LLaMA-Adapter V2: Parameter-Efficient Visual Instruction Model

Peng Gao[*‡†1], Jiaming Han[*1], Renrui Zhang[*1,2], Ziyi Lin[*2], Shijie Geng[3], Aojun Zhou[2]
Wei Zhang[1], Pan Lu, Conghui He[1], Xiangyu Yue[2], Hongsheng Li[†2], Yu Qiao[†1]

[1]Shanghai Artificial Intelligence Laboratory   [2]CUHK MMLab
[3]Rutgers University

{gaopeng, hanjiaming, zhangrenrui, qiaoyu}@pjlab.org.cn

Abstract

How to efficiently transform large language models (LLMs) into instruction followers is recently a popular research direction, while training LLM for multi-modal reasoning remains less explored. Although the recent LLaMA-Adapter demonstrates the potential to handle visual inputs with LLMs, it still cannot generalize well to open-ended visual instructions and lags behind GPT-4. In this paper, we present LLaMA-Adapter V2, a parameter-efficient visual instruction model. Specifically, we first augment LLaMA-Adapter by unlocking more learnable parameters (e.g., norm, bias and scale), which distribute the instruction-following ability across the entire LLaMA model besides adapters. Secondly, we propose an early fusion strategy to feed visual tokens only into the early LLM layers, contributing to better visual knowledge incorporation. Thirdly, a joint training paradigm of image-text pairs and instruction-following data is introduced by optimizing disjoint groups of learnable parameters. This strategy effectively alleviates the interference between the two tasks of image-text alignment and instruction following and achieves strong multi-modal reasoning with only a small-scale image-text and instruction dataset. During inference, we incorporate additional expert models (e.g. captioning/OCR systems) into LLaMA-Adapter to further enhance its image understanding capability without incurring training costs. Compared to the original LLaMA-Adapter, our LLaMA-Adapter V2 can perform open-ended multi-modal instructions by merely introducing 14M parameters over LLaMA. The newly designed framework also exhibits stronger language-only instruction-following capabilities and even excels in chat interactions. Our code and models are available at https://github.com/ZrrSkywalker/LLaMA-Adapter.

---

‡Project leader, * Equal contribution, † Corresponding author
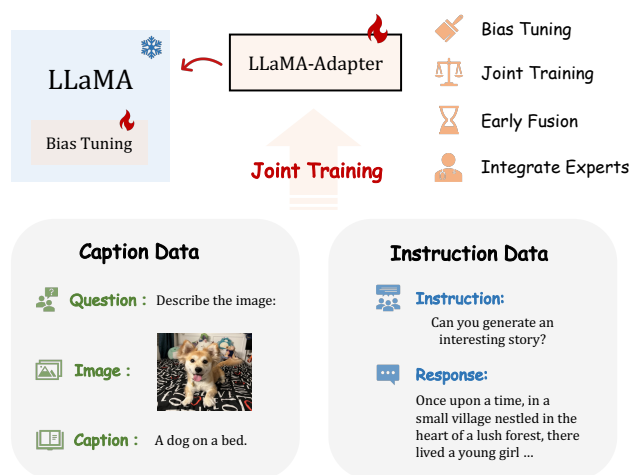
LLaMA-Adapter V2 :

Figure 1. **Training Pipeline of LLaMA-Adapter V2.** We introduce several strategies to enhance the capability of LLaMA-Adapter [72], which enable a parameter-efficient visual instruction model with superior multi-modal reasoning.

1. Introduction

Large Language Models (LLMs) [75] have garnered significant attention in the AI community because of their exceptional ability to comprehend, reason with, and generate human language. To make LLM's responses more vivid and convincing, recent works [7, 19, 50, 62, 68] have explored transforming LLMs into instruction-following models. For example, Stanford Alpaca [62] fine-tunes LLaMA [63] into an instruction-following model using instruction examples generated from OpenAI's Instruct-GPT model [48]. Follow-up works of Alpaca further extend LLaMA by utilizing higher-quality instruction data, such as ShareGPT [1] and those generated by GPT-4 [50]. Compared with the full fine-tuning paradigm adopted by Alpaca and Vicuna [7], LLaMA-Adapter [72] introduces

lightweight adapters with zero-initialized attention into the frozen LLaMA for parameter-efficient fine-tuning, along with multi-modality knowledge injection. Despite the significant progress, the aforementioned approaches are still unable to perform more advanced multi-modal instructions, *e.g.*, visual understanding like GPT-4 [47].

Most recently, studies such as MiniGPT-4 [78] and LLaVA [38] have sparked a new wave of research on extending language-only instruction models into multi-modal ones to empower LLMs with visual reasoning ability, in a similar way to LLaMA-Adapter. MiniGPT-4 connects a frozen visual encoder and an LLM by pre-training on 134 million image-text pairs, and then improves the model's performance by further fine-tuning on a well-aligned image-text dataset. LLaVA also leverages image-text pairs to align the visual model and the LLM. Different from MiniGPT-4, LLaVA fine-tunes the whole LLM on 150K high-quality multi-modal instruction data generated by GPT-4. While these methods demonstrate impressive multi-modal understanding capabilities, they require updating billions of model parameters and meticulously collecting a substantial amount of multi-modal training data, which are either annotated by humans or distilled from responses of OpenAI API.

In this paper, we aim to design a parameter-efficient visual instruction model. We build on the popular parameter-efficient LLaMA-Adapter to develop our new method, which we refer to as LLaMA-Adapter V2. LLaMA-Adapter was initially developed as an instruction-following model and can be easily transformed into a visual instruction model by incorporating visual features into the adaptation prompts. However, due to the lack of multi-modal instruction tuning data, the multi-modal version of LLaMA-Adapter is restricted as a traditional vision-language model. For instance, LLaMA-Adapter trained on COCO Caption [6] can only produce short image captions when given a specific prompt, such as "Generate caption for this image". The model cannot adapt to open-ended multi-modal instructions, such as complex visual reasoning and visual question answering tasks.

Although we do not currently leverage multi-modal instruction data, it is still possible to perform multi-modal instruction tuning for LLaMA-Adapter. We begin by using a frozen instruction-following LLaMA-Adapter model as the starting point and refine it by optimizing the visual projection layers on image-text pairs to ensure proper vision-language alignment. However, we have observed that the visual features tend to dominate the adaptation prompts, causing the instruction-following ability to deteriorate rapidly.

To address this challenge, we propose a simple ***early fusion of visual knowledge*** strategy that resolves the interference between the two tasks of image-text alignment and language instruction tuning. In LLaMA-Adapter, dynamic visual prompts are incorporated into the static adaptation prompts at the last L layers. However, in LLaMA-Adapter V2, we distribute the dynamic visual prompts to only the first K layers, where $K < N - L$ and N denote the total number of Transformer layers. As a result, image-text alignment no longer disrupts the model's instruction-following capability. With this strategy, we can achieve superior visual instruction learning through the ***joint training with disjoint parameters*** with image captioning data and instruction-following data, even in the absence of high-quality multi-modal instruction data. Additionally, we augment LLaMA-Adapter by unlocking more learnable parameters such as normalization, layer bias and scale, which we refer to as ***bias tuning of linear layers***. By increasing the model's tunable capacity, we can spread the instruction-following knowledge across the entire LLM. It is worth noting that these parameters only account for approximately 0.04% of the entire model, ensuring that LLaMA-Adapter V2 remains a parameter-efficient approach.

