

Enhancing Visual Classification using Comparative Descriptors

Hankyeol Lee¹, Gawon Seo¹, Wonseok Choi¹, Geunyoung Jung¹, Kyungwoo Song², Jiyoung Jung¹

¹Department of Artificial Intelligence, University of Seoul

²Department of Applied Statistics, Yonsei University

{leehk, gawon1224, wonseok.choi, gyjung975, jyjung}@uos.ac.kr, kyungwoo.song@yonsei.ac.kr

Abstract

The performance of vision-language models (VLMs), such as CLIP, in visual classification tasks, has been enhanced by leveraging semantic knowledge from large language models (LLMs), including GPT. Recent studies have shown that in zero-shot classification tasks, descriptors incorporating additional cues, high-level concepts, or even random characters often outperform those using only category names. In many classification tasks, while the top-1 accuracy may be relatively low, the top-5 accuracy is often significantly higher. This gap implies that most misclassifications occur among a few similar classes, highlighting the model’s difficulty in distinguishing between classes with subtle differences. To address this challenge, we introduce a novel concept of comparative descriptors. These descriptors emphasize the unique features of a target class against its most similar classes, enhancing differentiation. By generating and integrating these comparative descriptors into the classification framework, we refine the semantic focus and improve classification accuracy. An additional filtering process ensures that these descriptors are closer to the image embeddings in the CLIP space, further enhancing performance. Our approach demonstrates improved accuracy and robustness in visual classification tasks by addressing the specific challenge of subtle inter-class differences. Code is available at <https://github.com/hklee/Comparative-CLIP>

1. Introduction

The emergence of vision-language models (VLMs), such as CLIP [26], has contributed significantly to the vast progress in the field of visual classification. The standard zero-shot visual classification procedure—computing similarity between the query image and the embedded words for each category [21], then choosing the highest—has shown impressive performance on many popular benchmarks, including ImageNet [5]. Comparing against the word that names a category was a reasonable place to start [26].

While numerous studies have explored VLMs, recent research has shifted towards an innovative visual classification strategy [8, 20, 21, 23, 33] that enhances class representation by leveraging the semantic knowledge from large language models (LLMs) (e.g., GPT-3 [3]). Specifically, GPT-descriptor-extended CLIP (DCLIP) [21] diverges from the existing method that relied solely on class names, e.g., “A photo of a **Golden Retriever**.” by incorporating additional information, e.g., “A photo of a Golden Retriever, **which has golden fur**.” termed a descriptor. Employing such descriptors with additional cues in the class has led to improvements in classification performance as well as some explainability of the model. One can understand how the model classifies an image by checking the similarity of the image embedding with each descriptor. However, this approach has limitations mainly due to failure in descriptor creation. For example, GPT-3 occasionally produces descriptors about taste and smell in addition to vision, as well as descriptors with ambiguous words having multiple meanings, synthetic errors, or repetitive words [21]. These descriptors diminish the image classification performance.

An interesting approach named WaffleCLIP [32] substitutes these LLM-generated descriptors with either random words e.g., “foot loud” or sequence of random characters e.g., “jmhj, !J#m”, which have absolutely no semantic relation to class names. Without any external models, classification accuracy remained comparable to that of semantically meaningful descriptors. Moreover, when high-level concepts are added, e.g., “A photo of **an animal**: a Golden Retriever, which has jmhj, !J#m.”, the performance substantially exceeds the previous descriptor-based approach (DCLIP) [21]. This progression consequently triggered a reconsideration of the benefits of additional semantics introduced by LLM-generated descriptors [32]. However, this approach has an inherent downside: the complete loss of explainability. While using LLM-generated descriptors improved classification accuracy and provided insights into the model’s decision-making process [21], employing random descriptors to enhance performance significantly diminishes the interpretability of the model’s decisions.

