

Exploring Task-Level Optimal Prompts for Visual In-Context Learning

Yan Zhu^{1*}, Huan Ma^{1*}, Changqing Zhang^{1†}

¹College of Intelligence and Computing, Tianjin University, Tianjin, China

Abstract

With the development of Vision Foundation Models (VFMs) in recent years, Visual In-Context Learning (VICL) has become a better choice compared to modifying models in most scenarios. Different from retraining or fine-tuning model, VICL does not require modifications to the model’s weights or architecture, and only needs a prompt with demonstrations to teach VFM how to solve tasks. Currently, significant computational cost for finding optimal prompts for every test sample hinders the deployment of VICL, as determining which demonstrations to use for constructing prompts is very costly. In this paper, however, we find a counterintuitive phenomenon that most test samples actually achieve optimal performance under the same prompts, and searching for sample-level prompts only costs more time but results in completely identical prompts. Therefore, we propose task-level prompting to reduce the cost of searching for prompts during the inference stage and introduce two time-saving yet effective task-level prompt search strategies. Extensive experimental results show that our proposed method can identify near-optimal prompts and reach the best VICL performance with a minimal cost that prior work has never achieved.

Introduction

With the development of Vision Foundation Models (VFMs), many tasks in visual scenes no longer require training new models but can be solved by VLFs provided by model service providers, which are more affordable and convenient (Li et al. 2023; Wu, Sun, and Ouyang 2023). However, the performance of directly deploying VFMs for specific tasks is often unsatisfactory, and task-specific adaptation is necessary for better performance. In comparison to modifying model weights, Visual In-Context Learning (VICL) (Yang et al. 2024; Chen et al. 2024) is a better choice. Unlike retraining or fine-tuning models, VICL does not necessitate modifying the model’s weights or architecture, but only needs to teach VFMs how to solve problems using prompts with a set of demonstrations. However, VICL does not always achieve good results under arbitrary prompts, so determining the appropriate demonstrations to

construct prompts is the key challenge in enhancing VICL performance (Li et al. 2024; Dedhia et al. 2024).

Currently, there are primarily two ways to select demonstrations for constructing prompts: the rule-guided strategy and the reward-model-based strategy. The former strategy involves selecting demonstrations based on rules, like Un-supPR (Zhang, Zhou, and Liu 2023) and prompt-Self (Sun et al. 2023), which calculate the similarity between the query samples and the training samples using features extracted by the pretrained CLIP (Radford et al. 2021) model to select prompts. Although this type of method is very simple, these heuristic methods struggle to guarantee performance. To solve the issue, another strategy is to train a reward model that can predict performance under different prompts. For example, both SupPR (Zhang, Zhou, and Liu 2023) and In-MeMo (Zhang et al. 2024) train an additional reward model on the validation set to predict the compatibility between prompts and query samples. However, to obtain such a reward model, it requires a large amount of labeled data and incurs a high computational cost, or the reward model will suffer from severe overfitting. Unfortunately, this high computational cost severely hinders the deployment of VICL.

Is it really necessary to spend a significant amount of computation to find the optimal prompt for each individual sample? Perhaps it is not actually needed. As shown in Fig. 1(d), *we find that under a specific task, the prompts resulting in the best performance for different samples are always the same*. The prompt that performs best on the task, even if it cannot achieve optimal performance for some samples, can still obtain relatively good results compared with most prompts. In other words, there is no need to spend a massive amount of computational effort at the risk of overfitting to select different demonstrations to form a prompt for distinct samples (experimental results show that using a single prompt for all samples performs even better than those costly sample-level methods).

In this paper, we propose a reward-based, training-free approach to find the optimal task-level prompt. This strategy not only lessens the time taken during inference but also sustains performance when compared to more time-consuming methods. Specifically, we present two new strategies for effectively searching for the best demonstrations: (1) Top- K strategy and (2) Greedy search strategy. We concentrate on a general setting where a labeled set of size N is provided. The

*These authors contributed equally.

†Corresponding to zhangchangqing@tju.edu.cn.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

