and the instruction “Which region does the text x^t describe?” where x^t refers to the region caption. GC is an inverse task of VG. The model learns to generate a description based on the input image x^i and the instruction “What does the region describe? region: $\langle x_1, y_1, x_2, y_2 \rangle$ ”. For ITM, we use each original image-text pair as the positive sample and construct a negative sample by pairing the image with a randomly substituted caption. The model learns to discriminate whether the given image and text are paired by learning to generate “Yes” or “No” based on the input image x^i and the instruction “Does the image describe x^t ?”. As to image captioning, this task can naturally adapt to the sequence-to-sequence format. The model learns to generate the caption based on the given image and the instruction “What does the image describe?”. For VQA, the model learns to generate the correct answer given an image and a question.

For uni-modal representation learning, we design 2 tasks for vision and 1 task for language, respectively. The model is pretrained with image infilling and object detection for visual representation learning. Recent advances in generative self-supervised learning for computer vision show that masked image modeling is an effective pretraining task (Bao et al., 2021; He et al., 2021a). In practice, we mask the middle part of the images as the input. The model learns to generate the sparse codes for the central part of the image based on the corrupted input and the specified instruction “What is the image in the middle part?”. We additionally add object detection to pretraining following Xu et al. (2021). The model learns to generate human-annotated object representations, i.e., the sequence of object position and label, based on the input image and the text “What are the objects in the image?” as the instruction. Both tasks strengthen the representation learning on both pixel and object levels. For language representation learning, following the practice of Lewis et al. (2020), we pretrain the unified model on plain text data with text infilling.

Table 1. Experimental results on cross-modal understanding tasks including VQA and visual entailment. Note that we report the best results from the previous SOTAs, and specifically SimVLM is a huge-size model comparable to ViT-Huge pretrained on 1.8B image-text pairs, and Florence is built with CoSwin-H and RoBERTa and it is pretrained on 900M image-text pairs.

Model	VQA		SNLI-VE	
	test-dev	test-std	dev	test
UNITER	73.8	74.0	79.4	79.4
OSCAR	73.6	73.8	-	-
VILLA	74.7	74.9	80.2	80.0
VL-T5	-	70.3	-	-
VinVL	76.5	76.6	-	-
UNIMO	75.0	75.3	81.1	80.6
ALBEF	75.8	76.0	80.8	80.9
METER	77.7	77.6	80.9	81.2
VLMo	79.9	80.0	-	-
SimVLM	80.0	80.3	86.2	86.3
Florence	80.2	80.4	-	-
OFA _{Tiny}	70.3	70.4	85.3	85.2
OFA _{Medium}	75.4	75.5	86.6	87.0
OFA _{Base}	78.0	78.1	89.3	89.2
OFA _{Large}	80.3	80.5	90.3	90.2
OFA	82.0	82.0	91.0	91.2

In this way, we unify multiple modalities and multiple tasks to a single model and a pretraining paradigm. OFA is pre-trained jointly with those tasks and data. Thus, it can perform different tasks concerning natural language, vision, and cross-modality.

3.3. Pretraining Datasets

We construct pretraining datasets by incorporating vision & language data (i.e., image-text pairs), vision data (i.e., raw image data, object-labeled data), and language data (i.e., plain texts). For replication, we only use datasets that are publicly available. We carefully filter our pretraining data and exclude images that appear in the validation and test sets of downstream tasks to avoid data leakage. We provide more details about pretraining datasets in Appendix A.1.

3.4. Training & Inference

We optimize the model with the cross-entropy loss. Given an input x , an instruction s and an output y , we train OFA by minimizing $\mathcal{L} = -\sum_{i=1}^{|y|} \log P_{\theta}(y_i | y_{<i}, x, s)$, where θ refers to the model parameters. For inference, we apply the decoding strategies, e.g., beam search, to enhance the quality of generation. However, this paradigm has several problems in classification tasks: 1. optimizing on the entire vocabulary is unnecessary and inefficient; 2. the model may generate invalid labels out of the closed label set during inference. To overcome these issues, we introduce a search strategy based on prefix tree (Trie (Cormen et al., 2009)). Experimental results show that the Trie-based search can enhance the performance of OFA on classification tasks.

Table 2. Experimental results on MSCOCO Image Captioning. We report the results on the Karpathy test split. Note that SimVLM and LEMON are huge-size models.

Model	Cross-Entropy Optimization				CIDEr Optimization			
	B	M	C	S	B	M	C	S
VL-T5	34.5	28.7	116.5	21.9	-	-	-	-
OSCAR	37.4	30.7	127.8	23.5	41.7	30.6	140.0	24.5
UNICORN	35.8	28.4	119.1	21.5	-	-	-	-
VinVL	38.5	30.4	130.8	23.4	41.0	31.1	140.9	25.2
UNIMO	39.6	-	127.7	-	-	-	-	-
LEMON	41.5	30.8	139.1	24.1	42.6	31.4	145.5	25.5
SimVLM	40.6	33.7	143.3	25.4	-	-	-	-
OFA _{Tiny}	35.9	28.1	119.0	21.6	38.1	29.2	128.7	23.1
OFA _{Medium}	39.1	30.0	130.4	23.2	41.4	30.8	140.7	24.8
OFA _{Base}	41.0	30.9	138.2	24.2	42.8	31.7	146.7	25.8
OFA _{Large}	42.4	31.5	142.2	24.5	43.6	32.2	150.7	26.2
OFA	43.9	31.8	145.3	24.8	44.9	32.5	154.9	26.6

Table 3. Experimental results on the 3 datasets of referring expression comprehension, namely RefCOCO, RefCOCO+, and RefCOCOg. We report the Acc@0.5 on different test splits of the datasets.

Model	RefCOCO			RefCOCO+			RefCOCOg	
	val	testA	testB	val	testA	testB	val-u	test-u
VL-T5	-	-	-	-	-	-	-	71.3
UNITER	81.41	87.04	74.17	75.90	81.45	66.70	74.86	75.77
VILLA	82.39	87.48	74.84	76.17	81.54	66.84	76.18	76.71
MDETR	86.75	89.58	81.41	79.52	84.09	70.62	81.64	80.89
UNICORN	88.29	90.42	83.06	80.30	85.05	71.88	83.44	83.93
OFA _{Tiny}	80.20	84.07	75.00	68.22	75.13	57.66	72.02	69.74
OFA _{Medium}	85.34	87.68	77.92	76.09	83.04	66.25	78.76	78.58
OFA _{Base}	88.48	90.67	83.30	81.39	87.15	74.29	82.29	82.31
OFA _{Large}	90.05	92.93	85.26	85.80	89.87	79.22	85.89	86.55
OFA	92.04	94.03	88.44	87.86	91.70	80.71	88.07	88.78

See Appendix B for more details.

3.5. Scaling Models

In order to investigate how OFA of different model sizes perform in downstream tasks, we have developed 5 versions of OFA models, scaling from 33M to 940M parameters, and we list their detailed hyperparameters in Table 11 in Appendix A.2.

To be more specific, we have built basic models of Base and Large sizes, OFA_{Base} and OFA_{Large}. As our network configuration is similar to BART (Lewis et al., 2020), their sizes are similar to those of BART_{Base} and BART_{Large}. Additionally, we have developed OFA of a larger size, which we name it OFA_{Huge}, or OFA without specific mentioning in the tables. Its size is comparable to that of SimVLM_{Huge} and ViT_{Huge}. To investigate whether smaller OFA can still reach satisfactory performance, we have developed OFA_{Medium} and OFA_{Tiny}, which are solely around half and less than 20% as large as OFA_{Base}.

4. Experiments

This section provides experimental details and analyses to demonstrate the effectiveness of OFA. See Appendix A for implementation details.

4.1. Results on Cross-modal Tasks

We evaluate our models on different cross-modal downstream tasks, covering cross-modal understanding and generation. Specifically, we implement experiments on multimodal understanding datasets including VQA_{v2} for visual question answering and SNLI-VE (Xie et al., 2019) for visual entailment, and multimodal generation including MSCOCO Image Caption (Chen et al., 2015) for image captioning, RefCOCO / RefCOCO+ / RefCOCOg (Yu et al., 2016; Mao et al., 2016) for referring expression comprehension as this task can be viewed as bounding box generation, and MSCOCO Image Caption for text-to-image generation. More details are provided in Appendix A.4.

