

# IFCap: Image-like Retrieval and Frequency-based Entity Filtering for Zero-shot Captioning

Soeun Lee\* Si-Woo Kim\* Taewhan Kim Dong-Jin Kim†

Hanyang University, South Korea.

{soeun, boreng0817, taewhan, djdkim}@hanyang.ac.kr

## Abstract

Recent advancements in image captioning have explored text-only training methods to overcome the limitations of paired image-text data. However, existing text-only training methods often overlook the modality gap between using text data during training and employing images during inference. To address this issue, we propose a novel approach called Image-like Retrieval, which aligns text features with visually relevant features to mitigate the modality gap. Our method further enhances the accuracy of generated captions by designing a Fusion Module that integrates retrieved captions with input features. Additionally, we introduce a Frequency-based Entity Filtering technique that significantly improves caption quality. We integrate these methods into a unified framework, which we refer to as IFCap (Image-like Retrieval and Frequency-based Entity Filtering for Zero-shot Captioning). Through extensive experimentation, our straightforward yet powerful approach has demonstrated its efficacy, outperforming the state-of-the-art methods by a significant margin in both image captioning and video captioning compared to zero-shot captioning based on text-only training.<sup>1</sup>

## 1 Introduction

The task of image captioning generates appropriate textual descriptions for images by combining computer vision (CV) and natural language processing (NLP). With the emergence of Large Language Models (LLMs) and Vision and Language Models (VLMs), various works have studied efficient training methods for image captioning (Mokady et al., 2021; Luo et al., 2023; Ramos et al., 2023). These approaches develop effective captioning by using pre-trained models with few parameters or lightweight networks. However, these works rely on paired image-text data, which is costly (Kim

et al., 2019b, 2024). To overcome this, recent studies have explored text-only training methods for image captioning, aiming to solve the problem using only textual data (Nukrai et al., 2022; Li et al., 2023; Fei et al., 2023; Zeng et al., 2024; Wang et al., 2023; Liu et al., 2023; Ma et al., 2023).

Text-only training introduces a new direction in which models are trained solely using text data. Recent existing works have studied what to use as extra cues, such as extracted nouns (Fei et al., 2023), generated synthetic images (Liu et al., 2023; Ma et al., 2023) for training, and extracted tags from object detectors (Liu et al., 2023). However, existing methods that rely on object information are sensitive to incorrect data, and utilizing large external models (e.g., stable diffusion Rombach et al., 2022 or object detectors Carion et al., 2020) incurs additional costs. Thus, we aim to address the problem by acquiring diverse information cost-effectively without additional models.

The retrieval task involves finding relevant information in a database for a given query. Initially rooted in NLP (Lewis et al., 2020), the field has expanded into CV and into multi-modal retrieval. Depending on the input data and database, various retrieval methods are possible, such as image-to-text (Ramos et al., 2023) and text-to-text retrieval (Wang et al., 2023). In the existing text-only training study, there have been attempts to use the text-to-text retrieval method. However, existing works can't address the modality gap inherent in text-only training settings, where training is performed with text and inference with images. In addition, such works rely too much on retrieved captions without considering visual information. This modality gap and the use of a narrow scope of information may lead to performance degradation.

To verify this, we visualize the analysis result of the CLIP embedding feature of retrieved captions that the model uses in training via t-SNE in Fig. 2. The analysis is done on the COCO (Chen et al.,

\*Equal contribution. †Corresponding author.