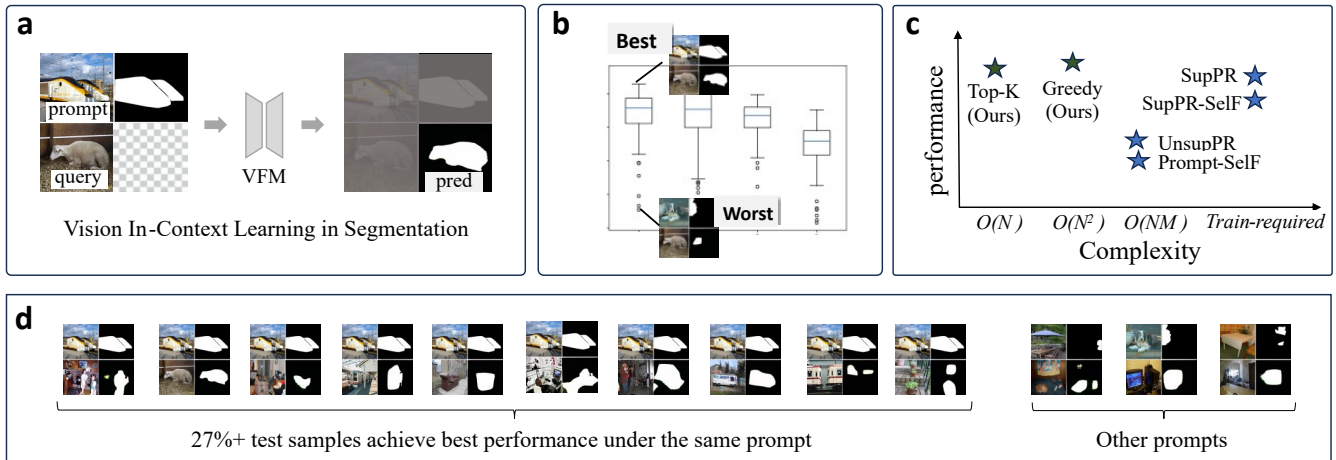


Figure 1: (a) The deployment example of VICL in segmentation task. The one-shot case is presented in the figure, and the few-shot prediction is the average prediction under multiple different demonstrations. (b) Selecting different prompts makes a significant impact on the tasks. The variance across different prompts is large, even resulting in cases where the metric approaches zero. (c) Comparison of time complexity and performance. Our methods (task-level prompting) significantly reduce complexity while ensuring that the performance is not worse than that of more complex methods. (d) Motivation for task-level prompting. We find that during the testing phase, most samples achieve optimal performance under the same prompt. As shown in the figure, more than 27% of the samples achieve their best performance under the same prompt, which means that finding the optimal task-level prompt ensures that at least 27% of the samples obtain the best prompt. In contrast, the sample-level prompt searching strategy only finds the optimal prompt for 15.03% of the samples (for details please refer to results section).

aim of our strategies is to carry out combinatorial optimization over this set to discover optimal demonstrations. In particular, Top- K strategy uses a straightforward method that first measures the performance of each individual demonstration (i.e., one-shot prompting) and then picks the top K best demonstrations to create the final prompts. Note that Top- K strategy assumes that the optimal prompt is typically built from demonstrations that perform well when used on their own. The Greedy search strategy follows the standard greedy search procedure, finding the best solution by making the best local choices at each step. At each step of the algorithm, the chosen demonstration is the one that allows the updated prompts to achieve the best performance.

To evaluate the effectiveness of our strategies, we conduct extensive experiments on various downstream tasks, such as foreground segmentation, single object detection, and colorization. Our results indicate that our method can significantly enhance the VFM’s in-context learning performance in an effective and interpretable manner. The overall contribution is summarized as follows:

- We introduce a task-level prompt to avoid the significant computational cost and the risk of overfitting associated with sample-level methods, which select different demonstrations for different samples.
- We propose two time-saving and effective prompt search strategies to identify near-optimal prompts and achieve SOTA performance with minimal cost, which has not been achieved in prior work.
- The effectiveness of these two strategies is demonstrated in various tasks. While saving more than 98% of the

prompt searching time, consistent relative improvements of over 6.2% are observed across different downstream tasks compared to state-of-the-art methods.

Related Work

Visual In-Context Learning

The emergence of large language models (LLMs) like GPT-3 (Brown et al. 2020), BLOOM (Workshop et al. 2022), and LLaMA (Touvron et al. 2023) introduces a new learning paradigm, In-Context Learning (ICL) (Li, Tang, and Mei 2019; Li et al. 2022; Wang et al. 2024; Baldassini et al. 2024), which refers to the process of conditioning an LLM to solve various downstream NLP tasks using prompts constructed from a few demonstration input-output pairs (Liu et al. 2022; Cho et al. 2023; Ma et al. 2024a) (i.e., few-shot prompting). In visual domain, MAE-VQGAN (Bar et al. 2022) utilizes the model’s grid inpainting capability to propose the first visual ICL model. Similarly, Painter (Wang et al. 2023a) performs standard masked image modeling on the stitch of input and output image pairs to train an ICL model. Additionally, SegGPT (Wang et al. 2023b), a variant of Painter, enhances the segmentation ability of the ICL model by randomly coloring similar semantic categories or objects. Previous works demonstrate the effectiveness of visual ICL models, and that selecting appropriate prompts ensures the model’s understanding of task knowledge.

Visual Prompt Selection

To fully leverage the powerful reasoning capabilities of visual context learning models like MAE-VQGAN (Bar et al.

