

ViPCap: Retrieval Text-Based Visual Prompts for Lightweight Image Captioning

Taewhan Kim, Soeun Lee, Si-Woo Kim, Dong-Jin Kim*

Hanyang University, South Korea
{taewhan, soeun, boreng0817, djdkim}@hanyang.ac.kr

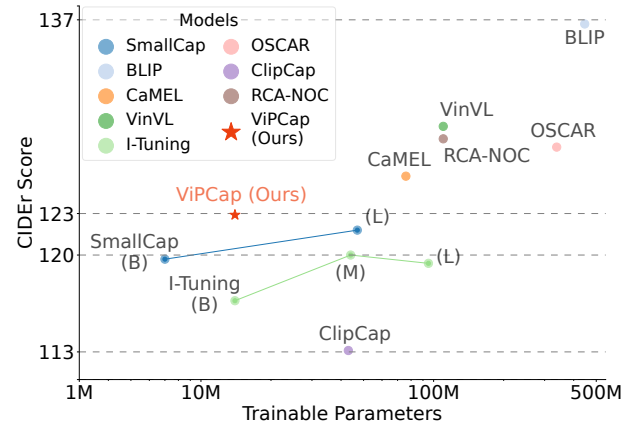
Abstract

Recent lightweight image captioning models using retrieved data mainly focus on text prompts. However, previous works only utilize the retrieved text as text prompts, and the visual information relies only on the CLIP visual embedding. Because of this issue, there is a limitation that the image descriptions inherent in the prompt are not sufficiently reflected in the visual embedding space. To tackle this issue, we propose ViPCap, a novel retrieval text-based visual prompt for lightweight image captioning. ViPCap leverages the retrieved text with image information as visual prompts to enhance the ability of the model to capture relevant visual information. By mapping text prompts into the CLIP space and generating multiple randomized Gaussian distributions, our method leverages sampling to explore randomly augmented distributions and effectively retrieves the semantic features that contain image information. These retrieved features are integrated into the image and designated as the visual prompt, leading to performance improvements on the datasets such as COCO, Flickr30k, and NoCaps. Experimental results demonstrate that ViPCap significantly outperforms prior lightweight captioning models in efficiency and effectiveness, demonstrating the potential for a plug-and-play solution.


Introduction

Vision and language (V&L) tasks, such as image captioning, have advanced with large-scale models like SimVLM (1.4B), PaLi (3B), and REVEAL (2.1B) (Wang et al. 2022a; Chen et al. 2023; Hu et al. 2023). However, these advanced multimodal models require a large number of parameters, resulting in high computational and dataset construction costs.

To enhance training efficiency, recent works suggest focusing on learning a mapping network, such as MAPL and BLIP-2 (Mañas et al. 2023; Li et al. 2023a), that bridges the modality gap between images and text or training learnable tokens like EVCap (Li et al. 2024). For example, BLIP-2 introduces the Qformer, which aligns the two modalities while keeping a V&L models frozen. Nevertheless, despite the efficiency of this approach, BLIP-2 still requires over 1B trainable parameters. Although EVCap trains learnable tokens with few trainable parameters, EVCap and MAPL



(a) Comparison of trainable parameter sizes and CIDEr scores for each model. B, M, and L denote the Base, Medium, and Large models, respectively.



(b) Comparison of total parameters between ViPCap (Ours) and other models. SC-B denotes SmallCap_{Base} model, SC-L represents the SmallCap_{Large} model.

Figure 1: (a) ViPCap shows the best efficiency among the lightweight captioning models. (b) EVCap and MAPL require large models such as EVA-CLIP (Sun et al. 2023), Vicuna (Chiang et al. 2023), and GPT-J (Wang and Komatsuzaki 2021). With under 0.3B total parameters, ViPCap achieves competitive performance despite its small size.

*Corresponding Author.



Retrieval: a brown dog with big eyes on a **chair**
 a brown dog peeks over the edge of a table
 a beagle dog looking innocent standing by a fence

GT : a dog is peaking its head behind a **chair**

SmallCap: a beagle is peeking out of a **window**

ViPCap : a dog peeking out from under a wooden **chair**

Figure 2: SmallCap (Ramos et al. 2023) fails to accurately represent visual information, such as a local object, in the ground truth (GT) or retrieval text (Retrieval). In contrast, our ViPCap accurately captures visual information found in GT or retrieval text.

total parameters as shown in Fig. 1b.

Recent models like SmallCap, EXTRA, and Re-ViLM (Ramos et al. 2023; Ramos, Elliott, and Martins 2023; Yang et al. 2023) are emerging to reduce computational costs using external knowledge. These models retrieve texts semantically similar to input images and use them as only text prompts. However, there is a limitation in that visual information relies only on the CLIP vision encoder. As depicted in Fig. 2, SmallCap, which uses retrieval caption as a text prompt without visual prompts, encounters an issue where it can not contain detailed visual descriptions in a caption. We suspect this is because the image description in text prompt is not utilized as visual information.

In this paper, we propose a retrieval text-based visual prompt for lightweight image captioning (**ViPCap**), leveraging retrieved texts describing image information as *visual prompts*. First, given that the retrieved text provides a comprehensive image description, we encode the text prompt into the CLIP embedding space and transform it into patch-level hidden representations to extract semantic information. To effectively enhance local visual representations using a global text representation, we randomly augment the semantic representation from the text prompt. In particular, we assume that the embedding vector follows a randomized Gaussian distribution and extract M semantic features from this distribution as the basis for a visual prompt.