<sup>1</sup>Code: <https://github.com/boreng0817/IFCap>

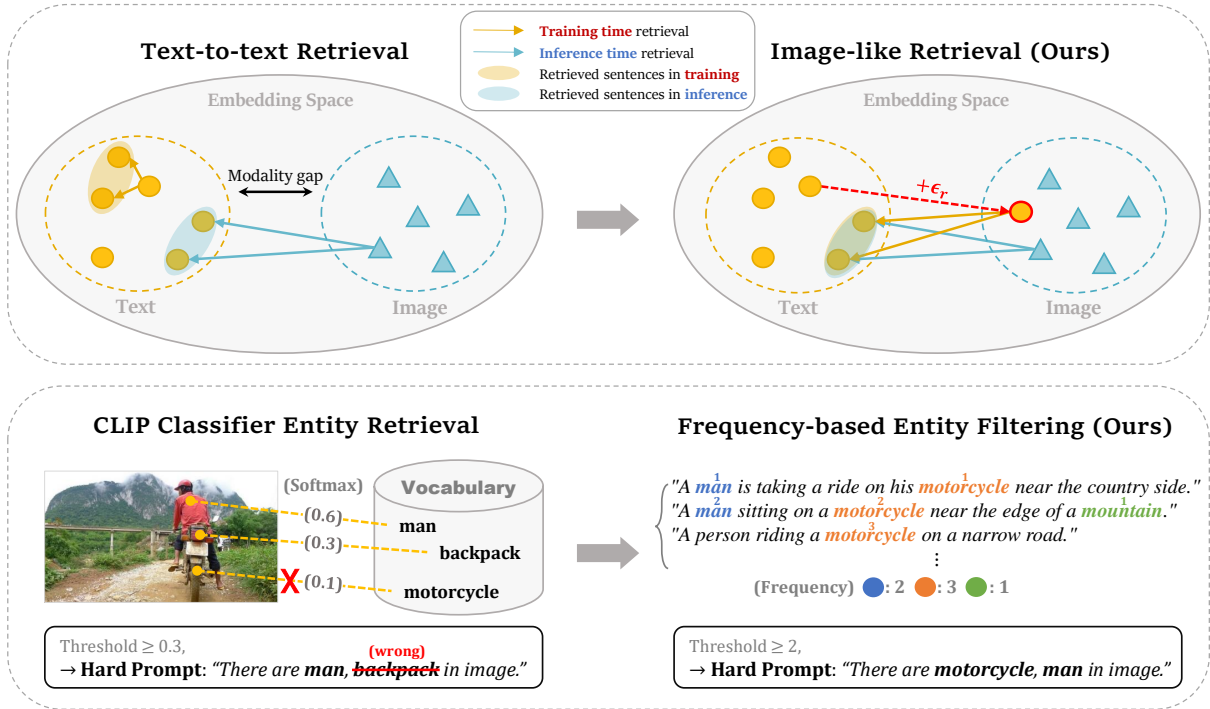


Figure 1: (Top) The previous text-to-text retrieval approach overlooks the modality gap, leading to different information use between training and inference. Our approach addresses this by aligning text features with the image embedding space during retrieval. (Bottom) The traditional CLIP classifier-based entity retrieval method struggles with entity detection as vocabulary size grows. Our approach detects frequently occurring words in retrieved captions, extracting entities more accurately without relying on a limited vocabulary.

2015) validation split, and the CLIP similarity-based KNN algorithm is used for retrieval. In the figure, there is a large difference between the distribution of features used after image-to-text retrieval and text-to-text retrieval, which shows that a modality gap exists between image and text.

To tackle this issue, we propose a novel approach called “Image-like Retrieval,” that addresses the modality gap between image and text data. We inject a noise into the CLIP text feature to act as a query in image feature distribution. Visualization results for this approach are shown in Fig. 2 right, demonstrating that our method exhibits a distribution highly similar to that of image-to-retrieval results and ground truth captions, unlike traditional text-to-text retrieval methods. Indeed, when our method is applied to the existing research (Wang et al., 2023), performance improvements are observed, as shown in Table 12.

Prior research (Wang et al., 2023) relies solely on retrieved captions, which may include wrong information in the input caption, potentially leading to inaccurate outputs. To address this, we design a *Fusion Module* that effectively integrates both

the original input and additional representations. Additionally, as shown by numerous studies (Fei et al., 2023; Ramos et al., 2023), prompts can clarify the information provided to the language model. We extract keywords from the input caption to construct a hard prompt, which is fed to the LLM, offering explicit guidance. This approach maximizes the utility of text data, guiding the model to generate accurate and relevant captions.

Guiding caption decoder with extracted entities from an image helps the model generate an accurate description of the image. However, we find that the previous works (Fei et al., 2023; Liu et al., 2023) show low entity detection precision, especially when the vocabulary is large as shown in Fig. 3. Therefore, we propose a Frequency-based Entity Filtering technique precisely utilizing entity information without relying on the vocabulary. During inference, we utilize retrieved sentences from images, parsing them into nouns and calculating their frequency. Then, we filter nouns with pre-defined thresholds and curate hard prompts for the text decoder. This simple method yields remarkable performance improvements.




Figure 2: The distribution of CLIP embedding features corresponding to images ■, paired captions ●, retrieved captions ● for a specific image, and the result of text-to-text retrieval ● and our Image-like Retrieval ●.

In summary, our contributions are as follows:

- We propose a novel approach, *Image-like Retrieval*, which achieves effects similar to image-to-text retrieval in text-only training. Then, we introduce a *Fusion Module* for interaction between existing and additional representations.
- We propose an entity filtering technique in inference, *Frequency-based Entity Filtering*, enhancing the language model by filtering frequently appearing entities in retrieved captions.
- Extensive evaluations show IFCap achieves state-of-the-art performance in various benchmarks, including video captioning.