2022) and Painter (Wang et al. 2023a), researchers are dedicated to exploring new algorithms for selecting appropriate prompts for different query samples/tasks (Huang et al. 2023; Ma et al. 2024b). These visual prompt selection methods can be divided into two categories, which are reward-model-based strategy and the rule-guided strategy:

- **Rule-guided** strategy methods include UnsupPR (Zhang, Zhou, and Liu 2023) and prompt-SELF (Sun et al. 2023). Both algorithms calculate the similarity between the query samples and the training samples using features extracted by the pretrained CLIP (Radford et al. 2021) model to select demonstrations. The latter enhances the selected prompts to make full use of prompt information. However, such unsupervised algorithms do not utilize the labeled information of the training set and do not introduce task content constraints in the selection process, thus limiting the performance of the algorithms.
- **Reward-model-based** strategy methods include SupPR (Zhang, Zhou, and Liu 2023) and In-MeMo (Zhang et al. 2024). Both algorithms train an additional scoring model on the training set to predict the compatibility between prompts and query samples. The latter proposes a prompt enhancer to improve the prompts and obtain a higher-performing scoring model. However, such supervised algorithms are costly to train, require a large amount of labeled data, and are prone to overfitting when the training samples are too few.

Therefore, to fully utilize the labeled information under few-shot supervision data, this paper proposes two simple greedy task-level prompt selection methods, which are Top- K prompt selection method and Greedy prompt selection method. The former has a time complexity of $O(N)$, and the latter has a time complexity of $O(N^2)$ in the worst case.

Methods

Problem Setup

Let $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^N$ be a validation set consisting of N image-label pairs, where x_i denotes an image, and y_i is the corresponding label (e.g., 0/1 masks in a segmentation task). The target of VICL is to select a subset of samples from \mathcal{S} to create a demonstration set, denoted as $\mathcal{P} = \{(x_i, y_i)\}_{i=1}^K$, which is used to prompt a pretrained foundation model f . This prompt aims to achieve the following goal: given a new query sample (x_q, y_q) , the foundation model should generate a prediction $\hat{y}_q = f(\mathcal{P}, x_q)$ for x_q that is as close as possible to the ground-truth y_q , which can be formally represented as minimizing $\mathcal{L}(\hat{y}_q, y_q)$, where $\mathcal{L}(\cdot)$ denotes the loss metric used for various specific tasks.

Sample-level Prompt

Sample-level prompt selection methods aim to find a \mathcal{P}^* for each query sample:

$$\mathcal{P}^* = \operatorname{argmin}_{\mathcal{P} \subseteq \mathcal{S}} \mathcal{L}(f(\mathcal{P}, x_q), y_q), \quad (1)$$

Since y_q is unknown, previous works (Zhang, Zhou, and Liu 2023; Zhang et al. 2024; Sun et al. 2023) focus on constructing a scoring function g (which can be a manually set rule,

like IOU, accuracy and so on, or parameters obtained from supervised learning) to automatically select the most suitable example(s) from the validation dataset \mathcal{S} for a query sample x_q . The prompt selection strategy is:

$$x^* = \operatorname{argmax}_{x \in \mathcal{S}} g(x, x_q). \quad (2)$$

These methods rank the training examples based on their scores and choose the top- K example pairs. When $K = 1$, they choose the optimal example pair as the prompt, $\mathcal{P}_q = \{(x^*, y^*)\}$.

Sample-level requires a significant amount of time to obtain a reward model (scoring model) and evaluate each query sample before VICL, which hinders the flexible deployment of VICL. However, these efforts seem to yield little benefit. In previous works, researchers find that differences in prompts can lead to significant variations in final performance, and intuitively search suitable prompts for different query samples. However, we discover a counter-intuitive phenomenon: although different prompts can cause dramatic changes in VICL’s performance (as shown in Fig. 1(b)), the prompts that perform well on every different query sample are always the same (as shown in Fig. 5). Therefore, we only need to spend a small amount of computational cost to find the best task-level prompt, rather than spending a large amount of computational cost to determine the same prompt for all samples individually.

Task-level Prompt

As the prompts that perform well on every different query sample are always the same, we propose the objective of task-level prompt selection methods to construct a demonstration set \mathcal{P} from labeled data \mathcal{S} shared across different query samples x_q :

$$\mathcal{P}^* = \operatorname{argmin}_{\mathcal{P} \subseteq \mathcal{S}} \sum_{\mathcal{D}} \mathcal{L}(f(\mathcal{P}, x_q), y_q), \quad (3)$$

where \mathcal{D} includes all the unseen query samples, namely $(x_q, y_q) \in \mathcal{D}$. Since obtaining \mathcal{D} is difficult, we substitute \mathcal{D} with known data \mathcal{S} based on the generalization of dataset:

$$\mathcal{P}^* = \operatorname{argmin}_{\mathcal{P} \subseteq \mathcal{S}} \sum_{\mathcal{S} - \mathcal{P}} \mathcal{L}(f(\mathcal{P}, x_q), y_q), \quad (4)$$

where $\mathcal{S} \neq \mathcal{P}$, and $(x_q, y_q) \in \mathcal{S} - \mathcal{P}$. The most straightforward approach in this method is to obtain all possible combinations of \mathcal{P} from \mathcal{S} and use the performance of \mathcal{P} on $\mathcal{S} - \mathcal{P}$ to obtain loss.

However, this approach still has a time complexity that is excessively high, with worst-case and best-case scenarios both being $O(2^N)$ (we integrate the features of prompt set by summation pooling without considering the effect of order, and it will be $O(N!)$ when considering order), making it difficult to achieve the optimal solution. Inspired by the method proposed in (Ma et al. 2024a), which aims to address the prompt selection in LLM, this paper extends two simple task-level prompt selection methods, which are Top- K prompt selection method and Greedy prompt selection method. The former has a time complexity of $O(N)$, and the latter has a time complexity of $O(N^2)$ in the worst case.

