

# Visual In-Context Learning for Large Vision-Language Models

Yucheng Zhou<sup>1</sup>, Xiang Li<sup>2</sup>, Qianning Wang<sup>3</sup>, Jianbing Shen<sup>1\*</sup>

<sup>1</sup> SKL-IOTSC, CIS, University of Macau

<sup>2</sup> Tianjin University, <sup>3</sup> Nanjing Audit University

yucheng.zhou@connect.um.edu.mo, jianbingshen@um.edu.mo

## Abstract

In Large Visual Language Models (LVLMs), the efficacy of In-Context Learning (ICL) remains limited by challenges in cross-modal interactions and representation disparities. To overcome these challenges, we introduce a novel Visual In-Context Learning (VICL) method comprising Visual Demonstration Retrieval, Intent-Oriented Image Summarization, and Intent-Oriented Demonstration Composition. Our approach retrieves images via “Retrieval & Rerank” paradigm, summarises images with task intent and task-specific visual parsing, and composes language-based demonstrations that reduce token count and alleviate cross-modal interaction problem. Experimental evaluations on five visual reasoning datasets demonstrate the effectiveness of our method. Moreover, our extensive experiments leverage information flow analysis to elucidate the effectiveness of our method, and investigate the impact of length and position of demonstrations for LVLM. The use of in-context unlearning further shows promise in resetting specific model knowledge without retraining.

## 1 Introduction

Large Language Models (LLMs) exhibit impressive reasoning abilities across various natural language tasks (Brown et al., 2020; Touvron et al., 2023). Researchers are actively investigating the extension of LLM capabilities to address challenges in the visual domain by integrating LLMs with vision models (Zhu et al., 2023; Bai et al., 2023). This endeavor has given rise to the development of Large Visual Language Models (LVLMs). LVLMs are designed to seamlessly fuse information from both images and text, enabling them to tackle more intricate tasks that demand a profound

\* Corresponding author. This work was supported in part by the FDCT grants 0102/2023/RIA2, 0154/2022/A3, and 001/2024/SKL, the MYRG-CRG2022-00013-IOTSC-ICI grant and the SRG2022-00023-IOTSC grant.

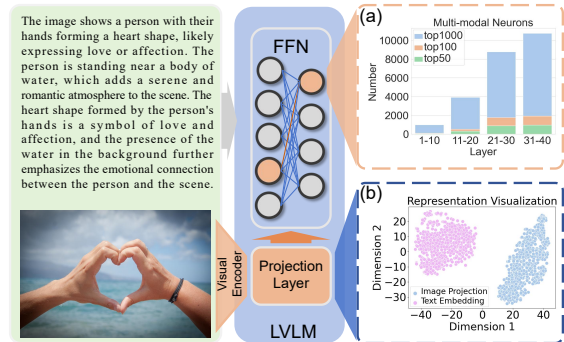


Figure 1: Illustrating Cross-Modal Challenges in LVLMs: (a) Distribution of multi-modal interaction neurons. (b) The distinct spaces are occupied by visual features and text embeddings before passing into LVLMs.

comprehension of both modalities (OpenAi, 2023; Zhu et al., 2023).

The LLM has a remarkable capability, known as In-Context Learning (ICL), which involves providing LLMs with a limited amount of labeled data as demonstrations to improve their reasoning ability (Brown et al., 2020; Dong et al., 2023; Zhang et al., 2023d). This approach can significantly enhance the performance of LLMs in various NLP tasks, such as translation (Garcia et al., 2023; Moslem et al., 2023), sentiment classification (Qin et al., 2023), and question-answering (Qin et al., 2023; Li et al., 2023). In the ICL, LLMs can flexibly adjust their behavior based on the provided context, allowing them to understand and perform tasks with few labeled data. The success of ICL has motivated research into extending the ICL capabilities to LVLMs. However, some studies (Chen et al., 2023; Peng et al., 2023) find that while LVLMs have ICL capabilities, they are not as pronounced as those observed in LLMs. Two factors lead to this difference: (1) As shown in Figure 1(a), as observed in previous research (Pan et al., 2023; Schwettmann et al., 2023), visual-language interactions occur at deeper layers in LVLMs, which

highlights the difficulty of cross-modal interactions. In ICL, label words aggregate information in shallow layers and subsequently distribute it in deeper layers (Wang et al., 2023). Consequently, the issue of cross-modal interactions significantly impacts the ICL capabilities of LVLMs. (2) From Figure 1(b), our analysis reveals that visual features and LLM embeddings occupy distinct spaces in LVLMs. This observation underscores the inherent cross-modal gap present in LVLMs. While many works (Min et al., 2022; Lu et al., 2023) are dedicated to enhancing the ICL capabilities of LLMs, the challenges faced by LVLMs in this regard differ substantially. This discrepancy arises from the difficulty in cross-modal interactions and inherent disparities in representation spaces within LVLMs, which impose limitations on their ICL performance.

In this study, we present a novel Visual In-Context Learning (VICL) method to enhance the ICL capability of LVLMs. VICL comprises Visual Demonstration Retrieval, Intent-Oriented Image Summarization, and Intent-Oriented Demonstration Composition. For Visual Demonstration Retrieval, we employ a pre-trained image encoder as a retriever to search for relevant candidate images for the provided image. Subsequently, we rerank the retrieved candidates using textual descriptions of the provided image. Moreover, LVLMs perform Intent-Oriented Image Summarization, automatically extracting image summary with task intent and task-specific visual parsing from image-label pairs. In addition, Intent-Oriented Demonstration Composition uses language cues to create demonstrations that enhance ICL of LVLMs, replacing images with image summary in demonstrations. Our method not only boosts in-context learning but also introduces in-context unlearning, allowing models to discard or reset specific knowledge through demonstration. The substitution of images with visual summaries significantly reduces the token count, enabling the concatenation of more demonstrations without encountering token limitations. In comparison to conventional visual-language interactions in standard visual ICL approaches, our method solely relies on language interactions to facilitate effective demonstration understanding.