## 2 Related work

### 2.1 Text-only Captioning

The advantage of CLIP (Radford et al., 2021) has been utilized in a variety of tasks, such as image captioning, image generation, and object detection. In the realm of image captioning, text-only training research is emerging that uses only text data for learning without image data, taking advantage of the CLIP characteristic that image embeddings and text embeddings are learned to be close. DeCap (Li et al., 2023) trains a text decoder using only textual data and introduces a support memory mechanism to project input images into the text embedding space during inference, facilitating the generation of captions. ViECap (Fei et al., 2023) recognizes the main entity of text data that comes as input and configures it as a prompt, allowing LLM to perform object-agnostic learning based on open vocabulary retrieval using CLIP.

### 2.2 Modality Gap

Vision language models such as CLIP aim to embed images and text closely in a shared space. How-

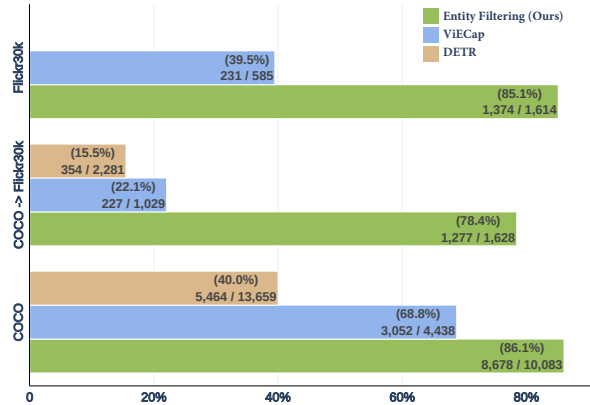


Figure 3: Precision of extracted entities in the COCO test set, total 5,000 images. If an extracted entity exists in the ground-truth caption, it counts as correct or else wrong. Three methods (Ours, ViECap2023, DETR2020) are compared with three different settings. Our method is illustrated in 3.3, and ViECap uses CLIP based classifier with the source domain’s vocabulary list. We follow the way SynTIC (Liu et al., 2023) uses DETR and employ the COCO vocabulary list. Due to the inaccessible vocabulary list of Flickr30k, DETR can’t be compared, and ViECap uses the VGOI (Zhang et al., 2021) vocabulary list in Flickr30k. Our method dominates the precision score and quantity of entities in every setting.

ever, it has been shown that these embeddings are located in two separate regions, with a significant gap between the modalities (Liang et al., 2022). This modality gap hinders the interaction between vision and text modalities and limits the quality of generated captions. Among the notable approaches addressing this issue, CapDec (Nukrai et al., 2022) assumes that the image embeddings paired with text embeddings are located within a small radius around the text embeddings and mitigates the gap with noise injection. CLOSE (Gu et al., 2022) highlights the low cosine similarity between images and their paired texts and uses a hyper-parameter-scaled noise injection technique to bridge the gap.

We focus on the modality gap for retrieval from a new perspective. Our goal is to perform text retrieval similar to image-to-text retrieval, considering the modality gap. The distinction from existing methods can be observed in Fig. 2 left.

### 2.3 Retrieval Augmented Generation

Retrieval has been used in diverse ways in NLP. Image captioning also benefits from retrieval modules by incorporating novel objects and new information into captions, allowing access to new domains without additional training. Retrieval is applied