Finally, we introduce additional expert models (*e.g.*, captioning, detection, and OCR systems) to enhance the image understanding capabilities of LLaMA-Adapter V2, setting our approach apart from others such as MiniGPT-4 and LLaVA that rely on massive image-text pair training data. By collaborating with these specialized experts, our framework gains increased flexibility and allows plugging in various experts for a wide variety of tasks without the necessity of pre-training on extensive vision-language data.

Fig. 1 and Fig. 4 illustrate the whole training and generation pipelines of our LLaMA-Adapter V2, respectively. We summarize our main contributions as follows:

- **Stronger Language Instruction Model.** With the parameter-efficient tuning strategy and high-quality language instruction data, LLaMA-Adapter V2 surpasses its predecessor LLaMA-Adapter in terms of language instruction-following performance. Moreover, LLaMA-Adapter V2 is capable of conducting multi-turn dialog, demonstrating its stronger ability as a language instruction model.

- **Balanced Visual Instruction Tuning.** We propose a simple early fusion strategy to solve the interference between image-text alignment and instruction-following learning targets. As a result, we transform LLaMA-Adapter V2 into a visual instruction model without the need of multi-modal instruction training data.

- **Integration of Expert Systems.** Instead of end-to-end pre-training on massive image-text pairs, we embrace modular design where different expert models can be integrated into our framework to enhance the image understanding abilities of LLMs.

2

Model	Language Instruction Data			Image-Text Data		Visual Instruction Data		Tuning Parameters
	Source	Type	Size	Source	Size	Source	Size	Size
MiniGPT-4 [78]	ShareGPT [1]	Conversation	70K	CC-VG-SBU-L400*	134M	CC+ChatGPT	5K	13B
LLaVA [38]	ShareGPT [1]	Conversation	70K	CC*	595K	COCO+GPT4	158K	13B
LLaMA-Adapter V2	GPT-4-LLM [50]	Single-turn	**52K**	COCO [6]	**567K**	-	**0**	**14M**

Table 1. **Training Comparison of Different Methods.** CC, VG and L400 represent Conceptual Caption [5, 57], Visual Genome [30] and LAION 400M [56], respectively. Note that we count all the data and tuning parameters needed to convert the pretrained vision model and LLM into a visual instruction model. ∗ denotes the filtered dataset.

2. Related Work

Instruction-following Language Models Large Language Models (LLMs) are pre-trained on extensive text corpora using autoregressive Transformer models to predict subsequent tokens. They have demonstrated strong capabilities as self-supervised [53], multi-task [54], and few-shot learners [4]. Recently, InstructGPT [48] and FLAN [8, 66] have shown that LLMs can be converted into instruction-following models by fine-tuning them on instructional datasets. To facilitate the generation of instruction-following examples, Self-Instruct [65] employs a semi-automated, iterative bootstrapping algorithm that expands a limited seed set of manually-written instructions and progressively expanding the collection of tasks using off-the-shelf LLMs. Alpaca [62] applies the Self-Instruct strategy to produce 52K high-quality instruction-following demonstrations and fine-tunes upon the open-source LLaMA [63] model, ultimately obtaining an instruction-following language model exhibiting many behaviors similar to OpenAI's text-davinci-003. Inspired by Alpaca's success, Vicuna [7] and GPT-4-LLM [50] further reveal that dialog and enhanced instruction-following capabilities can be ignited by fine-tuning on either user-shared ChatGPT conversations or instruction-following data generated by the GPT-4 API. However, Alpaca, Vicuna and GPT-4-LLM all fine-tune the full parameters of LLMs, resulting in unaffordable GPU memory usage and training costs. In contrast, LoRA [26] and LLaMA-Adapter [72] validate that parameter-efficient fine-tuning approaches can potentially replace full parameter updates during supervised fine-tuning of LLMs. In this paper, LLaMA-Adapter V2 goes a step further by constructing a parameter-efficient zero-shot visual instruction model that reuses the instruction-following ability of LLaMA-Adapter.

Visual Instruction Models Traditional image captioning [3, 9, 23, 32, 40, 64] and visual question answering (VQA) [2, 16, 20, 29, 34, 39, 41, 45, 77] approaches can be viewed as simplified versions of visual instruction models. However, their image understanding capabilities fall short compared to GPT-4. To be specific, image captioning models can only describe images using concise, short sentences, while VQA systems are capable of answering various visual-related questions but without providing explicit explanations. Therefore, these systems have yet to achieve the level of proficiency required to meet the demands of real-world applications. Recently, GPT-4 has showcased remarkable visual instruction-following abilities by handling mixed inputs of images and text for diverse tasks, ranging from meme explanation, document-level reasoning, exam problem-solving, and so on. In pursuit of developing visual instruction-following abilities akin to GPT-4, both LLaVA [38] and MiniGPT-4 [78] assemble high-quality multi-modal instruction-following data using Chat-GPT or GPT-4. They then integrate visual encoders and LLMs via a projection layer and fine-tune the models on the curated data. Diverging from these approaches, LLaMA-Adapter V2 strives to build a parameter-efficient visual instruction model without relying on multi-modal instruction data. Rather, LLaMA-Adapter V2 can function effectively using just language instruction data and image-text pairs.

Parameter-efficient Fine-tuning The pre-training-fine-tuning paradigm has been proven highly effective in various tasks such as visual recognition, language understanding, text generation, and image synthesis from textual descriptions. However, as model sizes continue to increase exponentially, fine-tuning every parameter in a colossal model becomes increasingly impractical. In contrast, parameter-efficient fine-tuning (PEFT) approaches [12,49] freeze most of the parameters of foundation models and only optimize a small portion of them. Many successful PEFT approaches [10, 17, 22, 28, 31, 59, 60, 71, 73] have been proposed for adopting popular pre-trained models such as BERT [11], GPT [4, 53, 54], ViT [13], CLIP [52], and Stable Diffusion [55] to various downstream tasks. In general, these PEFT approaches can be divided into three categories, namely prefix tuning (e.g. [35,76]), reparameterization (e.g. [10,26,43]), and adapters (e.g. [15,24,74]). In this paper, we present LLaMA-Adapter V2, an elegant and efficient marriage of prefix-tuning and adapter techniques. By utilizing an early fusion strategy and bias tuning, LLaMA-Adapter V2 injects visual features into large language models, yielding impressive multi-modal instruction-following performance with only 0.04% parameters of the entire LLaMA.