Our experiments across five image reasoning datasets evaluate our method’s effectiveness, comparing LVLM performance using our approach against a baseline method. Moreover, we employed information flow for interpretative analysis, to ver-

ify the effectiveness of our method. Furthermore, we examined the influence of demonstrations and their sequence length on LVLM’s ICL capability. We investigate the importance of intent-oriented image summaries and the impact of demonstrations order. In addition, we explore the application of in-context unlearning, demonstrating its feasibility for unlearning scenarios without additional training.

## 2 Related Work

### 2.1 Large Vision-Language Models

Large Vision-Language Models (LVLMs) are designed to comprehend and generate content across vision and language modalities, allowing them to perform tasks that involve understanding and generating not only text, but also information in visual forms (Yin et al., 2023).

LVLM can be broadly categorized into two main types, according to the output modalities: visual understanding and visual generation. Visual understanding models are capable of comprehending visual modality information provided to them and generating textual responses, enabling them to accomplish tasks such as image captioning, image question answering (Zhu et al., 2023; OpenAi, 2023; Alayrac et al., 2022), video understanding (Cho et al., 2021), video captioning (Bansal et al., 2023), etc. The typical structure of these models involves integrating the visual encoders based on transformer architecture (like Clip(Radford et al., 2021)) into a large language model.

On the other hand, visual generation models are equipped with visual decoders, enabling the decoding of feature vectors into images or videos. They have shown the ability to create high-quality outputs in generative tasks, such as generating text, images, and videos (Marcus et al., 2022; Saharia et al., 2022; Zhang et al., 2023b).

### 2.2 In-Context Learning

In-Context Learning exemplifies a paradigm where model weights require no optimization; rather, adjusting the model input (adding context) leads to correct output generation (Dong et al., 2023). An in-context learning prompt typically consists of two components: demonstration and new query. Demonstrations comprise multiple question-answer pairs, each presenting a complete question and its corresponding answer, while new queries involve inquiries posed to the model. Due to the emergent ability in large language models (Lu et al.,

2023), they can to some extent reference demonstrations to answer new questions (Min et al., 2022). With the advantage of not necessitating fine-tuning of model parameters, in-context learning has become a popular paradigm for applying large language models.

The inherent black-box nature of deep neural models renders the reasons behind the efficacy of in-context learning even more challenging to elucidate (Mao et al., 2024; Hahn and Goyal, 2023; von Oswald et al., 2023; Xie et al., 2022). One of the most widely accepted theoretical explanations at present is that when the pre-training text has long-range coherence, if the demonstrations in the prompt share potential concepts, in-context learning ability will emerge (Xie et al., 2022).

### 2.3 Machine Unlearning

In certain scenarios involving security or privacy concerns, it becomes imperative to selectively erase specific knowledge acquired by machine learning models (Goldsteen et al., 2022). Unlike conventional databases where information is explicitly stored in tabular forms, the entirety of a model’s acquired knowledge is implicitly embedded within its parameters. Consequently, the challenge arises of accurately expunging unwanted information without necessitating a complete retraining of the model, thereby minimizing interference with other retained knowledge. This intricate problem is addressed by a collective set of techniques known as machine unlearning (Bourtole et al., 2021; Nguyen et al., 2022; Zhang et al., 2023a; Koch and Soll, 2023).

Among these methodologies, the in-context unlearning (Pawelczyk et al., 2023) emerges as a promising solution, which specifically applies the in-context learning paradigm without updating any model parameters, making it suitable for large language models. The framework leverages a combination of incorrectly and correctly labeled examples from training datasets. By analyzing and understanding the nuances within these discrepancies, a unique prompt is constructed for each instance. This tailored prompt aims to highlight specific challenges posed by the labeling discrepancies, encouraging the model to refine its predictions during inference.

## 3 Background

Generative language models could self-supervised learn knowledge from pre-training corpora (Radford et al., 2018, 2019; Raffel et al., 2020). As the scale of model parameters and pre-training corpora expands, researchers have observed the emergent ability in large language models, enabling them to provide accurate answers merely by adjusting the input prompt without fine-tuning model parameters (Wei et al., 2022a,b; Fu et al., 2023).

The large language model is abstracted as a function denoted as  $\text{LLM}(\cdot)$ . Given an input prompt, denoted as  $p$ , it generates the corresponding output, denoted as  $o$ . In the most common case of question-answering, the prompt  $p$  is exactly the question  $q$  raised by the user. And the output  $o$  is expected to be the answer  $a$  to the question  $q$ .

$$\text{LLM}(o|p) = \text{LLM}(a|q) \quad (1)$$

However, under the in-context learning paradigm, the input prompt  $p_{\text{icl}}$  is carefully designed:

$$\begin{aligned} \text{LLM}(a|p_{\text{icl}}) &= \text{LLM}(a|\mathbb{D} \oplus q) \\ \mathbb{D} &= \mathbb{T}(\oplus_j^n (\hat{q}_j \oplus \hat{a}_j)), \end{aligned} \quad (2)$$

where the input prompt (denoted as  $p_{\text{icl}}$ ) is formed by concatenating (denoted as  $\oplus$ ) demonstrations  $\mathbb{D}$  and query  $q$ , the current question raised by the user.  $\mathbb{D}$  is composed of number  $n$  sets of complete questions ( $\hat{q}_j$ ) and answers ( $\hat{a}_j$ ), spliced together through a fixed template  $\mathbb{T}(\cdot)$ .

Considering the established efficacy of in-context learning on large language models, it was intuitive to extend the in-context learning approach to large visual-language models upon its emergence (Sun et al., 2023; Liu et al., 2023c; Xu et al., 2023; Zhang et al., 2023c). In the context of a visual question answering task, the formulation of in-context learning on large visual-language model can be delineated as follows:

$$\begin{aligned} \text{LVLM}(a|p_{\text{icl}}) &= \text{LVLM}(a|\mathbb{D} \oplus i \oplus q) \\ \mathbb{D} &= \mathbb{T}(\oplus_j^n (\hat{q}_j \oplus \hat{i}_j \oplus \hat{a}_j)), \end{aligned} \quad (3)$$

where the visual input prompt, denoted as  $p_{\text{icl}}$ , is created by concatenating demonstrations  $\mathbb{D}$ , query image  $i$ , and query text  $q$ .  $\mathbb{D}$  consists of  $n$  sets of complete questions ( $\hat{q}_j$ ), images ( $\hat{i}_j$ ), and answers ( $\hat{a}_j$ ), which are combined by a fixed template  $\mathbb{T}(\cdot)$ .

