

# Transfer Visual Prompt Generator across LLMs

Anonymous Author(s)

Affiliation

Address

email

## Abstract

Since developing a new Vision-Language LLM (VL-LLM) by pre-training on tremendous image-text pairs from scratch can be exceedingly resource-consuming, connecting an existing LLM with a comparatively lightweight visual prompt generator (VPG) becomes a feasible paradigm. However, further tuning the VPG component of the VL-LLM still incurs significant computational costs, such as thousands of GPU hours and millions of training data points. An alternative solution is transferring an existing VPG from one VL-LLM to the target VL-LLM. In this work, we investigate VPG transferability across LLMs for the first time, aiming to reduce the cost of VPG transfer. Specifically, we explore VPG transfer across different LLM sizes (*e.g.*, small-to-large) and types. We identify key factors to maximize transfer efficiency, based on which we develop a simple yet highly effective two-stage transfer framework, called **VPGTrans**. Notably, it enables VPG transfer from BLIP-2 OPT<sub>2.7B</sub> to BLIP-2 OPT<sub>6.7B</sub> with less than 10% of the GPU hours using only 10.7% of the training data compared to training a VPG for OPT<sub>6.7B</sub> from scratch. Furthermore, we provide a series of intriguing findings and discuss potential explanations behind them. Finally, we showcase the practical value of our VPGTrans approach, by customizing two novel VL-LLMs, including VL-LLaMA and VL-Vicuna, with recently released LLaMA and Vicuna LLMs.

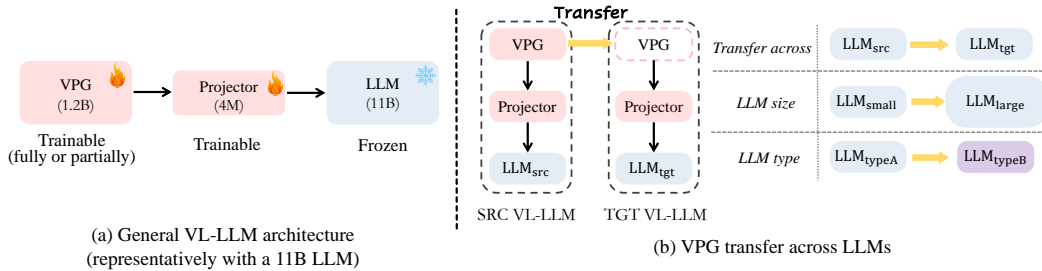


Figure 1: (a) The general architecture of VL-LLMs, *e.g.*, BLIP-2 [26] and PaLM-E [15], including a visual prompt generator (VPG), a linear projector and a backbone LLM. Typically, to tune the VL-LLM, only the VPG and the projector are updated, while the LLM is kept frozen. (b) This work investigates the VPG transferability across LLMs, including different LLM sizes and LLM types.

## 1 Introduction

**Background.** Recent years have witnessed a great rise in large-scale language models (LLMs) in ushering the human-like artificial intelligence. Text-based LLMs [35, 7, 37] are further enhanced by associating with other modalities such as vision, leading to the vision-language LLMs (VL-LLMs), *e.g.*, BLIP-2 [26], Flamingo [2], GPT-4 [8] for multimodal dialog system, and PaLM-E [15] for embodied AI system. To construct a VL-LLM, a visual prompt generator (VPG) module (*cf.* Fig. 1(a))

that produces soft prompts for the input images/videos is added<sup>1</sup> for bridging the gap between vision and language modalities. Currently, such architecture has been frequently adopted by many popular VL-LLMs. For example, BLIP-2 pre-trains a CLIP-ViT [36] combined with a Q-Former as VPG. To obtain the final VL-LLM, the VPG needs to be tuned. Ideally, the vast LLM backbone can remain untouched, leaving only the relatively lightweight VPG module to be fully or partially updated.<sup>2</sup>

**Motivation.** However, building a VL-LLM is inevitably computation-expensive, due to the huge overhead brought by the LLM. For example, training a BLIP-2 FlanT5<sub>XXL</sub> needs over 600 A100-GPU hours on over 100 million image-text pairs. Hopefully, transferring a pre-trained VPG (which is the main body of trainable parts) from an existing VL-LLM to a novel LLM instead of training from scratch,<sup>3</sup> offers a promising solution. Intuitively, all the VL-LLMs literally can share the same VPG infrastructure and utility,<sup>4</sup> which makes the VPG transfer theoretically feasible. In this work, we thus investigate the potential of transferring VPG across LLMs.

**Proposal.** Specifically, this paper examines the transferability of VPG across LLMs 1) with different sizes (in the same type), *i.e.*, *transfer across LLM sizes*, and 2) across different LLM types, *i.e.*, *transfer across LLM type*, as illustrated in Fig. 1(b).

- [*Transfer across LLM Sizes (TaS)*]. It has been a typical practice for LLM-related research [8] to validate the training strategy and the hyperparameter on smaller models (*e.g.*, OPT<sub>2.7B</sub>) and then scale up to larger ones (*e.g.*, OPT<sub>6.7B</sub>). It thus is worth exploring whether a VPG trained on a smaller LLM can be transferred to a larger LLM, resulting in reduced computational costs & data, and maintaining comparable performance.
- [*Transfer across LLM Types (TaT)*]. With a well-tuned VPG for a type of LLM, it is interesting to see if the VPG can be transferred to other types of LLMs even with different architectures (*e.g.*, decoder *v.s.* encoder-decoder). If the transfer can be achieved, how to make it more efficient?

We conduct a series of exploratory analyses (*cf.* §3.1) to identify the key factors for transfer efficiency. Based on our empirical study, we design a two-stage transfer learning framework (*cf.* §3.2), namely **VPGTrans**, that includes a projector warm-up (stage-1) and vanilla fine-tuning (stage-2). For stage-1, we find that the warming up projector before VPG tuning can effectively reduce the training step for adapting a pre-trained VPG to a new LLM, and avoid potential performance drop in the adaptation. To achieve an efficient warm-up, the projector will be well-initialized and then trained with an extremely large learning rate ( $5 \times lr$ ). For stage-2, there is a vanilla fine-tuning of both the VPG and projector. Despite its simplicity, VPGTrans is able to significantly speed up the VPG-transfer process without harming performance.

**Results and Findings.** Via extensive experiments on the transfer across LLM sizes and types (*cf.* §4 & §5), we gain the following key observations:

- VPGTrans helps avoid the performance drop caused by direct inheriting the VPG and achieves at most 10 times acceleration for the small-to-large transfer across LLMs in the same type.
- VPGTrans can also achieve comparable or better performance than training from scratch and achieve at most 5 times acceleration for the transfers between different model types.
- Notably, our VPGTrans helps achieve a **BLIP-2 ViT-G OPT<sub>2.7B</sub>→<sub>6.7B</sub>** transfer with **less than 10% of the GPU hours** and **10.7% training data** required for the original model training.
- Furthermore, our framework can even outperform the original BLIP-2 OPT<sub>6.7B</sub> on most of the evaluated datasets, with a **+2.9** improvement on VQAv2 and a **+3.4** improvement on OKVQA.

Our investigation further reveals some intriguing findings, for which we provide possible explanations:

- When conducting TaS from LLM<sub>src</sub> to LLM<sub>tgt</sub>, the size of LLM<sub>src</sub> is not the larger the better. The transfer sometimes even follows a counterintuitive principle of “*the smaller LLM<sub>src</sub>’s size the more speed-up and better performance*” (*cf.* §4.2).

<sup>1</sup>Also including a linear projector for dimension matching.

<sup>2</sup>Note that Flamingo also inserts some tunable parameters into the LLM part, but recent works [26, 15] found that freezing LLM can be more efficient.

<sup>3</sup>It is not a rigorous expression, because the VPG is typically a pre-trained model, like CLIP [36]. We use it for simplicity in this paper.

<sup>4</sup>while the projector can not be shared due to the dimension mismatch.




Figure 2: Comparing the cost between *training VPG from scratch* vs. *transferring VPG via our VPGTrans strategy*. Note the LLM via VPGTrans is FlanT5<sub>XL→XXL</sub> and OPT<sub>2.7B→6.7B</sub>, respectively.

- When conducting TaT, efficient VPG transfer can not be achieved between two small LLMs with our VPGTrans, due to the large gap between small LLMs’ embedding space (*cf.* §5.2).