Algorithm 1: Top- K Prompt Selection Method

- 1: **Given:** validation set $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^N$, pretrained model g , number K
 - 2: Initial demonstration set $\mathcal{P} = \{\}$
 - 3: **for** k in $1, \dots, K$ **do**
 - 4: $x^* = \operatorname{argmin}_{x \in \mathcal{S}} \sum_{\mathcal{S}-\{x\}} \mathcal{L}(f(\{x\}, x_q), y_q)$
 - 5: *Insert* the top sample x^* into \mathcal{P} and *remove* it from \mathcal{S}
 - 6: **end for**
 - 7: **return** \mathcal{P}
-

Top- K Prompt Selection Method The main idea of our Top- K Prompt Selection Method proposed in this section is to simplify the basic idea of task-level prompt selection by using prior human knowledge to constrain the demonstration sets obtained from \mathcal{S} , appropriately reducing the combination possibilities to achieve a reduction in time complexity. Specifically, this means moving from using $O(2^N)$ time to obtain the optimal demonstration set to using $O(N)$ time to obtain the optimal single prompt:

$$x^* = \operatorname{argmin}_{x \in \mathcal{S}} \sum_{\mathcal{S}-\{x\}} \mathcal{L}(f(\{x\}, x_q), y_q), \quad (5)$$

where $(x_q, y_q) \in \mathcal{S}-\{x\}$. However, for a given task, a single prompt sometimes cannot provide enough information to complete the task, and multiple prompts are needed. Therefore, to obtain a demonstration set composed of multiple prompts, we sort the scores of the single prompts and select the top K prompts to form the demonstration set. The specific algorithm is shown in Algorithm 1.

Nevertheless, the Top- K prompt selection method heavily relies on the choice of the hyperparameter K . When K is too small, the information may be insufficient; when K is too large, it may have a negative effect (as shown in Fig. 3). To automate the determination of the length of the demonstration set, this paper proposes a greedy prompt selection method.

Greedy Prompt Selection Method Our Greedy Prompt Selection Method proposed in this section follows the basic greedy approach, where at each stage, the current optimal prompt set is identified to guide the next stage’s operation. This involves selecting the most promising new sample to add to the existing demonstration set based on:

$$x^{greedy} = \operatorname{argmin}_{x \in (\mathcal{S}-\mathcal{P})} \sum_{\mathcal{S}-\mathcal{P}-\{x\}} \mathcal{L}(f(\mathcal{P} + \{x\}, x_q), y_q), \quad (6)$$

where $(x_q, y_q) \in \mathcal{S} - \mathcal{P} - \{x\}$. Although the new sample represents the optimal solution for the current demonstration set, simply increasing the quantity of prompts does not necessarily improve performance and can even have adverse effects (as shown in Fig. 3). Therefore, it’s crucial to compare the score $a_{new} = g(\mathcal{P} + \{x\}, \mathcal{S} - \mathcal{P} - \{x\})$ obtained from adding the new sample with the score $a_{ori} = g(\mathcal{P}, \mathcal{S} - \mathcal{P})$ of the original demonstration set. This comparison facilitates early stopping and pruning of the algorithm, aiming

Algorithm 2: Greedy Prompt Selection Method

- 1: **Given:** validation set $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^N$, pretrained model g
 - 2: Initial demonstration set $\mathcal{P} = \{\}$
 - 3: **while** \mathcal{S} is not null **do**
 - 4: Calculate the x^{greedy} according to Eq. 6
 - 5: *Insert* the greedy sample into \mathcal{P} that can improve score best and *remove* it from \mathcal{S}
 - 6: *Stop* searching when score can’t be improved
 - 7: **end while**
 - 8: **return** \mathcal{P}
-

to achieve global optimal solution through local optimal solutions. When $a_{ori} \leq a_{new}$, it indicates that the new demonstration set performs better on the validation set, allowing the new sample to be included in demonstration set for the next iteration:

$$\mathcal{P} = \mathcal{P} + \{x^{greedy}\}, \quad (7)$$

However, when $a_{ori} > a_{new}$, it suggests that adding the locally optimal sample no longer enhances the performance on the validation set, prompting the algorithm to halt the iteration and output the current demonstration set as the solution, $\mathcal{P}^* = \mathcal{P}$. The specific algorithm is illustrated in Algorithm 2.

Experiments

In this section, we discuss the performance comparison between the two proposed task-level prompt selection methods and other unsupervised prompt selection methods, as well as the performance impact of different methods for constructing \mathcal{S} . Additionally, we explore the necessity of using a greedy strategy to achieve a global solution, and conduct fine-grained, in-depth experiments starting from individual samples.