Integration of Expert Systems Collaboration among multiple expert models has proven to be a successful practice in the field of AI, as it often leads to improved performance and robustness. This is especially true in the realm of computer vision tasks, where ensembles of multiple models consistently achieve top positions on the challenge leaderboards. In recent years, with the continual expansion of LLMs' capabilities, there has been a growing trend [14, 18, 21, 27, 37, 42, 51, 58, 61, 67, 69, 79] to combine them with visual foundation models and leverage their combined strengths to tackle more complex vision-language tasks. By utilizing LLMs as a core controller for external visual models, these experts in turn support LLMs to perform a wider range of tasks that demand a deeper visual understanding. For example, recent studies such as HuggingGPT [58], Visual ChatGPT [67], Chameleon [42], MMReACT [69], and ViperGPT [61] utilize LLMs as the central manager to perform compositional task planning and call upon off-the-shelf expert models/tools to aid in various complex multi-modal tasks, including understanding, generation, search, reasoning, and programming, etc. In addition, PaLM-E [14], Inner Monologue [27], and Text2Motion [37] further extend the capabilities of LLMs to robotics by incorporating real-world sensor modalities. These embodied LLMs possess the power to comprehend natural language instructions and execute sequential manipulation planning in the real world. As a result, the aforementioned approaches facilitate the seamless integration of different expert systems and boost the overall performance and capabilities of LLMs. LLaMA-Adapter V2 stands out by integrating short, yet precise descriptions generated by LLaMA-Adapter during inference time for efficient zero-shot and training-free visual instruction understanding, setting it apart from other methods that require extensive multimodal data during training. In the future, more expert visual systems will be integrated into LLaMA-Adapter V2 for building stronger visual instruction models.

3. A Revisit of LLaMA-Adapter

Zero-initialized Attention. As a parameter-efficient fine-tuning solution for adapting LLaMA to acquire instruction-following capability, LLaMA-Adapter [72] freezes the entire LLaMA model [63] and introduces only an extra lightweight adapter module with 1.2M parameters. The adapter layers are employed at the higher Transformer layers of LLaMA and concatenate a set of learnable soft prompts as the prefix to the word tokens. To incorporate newly adapted knowledge into the frozen LLaMA, LLaMA-Adapter proposes a zero-initialized attention mechanism, which enables adaptively controlling the contribution of adaptation prompts to the word tokens by learning a gating factor initialized by zero. The gating magnitude progressively increases during training, thereby gradually injecting

the instruction-following ability into the frozen LLaMA. This strategy not only preserves LLaMA's language generation ability during the early training stages but also continuously incorporates new knowledge to enable a powerful instruction follower.

Simple Multi-modal Variant. Besides fine-tuning using language-only instructions, LLaMA-Adapter can also incorporate image and video inputs for multi-modal reasoning. For instance, when dealing with images, LLaMA-Adapter employs a pre-trained visual encoder such as CLIP [52] to extract multi-scale visual features. These features are then aggregated into a global feature and passed through a learnable projection layer to align the visual semantics with linguistic embedding space. Afterward, the global visual feature is added element-wisely to every adaptation prompt at the higher layers of the Transformer. This allows LLaMA-Adapter to generate responses based on both textual and visual inputs, resulting in competitive performance on the ScienceQA benchmark [41].

Open-ended Multi-modal Reasoning. While LLaMA-Adapter is capable of handling relatively simple tasks such as ScienceQA, it is still unclear whether it can generate open-ended responses, such as those required for general-purpose visual question answering. To investigate this, we first start with a LLaMA-Adapter pre-trained on language instruction data for leveraging its existing instruction-following capabilities. We then conduct experiments by fine-tuning its adapter modules and visual projection layers on the COCO Caption [6] dataset. However, we found that the newly learned visual cues tend to dominate the adaptation prompts, overriding the inherent instruction-following characteristics. Therefore, we propose LLaMA-Adapter V2, a parameter-efficient visual instruction model, to fully unleash the multi-modal potential of LLaMA.

4. LLaMA-Adapter V2

In this section, we present the technical details of LLaMA-Adapter V2, including ***bias tuning of linear layers*** (Sec. 4.1) to enhance its language instruction-following ability, ***joint training with disjoint parameters*** (Sec. 4.2) for balanced visual instruction tuning, ***early fusion of visual knowledge*** (Sec. 4.3) to balance textual and visual understanding, and ***integration with experts*** (Sec. 4.4) to boost zero-shot multi-modal reasoning.

4.1. Bias Tuning of Linear Layers

LLaMA-Adapter employs learnable adaptation prompts with the zero-initialized attention mechanism (Sec. 3) on the frozen LLaMA model, which allows for efficient incorporation of new knowledge. However, the parameter updates
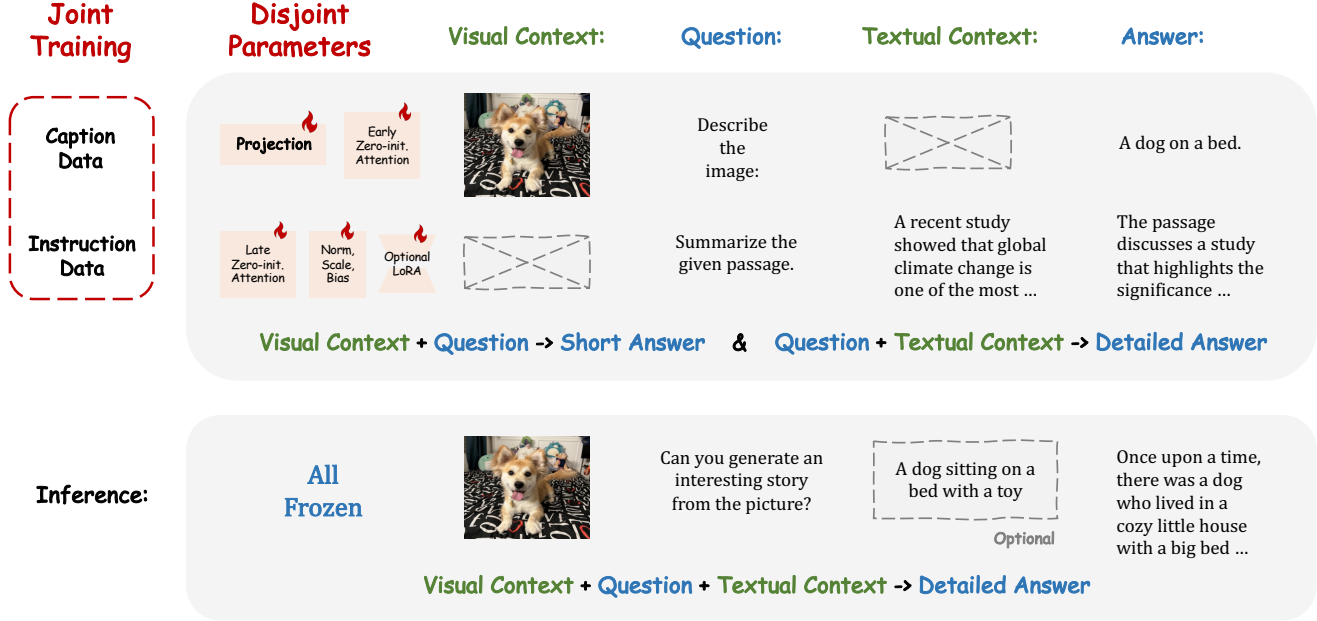
Figure 2. **Joint Training Paradigm in LLaMA-Adapter V2.** We utilize both image-text caption and language-only instruction data to jointly train LLaMA-Adapter V2, optimizing disjoint groups of learnable parameters.

are limited to the adaptation prompts and the gating factor, without modifying the internal parameters of LLMs, which restricts its ability to perform deep fine-tuning. In light of this, we propose a bias tuning strategy to further fuse instruction cues into LLaMA besides the adaptation prompts and the gating factor. Specifically, to adaptively handle the tasks of instruction-following data, we first unfreeze all the normalization layers in LLaMA. For each linear layer in the Transformer, we add a bias and a scale factor as two learnable parameters. We denote the input and pre-trained weights of a certain linear layer as $\mathbf{x}$ and $\mathbf{W}$, respectively. In LLaMA-Adapter V2, we modify the linear layer using the bias b and scale s as

$$\mathbf{y} = \mathbf{W} \cdot \mathbf{x} \quad \rightarrow \quad \mathbf{y} = s \cdot (\mathbf{W} \cdot \mathbf{x} + b), \quad (1)$$

$$\text{where} \quad b = \text{Init}(0), \quad s = \text{Init}(1). \quad (2)$$

Similar to zero-initialized attention, we initialize the bias and scale factors with zeros and ones, respectively, to stabilize the training process at early stages. With the incorporation of the bias tuning strategy and high-quality instruction data [50], LLaMA-Adapter V2 acquires superior instruction-following capabilities. Notably, the number of newly added parameters only accounts for 0.04% ($\sim$5M) of the entire LLaMA, showing that LLaMA-Adapter V2 is still a highly parameter-efficient approach.