**Contributions.** In this study, we for the first time show that effective VPG transfer across LLMs can be achieved under most conditions, suggesting that it is possible to build a new VL-LLM with considerably lower computational cost, as seen in Fig. 2. To summarize, we make the following key contributions:

- *Effective approach.* We investigate the key factors for VPG-transfer efficiency and propose a two-stage transfer framework VPGTrans. The approach helps achieve highly-efficient VPG transfer across LLMs meanwhile with less training data and even task improvements.
- *Intriguing findings.* By exploring the VPG transfer across LLMs, we reveal several intriguing findings and provide potential explanations that will shed light on further research.
- *Open source.* We showcase how to customize a novel GPT-4-like VL-LLM with our VPGTrans (*cf.* §6), and release two multimodal-version VL-LLMs: VL-LLaMA and VL-Vicuna. All codes and models will be released later.

## 2 Preliminary

This section will first give a briefing on the existing prevailing VL-LLMs, and then elaborate on the settings of the exploratory analyses of these VL-LLMs.

### 2.1 VL-LLM

**Architecture.** As illustrated in Fig. 1(a), current VL-LLMs mostly adopt a common architecture, including a visual prompt generator (VPG), a projector, and a backbone LLM. Typically, VPG takes images/videos as inputs, and encodes the visual input into a fixed length of soft prompts. Then, a linear projector is employed to align the soft prompt’s dimension to LLM’s word embedding dimension. Finally, the LLM will generate sentences based on the information from the soft prompt. We list some of the recent representative VL-LLMs in Table 1.

Table 1: VL-LLMs architectures and pre-training paradigm. †: it is a GPT-2-like LLM with relative position embeddings.

VL-LLMs	VPG	VPG Trainable	LLM	LLM Trainable
KOSMOS-1 [17]	CLIP [36]	All	Rand. Init. LM	All
Frozen [47]	NF-ResNet-50 [6]	NF-ResNet-50	GPT-2-like† [35]	No
Flamingo [2]	NFNet-F6 [6]+Resampler [19]	Resampler	Chinchilla [16]	Xattn-Dense
PaLM-E [15]	ViT [14] / OSRT [42]	All	PaLM [11]	No
BLIP-2 [26]	EVA-CLIP [45] + Q-Former [26]	Q-Former	OPT [52] / Flan-T5 [12]	No

**Training Paradigm.** Given a VL-LLM, typically the VPG and linear projector will be trained, fully or partially. For example, PaLM-E updates all of the parameters of VPG in the pre-training stage, while BLIP-2 and Flamingo freeze the ViTs and tune their Q-Former and Resampler, respectively. As the main part of the whole architecture, the LLM is usually frozen during the training or tuned only a small portion (*e.g.*, 10B for Flamingo-80B). KOSMOS-1 is an exception, which does not use a pre-trained LLM but trains the LLM from scratch. Such a training paradigm typically results in much longer training time and data (both multimodal and pure text corpus). Recent works [26, 15] show that adopting an existing LLM and freezing all of its parameters can also achieve excellent performance with significantly reduced computational cost, which leads to the trend into frozen pre-trained LLM. For example, BLIP-2 FlanT5<sub>XXL</sub> (12.1B) can achieve better zero-shot VQAv2




Figure 3: Comparisons between i) inheriting VPG from OPT<sub>125M</sub> and training it with randomly initialized projector for OPT<sub>350M</sub> and ii) training VPG and randomly initialized projector for OPT<sub>350M</sub> from scratch.

performance (65.0% in Acc.) compared with KOSMOS-1 (51.0% in Acc.) and Flamingo-80B (56.3% in Acc.). Thus, in this paper, we mainly focus on VPG transfer across frozen LLMs.

## 2.2 Experiment Settings

**Architecture.** We adopt BLIP-2’s architecture and training paradigm. In our exploration experiments, we consider using the VPG that consists of a CLIP ViT-L/14 [36], and a Q-Former that has already undergone a BLIP-like pre-training (the 1st stage pre-training in BLIP-2’s paper [26]).

**Training Data.** For all of the exploration experiments, we adopt human-annotated COCO caption dataset [28] and web image-text pairs SBU dataset [33], which results in 1.4 million image-text pairs.

**Transfer Direction.** For the small-to-large model transfer among the same type LLMs, we investigate: 1) OPT [52] (decoder-only) series including 125M, 350M, 1.3B, and 2.7B, and 2) FlanT5 [12] (encoder-decoder) ranging *base*, *large*, and *XL*. For the transfer across different types of LLMs, we consider the ones of OPT and FlanT5 with similar sizes.

**Evaluation.** To evaluate the performance of VL-LLMs, we choose two caption datasets: (1) COCO caption [28] (2) NoCaps [1], and three VQA datasets: (3) VQAv2 [4] (4) GQA [18] (5) OKVQA [30]. We make evaluations after the pre-training without task-specific fine-tuning and report the CIDEr [48] for all caption tasks and accuracy for all VQA tasks.

**Implementation Details.** We follow the same implementation details of BLIP-2, via the open code.<sup>5</sup> Concretely, we use FP16 and BFloat16 for OPT and FlanT5 respectively in the model training. For the learning rate, we first conduct a linear warm-up from 1e-6 to 1e-4, and then use a cosine learning rate schedule with the minimal  $lr=1e-5$  for 10 epochs. Due to the limited data amount, we slightly decrease the batch size, which we find beneficial for the final performance. Specifically, we set the batch size of 1,728 and 1,152 for OPT and FlanT5-based models, respectively.

## 3 Maximizing the Transfer Efficiency with a Two-stage Transfer Strategy

In this section, we first identify the key factors for maximizing transfer efficiency, based on which we then motivate our solution for better transfer.

### 3.1 Exploratory Analysis: Identifying Key Factors for VPG Transfer

Via selected experiments of small-to-large transfer among OPT models, we can obtain the following key observations. More systematical comparisons are conducted in the later section (*cf.* §4).

• **Inheriting the trained VPG can accelerate training.** To demonstrate this, we compare the convergence rates of VPG training on OPT<sub>350M</sub> from scratch, and inheriting VPG trained on OPT<sub>125M</sub>. The patterns are shown in Fig. 3. Overall, we find that inheriting VPG trained on OPT<sub>125M</sub> accelerates convergence, particularly for two caption tasks. However, for datasets requiring fine-grained visual perception such as VQAv2 and GQA, directly conduct continue training with an inherited VPG will harm the performance. We hypothesize that tuning VPG with a randomly initialized projector will compromise the existing fine-grained visual perception ability of VPG. The possible reason can be that, the VPG is typically a pre-trained model with powerful visual perception ability, and thus updating based on the gradient passed through a random projector will mislead the VPG at the initial steps [2, 27, 21].

<sup>5</sup><https://github.com/salesforce/lavis>

144 • **Warming up the linear projector can prevent performance drop and expedite VPG training.**

145 To verify this, we first conduct a warm-up training of the linear projector for 3 epochs, during which VPG and LLM are frozen. Subsequently, we jointly train VPG and the projector and plot the performance curve in Fig. 4 (the warm-up process is not included in this figure). The results show that the performance drop observed in Fig.3 can be avoided in Fig. 4. Additionally, we observe that the warm-up training leads to fewer training steps required for VPG and projector joint training. However, we must emphasize that warming up is a costly step. In the case of a large LLM, such as 6.7B, the trainable parameters of BLIP-2’s VPG will account for less than 10% of the total parameters, where freezing VPG can only lead to a reduction of 5.4% A100 hours (36.9 out of 684.0 A100 hours). We will elaborate on how to accelerate the linear projector warm-up in our later discussion (cf. 3.1).




Figure 4: First warming-up then transferring can avoid performance drop on VQAv2 and accelerate convergence for COCO caption.

160 • **Initializing LLM<sub>tgt</sub>’s projector with the help of the word converter can accelerate the linear projector warm-up.**

161 In fact, the VPG and projector trained on LLM<sub>src</sub> have already learned how to map the visual content to LLM<sub>src</sub>’s understandable soft prompt [32]. If we can convert the LLM<sub>src</sub>’s soft prompt to LLM<sub>tgt</sub>’s soft prompt, we can directly get a VPG suitable for LLM<sub>tgt</sub>. One natural idea is to leverage the word embeddings of both models as a proxy for the soft prompt [23]. The intuition behind the scene is that, the soft prompt works in the same format as normal words.

166 To validate our hypothesis, we conduct an experiment on the transfer from OPT<sub>125M</sub> to OPT<sub>1.3B</sub>. After training a linear word embedding converter (cf. §3.2(b)), we initialize the projector for OPT<sub>1.3B</sub> with the merged linear operation of the projector for OPT<sub>125M</sub> and converter. As shown in Table 2, we observe that the initialization can reduce the 3 epochs’ warm-up to 2 epochs.