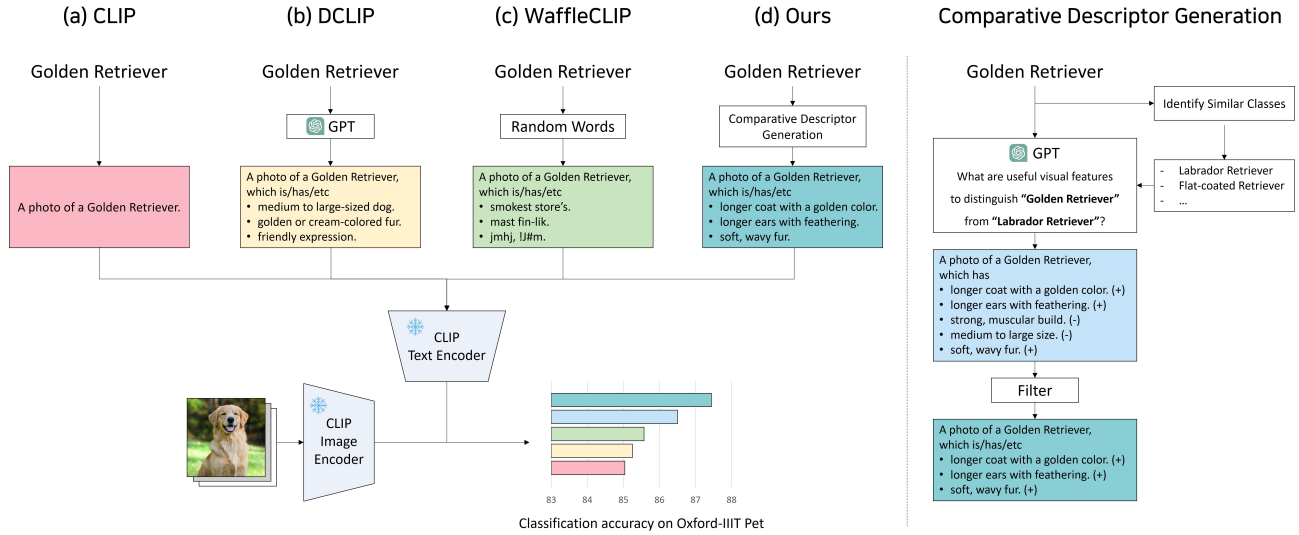


Figure 1. **Overview of our method and the baselines**(CLIP [26], GPT-descriptor-extended CLIP (DCLIP) [21], and WaffleCLIP [32]). (left) Comparison of our method with the baselines. Our proposed method outperforms the baselines on CLIP’s image classification task. (right) A detailed overview of our method. Our method generates descriptors through comparison with semantically similar classes. Then the filtering process is applied to retain only descriptors that are useful for classification, significantly increasing the classification accuracy.

During our analysis of existing experiments across various visual classification tasks, we found an inspiring phenomenon: as shown in Table 1, while top-1 accuracy may be relatively low, top-5 accuracy is often significantly higher. For example, in zero-shot classification on the Caltech-UCSD Birds 200 dataset [39], a notable difference is observed: top-1 accuracy stands at 51.36%, whereas top-5 accuracy is 83.48%. This gap implies that most misclassifications occur among a few similar classes out of the hundreds of classes in the dataset. It also indicates that the model has difficulty distinguishing between classes with subtle differences, often confusing classes with common characteristics. Based on this observation, we hypothesized that generating descriptors that emphasize the differences between the target class and its similar, frequently confused classes could enhance classification accuracy.

In this paper, we propose a novel approach to enhance the image classification performance of VLMs by generating *comparative descriptors* using LLMs. Our approach consists of two steps. Rather than directly prompting LLMs to generate descriptions for a specific class, we identify n classes that are semantically similar in advance, by measuring the distances between the text features of each class name. We then request the LLMs to generate descriptions for the target class comparatively, highlighting its distinct features relative to these semantically similar classes. For example, we query GPT, “What are useful features for distinguishing a {target class} from a {similar class} in a photo?” to generate comparative descriptors.

Next, we propose the following descriptor selection process, which we call the filtering process. This process is motivated by the observation that some descriptors generated by LLMs do not contribute to the classification and can even harm performance. Since VLMs and LLMs have different knowledge bases, some descriptors generated by LLMs may not be useful for classification in VLMs. Intuitively, removing these unhelpful descriptors would enhance classification accuracy. Specifically, we retain only the top- k descriptors that exceed a certain amount of similarity to the mean image feature of each class. These remaining k comparative descriptors effectively represent the target classes, distinguishing them well from their similar classes, thereby significantly enhancing the classification performance of VLMs.

In summary, our contributions are as follows: 1) We present a novel concept of comparative descriptors, which are visual descriptions emphasizing differences and unique features of the target class compared to semantically similar classes. This approach has reduced issues like ambiguity and the difficulty of understanding modalities, which are common problems associated with existing LLM-based descriptor generation methods. 2) We propose a simple yet effective filtering process that retains only class-specific descriptors contributing to image classification. By measuring the similarity of each descriptor to the mean image feature of its respective class, we ensure that only the most relevant and distinguishing descriptors are preserved. 3) Our method significantly improves the image classification performance of VLMs (particularly CLIP) across various datasets while

still preserving the inherent interpretability of the model’s decision.

Note that the classification leveraging comparative descriptors operates in a zero-shot setting, whereas the additional filtering process applies to a few-shot setting, as it involves calculating the mean image feature using a few examples for each class. In the experiments section, we independently evaluate the impact of both the comparative descriptors and the filtering process on enhancing the classification performance of VLMs, demonstrating their respective contributions.