Discussion. Our bias tuning strategy bears similarity to prior parameter-efficient methodologies such as BitFit [70]

for BERT fine-tuning and SSF [36] for visual prompt tuning [28]. However, both BitFit and SSF are designed for comprehension tasks with an 80-million parameter scale, whereas our bias tuning demonstrates its efficiency on large language models ranging from 7 billion to 65 billion parameters, such as LLaMA and GPT-3. Moreover, our bias tuning strategy is input-agnostic, unlike Low-Rank Adaptation (LoRA) that adds an input-aware bias using low-rank transformation, further reducing the fine-tuning cost.

4.2. Joint Training with Disjoint Parameters

Our goal is to simultaneously endow LLaMA-Adapter V2 with the capabilities of generating long language responses and multi-modal understanding. As shown in Fig. 2, we propose a joint training paradigm for LLaMA-Adapter V2 to leverage both image-text captioning data and language-only instruction examples. Due to the data volume difference between 500K image-text pairs and 50K instruction data, naively combining them for optimization can severely harm LLaMA-Adapter's instruction-following ability, just as discussed in Sec. 3. Therefore, our joint training strategy optimizes disjoint groups of parameters in LLaMA-Adapter V2 for image-text alignment and instruction-following respectively. Specifically, only the visual projection layers and early zero-initialized attention with gating are trained for image-text captioning data, while the late adaptation prompts together with zero gating, the unfrozen norm, newly added bias and scale factors (or optional low-rank adaption [25]) are utilized for learn-
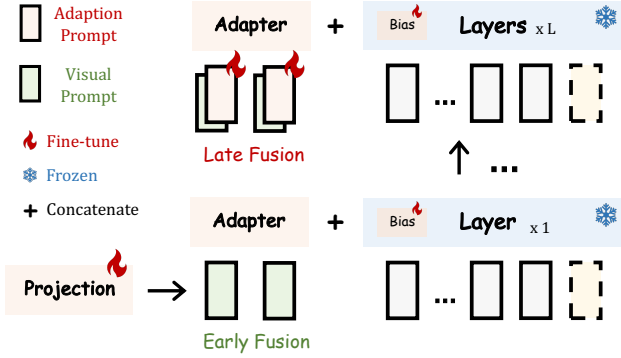
Figure 3. **Early Fusion of Visual Knowledge.** Following LLaMA-Adapter, we insert static adaptation prompts into the last L layers. For visual prompts, we insert them in the early stage of LLM, disjointing with adaptation prompts.
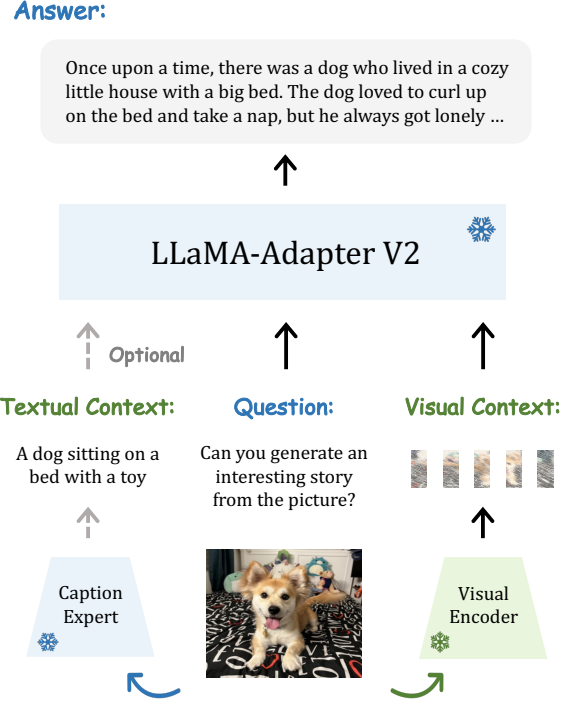

Figure 4. **Generation Pipeline of LLaMA-Adapter V2.** During inference, we introduce additional caption experts to generate a textual context for the input image, demonstrating strong visual understanding capacity of LLaMA-Adapter V2.

ing from the instruction-following data. Such disjoint parameter optimization naturally solves the interference issue between image-text understanding and instruction following, which contributes to the emergent visual instruction-following ability of LLaMA-Adapter V2.

Discussion. Aided by our joint training strategy, LLaMA-Adapter V2 requires no high-quality multi-modal instruction data like MiniGPT-4 [78] and LLaVA [38], but only image-text pairs and instruction-following data, as compared in Tab. 1. The captioning data extends LLMs for image understanding with short answers as shown in Fig. 2, while the language-only instruction data is utilized to preserve LLaMA's capacity to generate long detailed sentences. With such complementarity, LLaMA-Adapter V2 achieves superior multi-modal reasoning by solely small-scale image-text and instruction-following data, without the need of high-quality multi-modal instruction data.

4.3. Early Fusion of Visual Knowledge

To avoid interference between visual and language fine-tuning, we propose a simple early fusion strategy to prevent direct interactions between the input visual prompts and adaptation prompts. In LLaMA-Adapter, the input visual prompts are sequentially encoded by a frozen visual encoder with a learnable visual projection layer, and then added to the adaptation prompts at every inserted layer. In LLaMA-Adapter V2, we instead inject the encoded visual tokens and adaptation prompts to different Transformer layers without fusing them together, which is shown in Fig. 3. For the dataset-shared adaptation prompts, we follow LLaMA-Adapter by inserting them at the last L layers (*e.g.*, L=30). For the input visual prompts, we directly concatenate them with the word tokens at the first

Transformer layer with the zero-initialized attention, other than adding them to the adaptation prompts. Together with the proposed joint training, this simple early fusion strategy of visual tokens can effectively resolve the conflict between the two types of fine-tuning targets. This contributes to a parameter-efficient LLaMA-Adapter V2 with superior multi-modal reasoning capabilities.

4.4. Integration with Experts

Recent visual instruction models such as MiniGPT-4 [78] and LLaVA [38] require massive-scale image-text training to connect visual models and LLMs. In contrast, our LLaMA-Adapter V2 fine-tunes on much smaller-scale common image captioning data [6], making it more data-efficient. However, the image understanding ability of our approach is relatively weak, leading to occasional inaccurate or unrelated responses. Rather than collecting more image-text data or adopting stronger multi-modal modules, we propose integrating expert systems, such as captioning, OCR, and search engines, to supplement LLaMA-Adapter V2 with additional visual reasoning proficiency.