170 • **Linear projector warm-up enables faster convergence with an extremely large learning rate.**

171 To determine the most efficient transfer practice, we experiment with training the projector using different learning rates. Surprisingly, we find that the linear projector enables fast and stable convergence with an extremely large learning rate. Specifically, by setting the learning rate to 5 times of the original value, the COCO caption’s CIDEr score can reach 133.1 with **1 epoch** training, which is **higher than the 3 epochs** results of *w/o init.* as shown in Table 2.

Table 2: Comparison between linear projector warm-up with/without word embedding initialization. The metric is COCO caption’s CIDEr.

Epoch	w/ init.	w/o init.
1	130.2	126.1
2	132.7	131.6
3	133.4	132.8

180 **3.2 A Two-stage VPG Transfer Framework**

181 By connecting all the dots as discussed above in §3.1, we now design our two-stage VPGTrans framework for more efficient VPG transfer. As shown in Fig. 5, the stage-1 of VPGTrans performs projector warm-up and the stage-2 carries out a vanilla fine-tuning. Our results demonstrate that the VPGTrans is simple yet effective that can significantly speed up the transfer without compromising performance. Detailed results are given in the later sections (cf. §4 & 5).

186 ► **Stage-1: Projector Warm-up.**

187 (a) *Inherit VPG.* We first initialize the VPG for LLM<sub>tgt</sub> with the VPG trained on LLM<sub>src</sub>.

188 (b) *Projector Initialization.* Then, we initialize the projector for LLM<sub>tgt</sub> merged from the projector of LLM<sub>src</sub> and a linear word converter. Formally, we define the linear projector of LLM<sub>src</sub> as  $f_s(x) = W_s x + b_s$ , the linear projector for LLM<sub>tgt</sub> as  $f_t(x) = W_t x + b_t$ , and the word converter as  $g_c(x) = W_c x + b_c$ .

192 The word converter is a linear layer trained with text-only caption data to convert the LLM<sub>src</sub>’s word embeddings to LLM<sub>tgt</sub>’s word embeddings. We experiment with optimizing losses based on cosine similarity or Euclidean distance, and observe no significant difference between the two losses. Thus we simply use cosine similarity in our experiments. In cases where LLM<sub>src</sub> and LLM<sub>tgt</sub> use different tokenization methods, we optimize based on the overlapped tokens. Formally, for every given token  $k$ , we denote its word embeddings of LLM<sub>src</sub> and LLM<sub>tgt</sub> as  $x_s$  and  $x_t$ . Then, we minimize the loss:

$$\mathcal{L} = 1 - \text{sim}(g_c(x_s), x_t). \quad (1)$$

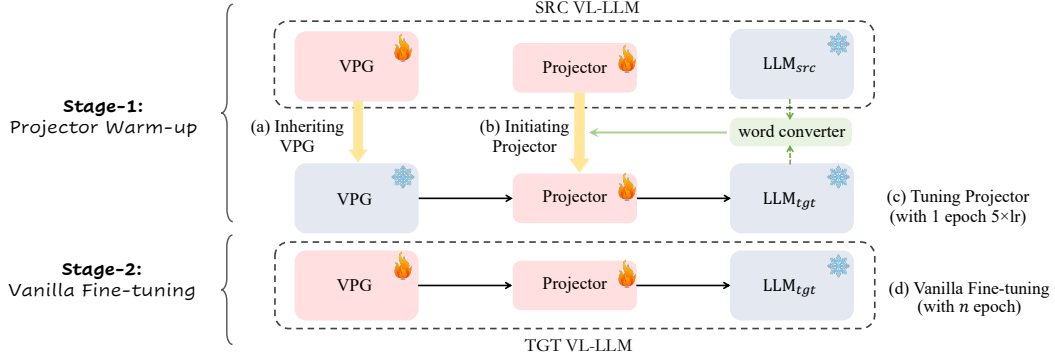


Figure 5: Our two-stage VPGTrans framework. **Stage-1** is to first (a) inherit the VPG of LLM<sub>src</sub> and (b) initialize the projector by merging the projector of LLM<sub>src</sub> and word converter. (c) Then the projector will be warmed up for 1 epoch with a large learning rate. **Stage-2** is to (d) conduct a vanilla fine-tuning for the VPG and projector for  $n$  epochs.

Once we obtain the word converter  $g_c(\cdot)$ , we can easily merge it with the projector of LLM<sub>src</sub> as:

$$f_t(x) = f_s(g_c(x)) = W_s(W_c x + b_c) + b_s, \quad (2)$$

resulting in  $f_t$ 's weight and bias as  $W_t = W_s W_c$  and  $b_t = W_s b_c + b_s$ .

**(c) Warm-up Training.** Then, we only train the projector in this stage with a frozen VPG and LLM. Specifically, we train the projector for 1 epoch with 5 times of the normal learning rate.

► **Stage-2: Vanilla Fine-tuning.**

**(d) Vanilla Fine-tuning.** In the final step, we conduct a joint training of VPG and projector for  $n$  epochs with a normal learning rate.

## 4 Exp-I: Transfer across Different Model Sizes

In this section, we conduct experiments to systematically illustrate the effectiveness of our VPGTrans and analyze the relationship between transfer efficiency and model size. For simplicity, we use **TaS** to represent the transfer across different model sizes.

### 4.1 Experimental Settings

In this part, we introduce baselines and transfer variants. For details about training data and implementation details, please refer to the experiment settings in the Preliminary (cf. 2.2).

**Baselines.** We mainly compare our VPGTrans with *training from scratch* (TFS) and *VPG inheritance* (VPG Inherit), where we report their performance on the aforementioned 5 tasks without further task-specific fine-tuning. For our VPGTrans, the word converter training only requires updating a linear layer on tokenized text data and typically takes less than 10 minutes on 1 A100 GPU with less than 15G GPU memory. Meanwhile, freezing the VPG can lead to at least 14 A100 minutes speed-up per epoch. **Therefore, we consider the whole stage-1 training as the 1st epoch for simplicity.**

**Transfer Variants.** We conducted experiments on transfer learning using 1) the OPT model across four different sizes: 125M, 350M, 1.3B, and 2.7B, and 2) the FlanT5 model across three sizes: *base*, *large*, and *XL*. However, we encountered significant instability during training with FlanT5<sub>large</sub>. As a result, we mainly present the transfer results between FlanT5<sub>base</sub> and FlanT5<sub>XL</sub>.

### 4.2 VPGTrans Enabling Faster Convergence without Performance Drop under TaS

First of all, as shown in Fig. 6, our VPGTrans can consistently accelerate the model convergence. For COCO caption and NoCaps that require more training steps to converge, our VPGTrans (green line) can be higher than the other two lines (blue and orange lines). To give a quantitative evaluation of the speed-up rate, we show the speed-up rate in Table 3. The speed-up rate is calculated by considering the number of epochs reduced to achieve the best TFS performance on a particular dataset. Formally, given a dataset  $D$ , TFS obtains the best performance  $p$  on  $D$  at epoch  $e_{tfs}$ , whereas VPGTrans first achieves a better performance than  $p$  at epoch  $e_{vt}$ . The speed-up rate on  $D$  is given by  $\frac{e_{tfs}}{e_{vt}}$ . According to Table 3, our VPGTrans can achieve at least 4 times speed-up on 40% of Transfer-Task



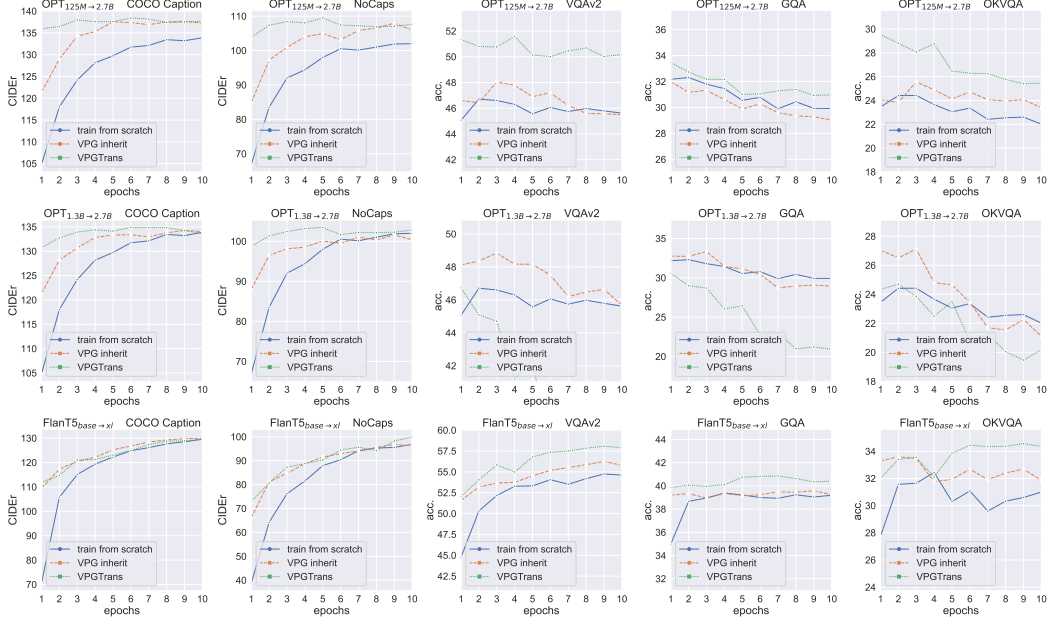


Figure 6: Comparison between different methods across 3 TaS variants on 5 tasks. Note that the model is directly evaluated after pre-training without further fine-tuning. Please refer to Appendix§C for other transfer variants.