Previous works, including CapDec and LinCIR (Nukrai, Mokady, and Globerson 2022; Gu et al. 2024), address the modality gap in V&L tasks using Gaussian distributions. Based on this approach, we propose modeling text features following a randomized Gaussian distribution. Unlike the heuristic approaches like CapDec, our approach generates semantic features sampled from a learnable distribution, aiming for a high correlation with visual features. We as-

sume these semantic features contain visual information and expect them to closely resemble the input image features. To achieve this, we employ a patch retrieval module that aligns semantic features with each input image patch.

The retrieved patch features are combined with the input image features to generate the *visual prompt* is added with the image features before decoder input. This approach aims to enhance the model’s ability to capture relevant visual representations. To the best of our knowledge, our work is the first to utilize retrieval text as a visual prompt for lightweight image captioning.

Our approach achieves superior performance on the COCO dataset (Lin et al. 2015) compared to our baseline model, SmallCap, and it significantly improves performance over previous lightweight models on the NoCaps dataset (Agrawal et al. 2019). In the experiments, we integrate our ViP module into retrieval-based models, text-only training models, and various prompts, resulting in consistent performance improvements.

Our contribution can be summarized as follows: (1) We propose a novel visual prompt for lightweight image captioning models named ViPCap, which leverages retrieved texts to generate visual prompts. (2) We introduce the ViP module, which retrieves semantic information from text features and combines it with image features to generate the visual prompt. (3) Extensive experiments demonstrate that our method is efficient and outperforms previous models across datasets like COCO and NoCaps, regardless of the text prompt types used.

Related Work

CLIP in captioning. With the advent of vision-language models (Radford et al. 2021; Jia et al. 2021), significant advancements have been made in V&L tasks. Notably, CLIP achieves multimodal alignment with 400M image-text pairs. The image captioning models that use a clip-based encoder such as BLIP-2, LLaVA, and MiniGPT-4 (Li et al. 2023a; Liu et al. 2023; Zhu et al. 2023) create a mapping network and pass input features to the decoder. In our study, we map the CLIP text encoder to the randomized Gaussian distribution, desiring to achieve a high correlation with visual features and leveraging it to generate visual prompts,

Prompt tuning. Initially used in NLP (Lester, Al-Rfou, and Constant 2021; Li and Liang 2021), this approach has extended to V&L and vision-only models (Jia et al. 2022; Wang et al. 2022b) to utilize pretrained knowledge for downstream without parameter training. Recent techniques focus on trainable tokens instead of handcrafted prompts (Gao, Fisch, and Chen 2021). Visual prompt tuning (Jia et al. 2022) adds learnable tokens to ViT for downstream tasks, while some methods (Bahng et al. 2022) integrate task-specific patches at the pixel level. Although previous methods achieved some success, they still face challenges with complex tasks like captioning. We introduce a visual prompt method suitable for captioning.

Retrieval for captioning. External knowledge helps reduce the cost of creating large image-text datasets for Multimodal LLM. Methods like SmallCap (Ramos et al. 2023), LMCap (Ramos, Martins, and Elliott 2023), Knight (Wang