Setup

According to (Bar et al. 2022; Zhang, Zhou, and Liu 2023; Sun et al. 2023), we conduct few-shot evaluations on three out-of-distribution (OOD) computer vision tasks:

- **Foreground Segmentation.** We use the Pascal-5i (Shaban et al. 2017) dataset, which is comprised of 4 different image splits, where each split contains data from 5 categories. For evaluation, we report the mean Intersection Over Union (mIOU) metric.
- **Single Object Detection.** We use the Pascal VOC 2012 (Everingham et al. 2015) dataset, which contains images and their associated detection boxes. The evaluation method is similar to Foreground Segmentation, and we report the mIOU metric.
- **Colorization.** We use a subset of the ImageNet (Russakovsky et al. 2015) dataset, which contains data from 1000 categories. We randomly sample 50,000 example pairs and image queries from the ImageNet validation set. For evaluation, we report the Mean Squared Error (MSE) loss (scaled up by a factor of 100).

	Method	Det.(mIOU) \uparrow	Color.(mse) \downarrow	Split-0	Split-1	Seg.(mIOU) \uparrow Split-2	Split-3	Avg.
Sample	UnsupPR	24.01 _{0.53}	64.95 _{0.27}	37.28 _{1.19} ¹	39.44 _{1.87}	33.32 _{3.04}	28.22 _{0.59}	34.57
	SupPR	-	-	36.97 _{1.25} ²	39.74 _{1.78}	34.24 _{1.88}	29.15 _{1.05}	35.03
	prompt-Self	15.00 _{0.57}	45.61 _{0.36} ¹	32.76 _{1.18}	38.45 _{0.92}	35.77 _{0.23}	29.38 _{0.20}	34.09
	SupPR-Self	-	-	33.42 _{0.96}	38.97 _{0.44}	35.64 _{0.46}	30.40 _{0.37} ²	34.61
	Oracle*	29.03 _{2.84}	61.56 _{1.27}	39.09 _{0.77}	44.37 _{0.98}	37.93 _{0.42}	32.40 _{1.06}	38.45
Task	Random	24.53 _{0.88}	66.39 _{0.17}	35.51 _{0.53}	39.83 _{2.39}	33.17 _{2.57}	25.30 _{1.55}	33.45
	Top-1	28.25 _{2.94} ¹	61.56 _{1.27} ²	32.41 _{2.31}	42.22 _{1.42} ¹	37.00 _{0.82} ²	30.20 _{2.44}	35.46 ²
	Greedy	28.25 _{2.94} ¹	61.56 _{1.27} ²	36.86 _{1.38}	42.22 _{1.42} ¹	37.11 _{0.87} ¹	30.84 _{1.84} ¹	36.76 ¹

Table 1: Main results. Our methods achieve the best performance and are close to the Oracle.

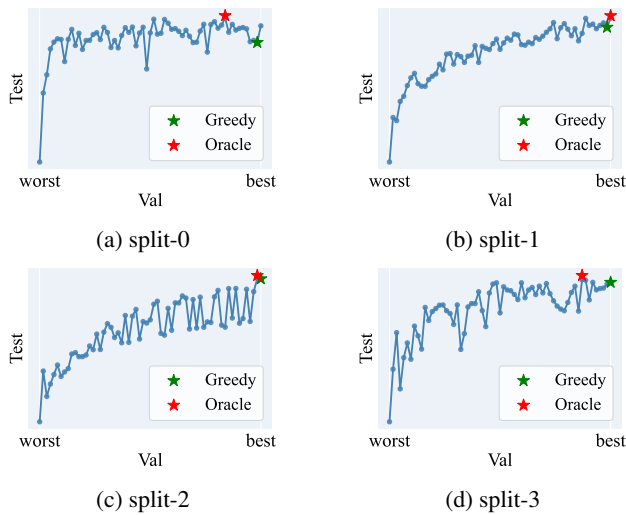


Figure 2: Performance of all prompts on test set, where “Greedy” and “Oracle” indicate performance under prompts selected by our strategy and the upper-bound performance across all prompts, respectively.

In each task, the few-shot in-context examples come from the training set, with a default size of $N = 16$. For all experiments, we perform evaluations on a pre-trained image inpainting model, MAE-VQGAN (Bar et al. 2022), which consists of an encoder and a decoder. When the prompt has K samples, each sample is combined with the query image into a grid of 2×2 sub-images. These are then passed through the encoder to obtain K features, which are summed to get fused features and then input to the decoder for the final output.

In the experiments, we mainly compare six unsupervised prompt selection methods:

- **Random.** Randomly select prompt combinations from the few-shot in-context examples. (The result of the random method is obtained by averaging the performance of all possible few-shot in-context examples combinations on the test set.)¹
- **Oracle.** The task-level best-performing combination on the test set from all possible few-shot in-context exam-




Figure 3: Performance of the demonstration set with a length of 2 and its subsets for each sample, where half of the combinations shows a performance drop when adding samples.

ples combinations (the upper bound performance of task-level prompting).