Table 3: The speed-up rate of our VPGTrans compared with *training from scratch* (TFS). The symbol "-" means VPGTrans can not achieve better performance than TFS.

Transfer	COCO Caption	NoCaps	VQAv2	GQA	OKVQA
OPT <sub>125M</sub> →350M	1.7	3.0	1.0	5.0	5.0
OPT <sub>125M</sub> →1.3B	9.0	10.0	9.0	-	2.0
OPT <sub>350M</sub> →1.3B	4.5	5.0	9.0	2.0	2.0
OPT <sub>125M</sub> →2.7B	10.0	10.0	2.0	2.0	3.0
OPT <sub>350M</sub> →2.7B	10.0	10.0	2.0	-	3.0
OPT <sub>1.3B</sub> →2.7B	3.3	3.3	2.0	-	1.5
FlanT5 <sub>base</sub> →XL	1.0	1.1	3.0	4.0	2.0
FlanT5 <sub>XL</sub> → OPT <sub>2.7B</sub>	5.0	5.0	2.0	2.0	3.0
OPT <sub>2.7B</sub> → FlanT5 <sub>XL</sub>	1.7	2.0	-	2.0	-

231 variants. Furthermore, for the two caption datasets, which take a long time to converge, our VPGTrans  
 232 OPT<sub>125M</sub>→2.7B delivers a 10 times speed-up.

233 Moreover, when compared with the *VPG inherit* in Fig. 6, our VPGTrans can achieve a higher  
 234 speed-up rate on all of the variants on caption tasks, and achieve better performance on most variants  
 235 except for OPT<sub>1.3B</sub>→2.7B. We refer the readers to Appendix§C.3 for more comparisons.

236 We provide interesting findings with respect to the efficiency transfer by VPGTrans in the following.

237 • **The smaller size of LLM<sub>src</sub>, the easier the transfer.** In our OPT based experiments, we notice an  
 238 interesting phenomenon: when transferring to a given LLM<sub>tgt</sub>, both the convergence rate and optimal  
 239 performance are roughly inversely proportional to the size of LLM<sub>src</sub>. For example, as shown in  
 240 Table 3, the OPT<sub>125M</sub>→2.7B and OPT<sub>350M</sub>→2.7B have much higher speed-up rate than OPT<sub>1.3B</sub>→2.7B on  
 241 all of the datasets. Meanwhile, as demonstrated in Fig. 6, the optimal performance of OPT<sub>125M</sub>→2.7B  
 242 is better than OPT<sub>1.3B</sub>→2.7B on 3 VQA tasks.

243 We hypothesize that training VPG on larger OPT will have a worse influence on VPG’s existing  
 244 fine-grained perception ability, which might be caused by the enlarging embedding dimensions. To  
 245 validate our hypothesis, we fix the VPG weight and only tune linear projectors to test VPGs trained  
 246 on different LLM<sub>src</sub> through cross-size transfer. The SPICE [3] metric on COCO caption is used to  
 247 evaluate the VPG’s visual perception ability, where SPICE is specifically designed for visual concept  
 248 perception in captions. As shown in Fig. 7, for each row, given the LLM<sub>tar</sub>, the performance of

Table 4: Comparison between models built with our VPGTrans and the original BLIP-2 ViT-G OPT<sub>6.7B</sub> and BLIP-2 ViT-G FlanT5<sub>XXL</sub>.

Models	VQAv2 val	GQA test-dev	OKVQA test	GPU hours	training data
BLIP-2 ViT-G OPT <sub>6.7B</sub>	54.3	36.4	36.4	631.5	129M
BLIP-2 ViT-G OPT <sub>2.7B</sub> → <sub>6.7B</sub> (ours)	57.2	36.2	<b>39.8</b>	<b>59.0</b>	<b>13.8M</b>
VL-LLaMA <sub>7B</sub> (ours)	<b>58.1</b>	<b>37.5</b>	37.4	67.1	13.8M
BLIP-2 ViT-G FlanT5 <sub>XXL</sub>	65.2	44.7	<b>45.9</b>	684.0	121.6M
BLIP-2 ViT-G FlanT5 <sub>XL</sub> → <sub>XXL</sub> (ours)	<b>65.2</b>	<b>45.0</b>	45.0	<b>32.4</b>	<b>5.3M</b>

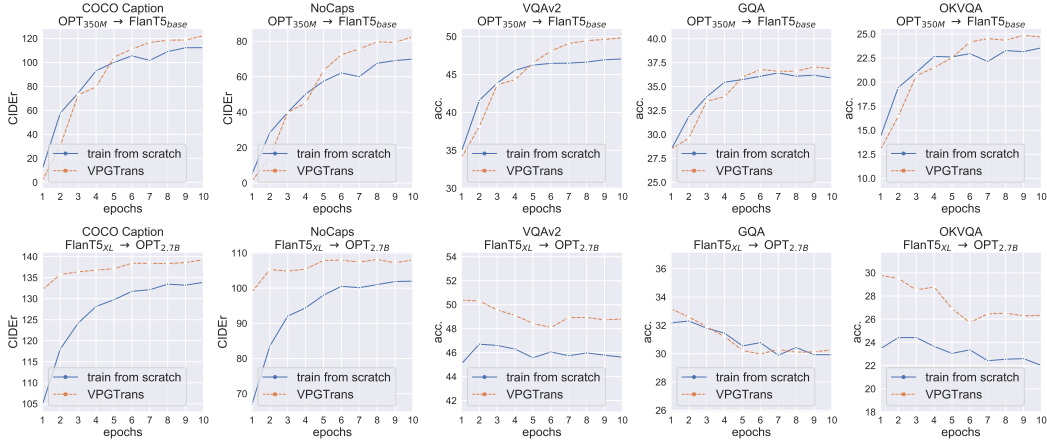


Figure 8: Comparison between different methods across 2 TaT variants on 5 tasks. Note that the model is directly evaluated after pre-training without further fine-tuning. Please refer to Appendix§D for other transfer variants.

VPG trained on smaller LLM<sub>src</sub> can outperform the larger ones in most conditions, which indicates a better visual perception ability of VPG trained on smaller LLM<sub>src</sub>. Therefore, **adapting a VPG from a smaller OPT model which is less affected, is helpful to take fewer steps to reach the TFS’s best performance and achieve even better performance.**

### 4.3 Scale-up Experiments

To validate the effectiveness of our VPGTrans on the real-world application level, we experiment on transferring from BLIP-2 ViT-G OPT<sub>2.7B</sub> to OPT<sub>6.7B</sub> and from BLIP-2 ViT-G FlanT5<sub>XL</sub> to FlanT5<sub>XXL</sub>. Please refer to Appendix§C.4 for implementation details.

**Speed-up with non-degenerated performances.** As shown in Table 4, we can see that (1) **OPT<sub>2.7B</sub>→<sub>6.7B</sub>**: our VPGTrans achieves a 10.7 times speed-up with only 10.7% training data, while the performance on VQAv2 and OKVQA have over 2 points improvement. (2) **FlanT5<sub>XL</sub>→<sub>XXL</sub>**: VPGTrans can achieve 21.1 times speed-up with less than 5% training data while achieving the same performance on VQAv2, higher performance on GQA and slightly lower performance on OKVQA. Note that continuing training the FlanT5<sub>XL</sub>→<sub>XXL</sub> only shows improvement on VQAv2 and GQA. Thus, we do not show a checkpoint with more training steps.

## 5 Exp-II: Transfer across Different Model Types

In this section, we further investigate the transfer across different model types. For simplicity, we mark this type of transfer as **TaT**.