Table 1 presents the performance of OFA and baseline models on VQA and SNLI-VE. In general, OFA achieves the best performance in both tasks with 82.0 on the VQA test set and 91.2 on the SNLI-VE test set. For smaller-size models, OFA_{Large} can outperform the recent SOTAs, e.g., VLMo and SimVLM, and OFA_{Base} can beat the SOTAs before the aforementioned two models in both tasks. This demonstrates that OFA can achieve superior performance on cross-modal understanding tasks and scaling up OFA can bring significant improvements, reflecting the strong potential of large-scale pretrained models.

Table 2 presents the performance of OFA and baseline models on the MSCOCO image captioning dataset. We report the results on the Karpathy test split, and we demonstrate the performance of models trained with Cross-Entropy optimization and additionally with CIDEr optimization based on reinforcement learning. In comparison with the previous SOTA SimVLM_{Huge} for Cross-Entropy optimization, OFA outperforms it by around 2 points in CIDEr evaluation. For CIDEr optimization, OFA of the 3 sizes all outperform the huge-size LEMON, and OFA demonstrates a new SOTA of 154.9 CIDEr score.

To evaluate the capability of visual grounding, we conduct experiments on RefCOCO, RefCOCO+, and RefCOCOg. While we unify locations to the vocabulary, visual grounding can be viewed as a sequence generation task. As there is only one target for each query, we limit the generation length to 4 in order to generate a bounding box by $\langle x_1, y_1, x_2, y_2 \rangle$. Experimental results in Table 3 show that OFA reaches the SOTA performance on the 3 datasets. Compared with the previous SOTA UNICORN (Yang et al., 2021), OFA achieves significant improvement with a gain of 3.61, 6.65 and 4.85 points on the testA sets of RefCOCO

Table 4. Experimental results on text-to-image generation. Models are evaluated on FID, CLIPSIM, and IS scores. OFA outperforms the baselines, including the concurrent SOTA NÜWA. We report the results of OFA_{Large}. Note that GLIDE additionally has 1.5B parameters for upsampling except for the 3.5B parameters.

Model	FID↓	CLIPSIM↑	IS↑
DALLE	27.5	-	17.9
CogView	27.1	33.3	18.2
GLIDE	12.2	-	-
Unifying	29.9	30.9	-
NÜWA	12.9	34.3	27.2
OFA	10.5	34.4	31.1

and RefCOCO+ as well as the test-u set of RefCOCOg.

Text-to-image generation is a challenging task even for pretrained models. As we pretrain OFA with the task “image-infilling”, i.e., recovering masked patches by generating the corresponding codes (Bao et al., 2021), and thus OFA is able to generate code. We thus directly finetune OFA on the MSCOCO Image Caption dataset for text-to-code generation. At the inference stage, we additionally transform the generated codes to an image with the code decoder. Specifically, we use the codes from VQGAN (Esser et al., 2021) following (Wu et al., 2021). Experimental results show that OFA outperforms the baselines in all the metrics. Note that increasing the sampling size during inference is expected to bring clear improvements on FID and IS. Compared with DALLE (Ramesh et al., 2021), CogView (Ding et al., 2021) and NÜWA (Wu et al., 2021), whose sampling sizes are 512, 60 and 60, respectively, OFA outperforms these SOTA methods on FID and IS with a much smaller sampling size 24. This illustrates that OFA has learned better correspondence among the query text, the image and the image codes.

We compare OFA with CogView and GLIDE on generation quality with normal and counterfactual queries.¹ Normal queries describe existing things in the real world, while counterfactual queries refer to those describing things that could only exist in our imagination. For normal queries, both CogView and OFA generate images semantically consistent with the given texts, in comparison with GLIDE. The generated examples from our model can provide more sophisticated details of objects, say the horse and the double-decker bus. For counterfactual queries, we find that OFA is the only one that can generate the three imaginary scenes, which indicates its imaginative power based on its strong capability to align text to the image. See Appendix C for more qualitative examples.

¹For more implementation details, please refer to Appendix A.4

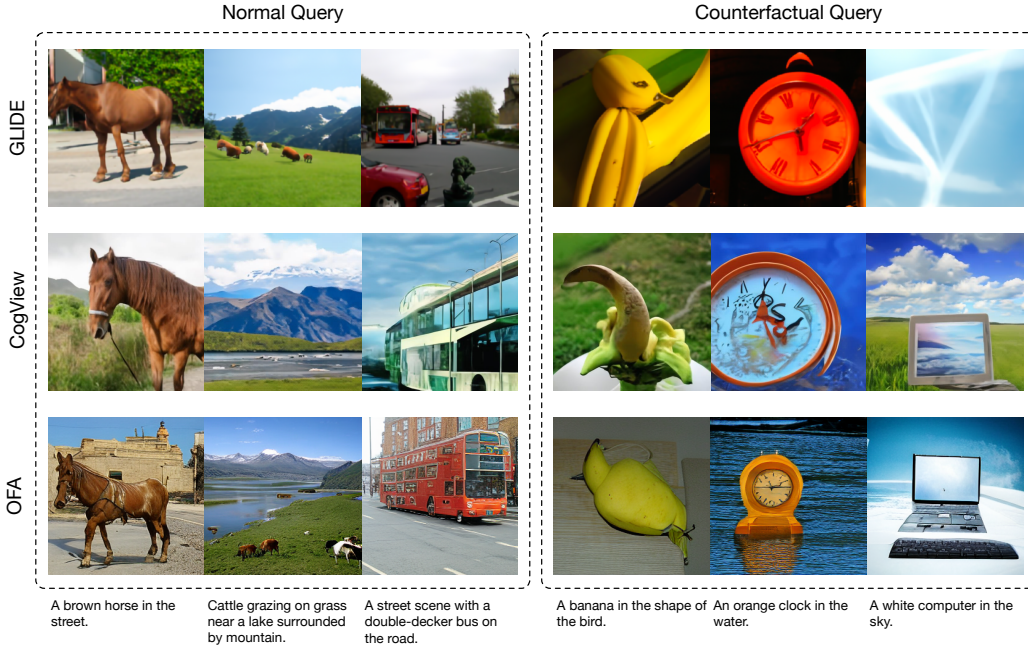


Figure 3. Qualitative comparison with state-of-the-art models for text-to-image generation. We present more qualitative examples of text-to-image generation for better demonstration in Appendix C.

4.2. Results on Uni-modal Tasks

As the design of OFA unifies different modalities, we evaluate its performance on unimodal tasks, namely tasks of natural language and computer vision. For natural language tasks, we evaluate OFA on 6 tasks of the GLUE benchmark (Wang et al., 2018) for natural language understanding and Gigaword abstractive summarization (Rush et al., 2015) for natural language generation. For computer vision, we evaluate OFA on the classic ImageNet-1K (Deng et al., 2009) dataset for image classification. More details are provided in Appendix A.4.

As OFA has been pretrained on the plain text data, it can be transferred to natural language downstream tasks. For natural language generation, it is essentially a sequence-to-sequence generation task, and for natural language understanding, typically text classification, we regard them as generation tasks where labels are essentially word sequences. Additionally, for each task, we design a manual instruction to indicate the model what types of questions it should answer. We list our instruction design in Appendix A.4.

We demonstrate that even a unified multimodal pretrained model can achieve highly competitive performance in natural language tasks. Specifically, in the evaluation of natural language understanding, OFA surpasses multimodal pretrained models by large margins in all tasks. In comparison with the state-of-the-art natural language pretrained models, including RoBERTa (Liu et al., 2019), XLNET (Yang et al.,

2019), ELECTRA (Clark et al., 2020), and DeBERTa (He et al., 2021b), OFA reaches a comparable performance. In the evaluation of natural language generation, OFA reaches a new state-of-the-art performance on the Gigaword dataset.

Also, OFA can reach a competitive performance in image classification. Table 7 shows the performance of OFA on image classification. OFA achieves higher accuracy than previous backbone models such as EfficientNet-B7 (Tan & Le, 2019) and ViT-L (Dosovitskiy et al., 2020). We also compare OFA with self-supervised pretraining models based on contrastive learning and masked image modeling. OFA outperforms contrastive-based models such as SimCLR (Chen et al., 2020a) and MoCo-v3 (Chen et al., 2020b; Chen & He, 2021) with similar parameters. Compared with pretrained models based on masked image modeling, e.g., BEiT-L (Bao et al., 2021) and MAE-L (He et al., 2021a), OFA can achieve similar performance.

These aforementioned results in both natural language and vision tasks indicate that a unified multimodal pretrained model is not only effective in multimodal tasks but also capable of tackling unimodal tasks, and in the future, it might be sufficient for such a model to solve complex tasks concerning different modality combinations.