As shown in Fig. 4, we utilize expert systems, such as captioning, detection, and OCR, to enhance the visual instruction-following capabilities of LLaMA-Adapter V2.

Instruction: How can I develop my critical thinking skills? Give me three suggestions.

LLaMA-Adapter V2:
1. Read widely: Reading books, articles, and other sources of information can help you develop your critical thinking skills. By exposing yourself to different perspectives and ideas, you can learn to analyze and evaluate information more effectively.
2. Practice mindfulness: Mindfulness is the practice of being present in the moment and paying attention to your thoughts and feelings. By being more aware of your thoughts and emotions, you can better understand and evaluate them. This can help you develop a more objective and rational approach to thinking.
3. Engage in open-minded discussions: Join a discussion group or debate club where you can engage in open-minded discussions with others. This can help you develop your critical thinking skills by challenging your assumptions and encouraging you to consider different perspectives.

LLaMA-Adapter: 1. Read widely and critically. 2. Engage in discussions with people who have different perspectives. 3. Ask questions and seek out answers.

Instruction: As a superhero, how would you explain your origin story to a curious child?

LLaMA-Adapter V2: As a superhero, I have a very interesting and exciting origin story to share with you. Once upon a time, there was a young boy named John who lived in a small town. John was always curious and wanted to help others, so he spent his days exploring the world around him and learning about different cultures and people. One day, while he was out exploring, he stumbled upon a mysterious cave. Inside the cave, John found a strange device that had the power to give him incredible abilities. With these abilities, John could fly, shoot lasers from his eyes, and even turn invisible.

John knew that he had to use these powers for good, so he decided to become a superhero and use his abilities to protect his town and the people who lived there. He started to train every day, learning how to use his powers and how to fight against evil. Eventually, John became the superhero known as "The Invisible Man" and he used his powers to help others and make the world a better place.

That's my origin story, and I'm sure you'll love it!

LLaMA-Adapter: Once upon a time, there was a young boy who had a special power. He used his power to help people in need and protect the innocent. He became known as the superhero, and he used his powers to fight evil and protect the world.

Instruction: Write a regex to remove all the HTML tags in a python string.

LLaMA-Adapter V2: Here is a simple regex that can be used to remove all HTML tags from a string in Python:

```python
import re

string = "This is a <strong>test</strong> string"

# Remove all HTML tags from the string
result = re.sub(r'<.*?>', '', string)

# Print the result
print(result)
```

This will output:

```
This is a test string
```

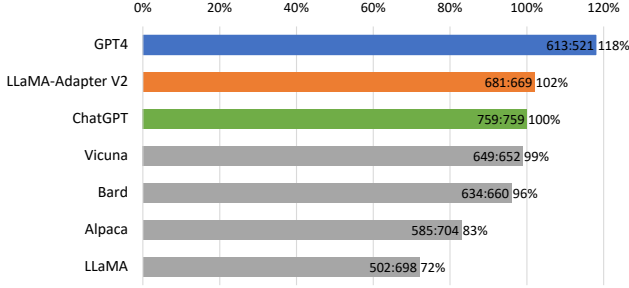**LLaMA-Adapter:**

```
re.sub(r'<[^>]*>', '', str)
```

Table 2. **Comparisons of Instruction-following Ability** between LLaMA-Adapter V2 and LLaMA-Adapter.

Given an input image, we encode its visual context using a pre-trained visual encoder and ask an expert system to produce a caption as the textual context. In our default implementation, we adopt LLaMA-Adapter pre-trained on COCO Caption [6] as the expert system, since it can generate short and accurate image descriptions. However, it is worth noting that any image-to-text model or even a search engine can serve as the expert system here. Our approach allows us to easily switch among different expert systems based on the specific downstream task at hand.

5. Experiments

5.1. Experimental Setups

Training Data. As shown in Tab. 1, we train LLaMA-Adapter V2 on 52K single-turn instruction data from GPT-4-LLM [50] and 567K captioning data from COCO Caption [6]. We do not use any visual instruction data like MiniGPT-4 [78] and LLaVA [38]. Besides, we also train a chatbot system using 80K conversation data collected by ShareGPT [1].

(a) The total quality scores compared with ChatGPT.
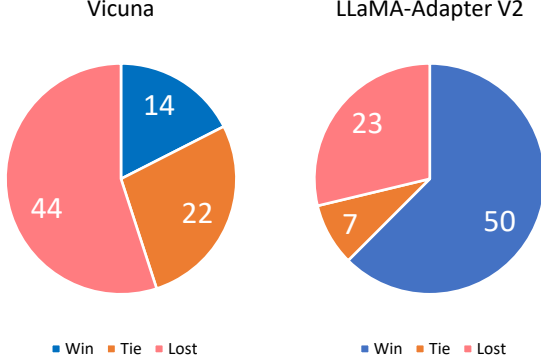
(b) Left: Vicuna *vs.* ChatGPT. Right: LLaMA-Adapter V2 *vs.* ChatGPT.

Figure 5. **Response Quality Comparisons** assessed by GPT4. The comparison baseline is ChatGPT. We use GPT4 to assess the response quality of different methods on 80 questions [7]. Note that our model is based on LLaMA-65B, while Vicuna is built on LLaMA-13B. But we only fine-tune **14M** parameters, while Vicuna requires fine-tuning the whole 13B parameters.

Implementation Details. For the LLaMA-7B model with 32 Transformer layers, we insert static adaptation prompts into the last 31 layers. Additionally, we append dynamic visual prompts to the first layer, with a prompt length set to 20. All the parameters in normalization layers, linear layer bias and scale are set to be updated during training, while the remaining parameters in LLaMA are kept frozen.

5.2. Stronger Language Instruction Model

With the proposed bias tuning strategy and high-quality instruction data [50], LLaMA-Adapter V2 was able to further enhance the instruction-following capability of LLaMA. As shown in Tab. 2, LLaMA-Adapter V2 can provide comprehensive answers to human instructions as well as the detailed explanations of the answers, while LLaMA-Adapter only delivers relatively short answers.

Given that bias tuning involves more learnable parameters for knowledge updating, it is possible to build a chatbot system that requires a deeper understanding of language context. By training LLaMA-Adapter V2 on 80K conversation data [1], we developed a stronger chatbot model.

Model	Data Scale		COCO Caption	
	PT	FT	BLEU@4	CIDEr
ClipCap [46]	0	0.6M	33.5	113.1
BLIP [33]	14M	0.6M	40.4	136.7
BLIP2 [32]	129M	0.6M	43.7	145.3
LLaMA-Adapter V2	0	0.6M	36.2	122.2

Table 3. **Comparisons on COCO Caption.** BLIP and BLIP2 require pretraining (PT) on COCO Caption [6], Visual Genome [30], Conceptual Captions [5, 57] and LAION [56], while ClipCap and our approach only fine-tune (FT) the model on COCO Caption.

Fig. 11 presents a chatbot example using a 7 billion model where the system is able to answer our questions, but its understanding of context is not very accurate. By scaling the model to 65 billion (Fig. 10), the chatbot becomes more powerful and answers our questions very well. In Fig 5, we also use GPT4 to assess the response quality, where LLaMA-Adapter V2 even beats ChatGPT on the total scores and 50/80 questions.