### 5.1 Experimental Settings

In this part, we introduce baselines and transfer variants. For details about training data and implementation details, please refer to the experiment settings in the Preliminary (cf. 2.2).




Figure 7: The confusion matrix. Only linear layers are trained for VPG evaluation. Models are tested on COCO caption with SPICE metric to compare the VPGs trained on different LLM<sub>src</sub>.



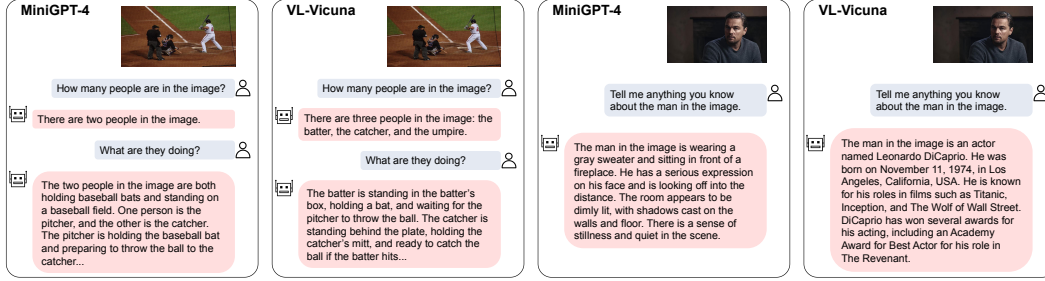


Figure 9: Comparison between MiniGPT-4 and our VL-Vicuna.

**Baselines.** We mainly compare our VPGTrans with *training from scratch* (TFS), and report the performance on the aforementioned 5 tasks. Other details (cf. 4.1) are totally the same with TaS experiments.

**Transfer Variants.** We conducted experiments on transfer learning between 1) OPT<sub>350M</sub> and FlanT5<sub>base</sub>, and 2) OPT<sub>2.7B</sub> and FlanT5<sub>XL</sub>.

## 5.2 VPGTrans Enabling Faster Convergence only on Large LLMs under TaT

• **There is no speed-up of TaT between two small LLMs.** A finding is that on TaT our VPGTrans does not show speed-up for small models, and even shows a degeneration of training speed in the initial several epochs. As shown in Fig. 8, when transferring from OPT<sub>350M</sub> to FlanT5<sub>base</sub>, the convergence speed of VPGTrans is even slower than TFS in the initial several epochs.

• **Speed-up of VPGTrans happens in large LLMs.** However, when moving to the large LLMs like OPT<sub>2.7B</sub> and FlanT5<sub>XL</sub>, there is an obvious speed-up. As shown in Table 3, we can see at least 2 times speed-up when transferring from FlanT5<sub>XL</sub> to OPT<sub>2.7B</sub>. We empirically find that **the soft prompts for larger LLM are more linear transferrable among different LLM types**. As shown in Fig. 8, when transferring between FlanT5<sub>base</sub> and OPT<sub>350M</sub>, the VGPTrans' 1st epoch results on two caption datasets are limited, where only a linear operation can be trained. The result of OPT<sub>350M</sub>→FlanT5<sub>base</sub> on the COCO caption is even near to zero. By contrast, the result of FlanT5<sub>XL</sub>→OPT<sub>2.7B</sub> with our VPGTrans are obviously higher than TFS. We hypothesize that larger LLM typically learned more generalizable text embeddings and share more similarity among relative word distances, which enables an easier VPG transfer.

## 6 Customizing New VL-LLMs with Any LLMs

Above, we thoroughly certify the efficacy of our proposed VPGTrans approach for higher efficient transfer of VPG. In this section, we illustrate how to apply the VPGTrans framework for VPG transfer to customize new VL-LLMs with any LLMs.

**VL-LLaMA.** By applying our VPGTrans, we can equip the recently released LLaMA [46] model with a VPG trained on BLIP-2 OPT<sub>6.7B</sub> to perceive the visual information. As shown in Table 4, we can see that our VL-LLaMA can outperform the original BLIP-2 OPT<sub>6.7B</sub> on all datasets.

**VL-Vicuna.** An exciting application of our VPGTrans is to build a GPT-4 [8] style multimodal conversation chatbot. To achieve our goal, we employ Vicuna [10] as our base LLM. Similarly, we transfer the VPG from BLIP-2 OPT<sub>6.7B</sub>, and add an extra instruction tuning using MiniGPT-4's self-instruct data [54]. We compare our model with MiniGPT-4 in Fig. 9. When compared to MiniGPT-4, our VL-Vicuna shows better visual perception ability. Please refer to Appendix E for more cases.

## 7 Conclusion

In this work, we conduct a comprehensive investigation to the problem of VPG transferability across LLMs. We first explore the key factors for maximizing the transfer efficiency under the VPG transfer across different LLM sizes and types. Based on the key findings, we propose a novel two-stage transfer framework, namely VPGTrans, which can help achieve comparable or better performance while significantly reducing training costs. Moreover, a list of important findings and possible reasons behind them are shown and discussed. Finally, we demonstrate the practical value of our VPGTrans, by customizing new VL-LLMs via VPG transfer from existing VL-LLMs.

## References

- [1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- [3] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 382–398. Springer, 2016.
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [5] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [6] Andy Brock, Soham De, Samuel L Smith, and Karen Simonyan. High-performance large-scale image recognition without normalization. In *International Conference on Machine Learning*, pages 1059–1071. PMLR, 2021.
- [7] 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.
- [8] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- [9] 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, pages 104–120. Springer, 2020.
- [10] 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. URL <https://vicuna.lmsys.org>.
- [11] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. PaLM: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [12] 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.
- [13] Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P Xing, and Zhiting Hu. Rlprompt: Optimizing discrete text prompts with reinforcement learning. *arXiv preprint arXiv:2205.12548*, 2022.
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [15] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- [16] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [17] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023.

- [18] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019.
- [19] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021.
- [20] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 709–727. Springer, 2022.
- [21] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054*, 2022.
- [22] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- [23] Brian Lester, Joshua Yurtsever, Siamak Shakeri, and Noah Constant. Reducing retraining by recycling parameter-efficient prompts. *arXiv preprint arXiv:2208.05577*, 2022.
- [24] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.
- [25] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022.
- [26] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [27] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. *arXiv preprint arXiv:2210.08823*, 2022.
- [28] 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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [29] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019.
- [30] 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.
- [31] Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. Teaching language models to support answers with verified quotes. *arXiv preprint arXiv:2203.11147*, 2022.
- [32] Ron Mokady, Amir Hertz, and Amit H Bermano. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021.
- [33] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24, 2011.
- [34] 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.
- [35] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [36] 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.

- [37] 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.
- [38] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- [39] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [40] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [42] Mehdi SM Sajjadi, Daniel Duckworth, Aravindh Mahendran, Sjoerd van Steenkiste, Filip Pavetic, Mario Lucic, Leonidas J Guibas, Klaus Greff, and Thomas Kipf. Object scene representation transformer. *Advances in Neural Information Processing Systems*, 35:9512–9524, 2022.
- [43] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.
- [44] Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, et al. On transferability of prompt tuning for natural language processing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3949–3969, 2022.
- [45] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- [46] 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.
- [47] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34: 200–212, 2021.
- [48] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [49] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022.
- [50] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- [51] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
- [52] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [53] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022.
- [54] 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.

## A Related Work

### A.1 Vision and Language Models

Vision and language models (VLMs) aim at understanding visual and textual information with a single model. Previously, the VLM mainly employs a pre-trained object detector as its feature extractor and conducts unsupervised training on a huge amount of image-text pairs. For example, VL-Bert [43], ViL-Bert [29] and Uniter [9] adopt Faster-RCNN [40] to extract image information into object features and take advantage of the masked language model as their pre-training task. Later, due to the prevalence of vision transformer (ViT), the VLM paradigm is turned into end-to-end training with a ViT as the visual encoder. The representative works include ALBEF [24], BLIP [25], and BEIT-v3 [49], which show the state-of-the-arts supervised training performance on a wide range of downstream tasks.

Recently, LLMs have shown their remarkable capability as zero/few-shot learners [7] and a series of emergent abilities [50] like in-context learning [7], and chain-of-thoughts reasoning [51]. A new paradigm, i.e., VL-LLMs is created by associating the VLM or pure vision encoders with LLMs. As we illustrated before, the VLM or visual encoders are typically able to convert the input vision signals into LLM-understandable soft prompts, and thus we call them VPG. The VL-LLMs advance in inheriting the great potentials of the backbone LLMs, and thus are capable of achieving excellent zero/few-shot performances [2, 26] on downstream tasks or be equipped with visual planning ability [15]. However, connecting the VPG to the existing LLMs with further tuning is costly. Even the BLIP-2 [26], targeted at efficient training, will take over 600 A100 GPU hours on over 100M image-text pairs for its largest model. With this regard, our proposed VPGTrans can effectively reduce the cost of building new VL-LLMs with the help of existing ones.