- **Top- K .** Select the top- K -performing samples under one-shot VICL on the validation set.
- **Greedy.** Use a greedy method to select from the few-shot in-context examples.
- **UnsupPR.** (Zhang, Zhou, and Liu 2023) Unsupervised prompt retrieval (UnsupPR), a method that uses off-the-shelf features for nearest example search. We use the cosine similarity between features extracted by CLIP’s vision encoder to search for the nearest example.
- **Prompt-Self.** (Sun et al. 2023) Use the nearest example obtained from UnsupPR as the query image and fuses the query image and prompt image using different arrangements to create 8 new fused images.

Also, we compare two supervised prompt selection methods:

- **SupPR.** (Zhang, Zhou, and Liu 2023) Supervised Prompt Retrieval(SupPR), is a supervised prompt retrieval method, which trains a neural network to choose

¹ $O(2^N)$ when $N = 16$, we randomly select $N' = 6$ samples for obtaining \mathcal{P} and use the remaining $N - N' = 10$ samples for validation set performance evaluation in the random, Oracle, and our methods.

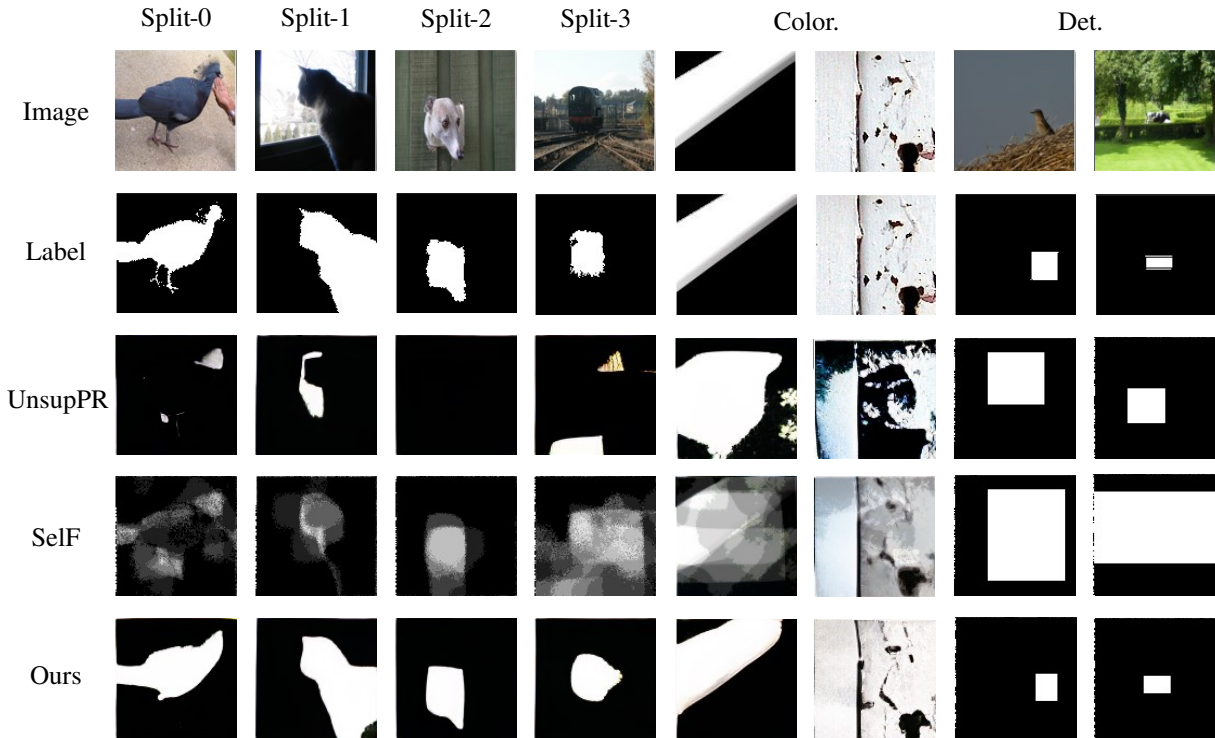


Figure 4: In-context results retrieved by several unsupervised prompt selection methods, where “Ours” is the Greedy Prompt Selection Method.




Figure 5: A comparison of the performance of prompts selected by different strategies among all possible prompts (the far right indicates that the selected prompt is the best-performing one among all possible prompts).

examples that directly maximize in-context learning performance.

- **SupPR-Self.** It is a combined method of SupPR and Prompt-Self, which uses SupPR to select prompt and Prompt-Self to do augment of prompt.

The eight methods are mainly divided into two categories: UnsupPR, Prompt-Self, SupPR and SupPR-Self belong to sample-level prompt selection methods, while Random, Or-

acle, and Ours (including Top- K and Greedy) belong to task-level prompt selection methods.

Results

We conduct experiments on different settings and report the results of three runs (with $seed = \{0, 1, 2\}$). In Tab. 1, it presents the main results of the experiments. The following findings can be observed: (1) Our methods achieve state-of-the-art performance across various datasets, both when compared to sample-level methods and task-level methods, demonstrating the effectiveness of our approaches. In Tab. 1, our methods achieve optimal results in detection and segmentation tasks, and even achieve a global optimal solution in the coloring task. (2) Under certain conditions, our methods can reach the global optimal solution, with overall results very close to Oracle. In the detection task, our Greedy method is less than 3% away from Oracle, and in the segmentation task, the average results across different splits are less than 6% away from Oracle. Additionally, in the coloring task, the results achieve the global optimal solution, consistent with Oracle (even if the results did not reach the optimal in the comparison of different methods). In Fig. 2, it shows that the results of the validation set and the test set are highly consistent, and our Greedy method can obtain results very close to the global optimal solution. (3) As shown in Tab. 3, our task-level methods not only achieve state-of-the-art performance, but also have the lowest worst-case complexity.