5.3. Visual Instruction Model

LLaMA-Adapter is primarily a language instruction model or a close-set vision-language model, while LLaMA-Adapter V2 is a powerful visual instruction model that is jointly trained on caption and language-only instruction data. In this section, we study LLaMA-Adapter V2's image captioning ability and demonstrate how it can be extended to general-purpose multi-modal understanding systems like GPT-4. In addition, we incorporate expert systems to further improve LLaMA-Adapter V2's visual instruction-following capability.

Image Captioning. As discussed in Sec. 3, LLaMA-Adapter supports multi-modal inputs by simply adding visual features to adaptation prompts. After fine-tuning on COCO Caption dataset, we transform LLaMA-Adapter into a strong image captioner. In Tab. 3, we can observe that LLaMA-Adapter achieves comparable results to BLIP [32, 33], even without pre-training on large-scale image-text data. However, LLaMA-Adapter cannot reuse the language modeling ability of the LLM since it is sensitive to specific prompts, such as "Generate caption for this image".

By employing early fusion and joint training, LLaMA-Adapter V2 has become a powerful visual instruction model that can simultaneously perform language instruction-following and image captioning. We provide some examples in Fig. 6 to compare the image captioning results between LLaMA-Adapter and LLaMA-Adapter V2. For an given image, LLaMA-Adapter can only produce a short, close-form image description, while LLaMA-Adapter V2 is able to generate natural and detailed image descriptions.

Success Case 1

V1: A bowl of cereal, a cup of coffee and a sandwich on a table.
V2: The image features a dining table set up with a variety of food items, including a bowl of fruit, a cup of coffee, and a plate of pastries. The table is set with a knife, fork, and a spoon, suggesting that the food is ready to be consumed. The dining table is situated in a kitchen, with a chair placed near the table. The chair is positioned next to a small table, which may be used for additional food items or drinks. The room is well-lit, with a lamp illuminating the area.

Success Case 2

V1: A group of people waiting to board a bus.
V2: In the image, a large bus is parked on the side of the road, with several people standing nearby. There are at least three people standing near the bus, with one person holding a backpack and another person holding a bag. A few more people can be seen in the background, walking towards the bus. The bus is painted white, and its front end is visible, with a few people standing near the front. A few cars are also visible in the background, adding to the busy street scene.

Failure Case

V1: A mouse and a duck in a body of water.
V2: In the image, two men are standing on a dock near a lake or river. One man is wearing a blue shirt and the other is wearing a white shirt. Both men are looking down and appear to be fishing. They are standing near a small rocky area and are in the water up to their knees. In the image, there is also a dog, a yellow labrador retriever, swimming in the water near the men. The dog appears to be swimming towards the men, perhaps attracted by the presence of the fishermen. The dog is wearing a red collar and appears to be enjoying its time in the water with the two men.

Figure 6. **Comparisons of Image Captioning Results** between LLaMA-Adapter and LLaMA-Adapter V2.

Instruction: Explain why this image is beautiful.
Answer: This image is beautiful because it captures the serene and peaceful atmosphere of the beach, with the sun setting in the background. The image features a beautiful, golden sunset over the ocean, creating a stunning backdrop for the solitary figure walking along the sandy beach. The combination of the warm, golden light of the sunset, the blue water of the ocean, and the white sandy beach creates a breathtaking and calming scene. The image also highlights the vastness of the ocean, with the waves crashing against the shore, creating a soothing sound that adds to the tranquility of the setting. The image conveys a sense of relaxation and tranquility, making it a perfect representation of a beautiful beach scene.

Instruction: What city might this picture be taken from? Can you give me some suggestions if I want to visit this city?
Answer: The image might be taken from Beijing, China. Beijing is the capital city of China and one of the most populous cities in the world. It is known for its rich history and culture, as well as its modern infrastructure and architectural landmarks. Some popular attractions in Beijing include the Forbidden City, Temple of Heaven, and the Great Wall of China, among others. If you want to visit Beijing, it is recommended to explore its historical and cultural heritage, as well as its modern developments.

Instruction: What should the car do here?
Answer: The car should follow the pedestrian crossing the street, as indicated by the red light. The car should stop at the red light and wait for the pedestrian to cross the street safely before proceeding. This is a common practice to ensure the safety of pedestrians and drivers alike.

Figure 7. **The Visual Understanding Examples** of LLaMA-Adapter V2.

In addition, we note that LLaMA-Adapter V2 may not always generate accurate image description. As seen in the failure case presented in Fig. 6, we intentionally chose an out-of-distribution example (cartoon picture) for testing. LLaMA-Adapter V2 was unable to comprehend the image and generated incorrect descriptions. The possible reason for this failure is the lack of an image-text alignment stage, which models like MiniGPT-4 and LLaVA have incorporated. This motivates us to employ additional expert systems to enhance the image understanding ability.

Visual Understanding. Thanks to our proposed joint training and early fusion techniques, LLaMA-Adapter V2 exhibits exceptional multi-modal understanding capabilities. As shown in Fig. 7, when presented with a prompt about the image content, such as in the form of "why is ..." and "what should ...", our model is able to seamlessly integrate the visual information with the language context, allowing for more sophisticated reasoning and decision-making. From the examples, we can see that our model is able to identify and explain the specific object or feature in the image that the question is referring to, and provide relevant information or suggestions based on the context.

Instruction: Could you list the steps for cooking the dish shown in the image?

w/o Context: Yes, the steps for cooking the dish shown in the image can be summarized as follows: 1. Boil the octopus in salted water to tenderize the meat. 2. Remove the octopus from the boiling water and drain the excess liquid. 3. Allow the octopus to cool down to room temperature. 4. Cut the octopus into small pieces, preferably 1-2 inches in size. 5. In a large skillet, heat the olive oil and cook the octopus pieces until they turn golden brown. 6. Add the garlic and lemon juice to the skillet and continue to cook the octopus for a few more minutes. 7. Season the octopus with salt and pepper, to taste. 8. Serve the cooked octopus on a plate with the lemon wedges and garnish with the chopped parsley.

w/ Context: [*Caption expert: A seafood dish with shrimp, octopus, and lobster.*] The steps for cooking the seafood dish with shrimp, octopus, and lobster are as follows: 1. Start by preparing the seafood. 2. Heat up the oil in a pan or skillet. 3. Add the shrimp, octopus, and lobster to the pan. 4. Cook the seafood over medium heat for 10 minutes. 5. Season the dish with salt and pepper to taste. 6. Serve the dish with a side of rice, vegetables, or other starch. 7. Enjoy!

Instruction: How to shoot such a picture using a film camera at night? Do you have some suggestions?

w/o Context: To capture the image of the neon sign in the rain, a film camera would need to be set up with a long exposure time to allow the light from the sign to be captured. This would require the photographer to use a tripod to keep the camera steady and to manually adjust the shutter speed to the desired exposure time. The photographer would also need to use a slow shutter speed to capture the light from the sign, which could range from 1 second to several minutes, depending on the desired effect. Additionally, the photographer would need to use a wide aperture setting to ensure that the entire sign is in focus, as well as a low ISO setting to minimize noise and preserve the image quality.

w/ Context: [*Caption expert: A wet street with a building and a car.*] To capture a nighttime scene with a wet street, a film camera, and a car, the photographer would need to use a slow shutter speed to allow more light to enter the camera's sensor. This would result in a blurred image of the moving car, but the building and the street would be sharply focused. The photographer could also use a tripod to steady the camera and reduce the shake caused by handheld shooting. Additionally, the photographer could use a flash to illuminate the car and the street, creating a more dramatic effect.