4.3. Zero-shot Learning & Task Transfer

The instruction-guided pretraining enables OFA to perform zero-shot inference. Following Uni-Perceiver (Zhu et al.,

Table 5. Experimental results on the GLUE benchmark datasets (Wang et al., 2018). For comparison, we list the performance of multimodal pretrained models as well the recent SOTA models that were pretrained on natural language data only. Following (Liu et al., 2019), we finetune RTE and MRPC starting from the checkpoint finetuned on MNLI.

Model	SST-2	RTE	MRPC	QQP	MNLI	QNLI
<i>Multimodal Pretrained Baseline Models</i>						
VisualBERT	89.4	56.6	71.9	89.4	81.6	87.0
UNITER	89.7	55.6	69.3	89.2	80.9	86.0
VL-BERT	89.8	55.7	70.6	89.0	81.2	86.3
VilBERT	90.4	53.7	69.0	88.6	79.9	83.8
LXMERT	90.2	57.2	69.8	75.3	80.4	84.2
Uni-Perceiver	90.2	64.3	86.6	87.1	81.7	89.9
SimVLM	90.9	63.9	75.2	90.4	83.4	88.6
FLAVA	90.9	57.8	81.4	90.4	80.3	87.3
UNIMO	96.8	-	-	-	89.8	-
<i>Natural-Language-Pretrained SOTA Models</i>						
BERT	93.2	70.4	88.0	91.3	86.6	92.3
RoBERTa	96.4	86.6	90.9	92.2	90.2	93.9
XLNet	97.0	85.9	90.8	92.3	90.8	94.9
ELECTRA	96.9	88.0	90.8	92.4	90.9	95.0
DeBERTa	96.8	88.3	91.9	92.3	91.1	95.3
<i>Ours</i>						
OFA	96.6	91.0	91.7	92.5	90.2	94.8

Table 6. Experimental results on Gigaword abstractive summarization. We report performance on the ROUGE evaluation (Lin, 2004).

Model	Gigaword		
	ROUGE-1	ROUGE-2	ROUGE-L
BERTSHARE	38.13	19.81	35.62
MASS	38.73	19.71	35.96
UniLM	38.45	19.45	35.75
PEGASUS	39.12	19.86	36.24
ProphetNet	39.55	20.27	36.57
UNIMO	39.71	20.37	36.88
OFA	39.81	20.66	37.11

2021), we evaluate our model on the 6 tasks of the GLUE benchmark, including single-sentence classification and sentence pair classification. Table 8 demonstrates that OFA generally outperforms Uni-Perceiver. However, both models do not achieve satisfactory performance in sentence-pair classification (with Acc. < 60%). We hypothesize that the missing sentence-pair data in the pretraining dataset attributes to the performance.

Also, we find that the model performance is highly sensitive to the design of instructions. To obtain the best result, one should search a proper instruction template possibly from a large pool of candidates. A slight change to manual prompts or model parameters may drastically influence the model performance, which is not robust. We leave this issue to the future work.

We observe that the model can transfer to unseen tasks well with new task instructions. We design a new task called grounded question answering and present examples in Fig-

Table 7. ImageNet-1K finetuning results. All the listed models do not use extra labeled image classification samples during training for fair comparison.

Model	Top-1 Acc.
EfficientNet-B7	84.3
ViT-L/16	82.5
DINO	82.8
SimCLR v2	82.9
MoCo v3	84.1
BEiT ₃₈₄ -L/16	86.3
MAE-L/16	85.9
OFA	85.6

Table 8. Zero-shot performance on 6 GLUE subtasks and SNLI-VE.

Model	SST-2	RTE	MRPC	QQP	QNLI	MNLI	SNLI-VE
	Acc.	Acc.	F1	F1	Acc.	Acc.	Acc. (dev/test)
Uni-Perceiver	70.6	55.6	76.1	53.6	51.0	49.6	-
OFA _{Base}	71.6	56.7	79.5	54.0	51.4	37.3	49.71 / 49.18

ure 4. In this scenario, given a question about a certain region on the image, the model should provide a correct answer. We find that the model can achieve a satisfactory performance in this new task, which reflects its strong transferability. Besides, OFA can solve tasks with the out-of-domain input data. For example, OFA without finetuning achieves satisfactory performance in VQA for the out-of-domain images. Examples are demonstrated in Figure 5. OFA can also perform accurate visual grounding on the out-of-domain images, e.g., anime pictures, synthetic images, etc., and we demonstrate more examples on Figure 11 in Appendix C.

4.4. Ablation on Multitask Pretraining

Thanks to the unified framework, OFA has been pretrained on multiple tasks and thus endowed with comprehensive capabilities. However, the effects of each task are still undiscovered. We verify their effects on multiple downstream tasks, including image captioning, VQA, image classification, and text-to-image generation.

We first evaluate how uni-modal pretraining tasks influence the performance in both cross-modal and uni-modal tasks. Table 9 demonstrates our experimental results. We observe some interesting phenomena about the effects of uni-modal pretraining tasks. Text infilling brings improvement on image caption (+0.8 CIDEr) and VQA (+0.4 Acc.). Natural language pretraining betters the contextualized representation of language and thus enhances performance in cross-modal tasks. However, it is noticed that the language pretraining task may degrade the performance in image classification, leading to the decrease in ImageNet-1K (-1.0



Figure 4. Qualitative results on an unseen task called grounded QA. In this task, the model should answer a question about a certain region in the image. More samples are provided in Figure 10 in Appendix C.




Figure 5. Qualitative results on the unseen-domain VQA. During pretraining, only real-world photographs are used for VQA. We present cases of VQA on out-of-domain images, i.e., the iconic and sci-fi images, and demonstrate their capability of transferring to unseen domains. More samples are provided in Figure 9 in Appendix C.

Acc.). Also, it is interesting to find that it does not encourage improvement in text-to-image generation (−0.1 CLIPSIM). It may attribute to the simplicity of text in this task, which indicates that improved representation of language does not affect the performance. As to image infilling, it significantly improves the performance in image classification (+1.0 Acc.) and text-to-image generation (+0.6 CLIPSIM). Learning to recover images is an effective self-supervised task for image representation, and it also encourages the decoder’s ability to generate image codes. However, it hurts the performance in image captioning and VQA. Both tasks require a strong capability in generating texts, and the decoder’s learning of image generation naturally brings performance degradation in captioning (−0.7 CIDEr) and VQA (−0.3 Acc.).

Furthermore, we evaluate how multimodal tasks impact the performance. Previous studies have provided evidence of the contribution of conventional pretraining tasks, e.g., MLM, MOC, ITM, VQA, image captioning, etc. (Chen et al., 2020c; Zhang et al., 2021a). However, they miss other tasks, e.g., detection and visual grounding & grounded captioning. We conduct experiments on these tasks and find that tasks predicting regions are crucial to multimodal tasks, with a performance increase in image captioning (+2.3 CIDEr &

Table 9. Ablation results of OFA. All models are pretrained for 250k step. *w/o ground.* represents the removal of both visual grounding and grounded captioning tasks. Note that all models are only finetuned with the cross-entropy loss in image captioning.

Model	Caption CIDEr	VQA Test-dev	ImageNet Top-1 Acc.	Image Generation FID / CLIPSIM / IS
OFA _{Base}	135.6	76.0	82.2	20.8 / 31.6 / 21.5
<i>w/o text infill.</i>	134.8	75.6	83.2	20.3 / 31.7 / 21.8
<i>w/o image infill.</i>	136.3	76.3	81.8	23.2 / 31.0 / 20.0
<i>w/o det.</i>	133.3	75.4	81.4	20.9 / 31.5 / 21.6
<i>w/o ground.</i>	134.2	75.5	82.0	21.2 / 31.5 / 21.5

+1.4 CIDEr) and VQA (+0.6 Acc. & +0.5 Acc.). It suggests that detection and visual grounding & grounded captioning help the model grasp fined-grained alignments between vision and language. Region information contributes little to text-to-image generation (+0.1 CLIPSIM & +0.1 CLIPSIM), as this task requires far less text-region alignment information. We surprisingly find that detection can encourage the performance in visual understanding (+0.8 Acc.). It indicates that incorporating region information might be essential to visual understanding, especially on images with complex objects.