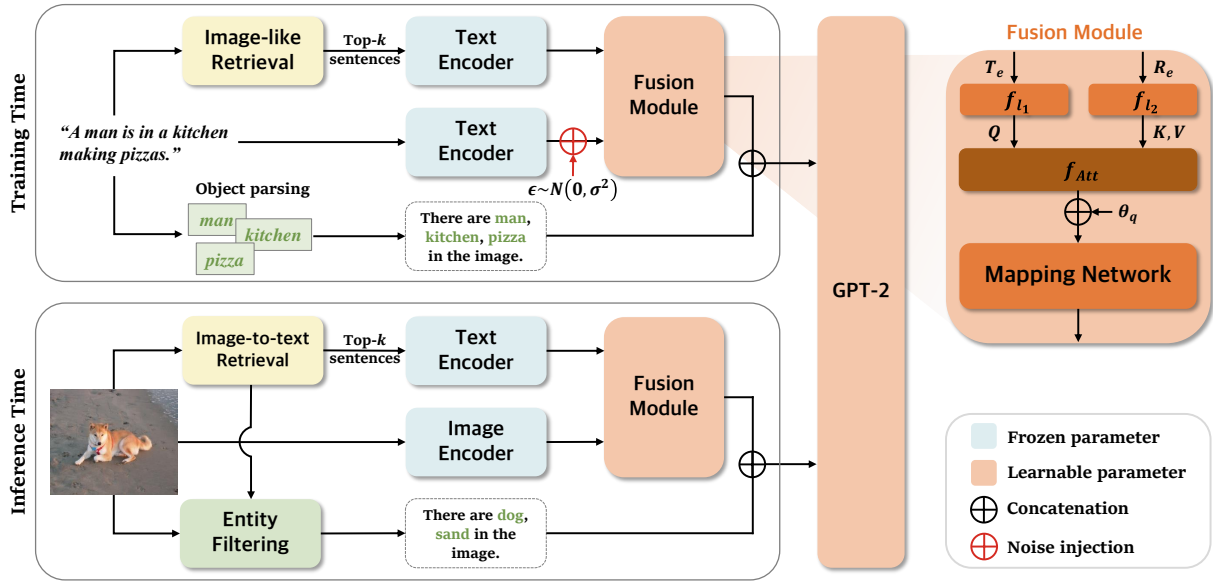


Figure 4: The overview of IFCap. During training, we extract nouns from the input text and retrieve  $k$  similar sentences using our Image-like Retrieval method. Extracted nouns are incorporated into a prompt template to form a hard prompt. Both the input text and retrieved sentences are encoded using the text encoder. These embeddings interact and combine through our Fusion Module before being fed into the LLM for sentence generation. During inference, we retrieve  $l$  sentences similar to the input image and construct a hard prompt by extracting entities via Frequency-based Entity Filtering from the retrieved sentences. The sentences are encoded using a text encoder, and the input image is encoded using an image encoder, followed by input into the Fusion Module. The subsequent process follows a procedure similar to the training phase.

in various ways in image captioning models. For instance, Smallcap (Ramos et al., 2023) retrieves captions relevant to the input image and uses them as instructions for the text decoder. In text-only image captioning, ViECap (Fei et al., 2023) retrieves novel objects from the input image and uses them as prompts, while Knight (Wang et al., 2023) uses retrieved captions as text features.

Most retrieval methods are based on image-to-text retrieval, but text-only captioning performs text-to-text retrieval. However, during inference, the modality gap caused by the input image leads to poor performance. Our method carefully addresses this issue to improve performance by considering the gap between image and text.

### 3 Methods

We propose a new text-only image captioning model, IFCap, which is illustrated in Fig. 4. During training, the model only utilizes text data, as is standard for text-only training models. First, we embed the input text using a text encoder. The text embeddings are then fed into a mapping network to close the gap between different modalities. Finally, the processed embeddings go through a caption decoder to generate the output caption.

Our IFCap utilizes a simple yet powerful retrieval mechanism and addresses the modality gap between image and text with Image-like Retrieval (Section 3.1). After performing Image-like Retrieval, we employ a Fusion Module (Section 3.2) to merge input embeddings with the retrieved features. During inference, we use the retrieved captions from the image to find accurate and detailed entities with Frequency-based Entity Filtering (Section 3.3).

#### 3.1 Image-like Retrieval (ILR)

While text-to-text retrieval can be effectively performed during training, it is likely to suffer from performance degradation during inference when an image is provided as input due to the modality gap. Therefore, Image-like Retrieval (ILR) aims to perform text-to-text retrieval in a manner that resembles image-to-text retrieval outcomes, given text input. For this, we propose an approach that inserts noise into the feature space of the input text, bringing it closer to the image feature space. The augmentation process is as follows:

First, we utilize the CLIP to embed the input text  $t_i$  and the text corpus  $\mathcal{T} = \{t_i\}_{i=1}^{N_c}$  with a text encoder  $\mathcal{E}_T$ . Then, we introduce noise  $\epsilon_r \sim N(0, \sigma_r^2)$  into



the embedding of input text  $T_i$ , aiming to adjust the text features to align more closely with the image feature space:

$$T_i = \mathcal{E}_T(t_i), \quad T_i^\epsilon = T_i + \epsilon_r. \quad (1)$$

Next, the retrieval step is performed using the noise-injected input text  $T_i^\epsilon$ . To identify the descriptions most relevant to  $T_i^\epsilon$ , the top- $k$  descriptions are retrieved by calculating the cosine similarity between  $T_i^\epsilon$  and all sentence embeddings in the text corpus. This process closely follows previous methods in image-to-text retrieval (Ramos et al., 2023), with the distinction that we perform retrieval based on  $T_i^\epsilon$  instead of images.

By utilizing this approach during training, we can enhance the ability of a model to provide image-like information even in a text-only training setting, thereby narrowing the modality gap and improving performance.

### 3.2 Fusion Module (FM)

In text-only image captioning, choosing which additional information to inject into the model and dealing with new representations with given data appropriately are important issues. To handle this problem, we use the attention mechanism (Vaswani et al., 2017) to fuse input text features and retrieved captions features to extract their meaningful interaction. The attention mechanism emphasizes certain important features, and due to its effectiveness, it has been widely utilized in the field of captioning (Xu et al., 2015).