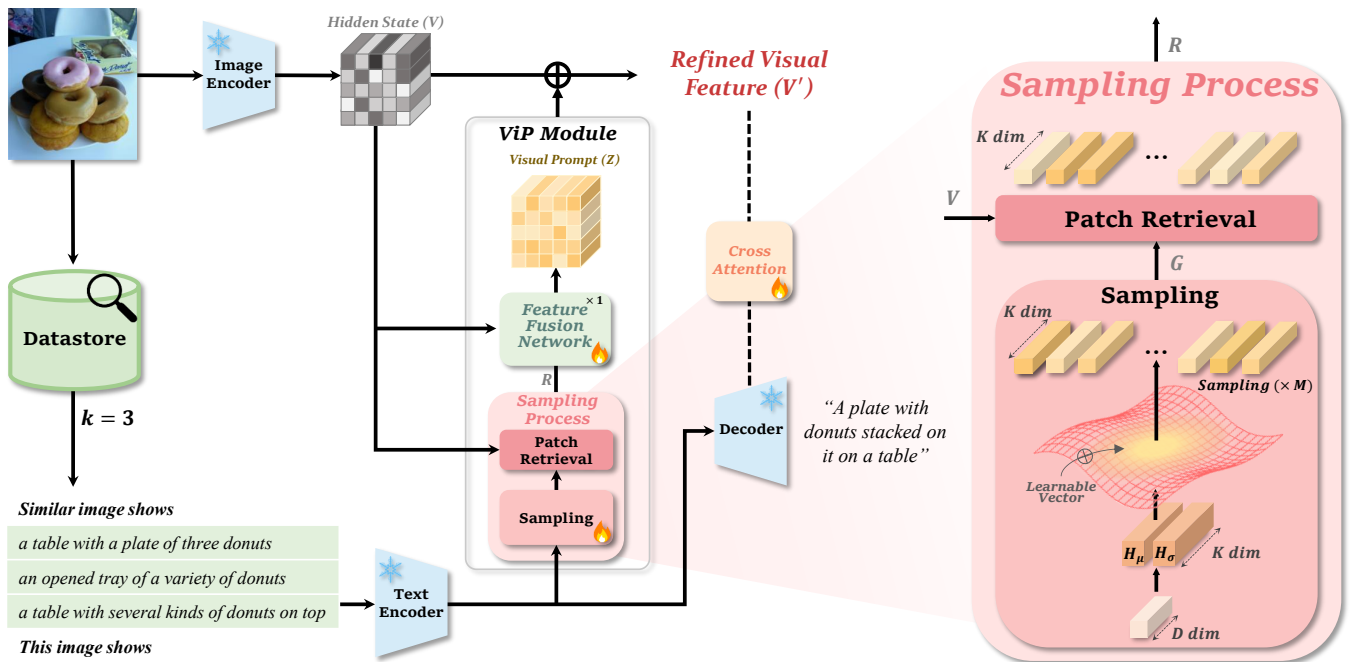


Figure 3: ViPCap leverages the CLIP text encoder to extract retrieved text features for visual prompts generation. The ViP module performs M sampling iterations from the text embedding distribution to extract semantic features G closely aligned with image patch features V . Then, in the patch retrieval module, we retrieve semantic vectors from G that are highly correlated with image patch features V . The retrieved semantic features R are fused with image features V within the Feature Fusion Network, and the resulting output is set as the visual prompt Z . Finally, the refined visual feature V' via summation with the visual prompt is fed to the decoder through the cross-attention layer.

et al. 2023), and MeaCap (Zeng et al. 2024) use text-based datasets and store captions in a datastore. SmallCap (Ramos et al. 2023), for example, generates captions using an input image and related text from a datastore with only 7 million trainable parameters, allowing fast and easy training, although it lacks visual knowledge. Our work enhances performance by generating visual prompts from retrieved text with detailed image descriptions.

Proposed Method

Our model adopts SmallCap (Ramos et al. 2023) as a baseline, which retrieves enriched expressions from an external datastore and integrates the pretrained CLIP encoder (Radford et al. 2021) with GPT-2 (Radford et al. 2018) through cross-attention layers. SmallCap consists of the retrieved texts with the following hard prompt and inputs it into the decoder. Similarly, we utilize this approach in encoding our prompts: Similar images show $\{\text{caption}_1\} \dots \{\text{caption}_k\}$. This image shows _____. However, SmallCap primarily focuses on text prompts, while visual representation relies solely on the performance of the vision encoder.

In this work, we introduce ViPCap, a novel approach that enhances the performance of lightweight image captioning by generating *visual prompts* based on the semantic information in text prompts as illustrated in Fig. 3. ViP module encodes retrieved texts into the CLIP embedding space

and converts them into patch-level representations. Then, assuming a multivariate Gaussian distribution, our method generates semantic features M times to obtain semantic features that are highly correlated to the local visual features. Unlike the heuristic Gaussian distribution approach in Capdec (Nukrai, Mokady, and Globerson 2022), our novel approach creates semantic features through a learnable distribution. Then, the semantic features are leveraged by the patch retrieval module to closely align with the input image patches. The matched semantic features are fused with the input image to generate the visual prompt. Finally, we train cross-attention layers to fuse the features in the decoder.

Randomized Gaussian Distribution Sampling

The ViP module aims to effectively sample the semantic information from the retrieved texts to enrich the input visual feature with semantic content in the form of visual prompts. Given retrieved texts T , our model encodes the retrieved text into D dimensional vector using pretrained CLIP text encoder $\phi(\cdot)$. Also, the CLIP image encoder embeds the input image into K dimensional visual features $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_N\} \in \mathbb{R}^{N \times K}$ representing N number of patch-level visual features.

When generating visual prompts, a single text feature might be insufficient to provide the necessary details to generate a visual prompt with complex patch-level local information. To address this, we employ a random augmentation techniques to sample semantic features from the Gaus-

sian distribution. Also, instead of learning multiple mapping functions for local regions individually, we empirically find that sampling random vectors helps better match with visual local representations.

We estimate the parameters of distribution of the text embedding $\vec{\mu}, \vec{\sigma} \in \mathbb{R}^K$ assuming it follows a multivariate Gaussian distribution $\mathcal{N}(\vec{\mu}, \vec{\sigma}^2 I)$. We design functions $\mathcal{H}_\mu(\cdot)$ and $\mathcal{H}_\sigma(\cdot)$ to map text features into the mean and standard deviation of multivariate Gaussian distribution. These functions are implemented as MLP layers to map from D dimensions to K dimensions while sampling the mean and standard deviation ($\mathcal{H} : \mathbb{R}^D \rightarrow \mathbb{R}^K$). We empirically find that adding an additional learnable vector ω_{add} with a hyperparameter α as a scaling factor to the MLP shows better performance and captures complex data structures more effectively. The α is used to expand the range of the learnable vector. Let $\vec{\mu}$ and $\vec{\sigma}$ are computed via $\vec{\mu} = \mathcal{H}_\mu(\phi(T)) + \alpha \cdot \omega_{add}$, and $\vec{\sigma} = \mathcal{H}_\sigma(\phi(T))$, respectively.

ViP module samples M number of semantic features from this learnable Gaussian distribution to obtain semantic features that are highly correlated to the local visual embedding. We define the set of semantic representation $G \in \mathbb{R}^{M \times K}$ obtained from the text features as:

$$G = \{\vec{g}_i \sim \mathcal{N}(\vec{\mu}, \vec{\sigma}^2 I; \phi(T))\}_{i=1}^M. \quad (1)$$

Additionally, increasing M allows semantic features G to deliver better fine-grained visual information to the input feature (See more details in Ablation.). Through the reparameterization trick (Kingma and Welling 2022), \vec{g} can be re-formulated as $\vec{g} = \vec{\mu} + \vec{\sigma} \cdot \vec{\epsilon}$, where $\vec{\epsilon} \sim N(0, I)$.