2. Related Work

2.1. Vision-Language Models (VLMs)

Vision-language models (VLMs) [14–16, 26, 30, 34] are multimodal architectures designed to process and integrate visual and linguistic information simultaneously. These models are capable of understanding both visual and textual data concurrently, learning the interactions between these two forms of information. Motivated by the achievements of large language models (LLMs) like BERT [6] and GPT [27], researchers have begun developing large-scale VLMs. These models, pre-trained on extensive datasets containing both images and text, show exceptional performance across various tasks. Specifically, CLIP [26] sets a significant milestone in the advancement of VLMs. CLIP learns the general correlations between images and text, showcasing remarkable flexibility and achieving superior performance in a wide array of visual tasks. Notably, it has reached state-of-the-art results in zero-shot classification, demonstrating its ability to accurately recognize images from unseen categories. This highlights CLIP’s proficiency in understanding and categorizing a broad spectrum of visual concepts, extending beyond predefined categories.

2.2. Large Language Models (LLMs)

Large language models (LLMs) [3, 6, 17, 27–29] have become foundational in the field of natural language processing. These powerful models excel at detecting patterns in vast amounts of text data, enabling them to tackle a diverse range of language-related tasks such as understanding language [6], generating text [28], translating languages [36], and summarizing texts [29]. The development of LLMs has significantly influenced VLM research. Leveraging the language comprehension capabilities and the knowledge learned from text data inherent in LLMs, VLMs have improved their ability to process and interpret the intricate relationships between images and text. This advancement has enhanced multimodal learning [18], facilitated the integration and expansion of knowledge across visual and textual domains [35], and led to the efficient integration and fine-tuning of VLMs using pre-trained LLMs [43].

2.3. Visual Classification of VLMs

Zero-shot classification [31, 40] is a task that allows a model to classify new classes or objects that were not seen during the training process. Recent studies leverage vision-language models (VLMs) for conducting zero-shot classification tasks. Specifically, there is research focused on utilizing both class names and additional information in termed descriptors. It is noteworthy that performance improves when using large language models (LLMs) to generate descriptors [8, 20, 21] or even employing random strings as descriptors [32]. On the other hand, few-shot classification [1, 37, 38, 41] is another branch that classifies new classes with minimal training data. This method is useful for real-world problems with limited data or labeling difficulties [25]. Particularly, training-free few-shot classification [42] is a type of few-shot learning technique that does not require training. Although it requires no training, it shows significant performance improvements compared to zero-shot approaches by leveraging a few image-label pairs from the dataset.

3. Method

We first discuss image classification in CLIP and classification when using descriptors in Section 3.1. Next, We describe how to generate comparative descriptors in Section 3.2. Specifically, to proactively identify classes likely to be misclassified, we propose an efficient method for detecting semantically similar classes. These similar classes can substitute for incorrect classes resulting from misclassification. We then generate a descriptor by querying the LLM with a comparison between the target class and similar classes. Lastly, Section 3.3 covers the filtering process to keep only descriptors that contribute to classification.

3.1. Image Classification with Descriptors

Given an image x and target categories C , the visual classification procedure of CLIP [26] is defined as,

$$\tilde{c} = \arg \max_{c \in C} s(\phi_I(x), \phi_T(f_c)) \quad (1)$$

where f_c denotes the prompt, e.g., “A photo of a Golden Retriever.”, ϕ_I and ϕ_T represent the image and text encoders respectively, and $s(\cdot, \cdot)$ represents the similarity score or the distance between two feature vectors calculated by the dot product.

To improve classification performance, we add a class descriptor [21] to the prompt f_c as, “A photo of a Golden Retriever, which has golden fur.” Given a set of descriptors D_c for each class c , classification procedure is reformulated as,

$$\tilde{c} = \arg \max_{c \in C} \frac{1}{|D_c|} \sum_{d \in D_c} s(\phi_I(x), \phi_T(d)) \quad (2)$$

ViT-B/32	IN1k	IN1k-V2	Caltech	CIFAR	CUB	SAT	Places	Food	Pets	DTD	Flowers	Aircraft	Cars	Avg
Top-1	62.04	54.67	78.39	64.31	51.36	40.89	39.11	82.57	85.04	43.19	62.97	24.96	58.62	57.55
Top-5	87.67	81.97	91.31	88.38	83.48	90.91	69.71	96.89	97.36	73.62	84.42	63.43	89.84	84.54

Table 1. **The top-1 and top-5 accuracy for image classification on various datasets using the CLIP ViT-B/32 model.** Even if the top-1 accuracy is relatively low, the top-5 accuracy is significantly high. This indicates that while the dataset contains many classes, the model can effectively capture coarse features but struggles to distinguish subtle differences between similar classes.