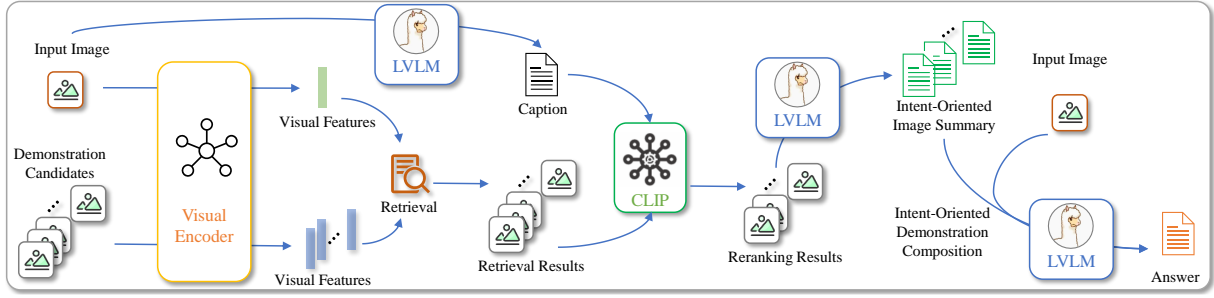


Figure 2: Overview of our Visual In-Context Learning (VICL) method. The Visual Encoder is used to encode images for retrieval, and CLIP is used for cross-modal reranking of images and caption; LVLM is used to generate caption for input images, generate intent-oriented image summaries, and predict answer based on the composed prompt.

## 4 Visual In-Context Learning

In this section, we elaborate on our approach VICL, which comprises three core components: Visual Demonstration Retrieval, Intent-Oriented Image Summarization, and Intent-Oriented Demonstration Composition. The pipeline of VICL is shown in Figure 2.

### 4.1 Visual Demonstration Retrieval

The Visual Demonstration Retrieval (VDR) component is the first step in our VICL method, designed to identify suitable samples for ICL as demonstrations. The goal of VDR is to discern and choose visual demonstrations that bear the utmost relevance to the current task. This process capitalizes on both the visual features of images and their accompanying textual descriptions. Following de facto “retrieval & rerank” paradigm (Zhou et al., 2023b), VDR comprises these two phases.

**Visual Retrieval.** Given an image  $I$ , our goal is to find a set of candidate demonstration images  $\mathcal{D} = \{I_1, I_2, \dots, I_n\}$  that are relevant to  $I$ . This is achieved by employing a pre-trained image encoder, Vision-Enc, that maps images into a high-dimensional feature space, i.e.,

$$\mathbf{f} = \text{Vision-Enc}(I), \quad (4)$$

$$\mathbf{f}_i = \text{Vision-Enc}(I_i), i \in \{1, 2, \dots, n\}, \quad (5)$$

where  $\mathbf{f}$  is the feature vector representing the embedding of the image  $I$ , and  $\mathbf{f}_i$  represents the embedding of the image  $I_i$  in the dataset. Vision-Enc denotes ViT (Dosovitskiy et al., 2021). The retrieval operation is defined as:

$$\mathcal{D}_q = \text{Retrieval}(I, \mathcal{D} \mid \text{Vision-Enc}), \quad (6)$$

The Retrieval function aims to select the top- $n$  images from  $\mathcal{D}$  whose embeddings are most similar

to the embedding of  $I$ , based on a similarity metric  $\text{Sim}$ . This approach ensures that the selected images  $\mathcal{D}_q$  are those that are most relevant to the query image in terms of visual features encoded within the high-dimensional feature space. The similarity metric is defined as:

$$\text{Sim}(\mathbf{f}, \mathbf{f}_i) = \frac{E(\mathbf{f}) \cdot E(\mathbf{f}_i)}{\|E(\mathbf{f})\| \|E(\mathbf{f}_i)\|}, \quad (7)$$

where  $(\cdot)$  is the dot product between two vectors.

**Cross-Modal Reranking.** After the visual retrieval, we obtain a set of candidate images  $\mathcal{D}_q$ . However, to ensure that the selected demonstrations are not only visually similar but also semantically relevant to  $I$ , we employ a reranking step using textual descriptions. We use a large vision language model, LVLM, to generate an image description  $T_q$  for the image  $I$ . The reranking process adjusts the initial rankings based on the semantic similarity between  $T_q$  and  $I_i$ ,  $I_i \in \mathcal{D}_q$ , which is computed by a pre-trained image-text model VL-Enc:

$$\hat{\mathcal{D}}_q = \text{Rerank}(\mathcal{D}_q, T_q \mid \text{VL-Enc}), \quad (8)$$

where  $\hat{\mathcal{D}}_q$  denotes the reranked set of demonstration images. VL-Enc refers to CLIP (Radford et al., 2021). This dual-stage approach allows us to harness the complementary strengths of visual and textual modalities, ensuring that the chosen visual demonstrations are not only visually pertinent but also contextually aligned with the query image’s task-specific requirements.

### 4.2 Intent-Oriented Image Summarization

Intent-Oriented Image Summarization (IOIS) aims to simplify LVLM’s ICL problem by generating a visual content summary from a task intent perspective. This summarization process focuses on

exploring the relationship of a given reference image, question and answer triplet, and generates an image summary encapsulating the task intent and the task-specific visual parsing.

**ICL in LLM.** For LLMs, given the reference question-answering pair  $\{\hat{Q}, \hat{A}\}$  and the input question  $Q$ , the ICL problem can be formalized as:

$$P(A | Q) = P(A | Q, T) \times P(T | \{\hat{Q}, \hat{A}\}), \quad (9)$$

where  $A$  is the predicted answer;  $T$  is task intent; the model needs to infer the task intent from given reference question-answer pairs to accurately respond to a new question.