### A.2 Prompt Transfer

In this paper, we investigate the VPG transfer, where the soft prompt is to represent the content of specific inputs like images and videos. In addition to the content prompt, the more explored soft prompt is the task prompt [22, 53, 20], where a sequence of soft prompts are tuned to assist the pre-trained models to achieve better performance on specific tasks. There have already been some works exploring the transferability of task prompts. For example, Su et al. [44] conducts a series of experiments to illustrate the transferability across tasks and models. Specifically, Su et al. [44] find that the transfer between similar tasks is beneficial for training speed-up and better performance. Lester et al. [23] proposes to recycle soft prompts across models with vocab-to-vocab transformations, or linear-combination transformations. Note that our word converter initialization is similar to the idea of vocab-to-vocab transformations. However, we do not observe a zero-shot transfer in our experiments like them, which indicates a potential difference between the content prompts and task prompts. Another way of soft prompt transfer [13] is to conduct prompt tuning on discrete prompts, and thus the discrete prompts can be directly shared across models. Different from these task prompts transfer works, our VPG transfer scenario actually suffers from fewer limitations. For example, the task soft prompts transfer suffers from the dimension change problem, where the main body of the trainable parameters should be processed. However, our VPG (the main trainable parameters) can naturally be shared among LLMs with different embedding dimensions and leave the dimension change problem to a simple projector with ignorable parameters.

## B Extended Findings in Exploratory Analysis

In this section, we show extended findings of exploratory analysis (*cf.* §3.1).

• **1. Merely tuning the projector can not achieve the best performance.** We want to clarify that merely tuning the projector is insufficient for achieving the best performance. Notably, as shown in Fig. 10, significant performance gaps are observed between the “*only linear*” (green curve) and “*train from scratch*” (orange curve) approaches for COCO caption and NoCaps. Therefore, if the goal is to build a multimodal conversation robot using carefully collected dialog data, training only the linear projector is insufficient to align with the provided data.




Figure 10: Comparisons between i) inheriting VPG from OPT<sub>125M</sub> and training it with randomly initialized projector for OPT<sub>350M</sub> and ii) training VPG and randomly initialized projector for OPT<sub>350M</sub> from scratch. iii) training only the projector.




Figure 11: Interpreting generated soft prompts for OPT<sub>125M</sub> with nearest words.

516 • **2. Word embedding converter can not replace a trained linear projector.** As demonstrated  
 517 in Fig. 11, we observe a common pattern: the last token of soft prompts is closest to EOS, while  
 518 the middle tokens represent the image content. Such a phenomenon indicates a similarity between  
 519 soft prompts and word embeddings. However, they are not identical. For instance, the norm of soft  
 520 prompts is typically around 10 times the average norm of word embeddings. It is important to note  
 521 that the linear projector initialization cannot replace the warm-up training. Using only the linear  
 522 projector initialization even yields a random performance. We believe that a better understanding of  
 523 how prompt works will further benefit the VPG’s transfer learning.

524 • **3. The projector warm-up is robust to a larger learning rate, while VPG can not.** The first  
 525 thing we want to clarify is that the 5 times normal learning rate will result in a training crash for  
 526 VPG. Additionally, we find that although increasing the learning rate to 10 times in the projector  
 527 warm-up does not yield any additional acceleration, the projector can converge without crashing  
 528 during training.

## 529 C Extended TaS Experiments

530 In this section, we first illustrate extending findings of TaS experiments (*cf.* §4). Then, we introduce  
 531 the implementation details of scale-up experiments. We also plot a more complete version of Fig. 6  
 532 in Fig. 13.

### 533 C.1 VPGTrans Enabling Stable Training under TaS

534 As we illustrated before, FlanT5<sub>large</sub> training is extremely unstable. Even when the learning rate  
 535 is adjusted to one-tenth of its original value, the model does not converge or shows a very slow  
 536 convergence rate after 4 epochs. However, we find that by lowering the learning rate for stage-2  
 537 training, our VPGTrans can achieve stable training on FlanT5<sub>large</sub>. We plot the performance curve of  
 538 the COCO caption in Fig. 12.

### 539 C.2 VPGTrans Enabling Training with Less Data under TaS

540 We empirically find that TaS can reduce the requirement for the amount of training data. By reducing  
 541 the training data to only COCO, we find no obvious performance drop. However, we want to stress  
 542 that **the retained data should be of high quality**, which means that if the same number of SBU data



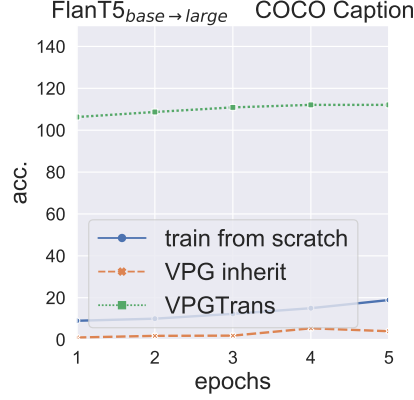


Figure 12: Comparison of training FlanT5<sub>large</sub> using different methods.

Table 5: The speed-up rate of our VPGTrans compared with *VPG inherit*. The symbol "-" means VPGTrans can not achieve better performance than *VPG inherit*.

Transfer	COCO Caption	NoCaps	VQAv2	GQA	OKVQA
OPT <sub>125M→350M</sub>	1.1	2.7	10.0	6.0	6.0
OPT <sub>125M→1.3B</sub>	-	2.7	-	-	-
OPT <sub>350M→1.3B</sub>	-	1.7	1.0	2.0	1.0
OPT <sub>125M→2.7B</sub>	3.0	3.0	3.0	1.0	3.0
OPT <sub>350M→2.7B</sub>	3.3	4.5	1.0	1.0	1.0
OPT <sub>1.3B→2.7B</sub>	2.3	3.0	-	-	-
FlanT5 <sub>base→XL</sub>	-	1.0	1.8	9.0	0.4

is retained, the performance especially for captioning will drop. The conclusion can also be found in Table 4. We refer the readers to the next subsection for more details.

### C.3 Comparison between VPGTrans and VPG Inherit

As shown in Fig. 13, the green line (our VPGTrans) can be higher than the orange line (VPG inherit) for the majority of various conditions. Especially when considering the best performance of different tasks, our VPGTrans can achieve better performance than VPG inherit (non "-" in Table 5) for over 74% Transfer-Task variants. Moreover, among the variants that our VPGTrans can achieve better performance, our VPGTrans can also achieve a speed-up on 69.2% conditions.

### C.4 Scale-up Experiment Implementation Details

The scale-up experiment refers to the results in Table 4. We try to imitate BLIP-2’s pre-training data composition. First of all, two human-annotated datasets **COCO** and **VG** are used. **SBU** is also used. Then, BLIP-2 uses BLIP to generate captions for the 115M web images and rank them with CLIP ViT-L/14. We also adopt similar synthetic data from **Laion-COCO**.<sup>6</sup> We report the concrete number of data we use in Table 4. For the stage-1 training, we keep the same as the previous validation experiments where COCO and SBU are used for warm-up with a 5 times the learning rate. Then, we use COCO, VG, and Laion-COCO for the stage-2 training. Note that we have tried to include Laion-COCO and VG for the stage-1 training, but found no obvious difference and thus use COCO and SBU for simplicity. For VL-Vicuna, to align with the conversation scenario, we further fine-tune our VL-Vicuna with MiniGPT-4’s self-instruct data, which is 3,439 image-text pairs.