3.2. Generating Comparative Descriptors

3.2.1 Identifying similar classes

To reduce misclassification, we identify classes similar to the target class by calculating the cosine similarity between the text features of their class names. The reasoning behind this approach is that CLIP has knowledge of both modalities, vision, and language. Therefore, we assume that, when given only text, it can leverage its inherent visual knowledge to find similar classes. The cosine similarity between the text features $\phi_T(f_i)$ and $\phi_T(f_j)$ of the classes f_i and f_j respectively is given by:

$$\cos(\phi_T(f_i), \phi_T(f_j)) = \frac{\phi_T(f_i) \cdot \phi_T(f_j)}{\|\phi_T(f_i)\| \|\phi_T(f_j)\|} \quad \forall i, j \in C \quad (3)$$

where $\|\cdot\|$ denotes the Euclidean norm of the vector. We select n similar classes having high cosine similarities with the target class.

3.2.2 Querying LLMs

Next, our goal is to generate comparative descriptors which include distinguishing attributes between the target class and its semantically similar classes. The descriptors are generated using GPT-4o [10] with the prompts for the queries as follows:

Q: What are useful features for distinguishing a {target class} from a {similar class} in the photo?

A: There are several useful visual features to tell the photo is a {target class}, not a {similar class}.

The proposed approach addresses the limitations associated with LLM-generated descriptors, such as challenges in modality comprehension and word ambiguity, effectively reducing their impact. This reduces the occurrence of misclassifications and thus improves the overall accuracy of classification tasks.

To effectively utilize in-context learning techniques with LLMs, we prepared ten sets of question-and-answer examples, following a consistent prompt template as shown above. For each instance of generating comparative descriptors, we randomly choose two sets from these ten. We arbitrarily select the target and similar classes from our datasets

to create comparative descriptors using ChatGPT. This dynamic selection process, rather than using a fixed set of examples, introduces a broader range of contexts to the model.

3.3. Filtering

In the final stage, we filter the generated descriptors to include only those that contribute to classification. Initially, we compute the mean image feature for each class by selecting a few images of the class from the dataset, making our method applicable even in scenarios with limited image availability. Note that the mean image feature is utilized solely during this filtering phase. We calculate the cosine similarity between the mean image feature of each class and every descriptor, deriving similarity scores. From these, we retain only the top- k descriptors based on their scores, discarding the rest.

We employ a lower bound for similarity scores in this process, set as the cosine similarity between the mean image feature and a CLIP-style text prompt (e.g., “A photo of a {class}.”). Descriptors falling below this threshold are excluded. If no descriptors for a class exceed this lower bound, classification proceeds using only the CLIP-style text prompt, without additional descriptors. For detailed steps of this filtering process, refer to the pseudo-code presented in Algorithm 1.

Algorithm 1 Filtering Process

```

1: for  $c \in \text{classes}$  do
2:    $Lower\ Bound \leftarrow \min(\cos(\bar{\phi}_{Ic}, \phi_T(f_c)), 0.3)$ 
3:   for  $d \in D_c$  do
4:      $score \leftarrow \cos(\bar{\phi}_{Ic}, \phi_T(d))$ 
5:     if  $score < Lower\ Bound$  then
6:       Discard( $d$ )
7:     end if
8:   end for
9:    $D_c \leftarrow \text{Select}(D_c, k)$ 
10:  if  $D_c$  is empty then
11:     $D_c \leftarrow f_c$ 
12:  end if
13: end for

```

ViT-B/32	IN1k	IN1k-V2	Caltech	CIFAR	CUB	SAT	Places	Food	Pets	DTD	Flowers	Aircraft	Cars	Avg
CLIP	62.04	54.67	78.39	64.31	51.36	40.89	39.11	82.57	85.04	43.19	62.97	24.96	58.62	57.55
DCLIP	63.66	56.28	81.23	64.94	53.61	41.37	41.64	83.06	85.25	44.26	66.63	26.67	59.08	59.05
WaffleCLIP	63.32	55.97	81.06	65.44	52.47	43.66	40.67	82.84	85.48	42.93	66.32	25.70	58.82	58.82
Ours	64.02	56.71	82.24	65.69	54.09	42.92	42.22	83.81	87.41	46.54	67.07	27.57	59.30	59.97
Ours + Filtering	65.66	57.56	83.55	64.81	56.41	51.08	44.18	84.62	87.04	53.78	73.43	28.61	60.15	62.37

Table 2. **Results of the image classification.** Comparison of classification results using our method with CLIP [26], DCLIP [21], and WaffleCLIP [32]. Our method outperforms the baseline performances across most evaluated datasets. Results for ViT-L/14 and ResNet50 are available in the supplementary material.

ViT-B/32	IN1k	IN1k-V2	Caltech	CIFAR	CUB	SAT	Places	Food	Pets	DTD	Flowers	Aircraft	Cars	Avg
DCLIP	22.53	19.69	48.92	35.04	4.45	38.43	26.22	37.03	17.06	25.48	12.67	7.98	13.17	23.74
Ours	31.57	27.54	60.60	41.07	18.62	46.17	33.37	51.54	35.76	35.69	25.09	14.97	28.73	34.67

Table 3. **Image classification results without class labels.** To evaluate the quality of the descriptor itself, we perform image classification after removing class labels and prefixes, using only the descriptor as text input. In this experiment, our method outperforms the existing methods. This means that our descriptors contribute significantly to the improvement of classification performance.