5. Conclusion

In this work, we propose **OFA**, a Task-Agnostic and Modality-Agnostic framework supporting Task Comprehensiveness. OFA achieves the unification in architecture, tasks and modalities, and thus is capable of multimodal & uni-modal understanding and generation, without specification in additional layers or tasks. Our experiments show that OFA creates new SOTAs in a series of tasks, including image captioning, VQA, visual entailment, and referring expression comprehension. OFA also demonstrates a comparable performance with language / vision pretrained SOTA models in uni-modal understanding and generation tasks, e.g., GLUE, abstractive summarization, and image classification. We provide a further analysis to demonstrate its capability in zero-shot learning and domain & task transfer, and we also verify the effectiveness of pretraining tasks.

In the future, we will continue exploring the issues discovered in this work. Also, we endeavor to figure out a reasonable solution to building an omni-model essentially generalizable to the complex real world.

Acknowledgments

We would like to thank Jie Zhang, Yong Li, Jiamang Wang, Shao Yuan, and Zheng Cao for their support to this project, and we would like to thank Guangxiang Zhao and Fei Sun for their insightful comments to our paper.

References

- Anderson, P., Fernando, B., Johnson, M., and Gould, S. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, pp. 382–398. Springer, 2016.
- Ba, L. J., Kiros, J. R., and Hinton, G. E. Layer normalization. *CoRR*, abs/1607.06450, 2016.
- Banerjee, S. and Lavie, A. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.
- Bao, H., Dong, L., and Wei, F. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Changpinyo, S., Sharma, P., Ding, N., and Soricut, R. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3558–3568, 2021.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020a.
- Chen, T., Saxena, S., Li, L., Fleet, D. J., and Hinton, G. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021.
- Chen, X. and He, K. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., and Zitnick, C. L. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- Chen, X., Fan, H., Girshick, R., and He, K. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- Chen, Y.-C., Li, L., Yu, L., Kholy, A. E., Ahmed, F., Gan, Z., Cheng, Y., and Liu, J. Uniter: Universal image-text representation learning. In *ECCV*, 2020c.
- Cho, J., Lei, J., Tan, H., and Bansal, M. Unifying vision-and-language tasks via text generation. In *ICML*, 2021.
- Clark, K., Luong, M., Le, Q. V., and Manning, C. D. ELECTRA: pre-training text encoders as discriminators rather than generators. In *8th International Conference on Learning Representations, ICLR 2020*. OpenReview.net, 2020.
- Cobbe, K., Kosaraju, V., Bavarian, M., Hilton, J., Nakano, R., Hesse, C., and Schulman, J. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to algorithms*. MIT press, 2009.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- Dai, Z., Liu, H., Le, Q. V., and Tan, M. Coatnet: Marrying convolution and attention for all data sizes. *arXiv preprint arXiv:2106.04803*, 2021.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T. (eds.), *NAACL-HLT 2019*, pp. 4171–4186. Association for Computational Linguistics, 2019.
- Ding, M., Yang, Z., Hong, W., Zheng, W., Zhou, C., Yin, D., Lin, J., Zou, X., Shao, Z., Yang, H., et al. Cogview: Mastering text-to-image generation via transformers. *arXiv preprint arXiv:2105.13290*, 2021.
- Dong, L., Yang, N., Wang, W., Wei, F., Liu, X., Wang, Y., Gao, J., Zhou, M., and Hon, H. Unified language model pre-training for natural language understanding and generation. In *NeurIPS 2019*, pp. 13042–13054, 2019.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Esser, P., Rombach, R., and Ommer, B. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12873–12883, 2021.

- Gan, Z., Chen, Y.-C., Li, L., Zhu, C., Cheng, Y., and Liu, J. Large-scale adversarial training for vision-and-language representation learning. [ArXiv](#), abs/2006.06195, 2020.
- Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. The pile: An 800gb dataset of diverse text for language modeling. [arXiv preprint arXiv:2101.00027](#), 2020.
- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., and Parikh, D. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In [Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition](#), pp. 6904–6913, 2017.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. [arXiv preprint arXiv:2006.07733](#), 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In [CVPR 2016](#), pp. 770–778, 2016.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. [arXiv preprint arXiv:2111.06377](#), 2021a.
- He, P., Liu, X., Gao, J., and Chen, W. Deberta: decoding-enhanced bert with disentangled attention. In [9th International Conference on Learning Representations, ICLR 2021](#). OpenReview.net, 2021b.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. [Advances in neural information processing systems](#), 30, 2017.
- Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In [International Conference on Machine Learning](#), pp. 2790–2799. PMLR, 2019.
- Hu, R. and Singh, A. Unit: Multimodal multitask learning with a unified transformer. [arXiv preprint arXiv:2102.10772](#), 2021.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., and Weinberger, K. Q. Deep networks with stochastic depth. In [ECCV](#), 2016.
- Huang, Y., Xue, H., Liu, B., and Lu, Y. Unifying multi-modal transformer for bi-directional image and text generation. In [Proceedings of the 29th ACM International Conference on Multimedia](#), pp. 1138–1147, 2021.
- Huang, Z., Zeng, Z., Liu, B., Fu, D., and Fu, J. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. [ArXiv](#), abs/2004.00849, 2020.
- Hudson, D. A. and Manning, C. D. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In [CVPR 2019](#), pp. 6700–6709, 2019.
- Jaegle, A., Borgeaud, S., Alayrac, J., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shihamer, E., Hénaff, O. J., Botvinick, M. M., Zisserman, A., Vinyals, O., and Carreira, J. Perceiver IO: A general architecture for structured inputs & outputs. [CoRR](#), abs/2107.14795, 2021a.
- Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., and Carreira, J. Perceiver: General perception with iterative attention. In Meila, M. and Zhang, T. (eds.), [ICML 2021](#), volume 139 of [Proceedings of Machine Learning Research](#), pp. 4651–4664. PMLR, 2021b.
- Kaiser, L., Gomez, A. N., Shazeer, N., Vaswani, A., Parmar, N., Jones, L., and Uszkoreit, J. One model to learn them all. [arXiv preprint arXiv:1706.05137](#), 2017.
- Karpathy, A. and Fei-Fei, L. Deep visual-semantic alignments for generating image descriptions. In [Proceedings of the IEEE conference on computer vision and pattern recognition](#), pp. 3128–3137, 2015.
- Ke, G., He, D., and Liu, T.-Y. Rethinking positional encoding in language pre-training. In [International Conference on Learning Representations](#), 2020.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L., Shamma, D. A., Bernstein, M. S., and Fei-Fei, L. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. [International Journal of Computer Vision](#), 123(1):32–73, 2017.
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al. The open images dataset v4. [International Journal of Computer Vision](#), 128(7):1956–1981, 2020.
- Lester, B., Al-Rfou, R., and Constant, N. The power of scale for parameter-efficient prompt tuning. [arXiv preprint arXiv:2104.08691](#), 2021.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In [ACL 2020](#), July 2020.

- Li, G., Duan, N., Fang, Y., Jiang, D., and Zhou, M. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. *CoRR*, abs/1908.06066, 2019a.
- Li, L. H., Yatskar, M., Yin, D., Hsieh, C.-J., and Chang, K.-W. Visualbert: A simple and performant baseline for vision and language. *ArXiv*, abs/1908.03557, 2019b.
- Li, W., Gao, C., Niu, G., Xiao, X., Liu, H., Liu, J., Wu, H., and Wang, H. UNIMO: towards unified-modal understanding and generation via cross-modal contrastive learning. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *ACL/IJCNLP 2021*, pp. 2592–2607. Association for Computational Linguistics, 2021.
- Li, X., Yin, X., Li, C., Hu, X., Zhang, P., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., Choi, Y., and Gao, J. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020.
- Liang, X., Wu, L., Li, J., Wang, Y., Meng, Q., Qin, T., Chen, W., Zhang, M., and Liu, T. R-drop: Regularized dropout for neural networks. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *NeurIPS 2021*, pp. 10890–10905, 2021.
- Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, Barcelona, Spain, July 2004. Association for Computational Linguistics.
- Lin, J., Yang, A., Zhang, Y., Liu, J., Zhou, J., and Yang, H. Interbert: Vision-and-language interaction for multi-modal pretraining. *arXiv preprint arXiv:2003.13198*, 2020.
- Lin, J., Men, R., Yang, A., Zhou, C., Ding, M., Zhang, Y., Wang, P., Wang, A., Jiang, L., Jia, X., et al. M6: A chinese multimodal pretrainer. *arXiv preprint arXiv:2103.00823*, 2021.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *ICLR 2019*, 2019.
- Lu, J., Batra, D., Parikh, D., and Lee, S. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- Lu, J., Goswami, V., Rohrbach, M., Parikh, D., and Lee, S. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10437–10446, 2020.
- Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A. L., and Murphy, K. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 11–20, 2016.
- Ordonez, V., Kulkarni, G., and Berg, T. L. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS 2011*, pp. 1143–1151, 2011.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision. In Meila, M. and Zhang, T. (eds.), *ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21 (140):1–67, 2020.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021.
- Rennie, S. J., Marcheret, E., Mroueh, Y., Ross, J., and Goel, V. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7008–7024, 2017.
- Rush, A. M., Chopra, S., and Weston, J. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 379–389, 2015.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. *Advances in neural information processing systems*, 29:2234–2242, 2016.