Task-level vs. Sample-level. As shown in Fig. 1(d), the task-level optimal prompt enables 27% of the samples to

Method	Det.(mIOU) \uparrow	Color.(mse) \downarrow	Seg.(mIOU) \uparrow				Avg.
			Split-0	Split-1	Split-2	Split-3	
Oracle*	29.03 _{2.84}	61.56 _{1.27}	39.09 _{0.77}	44.37 _{0.98}	37.93 _{0.42}	32.40 _{1.06}	38.45
Random	24.53 _{0.88}	66.39 _{0.17}	35.51 _{0.53}	39.83 _{2.39}	33.17 _{2.57}	25.30 _{1.55}	33.45
Top-1	28.25 _{2.94}	61.56 _{1.27}	32.41 _{2.31}	42.22 _{1.42}	37.00 _{0.82}	30.20 _{2.44}	35.46
Top-2	26.55 _{0.41}	63.70 _{0.95}	36.86 _{1.22}	43.07 _{1.13}	37.56 _{0.29}	29.74 _{1.15}	36.81
Top-4	26.72 _{2.22}	65.66 _{0.50}	37.71 _{0.68}	43.36 _{1.15}	36.43 _{1.17}	30.24 _{1.12}	36.94
Greedy	28.25 _{2.94}	61.56 _{1.27}	36.86 _{1.38}	42.22 _{1.42}	37.11 _{0.87}	30.84 _{1.84}	36.76

Table 2: Sort each sample based on its performance on the remaining validation set, and select the top- K as the target prompt set, where $K = \{1, 2, 4\}$.

	Method	Complexity
Sample	UnsupPR	$O(NM)$
	SupPR	Train-required
	prompt-Self	$O(NM)$
	SupPR-Self	Train-required
Task	Top-1	$O(N)$
	Greedy	$O(N^2)$

Table 3: A comparison of the worst-case complexity of different methods, where N and M represent the sizes of the demonstration set and test set, respectively ($N \ll M$).

achieve the optimal solution, while current sample-level prompt selection methods can only find the optimal solution for 15.03% of the samples. Moreover, we conduct a detailed study comparing the performance of prompts selected by the task-level method and the sample-level method on individual samples. As illustrated in Fig. 5, we compare the selected prompt with all possible demonstration sets on individual samples and plot the ranking frequency of these samples across the entire test dataset. It can be observed that both the task-level and sample-level methods exhibit long-tail distributions in their frequency curves, indicating that under a specific task, the prompts resulting in the best performance for different samples are always the same. However, the distribution of the task-level method is significantly more concentrated than that of the sample-level method, with frequency increasing more rapidly as performance improves. It demonstrates that the simple task-level method is more effective than the complex sample-level method. The prompt that performs best on the task, even if it cannot achieve optimal performance for some samples, can still obtain relatively good results compared with most prompts. In other words, there is no need to spend a massive amount of computational effort at the risk of overfitting to select different demonstrations to form a prompt for distinct samples.

Top- K vs. Greedy. Based on the simple assumption that the optimal prompt is typically built from demonstrations that perform well, the most straightforward approach is to sort each sample based on its performance on the remaining validation set and select the top K as the target demonstration set, which has a lower time complexity of $O(N)$.

Therefore, in this section, we compare the Greedy method with the Top- K method in Tab. 2. The Top- K method is also effective, and when $K = 1$, its results are very similar to Greedy (except for split-0 in the segmentation task). When comparing different K values, we can observe that for the segmentation task, a larger K value can lead to better results. However, simply choosing a higher K value does not necessarily improve performance shown in Figure 3. Particularly in the detection and colorization tasks, increasing the K value results in a significant performance decline. This illustrates a limitation of the Top- K method: it cannot reliably determine an appropriate hyperparameter K . Our proposed Greedy method, on the other hand, avoids this issue by adaptively determining the length of the demonstration set.

We visualize the in-context results of different datasets under various unsupervised prompt selection methods in Fig. 4, which are most distinguishable. Additionally, the association between UnsupPR and Self methods is noticeable. It is evident that some prompts selected by UnsupPR cause the model to find shortcuts, resulting in outputs that have low relevance to the query image but high relevance to the prompt image. For instance, in the coloring task, the outputs of UnsupPR are severely distorted. Consequently, since UnsupPR performs poorly on these samples, Self, which is an enhanced method based on UnsupPR, also performs poorly, resulting in scattered and messy integrated results. These sample results partially reflect the robustness of our method, but also reveal its limitations. The selected prompts do not handle some details well, such as bird legs and wall cracks.