4. Experiments

First, we provide experimental details in Section 4.1. We then assess the effectiveness of our method through image classification in Section 4.2, demonstrating its superior performance compared to existing methods. Next, in Section 4.3, we perform image classification using only descriptors as text input and show that our method generates high-quality descriptors. Additionally, we show that our method is not solely dependent on the number of descriptors by comparing it to the baseline after equalizing the number of descriptors in Section 4.4. We propose and evaluate a few-shot filtering process in Section 4.5, highlighting its practical utility even with limited data. Also, we apply our filtering process to the existing method to show the effectiveness of this process. We show the explainability, which is an advantage of using the LLM-generated descriptor in Section 4.6. Finally, we discuss some limitations of our approach in Section 4.7, providing a foundation for future research and improvements.

4.1. Experimental Details

In all experiments, we utilize CLIP [26] as the underlying VLM. Furthermore, since our method is training-free, all experiments were conducted using a single NVIDIA 3090 GPU.

4.1.1 Hyperparameters

In our experiment, we use two hyperparameters: the number of similar classes per class, denoted as n , and the maximum number of descriptors to retain after filtering, denoted as k . We empirically choose one of the values between 5 and 20 for both n and k . The hyperparameters used for each dataset are shown in supplementary material.

4.1.2 GPT

In the baselines [21, 32], GPT-3 [3] text-davinci-003 was employed as the LLM to generate descriptors, but this version has been deprecated. Therefore, to ensure a fair comparison, we re-conducted all experiments for both the baselines and our method using GPT-4o [10]. Note that we did not change any of the baselines’ settings, except for the GPT version. To demonstrate the fairness of our experiments, we disclose all responses.

4.1.3 Datasets

We conduct experiments on 13 datasets, including the 11 datasets tested in the baselines [21, 32]: ImageNet (IN1k) [5], ImageNetV2 (IN1k-V2) [11], CUB200-2011 (CUB) [39], EuroSAT (SAT) [9], Places365 (Places) [44], Food101 (Food) [2], Oxford-IIIT Pet (Pets) [24], Describable Textures (DTD) [4], Flowers102 (Flowers) [22], FGV-Caircraft (Aircraft) [19], and Stanford Cars (Cars) [12], along with 2 additional datasets: Caltech256 (Caltech) [7] and CIFAR100 (CIFAR) [13]. Note that ImageNetV2 is an expansion of the validation set designed to evaluate the generalization ability of ImageNet, sharing the same set of classes and not including a new training set. Therefore, we use the mean image feature of ImageNet during the filtering process for ImageNetV2.

4.2. Image Classification

We assess the effectiveness of our method by performing image classification on various datasets, each of which contains a wide range of domains. We then compare these results to those obtained by existing methods, as shown in Table 2.

Since our method generates descriptors by comparing similar classes, it is straightforward to perform well on domain-specific datasets such as Flowers [22] and Cars [12] which contain only images of flowers and cars, respectively. However, its performance on multi-domain datasets might be questionable. To address this concern, we present additional experimental results on two multi-domain datasets: Caltech [7] and CIFAR [13], which were not covered in the baselines [21, 32]. Our method surpasses the baselines across most evaluated datasets, irrespective of the domain. This superior performance demonstrates the validity and robustness of our approach.

In particular, the classification accuracy on the Describable Texture Dataset (DTD) [4] was significantly improved compared to existing methods. When we generate descriptors using the existing method [21], we frequently encounter failures due to the ambiguity of words. This leads to the production of descriptors irrelevant to the intended class and lacking sufficient detail. Specifically, with GPT-4o [10], descriptors are sometimes not created at all if the class is ambiguous. In contrast, we can make precise and relevant queries by generating descriptors using our comparison-based method. This approach reduces the occurrence of errors and produces context-rich descriptors, ultimately increasing classification accuracy. To offer further insight into these results, we compare the descriptors generated by DCLIP and our approach. See Figure 2 for examples.

4.3. Image Classification with Descriptors Only

To evaluate the quality of the descriptor, we modify the form of the input text. Instead of using the format “A photo of a {class}, which is/has/etc {descriptor}.” in Section 4.2, we remove the class name and prefix, using only the descriptor as the text input, *i.e.* “{descriptor}”, for classification. The results of this experiment are shown in Table 3. Our method achieves a much higher accuracy for all evaluated datasets than the existing method.