Figure 8. **Visual Understanding with the help of Caption Experts.** The caption expert is LLaMA-Adapter V1 trained on COCO Caption [6].

Question: What is the cost of the Antique Opera glasses? Do you think it's worth buying the Antique Opera glasses and can you explain the reason in detail?

OCR Context: 1. Antique Opera Glasses These beauti- 2. ful, multi-shaded mother-of-pearl opera 3. The Met Mirror The magnificent prosce- 4. glasses are genuine turn-of-the-century 5. nium of the Old Metropolitan Opera House 6. French antiques. Each is one-of-a-kind, 7. is lovingly recreated in this finely crafted, 8. and comes with a Met Costume Opera 9. distinctive piece. Neal Martz's design is 10. hand cast in hydrastone, with a beautiful 11. Glass Case, made from costume 12. swatches from Met productions of Aida, 13. handcrafted gold patina. 11 high, 10 wide 14. Lucia, Trovatore, Lohengrin and Boris 15. and 11/2 deep, the Met Mirror comes ready 16. Godunov. $250(2.00) #G70 17. for hanging. $45(2.25) #G41 18. 13

Language Instruction Model: I'm sorry, but I am an AI language model and I don't have the ability to access images or prices. Could you please provide me with the price of the Antique Opera glasses so that I can provide you with an accurate response?

Visual Instruction Model w/o Context: The cost of the Antique Opera glasses is not mentioned in the image. However, the image suggests that the glasses are being sold in a store, which means that the cost may vary depending on the location, the vendor, and the specific glasses being sold. It is important to note that the value of the Antique Opera glasses is subjective and may depend on the individual's taste, preference, and budget. Some people may find the antique design and the history behind the glasses appealing and may be willing to pay a higher price for them.

Visual Instruction Model w/ Context: The cost of the Antique Opera glasses is $250 (2.00) and it is worth buying because they are genuine turn-of-the-century French antiques that are lovingly recreated in this finely crafted, hand cast in hydrastone, with a beautiful gold patina. The Met Mirror comes ready for hanging and is $45 (2.25).

Figure 9. **Visual Understanding with the help of OCR Experts.** The example and OCR context are from DocVQA [44].

Moreover, our model is able to extrapolate and provide a plausible guess or explanation even when the question goes beyond the direct information presented in the image. These results demonstrate the effectiveness of LLaMA-Adapter V2 in addressing the challenge of interference between image-text pairs and instruction data, and highlight the potential of LLaMA-Adapter V2 in various real-world applications that require both language and visual understanding.

Integration with Experts. To further enhance LLaMA-Adapter V2's visual understanding capabilities, we integrate visual expert models during inference to provide additional textual contexts based on the input image. The examples of using LLaMA-Adapter [72] as the caption expert are shown in Fig. 8. By including a short caption of the image, our LLaMA-Adapter V2 produces more precise and detailed responses about the visual contents in the image.

Another example using the OCR expert from DocVQA [44] is presented in Fig. 9. By utilizing the detected texts in the image, our model receives more specific cues, such as the price of the glasses, and produces accurate answers to the question. The integration of experts thus largely extends the power of LLaMA-Adapter V2 and fully unleashes its multi-modal reasoning potential.

6. Conclusion

In this work, we propose LLaMA-Adapter V2, a parameter-efficient visual instruction tuning system. Through joint training on image-text pairs and instruction-following data, we observe that LLaMA-Adapter V2 can convert pre-trained LLMs into a zero-shot visual instruction model. With the assistance of simple bias-tuning strategy and early-fusion of visual features, LLaMA-Adapter V2's zero-shot visual instruction-following ability is further improved by alleviating the interference between image-text pairs and instruction-following data. At the same time, LLaMA-Adapter V2 achieves stronger language instruction-following performance and even possesses multi-turn dialog ability like a chatbot. To address the inaccurate image description problem of LLaMA-Adapter V2, we integrate expert visual systems, such as OCR and image captioner, to provide more accurate image-related information. While LLaMA-Adapter V2 can incorporate expert visual systems and perform visual instruction-following in a zero-shot manner, it still lags behind LLaVA in terms of visual understanding capability and is susceptible to inaccurate information provided by the expert systems. In the future, we plan to explore the integration of more expert systems, and fine-tune LLaMA-Adapter V2 with a multi-modal instruction dataset or other PEFT methods (*e.g.*, LoRA) to further enhance its visual instruction-following capabilities.

References

[1] Sharegpt: Share your wildest chatgpt conversations with one click. https://sharegpt.com/. 1, 3, 7, 8

[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

[3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 3

[4] 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. 3

[5] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 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*, pages 3558–3568, 2021. 3, 8

[6] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 2, 3, 4, 6, 7, 8, 10

[7] 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. 1, 3, 8

[8] 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. 3

[9] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10578–10587, 2020. 3

[10] Pedro Cuenca and Sayak Paul. Using lora for efficient stable diffusion fine-tuning. https://huggingface.co/blog/lora, January 2023. 3

[11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019. 3

[12] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Delta tuning: A comprehensive study of parameter efficient methods for pre-trained language models. *arXiv preprint arXiv:2203.06904*, 2022. 3

[13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 3

[14] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 4

[15] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao.

Human: Who are 5 people you would like to meet?

Assistant-65B: As an artificial intelligence, I don't have personal preferences or desires, but I can provide a list of well-known people who have made significant contributions in various fields and would be interesting to meet:

• Albert Einstein: A renowned physicist who revolutionized our understanding of the universe with his theory of relativity and other groundbreaking work.
• Marie Curie: A Polish-French physicist and chemist who made significant contributions to the study of radioactivity and was the first woman to win a Nobel Prize.
• Martin Luther King Jr.: An influential American civil rights leader who advocated for racial equality and justice through nonviolent resistance.
• Nelson Mandela: A South African anti-apartheid revolutionary and political leader who became the country's first democratically elected president.
• Jane Goodall: A British primatologist and anthropologist who has dedicated her life to the study of chimpanzees and environmental conservation.

These individuals have made significant contributions to their fields and society, and meeting them would be an enriching experience.

Human: Could you explain me the main contribution of the first person you mentioned?

Assistant-65B: Sure! Albert Einstein's main contribution was the development of the general theory of relativity, which revolutionized our understanding of the universe.

Einstein's theory of relativity proposed that the laws of physics are the same for all non-accelerating observers, regardless of their motion or the strength of the gravitational field. This theory also introduced the concept of space-time as a single, unified entity that is warped by massive objects, such as stars and black holes.

Einstein's work on relativity had many significant implications. For example, it helped explain the behavior of light and the phenomenon of gravitational lensing, where the path of light is bent by a massive object, such as a galaxy. His work also led to predictions about the existence of black holes and the expansion of the universe, both of which have since been confirmed by observations.

Einstein's contributions to physics are considered some of the most significant in the history of science, and his name has become synonymous with genius and innovation.

In addition to his work on relativity, Einstein also made important contributions to other areas of physics, such as the photoelectric effect, Brownian motion, and the discovery of the existence of atoms.