We first encode input text and retrieved captions using CLIP (Radford et al., 2021) text encoder, then inject a Gaussian noise  $\epsilon \sim N(0, \sigma^2)$  to input text feature for relieving the modality gap between image and text. Then we adjust the dimension of the input text feature and retrieved captions feature to the embedding space of caption decoder with linear layers  $f_{l_1}$  and  $f_{l_2}$  respectively, and apply cross-attention  $f_{Att}$  with  $T_e$  as query and  $R_e$  as key, then create fusion representation  $F_e$  containing input text and retrieved captions. Finally,  $F_e$  is fed into a trainable Mapping Network, which encodes the overall contents of the given input. We can summarize this process with equations.

$$T_e = T_i + \epsilon, \quad R_e = \mathcal{E}_T(\text{ILR}(T_i)), \quad (2)$$

$$F_e = f_{Att}(f_{l_1}(T_e), f_{l_2}(R_e)), \quad (3)$$

$$F = \text{Map}(F_e; \theta_q). \quad (4)$$

The noun implies intuitive and explicit information about objects in the image. For employing the property of nouns, we extract entities in each training text corpus and input images. We build a hard prompt  $h$  with extracted entities  $E = \{e_1, e_2, \dots, e_n\}$  to make the model aware of existing entities in the image. With retrieved captions and hard prompts with entities, the model can learn the ability to generate proper captions without images. We use auto-regressive loss to optimize our projector and caption decoder. (Details about the Fusion Module are in Sec. 4.1).

$$L_\theta = -\frac{1}{N} \sum_{i=1}^N \log(y_i | \mathbf{F}; \mathbf{h}; y_{<i}; \theta). \quad (5)$$

### 3.3 Frequency-based Entity Filtering (EF)

After retrieving  $l$  captions from an image, we use grammar parser tools (e.g., NLTK Bird and Loper, 2004) to extract nouns from the retrieved sentences and calculate the frequency of these extracted nouns as  $F = [f_1, f_2, \dots, f_n]$ . We then select nouns that have a frequency larger than a predefined threshold and place them into a hard prompt.

**Heuristic threshold:** Since frequency is discrete, we can manually find the best threshold by conducting experiments with every possible threshold. This allows us to determine the global optimal threshold.

**Adaptive threshold:** We can use a heuristic threshold, but these thresholds are often unsuitable for different environments, and performing extensive experiments incurs unnecessary costs. Instead, we can estimate the common distribution of noun frequencies as certain probability distributions. We can assume frequencies follow  $N(\mu_F, \sigma_F^2)$ .

$$\tau_{\text{adap}} = \mu_F + \sigma_F. \quad (6)$$

Any nouns with a frequency larger than  $\tau_{\text{adap}}$ , which places them in the upper 15%, can be considered outliers. Using this adaptive threshold, we can implement a flexible threshold that fits various settings. However, it does not guarantee global optima, leading to a trade-off relationship between heuristic thresholds and adaptive thresholds.

## 4 Experiments

### 4.1 Implementation Details

While verifying the state-of-the-art performance of our model, we use CLIP (ViT-B/32) as the image encoder and GPT2<sub>base</sub> (Radford et al., 2019)

Method	Image Encoder	Text Decoder	COCO				Flickr30k			
			B@4	M	C	S	B@4	M	C	S
CapDec (2022)	RN50x4	GPT-2 <sub>Large</sub>	26.4	25.1	91.8	11.9	17.7	20.0	39.1	9.9
DeCap (2023)	ViT-B/32	Transformer <sub>Base</sub>	24.7	25.0	91.2	18.7	21.2	21.8	56.7	15.2
CLOSE (2022)	ViT-L/14	T5 <sub>base</sub>	-	-	95.3	-	-	-	-	-
ViECap (2023)	ViT-B/32	GPT-2 <sub>Base</sub>	27.2	24.8	92.9	18.2	21.4	20.1	47.9	13.6
MeaCap <sub>InvLM</sub> (2024)	ViT-B/32	GPT-2 <sub>Base</sub>	27.2	25.3	95.4	19.0	22.3	22.3	<u>59.4</u>	15.6
Knight (2023)	RN50x64	GPT-2 <sub>Large</sub>	27.8	<u>26.4</u>	98.9	<u>19.6</u>	22.6	<b>24.0</b>	56.3	16.3
ICSD <sup>♣</sup> (2023)	ViT-B/32	BERT <sub>Base</sub>	<u>29.9</u>	25.4	96.6	-	<b>25.2</b>	20.6	54.3	-
SynTIC <sup>♣†</sup> (2023)	ViT-B/32	Transformer <sub>L=4, H=4</sub>	<u>29.9</u>	25.8	<u>101.1</u>	19.3	22.3	22.4	56.6	<u>16.6</u>
IFCap	ViT-B/32	GPT-2 <sub>Base</sub>	<b>30.8</b>	<b>26.7</b>	<b>108.0</b>	<b>20.3</b>	<u>23.5</u>	<u>23.0</u>	<b>64.4</b>	<b>17.0</b>