One of the reasons for these results can be found in the reduction of common attributes in the descriptors. The descriptors generated by the existing method are likely to contain common attributes. In contrast, our method generates descriptors by comparison, minimizing the generation of common attributes and allowing us to obtain distinct attributes. To be more specific, no matter how well a descriptor describes a class, if a semantically analogous descriptor exists for another class, this descriptor will not help distinguish between them. For better understanding, we present the examples in Figure 3. Consequently, this result demonstrates that the descriptors generated by our method have a much higher contribution to the classification.

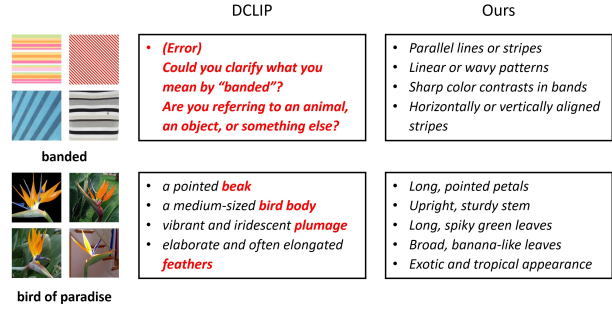


Figure 2. **Addressing ambiguity in descriptor generation.** We compare the descriptors generated by our method and the DCLIP [21] method on the (top) Describable Texture Dataset (DTD) [4] and (bottom) Flowers102 (Flowers) [22] dataset. DCLIP method failed to generate descriptors due to ambiguity (*e.g.* the word *banded* refers to both texture and a species of snake), and in some cases generated descriptors that were unrelated to the class. On the other hand, our method not only avoided failures but also enriched in context. This difference leads to a significant improvement in classification accuracy.

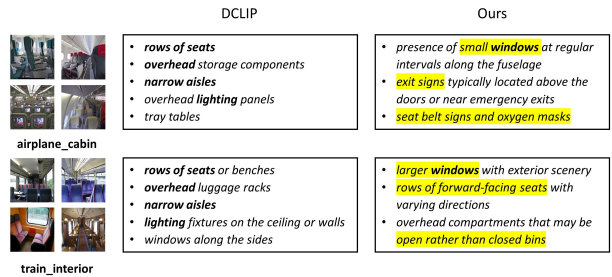


Figure 3. **Comparison of generated descriptors between similar classes.** Descriptors were generated for similar classes in the Places dataset. Bold text indicates an attribute that appears in both classes and highlighted text indicates a distinct attribute. The descriptors generated by DCLIP have semantically equivalent descriptors between the target class and its similar class. In contrast, our method minimizes semantically analogous descriptors, adds distinct features, and provides more detailed explanations. As a result, this makes it easier to distinguish between similar classes.

ViT-B/32	Random (k=5)	Filtering (k=5)
DCLIP	58.31	59.68
Ours	58.66	61.97

Table 4. **Image classification results with an equal number of descriptors.** We reduce the number of descriptors in two ways to make them equal, and both cases show better performance when using our descriptors. This demonstrates that our method does not depend solely on the number of descriptors.

Shot	IN1k	IN1k-V2	Caltech	CIFAR	CUB	SAT	Places	Food	Pets	DTD	Flowers	Aircraft	Cars	Avg
1	63.97	56.23	82.67	64.38	54.59	48.88	41.91	84.03	86.56	50.94	72.34	28.03	56.92	60.88
2	64.54	56.59	83.16	64.45	55.41	50.78	42.56	84.25	86.77	52.09	72.80	28.03	58.53	61.54
4	65.10	57.08	83.26	64.57	56.11	50.49	43.29	84.34	86.93	52.88	72.97	28.34	59.29	61.90
8	65.36	57.33	83.49	64.78	56.36	50.39	43.79	84.44	86.82	53.33	73.43	28.61	59.85	62.15
16	65.52	57.44	83.43	64.76	56.41	51.05	44.09	84.51	87.04	53.70	-	28.59	60.15	-
32	65.65	57.56	83.55	64.81	-	51.08	44.16	84.52	86.92	53.78	-	28.48	-	-
64	65.66	57.55	-	64.71	-	50.93	44.18	84.62	86.88	-	-	-	-	-

Table 5. **Image classification results with few-shot filtering applied.** Our filtering method shows superior results even when using a few images. This indicates that our method can be useful even in situations with limited data.

ViT-B/32	IN1k	IN1k-V2	Caltech	CIFAR	CUB	SAT	Places	Food	Pets	DTD	Flowers	Aircraft	Cars	Avg
DCLIP	63.66	56.28	81.23	64.94	53.61	41.37	41.64	83.06	85.25	44.26	66.63	26.67	59.08	59.05
+ Filtering	64.29	56.42	81.22	63.83	54.24	42.85	42.57	83.61	85.03	46.06	67.87	28.20	59.59	59.68

Table 6. **Image classification results when filtering is applied to the existing method [21].** Performance improved in most datasets. This demonstrates that our proposed filtering method is effective not only for descriptors generated by our method but also for those generated by various other methods.