Conclusion

In this paper, based on the observation that most test samples achieve optimal performance under the same prompt, we propose the task-level prompt strategy that significantly reduces inference computational costs. Furthermore, we introduce two train-free demonstration search strategy which can identify a near-optimal combination of demonstrations with minimal computational cost. Comprehensive experiments validate the effectiveness of our proposed method, demonstrating its ability to identify better demonstration combinations at a reduced cost compared to previous methods. These insights hold great promise for the further development and application of VACL, paving the way for more efficient and cost-effective paradigms.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No.62376193)

References

- Baldassini, F. B.; Shukor, M.; Cord, M.; Soulier, L.; and Piwowarski, B. 2024. What Makes Multimodal In-Context Learning Work? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1539–1550.
- Bar, A.; Gandelsman, Y.; Darrell, T.; Globerson, A.; and Efros, A. 2022. Visual prompting via image inpainting. *Advances in Neural Information Processing Systems*, 35: 25005–25017.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, Z.; Zhou, Q.; Shen, Y.; Hong, Y.; Sun, Z.; Gutfreund, D.; and Gan, C. 2024. Visual Chain-of-Thought Prompting for Knowledge-Based Visual Reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1254–1262.
- Cho, H.; Kim, H. J.; Kim, J.; Lee, S.-W.; Lee, S.-g.; Yoo, K. M.; and Kim, T. 2023. Prompt-augmented linear probing: Scaling beyond the limit of few-shot in-context learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12709–12718.
- Dedhia, B.; Chang, M.; Snell, J.; Griffiths, T.; and Jha, N. 2024. Im-promptu: in-context composition from image prompts. *Advances in Neural Information Processing Systems*, 36.
- Everingham, M.; Eslami, S. A.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111: 98–136.
- Huang, Q.; Dong, X.; Chen, D.; Zhang, W.; Wang, F.; Hua, G.; and Yu, N. 2023. Diversity-aware meta visual prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10878–10887.
- Li, F.; Jiang, Q.; Zhang, H.; Ren, T.; Liu, S.; Zou, X.; Xu, H.; Li, H.; Yang, J.; Li, C.; et al. 2024. Visual in-context prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12861–12871.
- Li, K.; Wang, Y.; Li, Y.; Wang, Y.; He, Y.; Wang, L.; and Qiao, Y. 2023. Unmasked teacher: Towards training-efficient video foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19948–19960.
- Li, Z.; Sun, Y.; Zhang, L.; and Tang, J. 2022. CTNet: Context-Based Tandem Network for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12): 9904–9917.
- Li, Z.; Tang, J.; and Mei, T. 2019. Deep Collaborative Embedding for Social Image Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(9): 2070–2083.
- Liu, H.; Tam, D.; Muqeeth, M.; Mohta, J.; Huang, T.; Bansal, M.; and Raffel, C. A. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35: 1950–1965.
- Ma, H.; Zhang, C.; Bian, Y.; Liu, L.; Zhang, Z.; Zhao, P.; Zhang, S.; Fu, H.; Hu, Q.; and Wu, B. 2024a. Fairness-guided few-shot prompting for large language models. *Advances in Neural Information Processing Systems*, 36.
- Ma, K.; Fu, Y.; Cao, C.; Hou, S.; Huang, Y.; and Zheng, D. 2024b. Learning Visual Prompt for Gait Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 593–603.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115: 211–252.
- Shaban, A.; Bansal, S.; Liu, Z.; Essa, I.; and Boots, B. 2017. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*.
- Sun, Y.; Chen, Q.; Wang, J.; Wang, J.; and Li, Z. 2023. Exploring effective factors for improving visual in-context learning. *arXiv preprint arXiv:2304.04748*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wang, X.; Wang, W.; Cao, Y.; Shen, C.; and Huang, T. 2023a. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6830–6839.
- Wang, X.; Zhang, X.; Cao, Y.; Wang, W.; Shen, C.; and Huang, T. 2023b. Seggpt: Segmenting everything in context. *arXiv preprint arXiv:2304.03284*.
- Wang, X.; Zhu, W.; Saxon, M.; Steyvers, M.; and Wang, W. Y. 2024. Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning. *Advances in Neural Information Processing Systems*, 36.
- Workshop, B.; Scao, T. L.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Wu, W.; Sun, Z.; and Ouyang, W. 2023. Revisiting classifier: Transferring vision-language models for video recognition. In *Proceedings of the AAAI conference on artificial intelligence*, 2847–2855.

Yang, Y.; Peng, H.; Shen, Y.; Yang, Y.; Hu, H.; Qiu, L.; Koike, H.; et al. 2024. Imagebrush: Learning visual in-context instructions for exemplar-based image manipulation. *Advances in Neural Information Processing Systems*, 36.

Zhang, J.; Wang, B.; Li, L.; Nakashima, Y.; and Nagahara, H. 2024. Instruct me more! random prompting for visual in-context learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2597–2606.

Zhang, Y.; Zhou, K.; and Liu, Z. 2023. What Makes Good Examples for Visual In-Context Learning? In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 17773–17794. Curran Associates, Inc.