**ICL in LVLM.** In LVLMs, the ICL problem extends to:

$$\begin{aligned} P(A | Q, I) &= P(A | I, Q, T, P) \\ &\times P(T | \{\hat{I}, \hat{Q}, \hat{A}\}) \\ &\times P(P | \{\hat{I}, \hat{Q}, \hat{A}\}), \end{aligned} \quad (10)$$

where the LVLM must first deduce the task intent  $T$  and image parsing strategy  $P$  from the reference image, question, and answer triplet  $\{\hat{I}, \hat{Q}, \hat{A}\}$  before analyzing the content of the given image based on these insights.

**VIOL.** Our IOIS method significantly simplifies this process by pre-generating a visual content summary that embodies both the task intent and the image parsing approach. This summary is produced by concatenating a carefully designed prompt with the given reference image, question, and answer, and then inputting this into the LVLM. The output is a summary that not only describes the image but also encapsulates task intent and task-specific visual parsing. Our approach can be represented as:

$$P(A|Q,I) = P(A|I,Q,S) \times P(S|\{\hat{I}, \hat{Q}, \hat{A}\}) \quad (11)$$

where  $S$  denotes the set of intent-oriented visual summarization for all reference images. This formulation demonstrates how our method modifies the LVLM’s ICL challenge by replacing the direct analysis of images with the interpretation of summarizations that are pre-aligned with the task’s intent and preferred image parsing methodology.

**Implementation of IOIS.** The specific procedure for generating the Intent-Oriented Image Summarization involves constructing a prompt that integrates the demonstration image with its corresponding label, underpinned by the task’s intent

and image parsing preferences. This prompt is then input into the LVLM to produce the summarization. The process can be mathematically function as:

$$\begin{aligned} S_i &= \text{LVLM}(\text{Prompt}(\hat{I}_i, \hat{Q}_i, \hat{A}_i), \\ &\text{where } S = \{S_1, S_2, \dots, S_l\}, \\ &\hat{I}_i, \hat{Q}_i, \hat{A}_i \in \{\hat{I}, \hat{Q}, \hat{A}\}, \end{aligned} \quad (12)$$

where  $l$  is the number of reference examples. Prompt is a function that formulates the input for the LVLM, encapsulating the demonstration image and label along with explicit cues about the task’s intent and the approach to image parsing.

This approach ensures that the LVLM’s ICL process is primed with a context that significantly lowers the cognitive load associated with cross-modal mapping, allowing the model to focus on reasoning within a linguistic framework that is inherently more aligned with its training. This strategic simplification not only enhances the efficiency of the LVLM’s ICL capabilities but also reduces the complexity associated with direct image analysis.

### 4.3 Intent-Oriented Demonstration Composition

Intent-Oriented Demonstration Composition (IODC) aims to effectively integrate the generated image summaries  $S_i$  with corresponding questions  $Q_i$  and answers  $A_i$  into a coherent demonstration for the LVLM.

**Composition of Demonstrations.** The IODC process involves the assembly of each image summary  $S_i$  with its corresponding question  $Q_i$  and answer  $A_i$  into a single, unified demonstration. This is achieved through the concatenation of these elements in a manner that preserves the logical and semantic coherence necessary for effective ICL. Formally, the process can be represented as:

$$\bar{D}_i = \text{Concat}(S_i, Q_i, A_i), \quad (13)$$

where Concat is the concatenation operation, and  $\bar{D}_i$  represents the  $i$ -th demonstration composed of the image summary, question, and answer triplet. This operation is performed for each set of  $S_i$ ,  $Q_i$ , and  $A_i$ , resulting in a collection of demonstrations:

$$\bar{D} = \{\bar{D}_1, \bar{D}_2, \dots, \bar{D}_n\}, \quad (14)$$

where  $\bar{D}$  denotes the complete set of demonstrations ready for presentation to the LVLM.

**Enhancing ICL with Demonstration Composition.** By replacing original images  $I_i$  with intent-oriented visual summaries  $S_i$ , we significantly reduce the complexity and token count inherent in direct image processing. This reduction allows for the inclusion of a larger number of demonstrations within the LVLM’s token limit, thereby enriching the context available for ICL without overwhelming the model with excessive information, i.e.,

$$\mathbf{A} = \text{LVLM}(\bar{\mathcal{D}}, \mathbf{I}, \mathbf{Q}). \quad (15)$$

The IODC methodology facilitates a shift in LVLM processing from a reliance on direct visual inputs to an emphasis on linguistic representations of visual content, grounded in the task’s intent. This shift not only streamlines the ICL process by minimizing the token count but also aligns the demonstrations more closely with the LVLM’s linguistic processing capabilities.

#### 4.4 Information Flow Analysis

Following Wang et al. (2023), we analyze the information flow of VICL in the LVLM. To calculate saliency score of each element in attention matrix, we employ Taylor expansion (Michel et al., 2019):

$$I_l = \sum_h \left| \mathbf{A}_{h,l}^\top \frac{\partial \mathcal{L}(x)}{\partial \mathbf{A}_{h,l}} \right| \quad (16)$$

where  $h$  and  $l$  represent different attention heads and transformer layers, respectively.  $\mathcal{L}(x)$  is the loss function for the task. Furthermore, we define four different information flow significance scores as follows:

$$S_{wp} = \frac{\sum_{(i,j) \in C_{wp}} I_l(i,j)}{|C_{wp}|}, \quad C_{wp} = \{(p_k, j) : k \in [1, C], j < p_k\} \quad (17)$$

$$S_{pq} = \frac{\sum_{(i,j) \in C_{pq}} I_l(i,j)}{|C_{pq}|}, \quad C_{pq} = \{(q, p_k) : k \in [1, C]\} \quad (18)$$

$$S_{vq} = \frac{\sum_{(i,j) \in C_{vq}} I_l(i,j)}{|C_{vq}|}, \quad C_{vq} = \{(q, v)\} \quad (19)$$

$$S_{ww} = \frac{\sum_{(i,j) \in C_{ww}} I_l(i,j)}{|C_{ww}|}, \quad C_{ww} = \{(i, j) : j < i\} - C_{wp} - C_{pq} - C_{vq} \quad (20)$$

where  $p_k$ ,  $C$ ,  $q$  and  $v$  represent the label words, the total number of label words, the target position and the input image, respectively.  $S_{wp}$  denotes the




Figure 3: Information flow results on the EmoSet.

significance of information flow from the image summaries to label words;  $S_{pq}$  represents the significance of information flow from label words to the target position;  $S_{vq}$  signifies the significance of information flow from label words to the input image part;  $S_{ww}$  indicates the significance of the information flow amongst all words, excluding influences represented by  $S_{wp}$ ,  $S_{pq}$ , and  $S_{vq}$ .