4.4. Evaluation with Equal Number of Descriptors

WaffleCLIP [32] pointed out that the number of descriptors can affect image classification accuracy. Our method identifies n similar classes for one target class and appends all the descriptors generated from n comparisons, it inevitably uses more descriptors than the baselines. Since it may be questionable that the performance improvement with our method is due to the increased number of descriptors, we perform an additional experiment with the same number of descriptors in the baseline and our approach.

We try both random selection and filtering to equalize the number of descriptors. The number of descriptors, k , is set to 5 for both approaches. We run both five times with different seeds or different images, and the average accuracy is shown in Table 4. Our method performs the baseline on average for both approaches. This demonstrates that the performance of our method is not solely dependent on the number of descriptors, but also indicates that it possesses better quality on its own.

4.5. Few-Shot Filtering

We propose a simple, training-free filtering process to retain only descriptors close to the given class and contribute to the classification. We present the classification accuracy obtained through this filtering process in the last row of Table 2, and show the variation in performance depending on the number of images used in Table 5. We repeat the experiment five times using different images to get reliable results and present the average accuracy. For most of the evaluated datasets, filtering using only one or two images also results in significant performance improvements. This demonstrates that our filtering can be used practically even in limited data situations.

Additionally, to demonstrate the validity of our proposed filtering process, we apply it to descriptors generated with DCLIP [21] and compare the results before and after filtering. As shown in Table 6, the use of filtering generally leads to an increase in accuracy.

However, the performance improvement is relatively low compared to applying filtering to the descriptors produced by our method, and for some datasets, we even observe a slight decrease in accuracy. This is because there are fewer descriptors available for filtering. Our filtering method aims to identify descriptors that are closely related to the class and useful for classification from a large pool of descriptors. If the number of these descriptors is limited, the filtering may be less effective.

4.6. Explainability

Our approach leverages descriptors generated by LLMs for image classification. Similar to DCLIP [21], our method not only allows the model to make precise decisions but also provides an explanation of those decisions and the reasons behind them. This capability significantly enhances the explainability of the model and provides deeper insights into its operational process.

As shown in Figure 4, we analyze by selecting a class and its similar counterpart from the dataset. Despite their resemblance, we find that the similarity scores of the target and similar classes show a huge gap. This gap enables the precise classification between the images of similar classes. Moreover, by comparing similarity scores, we can identify which features in the image contribute more to classification. This provides a clearer understanding of the reasoning process and enhances the explainability of our model.




Figure 4. **Examples of explainability.** (left) We show examples of decisions and justifications through our model. (right) We present the descriptors of each target class and the similar class within the datasets (from the top, Places365 [44], Food101 [2], Flowers102 [22]). The chart shows the similarity score between the descriptor and the image in the CLIP latent space. The higher the similarity score, the greater the influence of the descriptor on the decision-making process. The descriptor with the highest average similarity becomes the model’s classification prediction.

4.7. Limitations

In this section, we provide a detailed analysis of the limitations involved in our approach. This analysis will be useful in guiding future improvements and research directions.

We encounter an interesting problem while generating descriptors with LLMs for certain datasets. FGVC Aircraft (Aircraft) [19] and Stanford Cars (Cars) [12] are datasets that contain different types of aircraft and cars, respectively. When querying on these datasets, we find some attributes that contain visual information but are invisible. Think of the seating arrangement in aircraft or the engine in cars: these are visual components, but they are hard to observe due to the limited availability of images of aircraft and car interiors in the datasets. This resulted in less improvement in classification accuracy on these datasets compared to other datasets.

During our analysis of the experimental results, we discovered that the use of filtering can lead to a lack of diversity in the descriptors. More specifically, this phenomenon means that some filtered descriptors are semantically almost identical, but only slightly different in form (e.g., “short, rounded ear” and “small, rounded ears”). We assume that this phenomenon occurs because our filtering process only considers similarity scores. If an attribute has the highest similarity to the mean image feature of the class, it is reasonable to expect that a descriptor with a similar meaning and almost identical form would also have a high similarity. We will address this issue in future work by considering more factors in the filtering process.

5. Conclusion

In this work, we presented an innovative strategy for enhancing the image classification capabilities of vision-language models (VLMs) by generating comparative descriptors. By comparing the target class with semantically similar classes, our method effectively minimizes ambiguity and enhances the model’s ability to distinguish closely related classes. Coupled with a filtering process, our approach substantially boosts classification accuracy across various datasets, demonstrating notable improvements over existing baselines, even with limited image-label pairs. Integrating comparative descriptors with our filtering mechanism significantly advances image classification performance. This novel method not only improves the performance metrics of VLMs, particularly in models like CLIP [26] but also enriches the interpretability of the results, offering a more intuitive understanding of the model’s decision-making process.