Patch Retrieval Module for Semantic Features

We hypothesize that the semantic features $G = \{\vec{g}_1, \vec{g}_2, \dots, \vec{g}_M\}$ sampled from the Gaussian distribution contain the textual information describing the image. To effectively leverage random semantic features, we retrieve semantic features that contain useful visual information related to visual prompts based on feature similarity. In particular, we introduce a patch retrieval module, as depicted in Fig. 4, that compares the similarity between image patch-level representations V and the semantic features G . This module employs cosine similarity, denoted as $sim(\cdot, \cdot)$, to effectively identify the most relevant semantic information for each patch representation $\vec{v}_i \in V$. Consequently, ViPCap chooses N relevant vectors from M candidates, one responsible for each patch vector, with high similarity to the input feature, and generates R :

$$R = \{\vec{g}_{\mathcal{I}(j)}\}_{j=1}^N \in \mathbb{R}^{N \times K}, \quad (2)$$

where $\mathcal{I}(j) = \operatorname{argmax}_{i \in [1:M]} sim(\vec{g}_i, \vec{v}_j)$.

This simple calculation process extracts valuable information through R without any additional training. The extracted vectors have the potential to provide semantic representations that the vision encoder cannot offer.

Feature Fusion Network and Visual Prompts

After obtaining the representations R containing semantic information relevant to the visual features, we integrate them

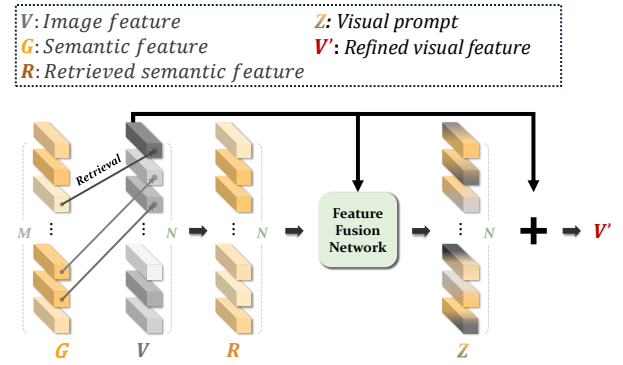


Figure 4: Calculating similarity between input feature V and the semantic features G . V retrieves essential semantic representations from G and combines these retrieved semantic features R through the fusion network. The fusion network generates a visual prompt Z by integrating image features V . After that, the visual prompt and image features are combined to create refined visual features V' .

with the input image to enhance the visual feature to generate the *visual prompt* Z . We introduce a Feature Fusion Network (FFN) designed to effectively fuse the visual representations V and the retrieved semantic features R , as shown in Fig. 4, by bridging the gap between them.

FFN is designed with a transformer architecture that incorporates both self-attention and cross-attention layers. Unlike previous mapping networks (Li et al. 2023a; Mokady, Hertz, and Bermano 2021) that typically consist of 8 to 12 layers, this module uses only a single layer ($l = 1$). This efficiency results from earlier stages where features are sampled to closely align with visual representations, making a single layer sufficient for effective fusion. This network uses the input feature V as the query and the retrieved semantic representation R as the key.

Note that a simple summation between V and R ignores distribution differences between features, and a simple concatenation cannot consider relational context between modalities (see Ablation. for more details).

Finally, we obtain refined visual features V' with a simple summation ($V' = V + Z$). Because the visual prompt closely aligns with the distribution of image features, we enable model to generate the refined visual features through simple summation. This method allows models like Small-Cap, which have a vision encoder and language decoder, to use refined visual feature V' instead of the input feature V , preserving the existing decoder design. This approach functions as a module that can be applied to various frameworks using encoders.

The decoder takes the input text embedding, while the refined visual feature V' is included conditionally when computing the loss function:

$$L_\theta = - \sum_{i=1}^Q \log P_\theta(y_i | y_{<i}, V'; \theta). \quad (3)$$

In Eq. (3), the main difference in our approach is the use of refined visual feature V' instead of V . In the cross-attention

Method	Training Param θ	COCO Test				Flickr30k Test		NoCaps Val			
		B@4	M	C	S	C	S	In	Near	Out	Entire
Large scale training models											
OSCAR _{Large} (2020)	338M	37.4	30.7	127.8	23.5	-	-	78.8	78.9	77.4	78.6
LEMON _{Huge} (2022)	675M	41.5	30.8	139.1	24.1	-	-	118.0	116.3	120.2	117.3
SimVLM _{Huge} (2022a)	632M	40.6	33.7	143.3	25.4	-	-	113.7	110.9	115.2	112.2
BLIP2 _{ViT-g OPT_{2.7B}} (2023a)	1.1B	43.7	-	145.8	-	-	-	123.0	117.8	123.4	119.7
CogVLM (2024)	1.5B	-	-	148.7	-	94.9	-	-	-	132.6	128.3
PaLI _{mT5-XXL} (2023)	1.6B	-	-	149.1	-	-	-	-	-	-	127.0
Lightweight models											
CaMEL (2022)	76M	39.1	29.4	125.7	22.2	-	-	-	-	-	-
I-Tuning _{Medium} (2023)	44M	35.5	<u>28.8</u>	120.0	<u>22.0</u>	72.3	19.0	<u>89.6</u>	77.4	58.8	75.4
ClipCap (2021)	43M	33.5	<u>27.5</u>	113.1	21.1	-	-	84.9	66.8	49.1	65.8
I-Tuning _{Base} (2023)	14M	34.8	28.3	116.7	21.8	61.5	16.9	83.9	70.3	48.1	67.8
SmallCap (2023)	7M	37.0	27.9	119.7	21.3	60.6	-	87.6	<u>78.6</u>	<u>68.9</u>	<u>77.9</u>
SmallCap _{d=16, Large} (2023)	47M	37.2	28.3	121.8	21.5	-	-	-	-	-	-
ViPCap (Ours)	14M	<u>37.7</u>	28.6	<u>122.9</u>	21.9	<u>66.8</u>	<u>17.2</u>	93.8	81.6	71.5	81.3