<sup>6</sup><https://laion.ai/blog/laion-coco>

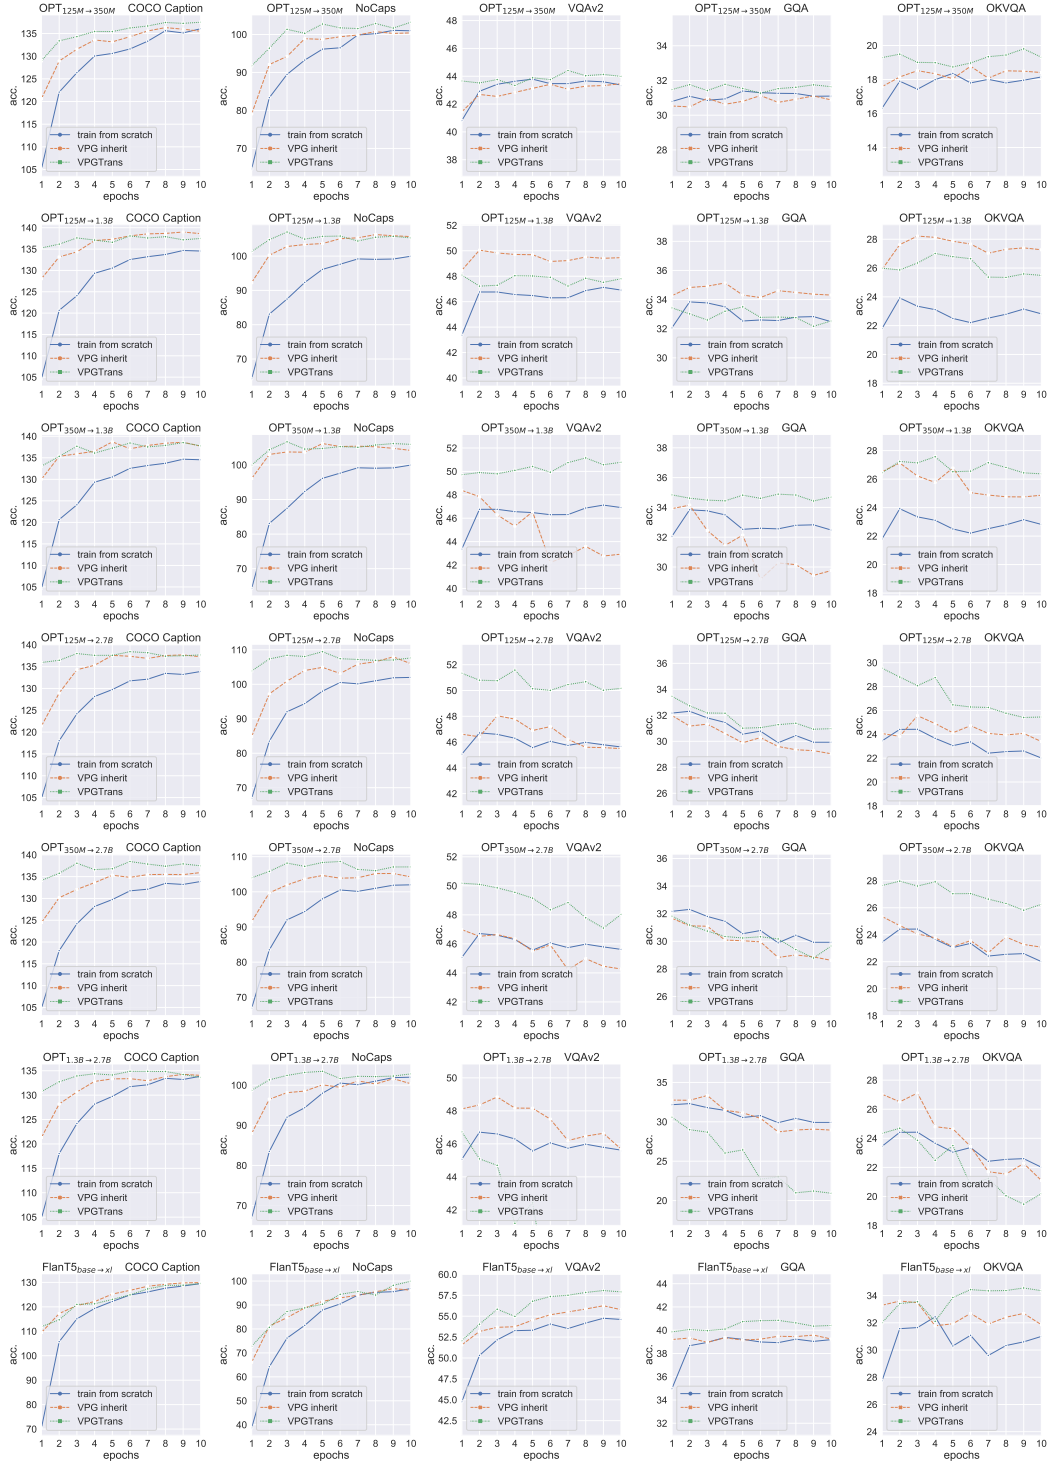


Figure 13: Comparison between different methods across 7 TaS variants on 5 tasks. Note that the model is directly evaluated after pre-training without further fine-tuning.

Table 6: Comparison between tuning only the projector (*i.e.* linear transfer) and the best results the VPGTrans achieve.

Transfer	COCO Caption		VQAv2	
	linear	best	linear	best
FlanT5 <sub>base</sub> $\rightarrow$ OPT <sub>350M</sub>	110.8	136.5	40.4	44.2
FlanT5 <sub>XL</sub> $\rightarrow$ OPT <sub>2.7B</sub>	132.1	139.3	50.4	50.3
OPT <sub>350M</sub> $\rightarrow$ FlanT5 <sub>base</sub>	1.1	122.1	34.0	49.9
OPT <sub>2.7B</sub> $\rightarrow$ FlanT5 <sub>XL</sub>	106.5	133.2	51.3	53.5

## D Extended TaT Experiments

In this section, we mainly illustrate extending findings of TaT experiments (*cf.* §5). We plot a more complete version of Fig. 8 in Fig. 14.

### D.1 Linear Transfer Gap between Different LLM’s Visual Prompts

As illustrated in Section 5.2, it is more difficult to transfer between two small LLMs with our VPGTrans due to the weaker linear transferability between two small LLMs’ visual prompts. To better support our results, we compare the results of tuning only the projector (*i.e.* linear transfer) and the best results the VPGTrans can achieve. As shown in Table 6, we can see that when conducting transfers between two large models (OPT<sub>350M</sub> and FlanT5<sub>base</sub>), the performance is typically far from the optimal results. However, when we transfer between OPT<sub>2.7B</sub> and FlanT5<sub>XL</sub>, the linear performance is near to the optimal performance. There is a 25.7 points gap between FlanT5<sub>base</sub>  $\rightarrow$  OPT<sub>350M</sub>’s *linear* and *best* on COCO caption datasets but only 7.2 points gap between FlanT5<sub>XL</sub>  $\rightarrow$  OPT<sub>2.7B</sub>’s *linear* and *best*. If considering the transfer between the small to large LLMs under TaT, both the VPG’s visual perception ability and transfer gap should be considered. We leave the systematical exploration for future works.

## E Extended Results for VL-LLM Customization

We show more comparisons between VL-Vicuna and MiniGPT-4 in Fig. 15. We can see that our VL-Vicuna has better visual perception ability. For example, when MiniGPT-4 falsely recognizes the three people in the image as two in the first example, our VL-Vicuna can not only recognize the number of people but also tell their roles. Moreover, our VL-Vicuna can successfully link the vision content with external knowledge. In the third example, our VL-Vicuna can recognize Leonardo and link the content with his films like Titanic.

## F Potential Impact and Limitations

Our VPGTrans is designed for building new VL-LLMs with lower computational cost, *i.e.*, shorter training time and less training data. With an already pre-trained VL-LLM, VPGTrans enables fast VPG transfer to build either a larger VL-LLM or a VL-LLM with a different type of LLM. We hope VPGTrans can facilitate teams in LLM community to customize their VL-LLMs with reduced cost. There are also possible limitations of the current version of VPGTrans. The first one is that our VPGTrans should rely on some already-aligned VPGs. The second potential limitation is that the VPGTrans-built VL-LLMs still suffer from the common problems of content generation AI systems [38, 39, 41, 7]. For example, the VL-Vicuna may make up some sentences with falsely recognized visual facts, like what is shown in Fig. 16. It is worth exploring associating our VPGTrans with training safer models [5, 34, 31].