Acknowledgments

This work was supported by the National Research Foundation of Korea Grant funded by the Korean Government (MSIT) (NRF-2020R1C1C1008726, NRF-2022R1A4A3033874). Also, this work was supported by the 2024 Advanced Facility Fund of the University of Seoul for Jiyoung Jung.

References

- [1] P. Bateni, R. Goyal, V. Masrani, F. Wood, and L. Sigal. Improved few-shot visual classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14481–14490. IEEE Computer Society, 2020. [3](#)
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 - mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. [5](#), [8](#)
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Matusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. [1](#), [3](#), [5](#)
- [4] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014. [5](#), [6](#)
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [1](#), [5](#)
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. [3](#)
- [7] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. Technical report, California Institute of Technology, 2007. [5](#), [6](#)
- [8] Songhao Han, Le Zhuo, Yue Liao, and Si Liu. LLMs as visual explainers: Advancing image classification with evolving visual descriptions. *ArXiv*, 2311.11904, 2024. [1](#), [3](#)
- [9] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. [5](#)
- [10] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. [4](#), [5](#), [6](#)
- [11] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do better imagenet models transfer better? *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2656–2666, 2018. [5](#)
- [12] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013. [5](#), [6](#), [8](#)
- [13] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. [5](#), [6](#)
- [14] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. Proceedings of Machine Learning Research, 2023. [3](#)
- [15] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 12888–12900. Proceedings of Machine Learning Research, 2022. [3](#)
- [16] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [3](#)
- [17] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *Computing Research Repository*, abs/1907.11692, 2019. [3](#)
- [18] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. VILBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. [3](#)
- [19] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *ArXiv*, abs/1306.5151, 2013. [5](#), [8](#)
- [20] Mayug Maniparambil, Chris Vorster, Derek Molloy, Noel Murphy, Kevin McGuinness, and Noel E. O’Connor. Enhancing clip with gpt-4: Harnessing visual descriptions as prompts. *2023 IEEE/CVF International Conference on Computer Vision Workshops*, pages 262–271, 2023. [1](#), [3](#)
- [21] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *The Eleventh International Conference on Learning Representations*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [22] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729, 2008. [5](#), [6](#), [8](#)
- [23] Shubham Parashar, Zhiqiu Lin, Yanan Li, and Shu Kong. Prompting scientific names for zero-shot species recognition. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9856–9861, Singapore, 2023. Association for Computational Linguistics. [1](#)
- [24] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. *2012 IEEE Conference*

- on *Computer Vision and Pattern Recognition*, pages 3498–3505, 2012. 5
- [25] Quan Quan, Qingsong Yao, Jun Li, and S. Kevin Zhou. Which images to label for few-shot medical landmark detection? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20606–20616, 2022. 3
- [26] 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. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. Proceedings of Machine Learning Research, 2021. 1, 2, 3, 5, 8
- [27] Alec Radford and Karthik Narasimhan. Improving language understanding by generative pre-training. 2018. 3
- [28] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. 3
- [29] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Computing Research Repository*, abs/1910.10683, 2019. 3
- [30] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. Proceedings of Machine Learning Research, 2021. 3
- [31] Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2152–2161. Proceedings of Machine Learning Research, 2015. 3
- [32] Karsten Roth, Jae-Myung Kim, A. Sophia Koepke, Oriol Vinyals, Cordelia Schmid, and Zeynep Akata. Waffling around for performance: Visual classification with random words and broad concepts. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 15700–15711. International Conference on Computer Vision, 2023. 1, 2, 3, 5, 6, 7
- [33] Oindrila Saha, Grant Van Horn, and Subhansu Maji. Improved zero-shot classification by adapting vlms with text descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1
- [34] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020. 3
- [35] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111. Association for Computational Linguistics, 2019. 3
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 3
- [37] Wai Keen Vong and Brenden M. Lake. Few-shot image classification by generating natural language rules. In *ACL Workshop on Learning with Natural Language Supervision*, 2022. 3
- [38] Bowen Wang, Liangzhi Li, Manisha Verma, Yuta Nakashima, Ryo Kawasaki, and Hajime Nagahara. Mtunet: Few-shot image classification with visual explanations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop*, 2021. 3
- [39] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. Technical Report CNS-TR-201, Caltech, 2010. 2, 5
- [40] Meng Ye and Yuhong Guo. Zero-shot classification with discriminative semantic representation learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5103–5111, 2017. 3
- [41] Chunjie Zhang, Chenghua Li, and Jian Cheng. Few-shot visual classification using image pairs with binary transformation. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(9):2867–2871, 2020. 3
- [42] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of CLIP for few-shot classification. In *Computer Vision – ECCV 2022*, pages 493–510. Springer Nature Switzerland, 2022. 3
- [43] Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. Adapting language models for zero-shot learning by meta-tuning on dataset and prompt collections. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2856–2878. Association for Computational Linguistics, 2021. 3
- [44] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2018. 5, 8