Table 1: Comparison with large pre-trained and lightweight models with existing methods on the COCO test, Flickr30k test, and NoCaps val set. CIDEr score is used for NoCaps evaluation. Our method shows the competitive performance in most metrics.

layers (θ), weights are optimized by reducing the cross-entropy loss associated with forecasting the Q tokens in the given reference y_1, \dots, y_q .

Experiments

In this section, we conduct experiments to show the effectiveness of ViPCap over existing methods. First, we evaluate the advantages of our proposed model over the baseline in Table 1 for in-domain and out-of-domain, and then we apply the ViP module with the text-only training model to identify the role of ViP module in Table 2, respectively. In Table 3, ViP module evaluates performance of sampling method from different probability distributions. Moreover, we test the plug-and-play solution with different models in Table 4, Table 5, and with different prompt styles in Table 6.

Experimental Setup

Training dataset. We conduct experiments on image captioning benchmarks, i.e., COCO dataset (Lin et al. 2015), NoCaps (Agrawal et al. 2019), Flickr30k (Plummer et al. 2016). For COCO and Flickr30k, we follow the Karpathy split (Karpathy and Fei-Fei 2015) used in the image captioning. We evaluate our model on the COCO and Flickr30k test set and NoCaps validation and test datasets, as well as the cross-domain experiments.

Training setup. ViPCap includes a vision encoder (ViT-B/32) and a language decoder (GPT-2_{Base}). Both are frozen during training. We train the ViP module and cross-attention layer. The FFN and cross-attention layer include a 12-head cross-attention layer with a single-layer block. To reduce the computational cost, we scale the dimension of cross-attention layers from 64 to 16. ViPCap requires 14M training parameters and is trained on a single NVIDIA 6000 GPU with a batch size of 128.

During training, the ViP module uses a patch size of $M=200$, and hyperparameter α is set to 5. The model selects three retrieved captions per image ($k=3$) from the COCO datastore due to the limitation of 77 context length size in CLIP. Captions are retrieved using CLIP ResNet-50x64 and processed with FAISS (Johnson, Douze, and Jégou 2017) for fast nearest neighbor search. Caption quality is evaluated using the metrics BLEU@4 (B@4) (Papineni et al. 2002), METEOR (M) (Denkowski and Lavie 2014), CIDEr (C) (Vedantam, Zitnick, and Parikh 2014), and SPICE (S) (Anderson et al. 2016).

Main Results

In-domain. We evaluate ViPCap on the COCO, Flickr30k, and NoCaps datasets in Table 1. Table 1 shows the evaluation results on the COCO dataset. The upper part of Table 1 refers to the performance of large models trained on large datasets. In particular, compared to OSCAR_{Large}, our model outperforms in B@4 score while using only 4% of the parameters. Our model, with 5 times fewer parameters than CaMEL, obtains the second-highest CIDEr score among lightweight models. In the COCO dataset, ViPCap exceeds the baseline, SmallCap, and outperforms SmallCap_{Large}, even though it has significantly fewer parameters (14M vs. 47M). The model also shows strong performance on the Flickr30k dataset.

Cross-domain. On the NoCaps validation dataset in the Table 1, it achieves a CIDEr score of 93.8 in in-domain data, surpassing all lightweight models. ViPCap exhibits superior performance in both in-domain and cross-domain, notably outperforming the previous SOTA models by over 3 points across most of metrics in the cross-domain. Additionally, it exceeds the large-scale training model, Oscar_{Large}, by more than 10 points on in-domain data, indicating further perfor-

Method	In-Domain								Cross-Domain											
	COCO				Flickr30k				COCO \Rightarrow Flickr30k				Flickr30k \Rightarrow COCO				COCO \Rightarrow NoCaps			
	B@4	M	C	S	B@4	M	C	S	B@4	M	C	S	B@4	M	C	S	I	N	O	E
CapDec (2022)	26.4	25.1	91.8	-	17.7	20.0	39.1	-	17.3	18.6	35.7	-	9.2	16.3	27.3	-	60.1	50.2	28.7	45.9
CapDec+ViP	27.0	25.6	94.2	18.8	18.6	20.1	44.4	14.4	15.7	18.0	35.8	11.8	9.5	16.3	30.7	9.2	60.2	50.9	33.7	47.8
Δ	0.6	0.5	2.4	-	0.9	0.1	5.3	-	-1.6	-0.6	0.1	-	0.3	-	3.4	-	0.1	0.7	5.0	1.9
ViECap (2023)	27.2	24.8	92.9	18.2	21.4	20.1	47.9	13.6	17.4	18.0	38.4	11.2	12.6	19.3	54.2	12.5	61.1	64.3	65.0	66.2
ViECap+ViP	27.3	25.1	93.6	18.4	21.2	20.2	48.8	13.9	17.4	18.1	40.2	11.1	13.6	19.3	55.2	12.7	62.2	64.9	67.1	67.2
Δ	0.1	0.3	0.7	0.2	-0.2	0.1	0.9	0.3	-	0.1	1.8	-0.1	1.0	-	1.0	0.2	1.1	0.6	2.1	1.0

Table 2: Results of applying the ViP module to the text-only training models: In-domain results on COCO and Flickr30k test sets; cross-domain results on Flickr30k, COCO test sets, and NoCaps validation set. CIDEr score is used for NoCaps evaluation. I is in-domain, N denotes near-domain, O represents out-of-domain, and E indicates entire dataset. Our method consistently improves performance across most metrics in both in-domain and cross-domain settings, regardless of the base models. Notably, a significant performance improvements in CIDEr is observed in cross-domain scenarios when combining ViP module.