- Sanh, V., Webson, A., Raffel, C., Bach, S. H., Sutawika, L., Alyafeai, Z., Chaffin, A., Stiegler, A., Scao, T. L., Raja, A., et al. Multitask prompted training enables zero-shot task generalization. [arXiv preprint arXiv:2110.08207](#), 2021.
- Schneider, S., Baevski, A., Collobert, R., and Auli, M. wav2vec: Unsupervised pre-training for speech recognition. [arXiv preprint arXiv:1904.05862](#), 2019.
- Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., and Komatsuzaki, A. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. [arXiv preprint arXiv:2111.02114](#), 2021.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In [Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pp. 1715–1725, 2016.
- Shao, S., Li, Z., Zhang, T., Peng, C., Yu, G., Zhang, X., Li, J., and Sun, J. Objects365: A large-scale, high-quality dataset for object detection. In [Proceedings of the IEEE/CVF International Conference on Computer Vision](#), pp. 8430–8439, 2019.
- Sharma, P., Ding, N., Goodman, S., and Soricut, R. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In [ACL 2018](#), pp. 2556–2565, 2018.
- Shleifer, S., Weston, J., and Ott, M. Normformer: Improved transformer pretraining with extra normalization. [arXiv preprint arXiv:2110.09456](#), 2021.
- Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., and Kiela, D. Flava: A foundational language and vision alignment model. [arXiv preprint arXiv:2112.04482](#), 2021.
- Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., and Dai, J. Vi-bert: Pre-training of generic visual-linguistic representations. In [International Conference on Learning Representations](#), 2019.
- Sun, Y., Wang, S., Li, Y., Feng, S., Chen, X., Zhang, H., Tian, X., Zhu, D., Tian, H., and Wu, H. ERNIE: enhanced representation through knowledge integration. [CoRR](#), abs/1904.09223, 2019a.
- Sun, Y., Wang, S., Li, Y., Feng, S., Tian, H., Wu, H., and Wang, H. ERNIE 2.0: A continual pre-training framework for language understanding. [CoRR](#), abs/1907.12412, 2019b.
- Tan, H. and Bansal, M. Lxmert: Learning cross-modality encoder representations from transformers. In [Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing \(EMNLP-IJCNLP\)](#), pp. 5100–5111, 2019.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In [International Conference on Machine Learning](#), pp. 6105–6114. PMLR, 2019.
- Thomee, B., Shamma, D. A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., and Li, L.-J. Yfcc100m: The new data in multimedia research. [Communications of the ACM](#), 59(2):64–73, 2016.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In [NIPS](#), 2017.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In [NeurIPS 2017](#), pp. 5998–6008, 2017.
- Vedantam, R., Lawrence Zitnick, C., and Parikh, D. Cider: Consensus-based image description evaluation. In [Proceedings of the IEEE conference on computer vision and pattern recognition](#), pp. 4566–4575, 2015.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Glue: A multi-task benchmark and analysis platform for natural language understanding. [arXiv preprint arXiv:1804.07461](#), 2018.
- Wang, W., Bao, H., Dong, L., and Wei, F. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. [ArXiv](#), abs/2111.02358, 2021a.
- Wang, Z., Yu, J., Yu, A. W., Dai, Z., Tsvetkov, Y., and Cao, Y. Simvlm: Simple visual language model pretraining with weak supervision. [ArXiv](#), abs/2108.10904, 2021b.
- Wei, J., Bosma, M., Zhao, V. Y., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners. [arXiv preprint arXiv:2109.01652](#), 2021.
- Wu, C., Liang, J., Ji, L., Yang, F., Fang, Y., Jiang, D., and Duan, N. N\” uwa: Visual synthesis pre-training for neural visual world creation. [arXiv preprint arXiv:2111.12417](#), 2021.
- Xie, N., Lai, F., Doran, D., and Kadav, A. Visual entailment: A novel task for fine-grained image understanding. [arXiv preprint arXiv:1901.06706](#), 2019.
- Xu, H., Yan, M., Li, C., Bi, B., Huang, S., Xiao, W., and Huang, F. E2e-vlp: End-to-end vision-language pre-training enhanced by visual learning. [arXiv preprint arXiv:2106.01804](#), 2021.

- Yang, Z., Dai, Z., Yang, Y., Carbonell, J. G., Salakhutdinov, R., and Le, Q. V. Xlnet: Generalized autoregressive pretraining for language understanding. In NeurIPS 2019, pp. 5754–5764, 2019.
- Yang, Z., Gan, Z., Wang, J., Hu, X., Ahmed, F., Liu, Z., Lu, Y., and Wang, L. Crossing the format boundary of text and boxes: Towards unified vision-language modeling. ArXiv, abs/2111.12085, 2021.
- Yu, F., Tang, J., Yin, W., Sun, Y., Tian, H., Wu, H., and Wang, H. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 35, pp. 3208–3216, 2021.
- Yu, L., Poirson, P., Yang, S., Berg, A. C., and Berg, T. L. Modeling context in referring expressions. In European Conference on Computer Vision, pp. 69–85. Springer, 2016.
- Yuan, L., Chen, D., Chen, Y.-L., Codella, N. C. F., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., Liu, C., Liu, M., Liu, Z., Lu, Y., Shi, Y., Wang, L., Wang, J., Xiao, B., Xiao, Z., Yang, J., Zeng, M., Zhou, L., and Zhang, P. Florence: A new foundation model for computer vision. ArXiv, abs/2111.11432, 2021.
- Yun, S., Han, D., Chun, S., Oh, S. J., Yoo, Y., and Choe, J. Cutmix: Regularization strategy to train strong classifiers with localizable features. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, pp. 6022–6031. IEEE, 2019.
- Zhang, H., Cissé, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018.
- Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., and Gao, J. Vinvl: Revisiting visual representations in vision-language models. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5575–5584, 2021a.
- Zhang, Z., Ma, J., Zhou, C., Men, R., Li, Z., Ding, M., Tang, J., Zhou, J., and Yang, H. M6-ufc: Unifying multi-modal controls for conditional image synthesis. arXiv preprint arXiv:2105.14211, 2021b.
- Zhao, G., Yang, W., Ren, X., Li, L., and Sun, X. Well-classified examples are underestimated in classification with deep neural networks. CoRR, abs/2110.06537, 2021.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. Random erasing data augmentation. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, pp. 13001–13008, 2020.
- Zhou, L., Palangi, H., Zhang, L., Hu, H., Corso, J. J., and Gao, J. Unified vision-language pre-training for image captioning and VQA. In AAAI 2020, pp. 13041–13049, 2020.
- Zhu, X., Zhu, J., Li, H., Wu, X., Wang, X., Li, H., Wang, X., and Dai, J. Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. arXiv preprint arXiv:2112.01522, 2021.

A. Implementation Details

A.1. Pretraining Datasets

We construct pretraining datasets by incorporating vision & language data (i.e., image-text pairs), vision data (i.e., raw image data, object-labeled data), and language data (i.e., plain texts). For replication, the pretraining datasets are publicly available. We carefully filter our pretraining data and exclude images that appear in the validation and test sets of downstream tasks to avoid data leakage. The statistics on the pretraining datasets are listed in Table 10.

Cross-modal Data For vision & language pretraining, we mainly apply image-text pairs, including image-caption pairs, image-QA pairs, and image-region pairs, as the pretraining data. For the pretraining tasks of image captioning and image-text matching, we collect Conceptual Caption 12M (CC12M) (Changpinyo et al., 2021), Conceptual Captions (CC3M) (Sharma et al., 2018), SBU (Ordonez et al., 2011), MSCOCO image captions (COCO) (Chen et al., 2015), and Visual Genome Captions (VG Captions) (Krishna et al., 2017). Specifically, the part of data from VG requires some additional processing. As texts in VG captions describe local regions on the images, we retrieve regions with area larger than 16,384 pixels and construct region-caption pairs. For visual question answering, we collect VQA_{v2} (Goyal et al., 2017), VG-QA (Krishna et al., 2017), as well as GQA (Hudson & Manning, 2019). VQA_{v2} is a visual question answering dataset with real-world photographs from COCO. VG-QA is also a visual question answering dataset with real-world photographs from VG. The questions of VG-QA are related to specific regions on the images. GQA is a large VQA dataset featuring compositional questions. The images of GQA are also collected from VG. For visual grounding and grounded captioning, we collect data from RefCOCO (Yu et al., 2016), RefCOCO+ (Yu et al., 2016), RefCOCOG (Mao et al., 2016) and VG captions. Additional processing is applied to VG Captions for this task. Specifically, we use the images of VG that contain regions with area smaller than 16,384 pixels for Visual Grounding, in order to encourage model to grasp fine-grained alignments between vision and language.