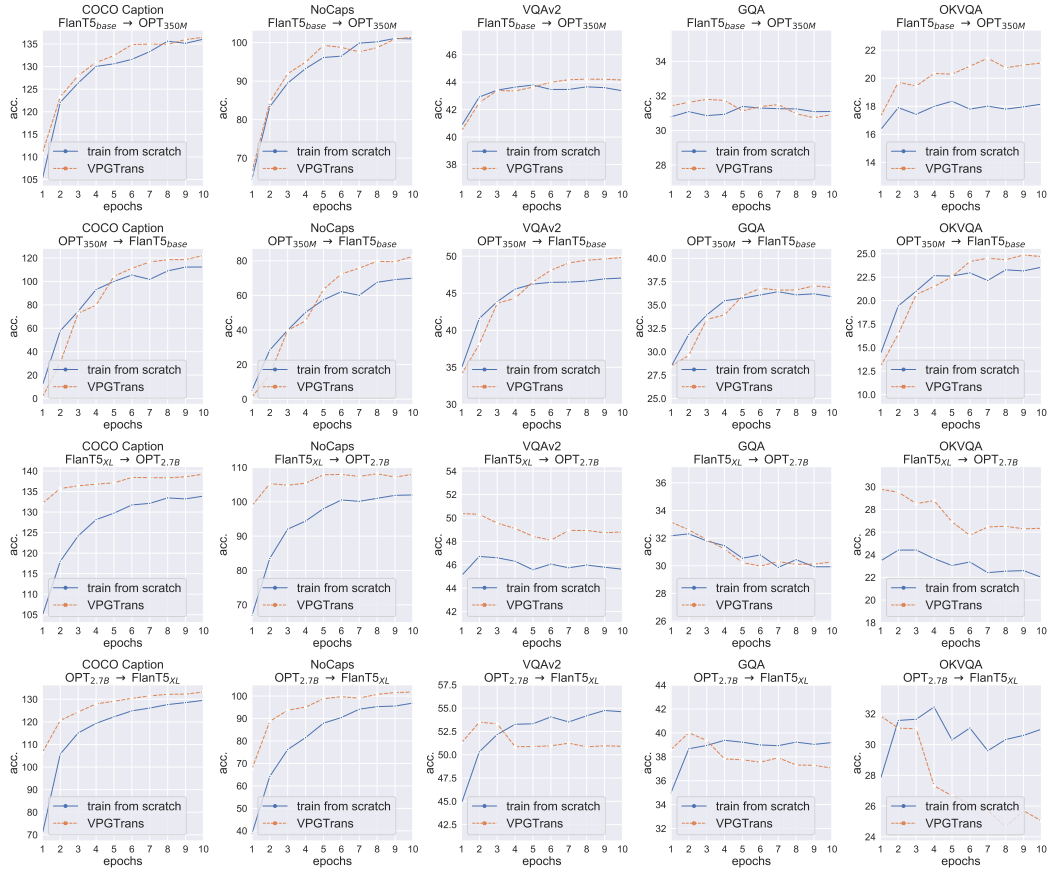


Figure 14: Comparison between different methods across 4 TaS variants on 5 tasks. Note that the model is directly evaluated after pre-training without further fine-tuning.

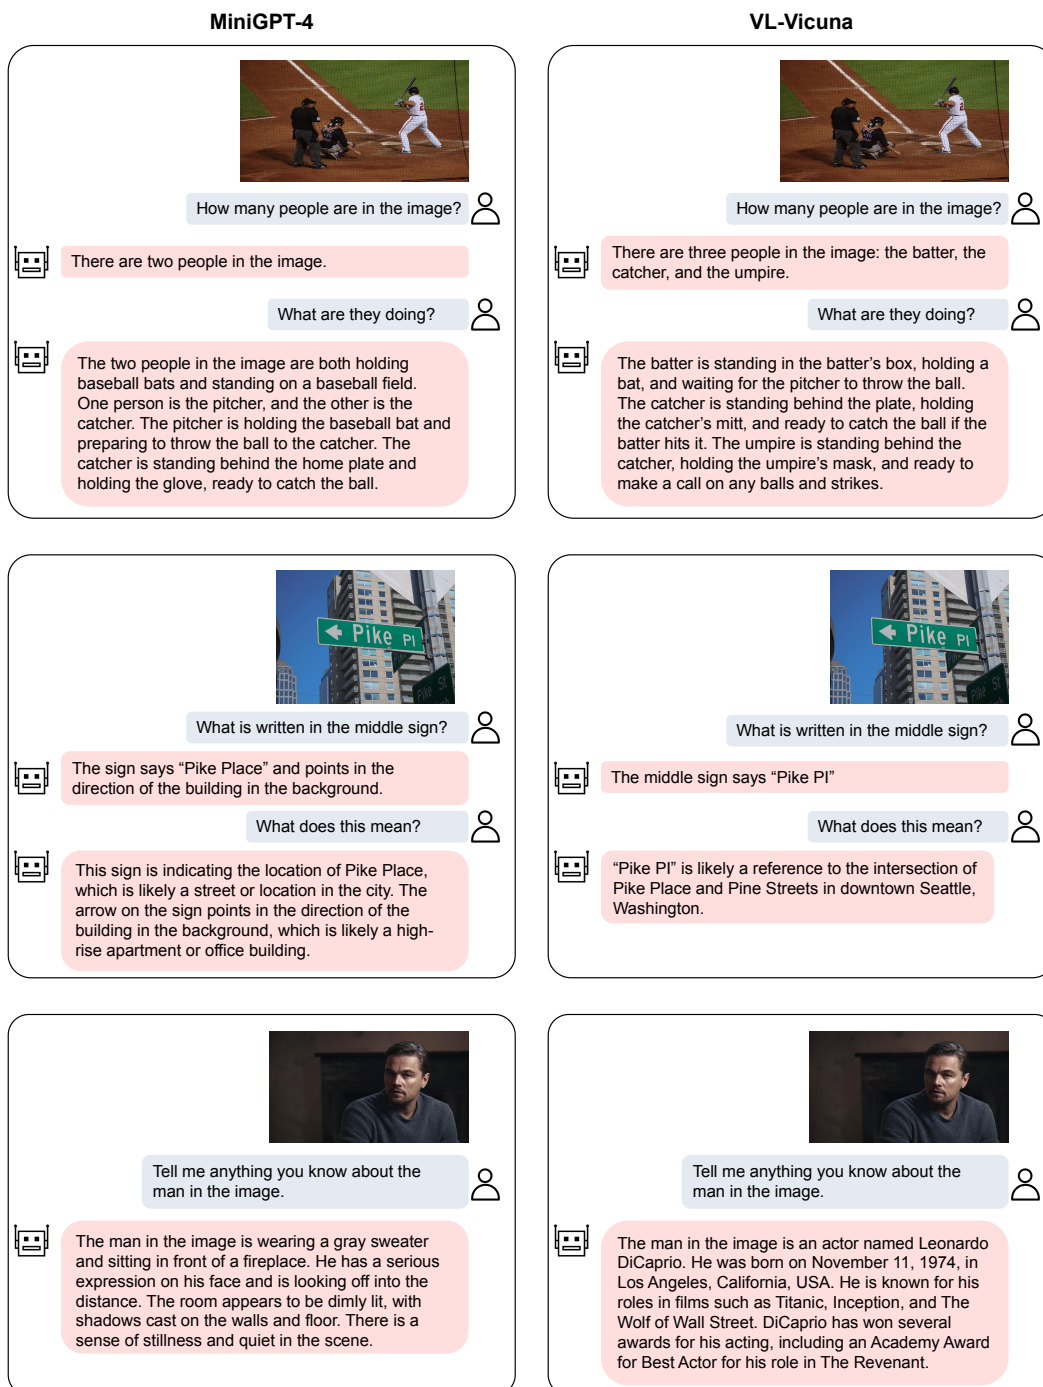


Figure 15: Comparison between MiniGPT-4 and our VL-Vicuna.

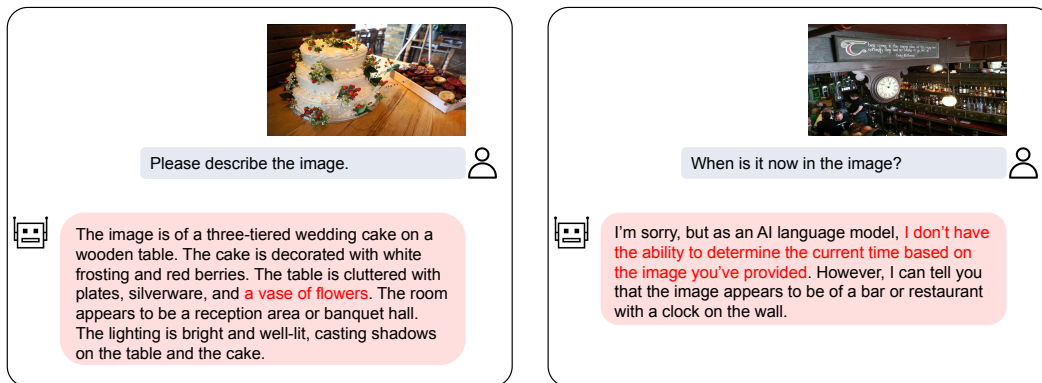


Figure 16: Failure cases of our VL-Vicuna.