	No Noise	$\mathcal{N}(0, 1)$	Unif(-1,1)	Unif(0,1) $\times \mathcal{N}(0, 1)$	$\alpha \cdot \omega_{add}$ (Ours)
CIDEr	121.1	121.4	121.9	122.0	122.9

Table 3: CIDEr scores estimated by adding vectors sampled from different probability distributions to the semantic features in the COCO test dataset. Uniform distribution denotes as Unif. Unif(0,1) $\times \mathcal{N}(0, 1)$ represents the method proposed by LinCIR (Gu et al. 2024).

Method	Enc.	Dec.	Ret.	B@4	CIDEr
MAPL	ViT	GPT-J 6B	\times	36.5	125.2
ViPCap (Ours)	-L/14	GPT2-L 0.7B	\checkmark	38.9 (2.4 \uparrow) 40.5 (4.0 \uparrow)	128.3 (3.1 \uparrow) 131.0 (5.8 \uparrow)

Table 5: CIDEr and BLEU@4 scores for different types of GPT decoders on the COCO test. Ret refers to the usage of the retrieved text.

Method	Enc.	Dec.	ViP	Ret	CIDEr
ViPCap	ViT	OPT -125M	\checkmark	\times \checkmark	122.0 122.5 (0.5 \uparrow)
(Ours)	-B/32	XGLM		\times \checkmark	116.8 121.2 (4.4 \uparrow)
EVCap	EVA-CLIP-g	Vicuna -13B	\times \checkmark	\checkmark \checkmark	140.1 141.3 (1.2 \uparrow)

Table 4: Performance improvements in CIDEr scores on COCO test across various decoders using the ViP module and the retrieved text. Ret refers to the usage of the retrieved text. It also maintains performance improvements when changing the encoder and decoder to larger models.

mance in entire data. Table 1 shows that ViPCap is highly suitable for zero-shot tasks and real-world scenarios.

Additional vector design. We add vectors to enable the model to capture more complex data structures when estimating the Gaussian distribution. In the Table 3, we experiment by adding a vector to the semantic features across various design cases. *No noise* is adding nothing. We empirically find that our method enhances performance compared to not adding a vector or adding a vector sampled from a different distribution during the sampling process.

ViP Module Capability

We explore different model sizes and prompt styles to evaluate the capabilities of our model.

Plug-and-Play manner. In Table 2, we evaluate the COCO, Flickr30k, and NoCaps test datasets to explicitly demonstrate the visual prompt ability by the ViP module. We combine our module with text-only training models.

As a result, applying the ViP module to CapDec and ViECap (Nukrai, Mokady, and Globerson 2022; Fei et al. 2023) improves performance, demonstrating that the ViP module can function as both a visual prompt and an image feature. The ViP module easily fuses with the vision encoders without modifying the framework. CapDec and ViECap with the ViP module result in an average increase of 3.5 points in CIDEr score on cross-domain in the NoCaps dataset. Our method shows the capability of ViP module in zero-shot tasks across real-world scenarios. We do not test with DeCap (Li et al. 2023b) due to its focus on memory efficiency, which does not align with our goals.

Model-agnostic. Table 4 reveals that combining the ViP module with OPT (Zhang et al. 2022) and XGLM (Lin et al. 2022) as decoders, along with using both ViP and the retrieved text, leads to a notable improvement in performance. This indicates the capability as a model-agnostic and flexible framework. Additionally, similar to SamllCap, combining EVCap (Li et al. 2024), which utilizes retrieval data, with the ViP module enhances performance. This means our ap-

Prompt	ViP	
	×	✓
“This image shows”	111.1	116.0 (4.9 ↑)
Retrieval prompt	117.3	119.9 (2.6 ↑)

Table 6: CIDEr results of the ViP module on COCO val dataset show the potential of our prompt-agnostic model.

<i>PR</i>	ω	α	B@4	CIDEr
×	✓	×	36.6	120.1
×	✓	✓	36.9	121.1
✓	×	×	37.0	121.1
✓	✓	×	37.3	121.9
✓	✓	✓	37.7	122.9

Table 7: Experimental results on the COCO test dataset. *PR* denotes the Patch Retrieval module, ω is the learnable vector, and α represents the scale factor used in the $\alpha \cdot \omega_{add}$ with sampling Gaussian distribution.

proach can be effectively applied to models leveraging retrieval data, as well as to large-scale models such as EVA-CLIP and Vicuna (Sun et al. 2023; Chiang et al. 2023).

We cannot apply our module to MAPL (Mañas et al. 2023) due to its limited access to the training code. Hence, in Table 5, we evaluate the same encoder with different types of GPT decoders. We use “This image shows” as the hard prompt. Notably, despite our decoder having approximately 9 times fewer parameters, it achieves a CIDEr score improvement of over 3.1 points and nearly 6 points when using the retrieved text. Therefore, ViP module presents a consistent performance improvement when combined with SOTA models regardless of model size.