Table 1: Result on the In-domain captioning including COCO test split and Flickr30k test split. Every result is copied from the original papers. ♣: Utilizes text-to-image generation model in the training time, †: Utilizes object detector during the training and inference time. IFCap achieves state-of-the-art in most metrics. The best number overall is in **bold** and the second best in underline.

Method	COCO $\Rightarrow$ Flickr				Flickr $\Rightarrow$ COCO			
	B@4	M	C	S	B@4	M	C	S
DeCap (2023)	16.3	17.9	35.7	11.1	12.1	18.0	44.4	10.9
ViECap (2023)	17.4	18.0	38.4	11.2	12.6	19.3	54.2	12.5
Knight (2023)	<u>21.1</u>	<b>22.0</b>	<u>48.9</u>	<u>14.2</u>	<u>19.0</u>	<u>22.8</u>	<u>64.4</u>	<u>15.1</u>
SynTIC (2023)	17.9	18.6	38.4	11.9	14.6	19.4	47.0	11.9
SynTIC- <i>TT</i>	19.4	20.2	43.2	13.9	<b>20.6</b>	21.3	<u>64.4</u>	14.3
IFCap <sup>♣</sup>	17.8	19.4	47.5	12.7	14.7	20.4	60.7	13.6
IFCap- <i>TT</i>	<b>21.2</b>	<u>21.8</u>	<b>59.2</b>	<b>15.6</b>	<u>19.0</u>	<b>23.0</b>	<b>76.3</b>	<b>17.3</b>

Table 2: Results on the Cross-domain captioning. *-TT*: models can access to target domain’s corpus during inference time. ♣: without Entity Filtering module in the inference time. IFCap achieves state-of-the-art in most metrics.

as the text decoder. Parameters in the image encoder are frozen during training, and the text decoder and Fusion Module are trained. We train a total of 5 epochs, learning rate of  $2 \times 10^{-5}$ , use scheduler for learning rate scheduler, AdamW optimizer (Kingma and Ba, 2014), and set batch size 80. We use a single NVIDIA RTX4090 with 24GB VRAM; it takes about an hour and uses 12GB of VRAM during training.

**Image-like Retrieval:** We first discover adequate  $\sigma_r$  for Image-like Retrieval. Based on our experiment (Fig. 5), we choose  $\sigma_r$  as 0.04 in most cases. We retrieve  $k$  sentences with noise-injected input text feature  $T_e$ .

**Fusion Module:** We project  $T_e \in \mathbb{R}^d$  and  $R_e \in \mathbb{R}^{d \times k}$  with  $f_{l_1}, f_{l_2}$  into  $\mathbb{R}^{d_{gpt}}, \mathbb{R}^{d_{gpt} \times k}$  respectively where  $d$  is the CLIP dimension and  $d_{gpt}$  is the dimension of GPT-2 embedding space. We use projected  $T_e$  as query and  $R_e$  as key in  $f_{Att}$  layer. Finally,  $F_e$  and  $\theta_q$  are concatenated and fed into

Method	COCO $\Rightarrow$ NoCaps Val							
	In		Near		Out		Entire	
	C	S	C	S	C	S	C	S
DeCap (2023)	65.2	-	47.8	-	25.8	-	45.9	-
CapDec (2022)	60.1	10.2	50.2	9.3	28.7	6.0	45.9	8.3
ViECap (2023)	61.1	10.4	64.3	9.9	65.0	8.6	66.2	9.5
IFCap <sup>♣</sup>	<b>70.1</b>	<b>11.2</b>	<b>72.5</b>	<b>10.9</b>	<b>72.1</b>	<b>9.6</b>	<b>74.0</b>	<b>10.5</b>

Table 3: Results on the NoCaps validation split. ♣: without Entity Filtering module in the inference time. IFCap achieves state of the art in every metric.

the Mapping Network, which consists of 8 layered transformers (Vaswani et al., 2017).

**Frequency-based Entity Filtering:** From the input image, we retrieve  $l$  sentences and extracted nouns to obtain frequency  $F$ . With the predefined threshold, we filter entities and build hard prompt  $h$ , providing more accurate and diverse entities to the caption decoder.