Uni-modal Data Uni-modal data include vision and language data. Vision data consists of raw images for image infilling and object-labeled images for object detection. For image infilling, we collect raw images from OpenImages, YFCC100M (Thomee et al., 2016) and ImageNet-21K (Deng et al., 2009), and exclude annotations. Thus the model is unable to access labels in the pretraining stage. For object detection, we collect OpenImages (Kuznetsova et al., 2020), Object365 (Shao et al., 2019), VG and COCO for object detection. Language data consist of plain texts, i.e., passages consisting of sentences. We use around 140GB of data from Pile (Gao et al., 2020) to leverage its diversity.

Table 10. Statistics on the datasets of pretraining tasks. For language data, 140G* represents the storage space of the plain texts.

Type	Pretraining Task	Source	#Image	#Label
Vision&Language	Image Captioning Image-Text Matching	CC12M, CC3M, SBU, COCO, VG-Cap	14.78M	15.25M
	Visual Question Answering	VQA _{v2} , VG-QA, GQA	178K	2.92M
	Visual Grounding Grounded Captioning	RefCOCO, RefCOCO+, RefCOCOG, VG-Cap	131K	3.20M
Vision	Detection	OpenImages, Object365, VG, COCO	2.98M	3.00M
	Image Infilling	OpenImages, YFCC100M, ImageNet-21K	36.27M	-
Language	Masked Language Modeling	Pile (Filter)	-	140G*

A.2. Model Configuration

We provide the details of our proposed OFA of 5 different sizes, ranging from 33M to 930M parameters in Table 11. The details include the number of total parameters, the choice of visual backbone, the hidden size, the number of layers in the encoder and decoder, the number of attention heads, etc.

Table 11. Detailed hyperparameters of OFA model configuration. We list the configuration for OFA of 5 different sizes.

Model	#Param.	Backbone	Hidden size	Intermediate Size	#Head	#Enc. Layers	#Dec. Layers
OFA _{Tiny}	33M	ResNet50	256	1024	4	4	4
OFA _{Medium}	93M	ResNet101	512	2048	8	4	4
OFA _{Base}	182M	ResNet101	768	3072	12	6	6
OFA _{Large}	472M	ResNet152	1024	4096	16	12	12
OFA _{Huge}	930M	ResNet152	1280	5120	16	24	12

A.3. Pretraining Details

For the image processing, we first resize and crop the images into different resolutions, 256×256 for OFA_{Tiny} and OFA_{Medium}, 384×384 for OFA_{Base}, 480×480 for OFA_{Large} and OFA_{Huge}, with a fixed patch size of 16×16 . Note that training OFA_{Large} and OFA_{Huge} are time and computation consuming, we first train them with images of the resolution of 384×384 and 256×256 , and continue pretraining with images of the resolution of 480×480 .

For each patch, we obtain its feature vector with the first three blocks of ResNet (He et al., 2016). The ResNet module is jointly trained along with the transformer module. Note that through extensive experiments we find that random sampling patches (Huang et al., 2020) does not bring additional benefits in our scenario. For the text processing, we tokenize the texts with the same BPE Tokenizer (Sennrich et al., 2016) as BART (Lewis et al., 2020). The maximum text sequence length of both encoder and decoder is set to 256. We share parameters between the embedding and the decoder softmax output layer.

From our preliminary experiments, we find that the initialization for Transformer plays an important role. For OFA_{Base} and OFA_{Large}, we initialize the transformer with most of the weights of BART_{Base} and BART_{Large} considering the slight difference between OFA Transformer and BART as described in Sec 3.1. For OFA of the other sizes, we pretrain language models with the same pretraining strategy with BART and use the pretrained weights to initialize the Transformer in OFA.

We use the AdamW (Loshchilov & Hutter, 2019) optimizer with $(\beta_1, \beta_2) = (0.9, 0.999)$ and $\epsilon = 1e-8$ to pretrain our models. We set the peak learning rate to $2e-4$, and apply a scheduler with linear decay with a warmup ratio of 0.01 to control the learning rate. For regulation, we set dropout to 0.1 and use weight decay with 0.01. We employ stochastic depth (Huang et al., 2016) with a 0.1 rate (applied to encoder and decoder except for convolution blocks). We mix all the pretraining data within each batch, which contains 2,048 vision&language samples, 256 object detection samples, 256 image-only samples and 512 text-only samples. All models are pretrained for at least 300K steps except the models used for ablation study.

A.4. Details of Downstream Tasks

We verify the capability of OFA on various downstream tasks in both finetuning and zero-shot settings. We design various task-specific instructions to transfer the knowledge learned from pretraining to downstream tasks effectively. The instructions of different tasks are listed in Table 12. For finetuning, if not specified, the input image resolution is set to 480×480 . The experimental details of different downstream tasks, including both multimodal and uni-modal tasks, are listed below:

Image Captioning Image captioning is a standard vision & language task that requires models to generate an appropriate and fluent caption for an image. We adopt the most widely used MSCOCO Image Caption dataset (Chen et al., 2015) to evaluate the multi-modal generation capability of OFA. We report BLEU-4 (Papineni et al., 2002), METEOR (Banerjee & Lavie, 2005), CIDEr (Vedantam et al., 2015), and SPICE (Anderson et al., 2016) scores on the Karpathy test split (Karpathy & Fei-Fei, 2015). Following the previous standard practice, we first finetune OFA with the cross-entropy loss for 2 epochs with a batch size of 128 and a learning rate of $1e-5$, and label smoothing is set to 0.1. We then finetune the model with CIDEr optimization for 3 epochs with a batch size of 64, and disable dropout and stochastic depth. We report scores at the two stages.

Visual Question Answering Visual question answering (VQA) is a cross-modal task that requires the models to answer the question given an image. Previous works such as VLMO (Wang et al., 2021a) or SimVLM (Wang et al., 2021b) define VQA as a classification task. They use a linear output layer to predict the probability of each candidate answer on a given set. In contrast with these studies, to adapt the generative OFA model to VQA benchmark, we use the Trie-based search strategy

Table 12. Instructions for downstream tasks.

Task	Dataset	Instruction	Target
Image Captioning	COCO	[Image] What does the image describe?	{Caption}
Visual Question Answering	VQA	[Image] {Question}	{Answer}
Visual Entailment	SNLI-VE	[Image] Can image and text1 “{Text1}” imply text2 “{Text2}”?	Yes/No/Maybe
Referring Expression Comprehension	RefCOCO, RefCOCO+, RefCOCOg	[Image] Which region does the text “{Text}” describe?	{Location}
Image Generation	COCO	What is the complete image? caption: {Caption}	{Image}
Image Classification	ImageNet-1K	[Image] What does the image describe?	{Label}
Single-Sentence Classification	SST-2	Is the sentiment of text “{Text}” positive or negative?	Positive/Negative
Sentence-Pair Classification	RTE	Can text1 “{Text1}” imply text2 “{Text2}”?	Yes/No
	MRPC	Does text1 “{Text1}” and text2 “{Text2}” have the same semantics?	Yes/No
	QQP	Is question “{Question1}” and question “{Question2}” equivalent?	Yes/No
	MNLI	Can text1 “{Text1}” imply text2 “{Text2}”?	Yes/No/Maybe
	QNLI	Does “{Text}” contain the answer to question “{Question}”?	Yes/No
Text Summarization	Gigaword	What is the summary of article “{Article}”?	{Summary}

mentioned in Sec. 3.4 to ensure that the answer generated by OFA is constrained in the candidate set. We evaluate our model with other baselines on the commonly used VQAv2 dataset (Goyal et al., 2017). Accuracy scores on both test-dev and test-std sets are reported. The OFA models of all the reported sizes are finetuned for 40,000 steps with a batch size of 512. The learning rate is $5e-5$ with the label smoothing of 0.1. When finetuning OFA_{Large} and OFA_{Huge}, we increase the image resolution from 480 to 640. We employ linear interpolation for the absolute positional embedding following Dosovitskiy et al. (2020) when transferring the pretrained OFA to VQA finetuning. During Trie-based searching, we constrain the generated answers over the most frequent 3,129 answer candidates. Exponential moving average (EMA) with decay rate 0.9999 is employed in finetuning.