Prompt-agnostic. Table 6 compares ViPCap performance with and without the retrieval module. SmallCap scores 111.1 without retrieval and 117.3 with retrieval. ViPCap scores 116.0 without retrieval, using simple prompts like “This image shows ...” and 119.9 with retrieval. Additionally, in CapDec, ViECap, and EVCap, we observe notable results by leveraging hard prompts such as “a photo of” and “There are $entity_1$, $entity_2$, ...”. This demonstrates that ViPCap addresses competitive performance even with simple hard prompts and suggests its potential as a flexible visual prompt module applicable to various prompt types.

Ablation Studies

We conduct ablation studies to evaluate the effects of various components in our work, including the design of ViP module, additional vector, and feature fusion network strategy.

Effect of ViP module components. We explore the effect of the patch retrieval module, learnable vector, and scale factor in Table 7. Experimental results show that the patch retrieval module achieves better performance, which can be attributed to its close alignment with visual features. Adding a learnable vector during sampling improves performance. Scaling

FFN	B@4	CIDEr
Sum	36.4	119.1
Concat	36.9	119.8
MLP→Sum	36.8	119.8
MLP→Concat	37.1	119.8
Ours	37.7	122.9

Table 8: CIDEr and BLEU@4 on COCO test dataset for different Feature Fusion Network (FFN) designs.

<i>M</i>	100	150	200	250	300	400	500
B@4	37.3	37.2	37.7	37.3	37.2	37.2	36.9
CIDEr	121.4	121.1	122.9	121.5	121.4	121.4	121.1

Table 9: CIDEr and BLEU@4 scores on the COCO test data depending on sampling *M* number of semantic feature.

this vector with the ω parameter provides the optimal framework, and all experiments in this paper are conducted with the same scaling value.

Feature Fusion Network design. In the Table 8, our FFN shows more effectiveness in mitigating the modality gap between image and text semantic features, and generating refined visual features with only a single layer. Compared to our approach, the summation, concatenation, and MLP methods struggle to capture distribution and relational context between modalities.

Effect of *M* number sampling. We extract semantic features through *M* number sampling. In Table 9, where *M* number of sampling ranges from 100 to 500, the highest score is observed at $M = 200$. When there are relatively too many patches ($M > 200$), the local representation G becomes scattered, and leads to no further improvement in performance.

Conclusion

In this work, we introduce ViPCap, a novel approach that generates visual prompts by leveraging semantic information from retrieved text embedding through the ViP module. ViPCap performs well across both in-domain and out-of-domain datasets. The ViP module proposes a plug-and-play method that generates visual prompts based on various models and prompt types. Future work will explore using learnable tokens as visual prompts for better flexibility.

Acknowledgements

This was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00245661).

References

Agrawal, H.; Desai, K.; Wang, Y.; Chen, X.; Jain, R.; Johnson, M.; Batra, D.; Parikh, D.; Lee, S.; and Anderson, P.