**Datasets, metrics** We evaluate our model in human-annotated datasets. For in-domain generalization, we test our model on MS-COCO (Chen et al., 2015), Flickr30k (Young et al., 2014), and utilize Karpathy split (Karpathy and Fei-Fei, 2015). Also, to check the model’s performance in the unseen scenarios, we use the NoCaps (Agrawal et al., 2019) validation set. For metrics, we use common image captioning metrics CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016), BLEU@ $n$  (Papineni et al., 2002), and METEOR (Banerjee and Lavie, 2005). More details about datasets and metrics are included in the appendix (Sec. A).

Method	MSR-VTT				MSVD			
	B@4	M	C	S	B@4	M	C	S
ZeroCap (2022b)	2.3	12.9	5.8	-	2.9	16.3	9.6	-
MAGIC (2022)	5.5	13.3	7.4	4.2	6.6	16.1	14.0	2.9
CLMs (2022)	6.2	17.8	10.1	6.5	7.0	16.4	20.0	3.1
CapDec (2022)	8.9	23.7	11.5	5.9	7.9	23.3	34.5	3.2
EPT (2022a)	3.0	14.6	11.3	-	3.0	17.8	17.4	-
Knight (2023)	25.4	<b>28.0</b>	31.9	<b>8.5</b>	37.7	<b>36.1</b>	63.8	5.0
IFCap	<b>27.1</b>	25.9	<b>38.9</b>	6.7	<b>40.6</b>	34.2	<b>83.9</b>	<b>6.3</b>

Table 4: Results on the Video captioning including MSR-VTT and MSVD. IFCap achieves state-of-the-art in most metrics.

Image-like Retrieval	Fusion Module	Entity Filtering	COCO			
			B@4	M	C	S
✓	✓	✓	<b>30.8</b>	<b>26.7</b>	<b>108.0</b>	<b>20.3</b>
✓		✓	29.2	26.0	104.0	19.9
✓	✓		28.5	26.0	102.0	20.0
		✓	27.2	24.7	97.3	18.5
✓			27.7	25.6	99.0	19.4
			27.2	24.8	92.9	18.2

Table 5: Ablation studies of the key components of IFCap.

## 4.2 Text-only Captioning

We compare our model with other state-of-the-art text-only image captioning models. CapDec (Nukrai et al., 2022) and ViECap (Fei et al., 2023) are based on Clipcap (Mokady et al., 2021). They use predefined Gaussian noise for aligning text and image features. Similarly, CLOSE (Gu et al., 2022) uses various noise settings, and De-Cap (Li et al., 2023) uses a memory bank. And a recent approach to text-only image captioning, Knight (Wang et al., 2023) only utilizes text features with a retrieval mechanism, also Mea-Cap (Zeng et al., 2024) processes retrieved sentences into Subject-Predicate-Object triplets and employs them as additional information. ICSD (Ma et al., 2023) and SynTIC (Liu et al., 2023) utilize text-to-image generation models like Stable Diffusion (Rombach et al., 2022) to close the gap.

## 4.3 In-domain Captioning

We benchmark our IFCap on in-domain settings in Table 1 including COCO and Flickr30k. We compare our methods with previous state-of-the-art in text-only image captioning. Our IFCap dominates every metric in the COCO dataset compared to models that utilize larger models (Gu et al., 2022; Wang et al., 2023) and have complex train-

Design Choice Reference	Pre- $\epsilon$	Post- $\epsilon$	Retrieval	COCO			
				B@4	M	C	S
ViECap				27.2	24.8	92.9	18.2
Smallcap			✓	23.5	24.2	88.5	18.2
Knight		✓	✓	26.0	24.6	92.9	18.3
Knight + ILR	✓	✓	✓	27.2	25.0	93.9	18.3
IFCap	✓		✓	<b>28.5</b>	<b>26.0</b>	<b>102.0</b>	<b>20.0</b>

Table 6: Importance of noise injection timing of **Image-like Retrieval**. **Pre- $\epsilon$**  refers to noise injection before retrieval, and **Post- $\epsilon$**  refers to noise injection to retrieved features.

$k$ retrieved sentences	COCO			
	B@4	M	C	S
3	28.1	25.7	100.0	19.5
<b>5</b>	<b>28.5</b>	<b>26.0</b>	<b>102.0</b>	<b>20.0</b>
7	28.2	<b>26.0</b>	101.7	19.8

Table 7: Ablation studies of the number of retrieved captions  $k$  for **Fusion Module**.

ing time (Ma et al., 2023; Liu et al., 2023). Also, in Flickr30k, IFCap shows decent performance in BLEU@4 and METEOR and achieves the best scores in CIDEr and SPICE.