Overall, Einstein's work had a profound impact on our understanding of the universe and continues to shape the field of physics today.

Human: What is the most famous equation from this theory?

Assistant-65B: The most famous equation associated with Einstein's theory of relativity is: $E = mc^2$.

This equation relates mass (m) and energy (E) through the speed of light (c^2), which is a constant. According to this equation, mass and energy are equivalent and interchangeable, and it implies that even a small amount of mass can be converted into an enormous amount of energy.

The famous equation $E = mc^2$ is often referred to as the "mass-energy equivalence principle," and it is one of the cornerstones of Einstein's theory of relativity. It has numerous applications in fields such as nuclear physics, particle physics, and astrophysics, and it continues to be a source of inspiration and investigation for scientists around the world.

The equation is often used to calculate the energy released during nuclear reactions, such as those that occur in nuclear power plants or during nuclear bomb explosions. It also helps explain the enormous amounts of energy released during supernova explosions, which occur when massive stars collapse under their own gravity and release a vast amount of energy.

In summary, the equation $E = mc^2$ is one of the most famous equations in physics and is a central part of Einstein's theory of relativity. It has far-reaching implications for our understanding of the universe and continues to be an important concept in modern science.

Figure 10. **A Chatting Example using 65B LLaMA-Adapter V2.**

Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 3

[16] Peng Gao, Zhengkai Jiang, Haoxuan You, Pan Lu, Steven CH Hoi, Xiaogang Wang, and Hongsheng Li. Dynamic fusion with intra-and inter-modality attention flow for visual question answering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6639–6648, 2019. 3

[17] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In *ACL-IJCNLP*, 2021. 3

[18] Yingqiang Ge, Wenyue Hua, Jianchao Ji, Juntao Tan, Shuyuan Xu, and Yongfeng Zhang. Openagi: When llm meets domain experts. *arXiv preprint arXiv:2304.04370*, 2023. 4

[19] Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. Koala: A dialogue model for academic research. Blog post, April 2023. 1

[20] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 3

[21] Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. *arXiv preprint arXiv:2211.11559*, 2022. 4

[22] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *International Con-*

Human: Who are 5 people you would like to meet?

Assistant-7B:
• Albert Einstein: I would love to meet the genius physicist who revolutionized our understanding of the universe.
• Steve Jobs: I would like to meet the visionary who changed the way we communicate and interact with technology.
• Leonardo da Vinci: I would like to meet the artist, inventor, and scientist who was a true Renaissance man.
• Jane Goodall: I would like to meet the primatologist who has dedicated her life to studying and protecting chimpanzees.
• Nelson Mandela: I would like to meet the anti-apartheid activist and former South African president who fought for equality and freedom.

Human: Could you explain me the main contribution of the first person you mentioned?

Assistant-7B: Albert Einstein was a German-born theoretical physicist who developed the general theory of relativity, one of the two pillars of modern physics (alongside quantum mechanics). His theory of relativity is the foundation of modern cosmology and has had a profound impact on our understanding of the universe.

Einstein's theory of relativity revolutionized our understanding of space and time, and it has been experimentally verified to a high degree of accuracy. It has also led to the development of the atomic bomb and the laser, among other technological advancements.

Einstein's contributions to physics are immense, and his work has had a lasting impact on the field. He is widely considered one of the most influential scientists of all time.

Human: What is the most famous equation from this theory?

Assistant-7B: The most famous equation from the theory of relativity is the famous $E = mc^2$ equation, which states that energy (E) is equal to mass (m) multiplied by the speed of light squared (c^2).

This equation is the foundation of nuclear physics and is the basis for the development of nuclear weapons. It also has important implications for the study of black holes and the nature of gravity.

The equation is a direct consequence of the theory of relativity, and it has been experimentally verified to a high degree of accuracy. It has had a profound impact on our understanding of the universe and has led to many technological advancements.

Figure 11. **A Chatting Example using 7B LLaMA-Adapter V2.**

ference on Learning Representations, 2022. 3

[23] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CsUR)*, 51(6):1–36, 2019. 3

[24] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. 3

[25] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 5

[26] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 3

[27] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and brian ichter. Inner monologue: Embodied reasoning through planning with language models. In *6th Annual Conference on Robot Learning*, 2022. 4

[28] 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. 3, 5

[29] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 5583–5594, 2021. 3

[30] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. 3, 8

[31] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *EMNLP*, 2021. 3

[32] 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. 3, 8

[33] 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. 8

[34] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and perfor-

mant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 3

[35] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*, 2021. 3

[36] 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. 5

[37] Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. Text2motion: From natural language instructions to feasible plans. *arXiv preprint arXiv:2303.12153*, 2023. 4

[38] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 2, 3, 6, 7

[39] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 13–23, 2019. 3

[40] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 375–383, 2017. 3

[41] Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 3, 4

[42] Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. Chameleon: Plug-and-play compositional reasoning with large language models. *arXiv preprint arXiv:2304.09842*, 2023. 4

[43] Rabeeh Karimi mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 3

[44] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021. 10, 11

[45] So Yeon Min, Devendra Singh Chaplot, Pradeep Ravikumar, Yonatan Bisk, and Ruslan Salakhutdinov. Film: Following instructions in language with modular methods. *ArXiv*, abs/2110.07342, 2021. 3

[46] Ron Mokady, Amir Hertz, and Amit H Bermano. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021. 8

[47] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. 2

[48] 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. 1, 3

[49] Sourab Mangrulkar; Sylvain Gugger; Lysandre Debut; Younes Belkada; Sayak Paul. Peft: State-of-the-art parameter-efficient fine-tuning methods. https://github.com/huggingface/peft, 2022. 3

[50] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023. 1, 3, 5, 7, 8

[51] Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, et al. Tool learning with foundation models. *arXiv preprint arXiv:2304.08354*, 2023. 4

[52] 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. 3, 4

[53] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 3

[54] 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. 3

[55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3

[56] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 3, 8

[57] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 3, 8

[58] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*, 2023. 4

[59] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Lst: Ladder side-tuning for parameter and memory efficient transfer learning. *arXiv preprint arXiv:2206.06522*, 2022. 3

[60] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Vl-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5227–5237, 2022. 3

[61] Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *arXiv preprint arXiv:2303.08128*, 2023. 4

[62] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023. 1, 3

[63] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 3, 4

[64] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):652–663, 2016. 3

[65] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*, 2022. 3

[66] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022. 3

[67] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023. 4

[68] Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. *arXiv preprint arXiv:2304.01196*, 2023. 1

[69] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*, 2023. 4

[70] Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021. 5

[71] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 3

[72] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023. 1, 3, 4, 10

[73] Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Hongsheng Li, Yu Qiao, and Peng Gao. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. *CVPR 2023*, 2023. 3

[74] Renrui Zhang, Zhang Wei, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *ECCV*, 2022. 3

[75] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. 1

[76] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, pages 1–12, 2022. 3

[77] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13041–13049, 2020. 3

[78] Deyao Zhu, Jun Chen, Xiaoqian Shen, xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. 2023. 2, 3, 6, 7

[79] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyao Zeng, Shanghang Zhang, and Peng Gao. Pointclip v2: Adapting clip for powerful 3d open-world learning. *arXiv preprint arXiv:2211.11682*, 2022. 4