2019. nocaps: novel object captioning at scale. In *Proceedings of the IEEE International Conference on Computer Vision*, 8948–8957.
- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. SPICE: Semantic Propositional Image Caption Evaluation. *ArXiv*, abs/1607.08822.
- Bahng, H.; Jahanian, A.; Sankaranarayanan, S.; and Isola, P. 2022. Exploring Visual Prompts for Adapting Large-Scale Models. *arXiv*:2203.17274.
- Barraco, M.; Stefanini, M.; Cornia, M.; Cascianelli, S.; Baraldi, L.; and Cucchiara, R. 2022. CaMEL: Mean Teacher Learning for Image Captioning. *arXiv*:2202.10492.
- Chen, X.; Wang, X.; Changpinyo, S.; Piergiovanni, A.; Padlewski, P.; Salz, D.; Goodman, S.; Grycner, A.; Mustafa, B.; Beyler, L.; Kolesnikov, A.; Puigcerver, J.; Ding, N.; Rong, K.; Akbari, H.; Mishra, G.; Xue, L.; Thapliyal, A.; Bradbury, J.; Kuo, W.; Seyedhosseini, M.; Jia, C.; Ayan, B. K.; Riquelme, C.; Steiner, A.; Angelova, A.; Zhai, X.; Houlsby, N.; and Soricut, R. 2023. PaLI: A Jointly-Scaled Multilingual Language-Image Model. *arXiv*:2209.06794.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Denkowski, M. J.; and Lavie, A. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *WMT@ACL*.
- Fei, J.; Wang, T.; Zhang, J.; He, Z.; Wang, C.; and Zheng, F. 2023. Transferable Decoding with Visual Entities for Zero-Shot Image Captioning. *arXiv*:2307.16525.
- Gao, T.; Fisch, A.; and Chen, D. 2021. Making Pre-trained Language Models Better Few-shot Learners. *arXiv*:2012.15723.
- Gu, G.; Chun, S.; Kim, W.; Kang, Y.; and Yun, S. 2024. Language-only Efficient Training of Zero-shot Composed Image Retrieval. *arXiv*:2312.01998.
- Hu, X.; Gan, Z.; Wang, J.; Yang, Z.; Liu, Z.; Lu, Y.; and Wang, L. 2022. Scaling Up Vision-Language Pre-training for Image Captioning. *arXiv*:2111.12233.
- Hu, Z.; Iscen, A.; Sun, C.; Wang, Z.; Chang, K.-W.; Sun, Y.; Schmid, C.; Ross, D. A.; and Fathi, A. 2023. REVEAL: Retrieval-Augmented Visual-Language Pre-Training with Multi-Source Multimodal Knowledge Memory. *arXiv*:2212.05221.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q. V.; Sung, Y.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. *arXiv*:2102.05918.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual Prompt Tuning. *arXiv*:2203.12119.
- Johnson, J.; Douze, M.; and Jégou, H. 2017. Billion-scale similarity search with GPUs. *arXiv*:1702.08734.
- Karpathy, A.; and Fei-Fei, L. 2015. Deep Visual-Semantic Alignments for Generating Image Descriptions. *arXiv*:1412.2306.
- Kingma, D. P.; and Welling, M. 2022. Auto-Encoding Variational Bayes. *arXiv*:1312.6114.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. *arXiv*:2104.08691.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. *arXiv*:2301.12597.
- Li, J.; Vo, D. M.; Sugimoto, A.; and Nakayama, H. 2024. EVCap: Retrieval-Augmented Image Captioning with External Visual-Name Memory for Open-World Comprehension. *arXiv*:2311.15879.
- Li, W.; Zhu, L.; Wen, L.; and Yang, Y. 2023b. DeCap: Decoding CLIP Latents for Zero-Shot Captioning via Text-Only Training. In *The Eleventh International Conference on Learning Representations*.
- Li, X.; Yin, X.; Li, C.; Zhang, P.; Hu, X.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; Choi, Y.; and Gao, J. 2020. Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. *arXiv*:2004.06165.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. *arXiv*:2101.00190.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C. L.; and Dollár, P. 2015. Microsoft COCO: Common Objects in Context. *arXiv*:1405.0312.
- Lin, X. V.; Mihaylov, T.; Artetxe, M.; Wang, T.; Chen, S.; Simig, D.; Ott, M.; Goyal, N.; Bhosale, S.; Du, J.; Pasunuru, R.; Shleifer, S.; Koura, P. S.; Chaudhary, V.; O’Horo, B.; Wang, J.; Zettlemoyer, L.; Kozareva, Z.; Diab, M.; Stoyanov, V.; and Li, X. 2022. Few-shot Learning with Multilingual Language Models. *arXiv*:2112.10668.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. *arXiv*:2304.08485.
- Luo, Z.; Hu, Z.; Xi, Y.; Zhang, R.; and Ma, J. 2023. I-Tuning: Tuning Frozen Language Models with Image for Lightweight Image Captioning. *arXiv*:2202.06574.
- Mañas, O.; Rodriguez, P.; Ahmadi, S.; Nematzadeh, A.; Goyal, Y.; and Agrawal, A. 2023. MAPL: Parameter-Efficient Adaptation of Unimodal Pre-Trained Models for Vision-Language Few-Shot Prompting. *arXiv*:2210.07179.
- Mokady, R.; Hertz, A.; and Bermano, A. H. 2021. ClipCap: CLIP Prefix for Image Captioning. *arXiv*:2111.09734.
- Nukrai, D.; Mokady, R.; and Globerson, A. 2022. Text-Only Training for Image Captioning using Noise-Injected CLIP.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Annual Meeting of the Association for Computational Linguistics*.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2016. Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models. *arXiv*:1505.04870.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.;

- Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2018. Language Models are Unsupervised Multitask Learners.
- Ramos, R.; Martins, B.; and Elliott, D. 2023. LMCap: Few-shot Multilingual Image Captioning by Retrieval Augmented Language Model Prompting. arXiv:2305.19821.
- Ramos, R.; Martins, B.; Elliott, D.; and Kementchedjhieva, Y. 2023. SmallCap: Lightweight Image Captioning Prompted with Retrieval Augmentation. arXiv:2209.15323.
- Ramos, R. P.; Elliott, D.; and Martins, B. 2023. Retrieval-augmented Image Captioning. In *Conference of the European Chapter of the Association for Computational Linguistics*.
- Sun, Q.; Fang, Y.; Wu, L.; Wang, X.; and Cao, Y. 2023. EVA-CLIP: Improved Training Techniques for CLIP at Scale. arXiv:2303.15389.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2014. CIDEr: Consensus-based image description evaluation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4566–4575.
- Wang, B.; and Komatsuzaki, A. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Wang, J.; Yan, M.; Zhang, Y.; and Sang, J. 2023. From Association to Generation: Text-only Captioning by Unsupervised Cross-modal Mapping. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 4326–4334. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; Song, X.; Xu, J.; Xu, B.; Li, J.; Dong, Y.; Ding, M.; and Tang, J. 2024. CogVLM: Visual Expert for Pretrained Language Models. arXiv:2311.03079.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2022a. SimVLM: Simple Visual Language Model Pre-training with Weak Supervision. arXiv:2108.10904.
- Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022b. DualPrompt: Complementary Prompting for Rehearsal-free Continual Learning. arXiv:2204.04799.
- Yang, Z.; Ping, W.; Liu, Z.; Korthikanti, V.; Nie, W.; Huang, D.-A.; Fan, L.; Yu, Z.; Lan, S.; Li, B.; Liu, M.-Y.; Zhu, Y.; Shoeybi, M.; Catanzaro, B.; Xiao, C.; and Anandkumar, A. 2023. Re-ViLM: Retrieval-Augmented Visual Language Model for Zero and Few-Shot Image Captioning. arXiv:2302.04858.
- Zeng, Z.; Xie, Y.; Zhang, H.; Chen, C.; Wang, Z.; and Chen, B. 2024. MeaCap: Memory-Augmented Zero-shot Image Captioning. arXiv:2403.03715.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; Mihaylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.;
- Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models. arXiv:2205.01068.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. arXiv:2304.10592.