As shown in Figure 3, the analysis reveals varying degrees of importance in information flow within VICL across different layers and attention heads. In the shallow layers,  $S_{wp}$  is highly important but diminishes as the layers progress. This suggests that image summaries are crucial in determining label words in the early stages, but their influence weakens with increasing model depth. Unlike  $S_{wp}$ , the importance of  $S_{pq}$  increases as the layers deepen. This implies that the influence of label words on determining the target position becomes more significant in the later stages of the model. Similar to  $S_{pq}$ ,  $S_{vq}$  shows some importance in the early stages but diminishes as the layers deepen. This suggests that the influence of label words on the input image weakens as the model progresses.  $S_{ww}$  remains relatively stable throughout the training process, showing no significant trend. Image summaries are crucial for label words aggregating information in the early stages, but the model increasingly emphasizes the relationship between label words and target position as it deepens.

## 5 Experiments

### 5.1 Experimental Settings

**Dataset.** In our experiments, we used five popular datasets related to image content reasoning: EmoSet (Yang et al., 2023), Emotion6 (Peng et al., 2015), UnBiasedEmo (Panda et al., 2018), CIFAR10 (Krizhevsky, 2009), and MNIST (Deng,

Model	Method	EmoSet	Emotion6	UnBiasedEmo	CIFAR10	MNIST
LLaVA-7B (Liu et al., 2023a)	Zero-Shot	0.23	0.31	0.31	0.75	0.85
	ICL	0.32	0.40	0.38	0.68	0.77
	VICL	0.69	0.70	0.76	0.84	0.88
MiniGPT-4 (Zhu et al., 2023)	Zero-Shot	0.21	0.27	0.27	0.61	0.84
	ICL	0.28	0.34	0.36	0.65	0.67
	VICL	0.61	0.61	0.74	0.76	0.85
Qwen-VL (Bai et al., 2023)	Zero-Shot	0.22	0.29	0.30	0.74	0.78
	ICL	0.31	0.39	0.37	0.66	0.75
	VICL	0.64	0.63	0.74	0.84	0.85
LLaVA-13B (Liu et al., 2023a)	Zero-Shot	0.32	0.34	0.38	0.79	0.87
	ICL	0.32	0.52	0.42	0.70	0.80
	VICL	<b>0.72</b>	<b>0.76</b>	<b>0.78</b>	<b>0.87</b>	<b>0.90</b>

Table 1: Comparison results of LVLM on five datasets.

2012). EmoSet, Emotion6, and UnBiasedEmo are datasets for image emotion classification, where the goal is to infer emotions based on the content of the images. EmoSet consists of 8 emotions, while Emotion6 and UnBiasedEmo consist of 6 emotions, each with multiple sub-categories. CIFAR10 and MNIST are classification datasets, where the task is to identify the object category in the images. For each of EmoSet, Emotion6, UnBiasedEmo, CIFAR10 and MNIST, we sampled 100 and 1000 samples as demonstration candidates and test sets, respectively. The metric for all test sets is accuracy.

**Prompts.** We have considered three distinct prompts for experimental comparison: (1) “Zero-shot” involves using the instruction and input image directly as the prompt without providing any demonstrations, formatted as “{instruction} {image}”. (2) “ICL” (In-Context Learning) includes first retrieving demonstrations using Visual Demonstration Retrieval, then integrating these demonstrations into the prompt, formatted as “{instruction} {demonstrations} {image}”. (3) “VICL” (Visual In-Context Learning) involves first retrieving demonstrations using Visual Demonstration Retrieval, then converting the images from the demonstrations into text using Intent-Oriented Image Summarization, concatenating this text back into the demonstrations, and finally integrating them into the prompt, formatted as “{instruction} {text demonstrations} {image}”. Detailed prompt specifications are provided in Appendix A.

**Large Vision-Language Models.** In this paper, we leverage four state-of-the-art LVLMs with various prompts to perform visual reasoning tasks. These models are LLaVA-7B (Liu et al., 2023a), MiniGPT-4 (Zhu et al., 2023), Qwen-VL (Bai et al.,

2023), and LLaVA-13B (Liu et al., 2023a). LLaVA-7B and LLaVA-13B are based on the visual instruction tuning (VIT) technique, which aligns a frozen visual encoder and a large language model (LLM) using one projection layer. MiniGPT-4 is an open-source chatbot that fine-tunes LLaMA/Vicuna on GPT-generated multi-modal instruction-following data. Qwen-VL is a versatile vision-language model that can perform understanding image and text. We compare and analyze the in-context learning performance and capabilities of these models.

## 5.2 In-Context Learning

We analyze the performance of LVLMs using Zero-Shot, ICL, and VICL approaches on five datasets. The results in Table 1 show the effectiveness of VICL across different LVLMs and datasets. VICL consistently outperforms both ICL and Zero-Shot across all models and datasets. This improvement underscores the effectiveness of VICL in enhancing the in-context learning capability of LVLMs by providing intent-oriented demonstrations. The method significantly bridges the cross-modal gap, allowing LVLMs to better understand and incorporate visual information within LLM reasoning processes. The performance increase is more pronounced in models like LLaVA-13B, where VICL boosts performance notably compared to the baseline Zero-Shot and ICL methods. This suggests that models with higher capacity or more parameters benefit more from the VICL approach due to their stronger ability to reason with multi-modal information.