Visual Entailment Visual entailment requires the model to evaluate how the given image and text are semantically correlated, i.e., entailment, neutral, or contradiction. We perform experiments on the SNLI-VE dataset (Xie et al., 2019). The image premise, text premise and text hypothesis are fed to the encoder, and the decoder generates appropriate labels. To transfer the knowledge learned by pretraining to this task, we convert the labels “entailment / neutral / contradiction” to “yes / maybe / no”. We also use the Trie-based search strategy to constrain the generated labels over the candidate set. We report accuracy on both dev and test sets. The OFA model is finetuned for 6 epochs with a learning rate of $2e-5$ and a batch size of 256.

Referring Expression Comprehension Referring expression comprehension requires models to locate an image region described by a language query. Different from the approach taken by most previous methods (Lu et al., 2019; Chen et al., 2020c) which ranks a set of candidate bounding boxes detected by a pretrained object detector, our method directly predicts the best matching bounding box without any proposals. We perform experiments on RefCOCO (Yu et al., 2016), RefCOCO+ (Yu et al., 2016), and RefCOCOg (Mao et al., 2016). Consistent with other downstream tasks, we formulate referring expression comprehension as a conditional sequence generation task. In detail, given an image and a language query, OFA generates the box sequence (e.g., $\langle x_1, y_1, x_2, y_2 \rangle$) in an autoregressive manner. We report the standard metric Acc@0.5 on the validation and test sets. For finetuning, the input image resolution is set to 512×512 . We finetune the OFA model on each dataset for about 10 epochs with a batch size of 128. The learning rate is $3e-5$ with the label smoothing of 0.1. Each query only corresponds to an image region, so we limit the maximum generated length to 4 during inference.

Image Generation Following the same setting with Wu et al. (2021), we train our model on the MS COCO train split and evaluate our model on the validation split by randomly sampling 30,000 images. We use Fréchet Inception Distance (FID) (Heusel et al., 2017) and Inception Score (IS) (Salimans et al., 2016) to evaluate the quality of the images. Following the

previous studies (Huang et al., 2021; Wu et al., 2021), we also compute CLIP Similarity Score (CLIPSIM) to evaluate the semantic similarity between the query text and the generated images. During finetuning, OFA learns to generate the image code sequence according to the given text query only. The model is first finetuned with the cross-entropy loss and then with CLIPSIM optimization following Rennie et al. (2017) and Huang et al. (2021). In the first stage, we finetune the OFA model for about 50 epochs with a batch size of 512 and a learning rate of $1e - 3$. In the second stage, the model is finetuned for extra 5000 steps with a batch size of 32 and a learning rate of $1e - 6$. During the evaluation, we sample 24 images with the resolution of 256×256 for each query and choose the best one using the pretrained CLIP model (Radford et al., 2021).

For the case study, we compare OFA with CogView and GLIDE. CogView provides an API website². Note that this API samples 8 images of resolution of 512×512 for each query. We select the first one of generated images and resize it to the resolution of 256×256 . GLIDE provides a Colab notebook.³ Note that the only publicly available GLIDE model is of *base* size ($\sim 385M$).

Image Classification We provide finetuning results on ImageNet-1K (Deng et al., 2009) following recent studies in self-supervised learning for computer vision. During finetuning and inference, a Trie-based search strategy is employed to constrain the generated text into the set of 1,000 candidate labels. We finetune OFA for 32 epochs and a batch size of 256. The learning rate is $5e - 5$. The ratio for label smoothing is 0.1. The encouraging loss proposed in (Zhao et al., 2021) is employed with the hyperparameter LE set to 0.75. Following Bao et al. (2021), we use the same random resize cropping, random flipping, RandAug (Cubuk et al., 2020) and random erasing (Zhong et al., 2020) transformations as data augmentation strategies. Mixup (Zhang et al., 2018) and CutMix (Yun et al., 2019) are used with overall 0.5 probability to be performed on each batch and alpha is 0.8 and 1.0, respectively. To adapt the mixed soft target of Mixup and CutMix into generation paradigm during finetuning, we run the decoder twice each with one of the target sequences to be mixed and sum the loss weighted by the mixing ratio.

Natural Language Understanding To verify the natural language understanding ability of OFA, we select 6 language understanding tasks from the GLUE benchmark (Wang et al., 2018), including single-sentence classification tasks and sentence-pair classification tasks. To adapt to sentence-pair classification, previous models (Devlin et al., 2019; Liu et al., 2019) usually use segment embeddings to distinguish different sentences. Unlike those models, OFA can apply the model to sentence-pair classification tasks by constructing appropriate instructions without introducing additional segment embeddings. For the hyper-parameters of finetuning, we tune the training epochs among $\{5, 7, 10\}$, learning rate among $\{3e - 5, 5e - 5, 6e - 5, 7e - 5, 1e - 4\}$, batch size among $\{32, 64, 128\}$, weight decay among $\{0.01, 0.05\}$, and dropout rate among $\{0.0, 0.1\}$. We report the best performance on the development set for each task.

Natural Language Generation We verify the natural language generation ability of OFA on the Gigaword dataset (Rush et al., 2015). We report ROUGE-1 / ROUGE-2 / ROUGE-L to evaluate the generation results following Rush et al. (2015). We finetune the OFA models for around 6 epochs with a batch size of 512. The learning rate is $1e - 4$ with the label smoothing of 0.1, and the maximum input text sequence length is set to 512. R-drop (Liang et al., 2021) is applied for finetuning. In the stage of inference, we set the length penalty to 0.7 and beam size to 6, and limit the maximum generated length to 32.

B. Trie-based Search

This section describes how to use Trie-based search to improve model performance on downstream classification tasks. When dealing with classification tasks, we first construct a Trie where nodes are annotated with tokens from the candidate label-set. During finetuning, the model computes the log-probabilities of the target tokens based on their positions on the Trie. As shown in Figure 6, when computing the log-probabilities of the target token “sky”, we only consider tokens in $\{\text{“sky”}, \text{“ocean”}\}$ and forcefully set the logits for all invalid tokens to $-\infty$. During inference, we constrain the generated labels over the candidate set. As shown in Table 13, Trie-based search strategy can boost the performance of OFA in various downstream classification tasks.

²<https://wudao.aminer.cn/CogView/index.html>

³<https://colab.research.google.com/drive/1q6tJ58UKod1eCokbaUNGzF3K5BbX1B5m>

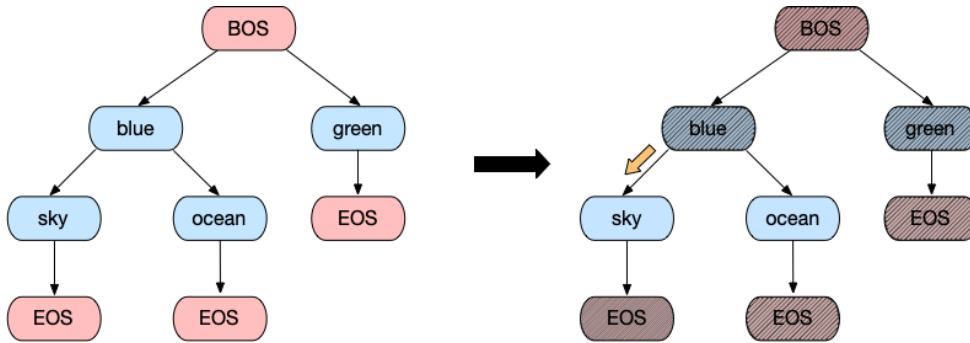


Figure 6. Example of the Trie-based search where the constraint labels are “blue sky”, “blue ocean” and “green”. When computing the log-prob of token “sky”, we only consider tokens in {“sky”, “ocean”} and forcefully set the logits for all invalid tokens to $-\infty$.

Table 13. Ablation results of Trie. The removal of Trie-based search degenerates the performance on downstream tasks. Note that the baseline OFA_{Base} is only pre-trained for 250k steps, which is also used in Table 9.