## 5.3 Analysis

**Visual Demonstration Retrieval.** As shown in Figure 4, our experiments compare retrieval and




Figure 4: Different retrieval method comparison. “*V-Ret + VL-Rank*” denotes the combination of ViT for retrieval and CLIP for reranking. “*VL-Ret + V-Rank*” refers to CLIP for retrieval and ViT for reranking. “*V-Ret*” and “*VL-Ret*” are ViT and CLIP alone for retrieval, respectively. “*Random*” is random sampling.




Figure 5: Impact of Image Demonstration Number. reranking strategies on three datasets. “*V-Ret + VL-Rank*”, outperforms others on the Emotion6 dataset, highlighting the benefits of broad retrieval by ViT complemented by CLIP’s nuanced understanding. “*VL-Ret + V-Rank*”, maintains consistent performance across datasets but falls short of “*V-Ret + VL-Rank*” on Emotion6 and UnBiasedEmo, suggesting ViT’s unique approach may not always enhance performance. “*V-Ret*”, shows consistency but lacks the leading performance of combined approaches. “*VL-Ret*”, performs better than random sampling but lags behind two-step methods. “*Random*” emphasizes the need for strategic retrieval and reranking.

**Impact of Image Demonstration Number.** Our experimental results reveal a clear impact of the number of image demonstrations on the performance of both ICL and VICL. As shown in Figure 5, the ICL method shows a modest increase in performance with an increasing number of demonstrations, particularly evident in the progression from one to three demonstrations. However, the performance tends to plateau or even slightly de-




Figure 6: Impact of Context Length.




Figure 7: Impact of Demonstration Order.

crease beyond three demonstrations, suggesting a diminishing return on additional demonstrations. In stark contrast, the VICL method exhibits a more pronounced improvement with the increase in the number of demonstrations. This indicates that the VICL method effectively leverages additional demonstrations, translating into substantial performance gains.

**Impact of Context Length.** The performance dynamics of the VICL method, as observed in Figure 6, exhibit a discernible correlation with the length of the context provided. The results demonstrate an initial increase in accuracy with the expansion of context length of VICL. Particularly, on the EmoSet dataset, the accuracy ascends from 0.3 to a peak of 0.69, followed by a tapering off and slight fluctuations thereafter. A similar trend is observable in the Emotion6 and UnBiasedEmo dataset. However, beyond certain context lengths, there is a general trend of diminishing gains, or even a slight decline in accuracy. This highlights the balance between providing sufficient context for the model to leverage and avoiding an excessive amount.

**Order of Demonstrations.** We delve into the influence of the position of examples with posi-



Method	EmoSet	Emotion6	UnBiasedEmo
Standard	0.61	0.62	0.65
Task Intent	0.64	0.65	0.71
Image Parsing	0.66	0.68	0.69
IOIS	<b>0.69</b>	<b>0.70</b>	<b>0.76</b>

Table 2: Impact of Visual Summarization.

tive labels – labels as same as the true category of the prediction sample – within the demonstration sequence. Following (Liu et al., 2023b; Zhou et al., 2023a), we split the positions into three distinct sections: head, middle, and tail. As depicted in Figure 7, the head position yields the highest accuracy across all datasets, and the tail position demonstrates the next best performance. The middle position shows the least favorable performance. This trend could suggest that the model’s predictions are more influenced by examples positioned at the beginning and end of the sequence. These observations underscore the significance of demonstration order in visual in-context learning.

**Impact of Visual Summarization Method.** To evaluate the effect of various visual summarization for VICL, we consider four strategies, i.e., Standard captioning, Task Intent Summarization, Image Parsing Summarization, and Intent-Oriented Image Summarization (IOIS). Details can be found in Appendix B. As Table 2 shown, task intent summarization yields a moderate increase in accuracy, demonstrating the benefit of aligning the image summary with the task. Image Parsing, including a detailed visual reasoning process, can significantly enhance performance. The IOIS method, leveraging the strengths of the previous two approaches, achieves the best performance with a notable margin. The improvement demonstrates the efficacy of integrating task intent with image parsing, suggesting that an understanding of both task and visual content is paramount.

#### 5.4 In-Context Unlearning

We evaluate the capability of models to unlearn specific information, as shown in Table 3. We randomly selected sub-classes from the dataset and replaced the class to build the Unlearning Set, while the entire dataset constitutes the All Set. The details can be found in the Appendix C. The Zero-Shot shows the lowest performance, indicating a limited ability to disregard incorrect sub-class information based on the model’s pre-existing knowledge. ICL exhibits a marked improvement in the Unlearn-

Method	Emotion6		UnBiasedEmo	
	Unlearning Set	All Set	Unlearning Set	All Set
Zero-Shot	0.1	0.26	0.08	0.24
ICL	0.57	0.35	0.49	0.36
VICL	<b>0.77</b>	<b>0.69</b>	<b>0.82</b>	<b>0.74</b>

Table 3: In-Context Unlearning.

ing Set, demonstrating its ability to adapt to the new context provided by the altered demonstrations. VICL method significantly outperforms the other approaches, achieving the highest unlearning accuracy. VICL also maintains superior performance in the All Set, indicating robustness in distinguishing between correctly and incorrectly labeled samples and adjusting its inferences accordingly.

## 6 Conclusion

This paper has introduced the integration of In-Context Learning (ICL) into Large Visual Language Models (LVLMs), addressing challenges in cross-modal interactions and the distinct representation spaces. Through Visual In-Context Learning (VICL), incorporating Visual Demonstration Retrieval, Intent-Oriented Image Summarization, and Demonstration Composition, LVLMs show enhanced performance in understanding visual and textual information. In VICL, we have not only streamlined the in-context learning process but also introduced the concept of in-context unlearning, allowing LVLMs to adjust their knowledge base dynamically without the need for retraining. Our method shows effective in improving LVLMs in processing multi-modal tasks. The extensive evaluations verify the effectiveness of our VICL method, highlighting its potential to bridge the gap between visual and linguistic modalities.