Model	VQA Test-dev Acc.	SNLI-VE Dev Acc.	ImageNet Top-1 Acc.	MRPC F1	QQP F1
OFA _{Base}	76.03	89.2	82.2	90.6	88.4
<i>w/o Trie</i>	75.86(-0.17)	89.0(-0.2)	81.9(-0.3)	90.1(-0.5)	88.2(-0.2)

C. Qualitative Examples


This section provides more qualitative examples of multiple tasks, including text-to-image generation, open-domain VQA, grounded question answering, and open-domain visual grounding, from the generation of OFA. By reading this section, we hope that readers can better perceive OFA.



An eagle view of a magic city.



A pathway to a temple with sakura trees in full bloom, HD.



A beautiful painting of native forest landscape photography, HD.



An art painting of a soldier, in the style of cyperpunk.




The golden palace of the land of clouds.



Rustic interior of an alchemy shop.



An art painting of a dog, in the style of steampunk, white background.



A strawberry splashing in the coffee in a mug under the starry sky.



Elf elk in the forest illustration, HD, fantasy art.

Figure 7. Examples of text-to-image generation. For better demonstration, we continue finetuning OFA on a subset of LAION-400M (Schuhmann et al., 2021).



An art painting of a city, in the style of cyberpunk.



A painting of the sunset cliffs in the style of fantasy art.



A painting of the superman.



An art painting of a city, in the style of steampunk.



A painting of the sunset cliffs in the style of dark fantasy art.



A painting of the superman, in the dark style.



An oil painting of a pizza on the beach.




An oil painting of a computer in the sky.



A painting of the ship in the style of etching, HD.

Figure 8. Examples of text-to-image generation.




Q: what is the man sitting on?

A: light bulb



Q: what is the dog doing now?

A: getting a bath



Q: what is the person in the right-bottom corner holding now?

A: computer




Q: what is the mood of the children in the picture?

A: happy



Q: what is the man doing?

A: walking



Q: what is the name of the largest planet in the picture?

A: sun

Figure 9. More samples of VQA task on unseen domains. The answers are generated by pretrained OFA without finetuning. The datasets used in VQA pretraining task only contain real-world photographs. We present more cases of VQA task on out-of-domain (non-photographic) images and demonstrate the capability of transferring OFA to these unseen domains.



Q: what color is the car in the region? region: <loc301> <loc495> <loc501> <loc596>

A: tan



Q: what color is the car in the region? region: <loc512> <loc483> <loc675> <loc576>

A: gray



Q: what color is the roof in the region? region: <loc521> <loc176> <loc689> <loc290>

A: brown



Q: what color is the house in the region? region: <loc295> <loc120> <loc524> <loc491>

A: light blue



Q: what color is the house in the region? region: <loc534> <loc172> <loc731> <loc516>

A: White



Q: what object is in the region? region: <loc571> <loc175> <loc598> <loc240>

A: chimney

Figure 10. Samples of the unseen grounded question answering task. In this task, the model should answer a question about a particular region in the image. This task is unseen in pretraining. We demonstrate that directly transferring pretrained OFA to this new task without finetuning works well.

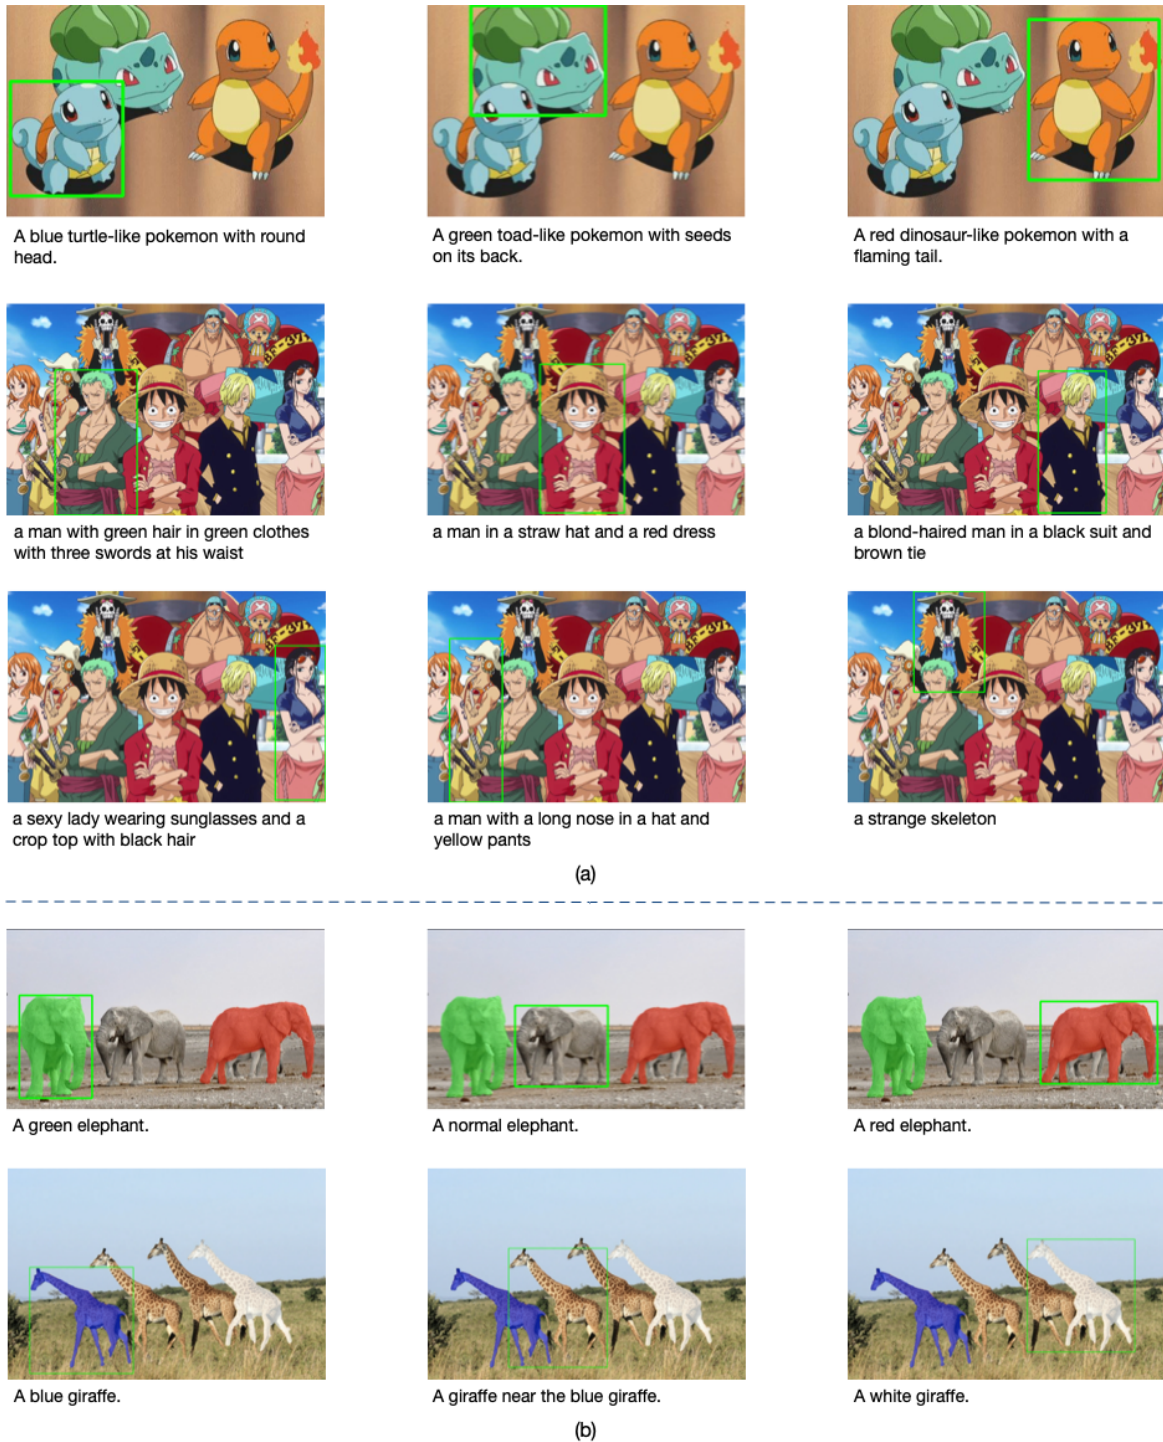


Figure 11. Samples of visual grounding task generated by OFA for various unseen domains: (a) anime (the corresponding animations are *Pokemon* and *One Piece*); (b) synthetic images with attribute combinations.