## Limitations

This study introduces the VICL method to advance LVLMs, yet acknowledges several limitations warranting further investigation: (1) The efficacy of VICL heavily depends on the performance of LVLMs, which in turn is highly reliant on the original parameter size and the scale of training data. Further validation with larger LVLMs requires more computational resources. (2) While VICL demonstrates promise in visual reasoning tasks, its broader applications can be explored. Some Difficult tasks may require improved strategies for VICL method.

## References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L. Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. 2022. [Flamingo: a visual language model for few-shot learning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. [Qwen-vl: A frontier large vision-language model with versatile abilities](#). *CoRR*, abs/2308.12966.
- Hritik Bansal, Yonatan Bitton, Idan Szpektor, Kai-Wei Chang, and Aditya Grover. 2023. [Videocon: Robust video-language alignment via contrast captions](#). *CoRR*, abs/2311.10111.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. [Machine unlearning](#). In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 141–159. IEEE.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Shuo Chen, Zhen Han, Bailan He, Mark Buckley, Philip H. S. Torr, Volker Tresp, and Jindong Gu. 2023. [Understanding and improving in-context learning on vision-language models](#). *CoRR*, abs/2311.18021.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 1931–1942. PMLR.
- Li Deng. 2012. [The MNIST database of handwritten digit images for machine learning research \[best of the web\]](#). *IEEE Signal Process. Mag.*, 29(6):141–142.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023. [A survey for in-context learning](#). *CoRR*, abs/2301.00234.
- 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. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. [Complexity-based prompting for multi-step reasoning](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Xavier Garcia, Yamini Bansal, Colin Cherry, George F. Foster, Maxim Krikun, Melvin Johnson, and Orhan Firat. 2023. [The unreasonable effectiveness of few-shot learning for machine translation](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10867–10878. PMLR.
- Abigail Goldstein, Gilad Ezov, Ron Shmelkin, Micha Moffie, and Ariel Farkash. 2022. [Data minimization for GDPR compliance in machine learning models](#). *AI Ethics*, 2(3):477–491.
- Michael Hahn and Navin Goyal. 2023. [A theory of emergent in-context learning as implicit structure induction](#). *CoRR*, abs/2303.07971.
- Korbinian Koch and Marcus Soll. 2023. [No matter how you slice it: Machine unlearning with sisa comes at the expense of minority classes](#). *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 622–637.
- Alex Krizhevsky. 2009. Learning multiple layers of features from tiny images.
- Tianle Li, Xueguang Ma, Alex Zhuang, Yu Gu, Yu Su, and Wenhua Chen. 2023. [Few-shot in-context learning on knowledge base question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 6966–6980. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. [Visual instruction tuning](#). *CoRR*, abs/2304.08485.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023b. [Lost in the middle: How language models use long contexts](#). *CoRR*, abs/2307.03172.
- Yihao Liu, Xiangyu Chen, Xianzheng Ma, Xintao Wang, Jiantao Zhou, Yu Qiao, and Chao Dong. 2023c. [Unifying image processing as visual prompting question answering](#). *CoRR*, abs/2310.10513.

- Sheng Lu, Irina Bigoulaeva, Rachneet Sachdeva, Harish Tayyar Madabushi, and Iryna Gurevych. 2023. [Are emergent abilities in large language models just in-context learning?](#) *CoRR*, abs/2309.01809.
- Haitao Mao, Guangliang Liu, Yao Ma, Rongrong Wang, and Jiliang Tang. 2024. [A data generation perspective to the mechanism of in-context learning.](#) *CoRR*, abs/2402.02212.
- Gary Marcus, Ernest Davis, and Scott Aaronson. 2022. [A very preliminary analysis of DALL-E 2.](#) *CoRR*, abs/2204.13807.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 14014–14024.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 11048–11064. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. [Adaptive machine translation with large language models.](#) In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation, EAMT 2023, Tampere, Finland, 12-15 June 2023*, pages 227–237. European Association for Machine Translation.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. 2022. [A survey of machine unlearning.](#) *CoRR*, abs/2209.02299.
- OpenAi. 2023. [Gpt-4v\(ision\) system card.](#)
- Haowen Pan, Yixin Cao, Xiaozhi Wang, and Xun Yang. 2023. [Finding and editing multi-modal neurons in pre-trained transformer.](#) *CoRR*, abs/2311.07470.
- Rameswar Panda, Jianming Zhang, Haoxiang Li, Joon-Young Lee, Xin Lu, and Amit K. Roy-Chowdhury. 2018. [Contemplating visual emotions: Understanding and overcoming dataset bias.](#) In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part II*, volume 11206 of *Lecture Notes in Computer Science*, pages 594–612. Springer.
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. 2023. [In-context unlearning: Language models as few shot unlearners.](#) *CoRR*, abs/2310.07579.
- Kuan-Chuan Peng, Tsuhan Chen, Amir Sadovnik, and Andrew C. Gallagher. 2015. [A mixed bag of emotions: Model, predict, and transfer emotion distributions.](#) In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 860–868. IEEE Computer Society.
- Yingzhe Peng, Xu Yang, Haoxuan Ma, Shuo Xu, Chi Zhang, Yucheng Han, and Hanwang Zhang. 2023. [ICD-LM: configuring vision-language in-context demonstrations by language modeling.](#) *CoRR*, abs/2312.10104.
- Chengwei Qin, Aston Zhang, Anirudh Dagar, and Wenming Ye. 2023. [In-context learning with iterative demonstration selection.](#) *CoRR*, abs/2310.09881.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision.](#) In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. [Improving language understanding by generative pre-training.](#)
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners.](#) *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer.](#) *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. 2022. [Photorealistic text-to-image diffusion models with deep language understanding.](#) In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Sarah Schwettmann, Neil Chowdhury, Samuel Klein, David Bau, and Antonio Torralba. 2023. [Multimodal neurons in pretrained text-only transformers.](#) In *IEEE/CVF International Conference on Computer Vision, ICCV 2023 - Workshops, Paris, France, October 2-6, 2023*, pages 2854–2859. IEEE.
- Yanpeng Sun, Qiang Chen, Jian Wang, Jingdong Wang, and Zechao Li. 2023. [Exploring effective factors for improving visual in-context learning.](#) *CoRR*, abs/2304.04748.

- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 35151–35174. PMLR.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. [Label words are anchors: An information flow perspective for understanding in-context learning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 9840–9855. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#). *Trans. Mach. Learn. Res.*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022b. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. [An explanation of in-context learning as implicit bayesian inference](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jiarui Xu, Yossi Gandelsman, Amir Bar, Jianwei Yang, Jianfeng Gao, Trevor Darrell, and Xiaolong Wang. 2023. [Improv: Inpainting-based multimodal prompting for computer vision tasks](#). *CoRR*, abs/2312.01771.
- Jingyuan Yang, Qirui Huang, Tingting Ding, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2023. [Emoset: A large-scale visual emotion dataset with rich attributes](#). In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 20326–20337. IEEE.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2023. [A survey on multimodal large language models](#). *CoRR*, abs/2306.13549.
- Haibo Zhang, Toru Nakamura, Takamasa Isohara, and Kouichi Sakurai. 2023a. [A review on machine unlearning](#). *SN Comput. Sci.*, 4(4):337.
- Pan Zhang, Xiaoyi Dong, Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuangrui Ding, Songyang Zhang, Haodong Duan, Wenwei Zhang, Hang Yan, Xinyue Zhang, Wei Li, Jingwen Li, Kai Chen, Conghui He, Xingcheng Zhang, Yu Qiao, Dahua Lin, and Jiaqi Wang. 2023b. [Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition](#). *CoRR*, abs/2309.15112.
- Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. 2023c. [What makes good examples for visual in-context learning?](#) *CoRR*, abs/2301.13670.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023d. [Automatic chain of thought prompting in large language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Yucheng Zhou, Xiubo Geng, Tao Shen, Chongyang Tao, Guodong Long, Jian-Guang Lou, and Jianbing Shen. 2023a. [Thread of thought unraveling chaotic contexts](#). *CoRR*, abs/2311.08734.
- Yucheng Zhou, Tao Shen, Xiubo Geng, Chongyang Tao, Can Xu, Guodong Long, Binxing Jiao, and Daxin Jiang. 2023b. [Towards robust ranker for text retrieval](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 5387–5401. Association for Computational Linguistics.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. [Minigt-4: Enhancing vision-language understanding with advanced large language models](#). *CoRR*, abs/2304.10592.

## A Prompts

There are prompts for the three methods, i.e., Zero-Shot, ICL, and VICL.

- For EmoSet, Emotion6 and UnBiasedEmo dataset, the prompt for Zero-Shot is “Question: Do you feel which emotion when seeing this image? There is an emotion category list: [{{Label List}}]. Image: {image}. Answer: ”.
- For CIFAR10 and MNIST dataset, the prompt for Zero-Shot is “Question: What you see in this image? There is a category list: [{{Label List}}]. Image: {image}. Answer: ”.
- For EmoSet, Emotion6 and UnBiasedEmo dataset, the prompt for ICL is “Question: Do you feel which emotion when seeing this image? There is an emotion category list: [{{Label List}}]. Image 1: {image-1}. Answer: {label-1}. Image 2: {image-2}. Answer: {label-2} ... Image N: {image-N}. Answer: ”.
- For CIFAR10 and MNIST dataset, the prompt for ICL is “Question: What you see in this image? There is a category list: [{{Label List}}]. Image 1: {image-1}. Answer: {label-1}. Image 2: {image-2}. Answer: {label-2} ... Image N: {image-N}. Answer: ”.
- For EmoSet, Emotion6 and UnBiasedEmo dataset, the prompt for VICL is “Question: Do you feel which emotion when seeing this image? There is an emotion category list: [{{Label List}}]. Image 1: {summary-1}. Answer: {label-1}. Image 2: {summary-2}. Answer: {label-2} ... Image N: {image-N}. Answer: ”.
- For CIFAR10 and MNIST dataset, the prompt for VICL is “Question: What you see in this image? There is a category list: [{{Label List}}]. Image 1: {summary-1}. Answer: {label-1}. Image 2: {summary-2}. Answer: {label-2} ... Image N: {image-N}. Answer: ”.

## B Visual Summarization Prompt

We investigate different visual summarization method for VICL, and the prompt for summarization as follows:

- The Standard captioning approach employs conventional captioning techniques, and its

prompt is “Generate a detailed description of the content depicted in the provided image.”.

- Task Intent method enriches image descriptions with task-specific intent, and its prompt is “Given an image and a corresponding label, generate a descriptive caption that not only describes the image content but also conveys the intention or purpose behind the depicted scene.”.
- Image Parsing goes further by incorporating descriptions of both image observations and the reasoning process, and its prompt is “You are presented with an image along with accompanying labels. Your task is to provide a detailed description of the image content while also explaining the observations and reasoning process behind your description.”.
- IOIS combines features of both Task Intent and Image Parsing to provide comprehensive summaries, and its prompt is “Generate a descriptive caption for the provided image and labels, elucidating both the visual content and the underlying purpose or intention depicted. Craft a clear and concise description that seamlessly integrates details from the image and labels, highlighting connections between visual cues and semantic meaning. Your caption should not only describe what is visible in the image but also convey the task-oriented aspect.”.

## C In-Context Unlearning

The “Unlearning Set” comprises samples from five randomly selected sub-classes with labels reassigned to alternate categories and incorporated into the demonstration set and test set. Specifically, we randomly select an example corresponding to the input image’s class and include it in the demonstration set. Other examples in the demonstration set are drawn from samples belonging to standard categories. This setup is designed to assess the model’s ability to discard previously learned sub-class information when exposed to intentionally mislabeled examples. The performance on this set directly reflects the unlearning accuracy. Meanwhile, the “All Set” includes the Unlearning Set combined with additional samples from standard categories.