## 4.4 Cross-domain Captioning

We validate IFCap’s transfer ability through diverse domains, including the NoCaps validation set and cross-domain from COCO  $\rightarrow$  Flickr30k and vice versa. In NoCaps, we use the same model trained in the COCO domain to test how the model recognizes unseen objects during training. In the NoCaps validation split, our IFCap performs the best in every metric and every domain compared to previous state-of-the-art text-only image captioning models (Li et al., 2023; Nukrai et al., 2022; Fei et al., 2023). Also, in cross-domain settings between COCO and Flickr30k, IFCap wins state-of-the-art in most metrics and the second best in some metrics.

## 4.5 Video Captioning

In video captioning, we train our model in the same manner as previous experiments. First, we perform Image-like Retrieval on the corpus from each video captioning dataset MSVD (Wu et al., 2017) and MSR-VTT (Xu et al., 2016). For inference time, we sample 5 images from input video and calculate the average of their CLIP image features. We also retrieve 5 sentences from each sampled image, 25 in total, and also calculate the average of CLIP text features per image. Most of the met-

Transformer # Layers	Cross-Attention # Layers	COCO			
		B@4	M	C	S
1	1	23.9	24.6	86.9	17.8
	4	26.2	24.4	92.8	18.0
2	1	27.4	24.9	95.0	18.5
	4	26.4	24.9	95.5	18.4
4	1	27.4	25.5	99.7	19.1
	4	27.9	25.8	99.1	19.4
8	1	28.3	<b>26.0</b>	<b>102.0</b>	<b>20.0</b>
	4	<b>28.4</b>	25.7	100.6	19.5

Table 8: Ablation studies of the number of transformer layers and cross-attention layers of the **Fusion Module**.

$l$ retrieved sentences	COCO				Flickr			
	B@4	M	C	S	B@4	M	C	S
5	29.9	26.4	106.1	20.2	<b>23.5</b>	22.2	61.9	16.0
7	30.3	26.5	107.2	<b>20.3</b>	<b>23.5</b>	<b>23.0</b>	<b>64.4</b>	<b>17.0</b>
9	<b>30.8</b>	<b>26.7</b>	<b>108.0</b>	<b>20.3</b>	23.4	22.6	62.9	16.6

Table 9: Ablation studies of the number of retrieved sentences  $l$  for **Entity Filtering**.

rics in both datasets, IFCap, fulfills state-of-the-art performance, except METEOR.

#### 4.6 Ablation Study

We conduct extensive experiments to identify the impact of each key component in IFCap, Image-like Retrieval (**ILR**), Fusion Module(**FM**), and Frequency-based Entity Filtering(**EF**). Also, for each component, we search the best hyperparameter in the COCO test split with an in-domain setting.

**Key Components:** We check the strength of each component by detaching from our best model, which consists of all 3 components Table 5. First, removing **FM**, we simply concatenate the input text feature and retrieved features after applying dimension mapping layers  $f_{i_1}$  and  $f_{i_2}$  and passing it to the caption decoder. Removing **EF** is simply applying entity extraction via CLIP classifier like (Fei et al., 2023) does. Demounting **ILR** makes inaccessible to retrieval features solely using input features; hence **FM** can't exist without **ILR**. Adding more components into the baseline, we can explicitly notice performance improvement. So, using all three key components constitutes a state-of-the-art model, which is IFCap. Note that IFCap has 2 variants, IFCap and IFCap\*, with **EF** and without **EF** respectively. To see a full comparison of various datasets, including in-domain, cross-domain, and

$\tau$	COCO				Flickr30k			
	B@4	M	C	S	B@4	M	C	S
1	6.5	18.7	6.4	17.0	6.8	18.9	3.9	15.4
2	21.4	26.5	80.3	21.0	18.9	<b>23.4</b>	52.2	<b>17.9</b>
3	28.1	26.8	103.6	<b>21.1</b>	23.5	23.0	<b>64.4</b>	17.0
4	30.2	<b>26.7</b>	107.7	20.7	<b>23.8</b>	22.3	61.1	15.9
5	<b>30.8</b>	<b>26.7</b>	<b>108.0</b>	20.3	<b>23.8</b>	21.9	59.1	15.3
6	30.4	26.4	106.2	19.9	23.6	21.7	57.3	15.0
7	30.0	26.1	104.6	19.6	23.6	21.6	56.5	14.8
8	29.8	26.0	103.4	19.4	23.7	21.6	55.9	14.7

Table 10: Ablation studies of heuristic threshold  $\tau$  of **Entity Filtering**.

$\tau_{\text{adap}}$	COCO				Flickr30k			
	B@4	M	C	S	B@4	M	C	S
<i>Lognormal</i> ( $\mu, \sigma^2$ )								
$\mu$	22.0	26.6	83.8	<b>21.1</b>	19.0	<b>23.4</b>	52.7	<b>17.9</b>
$\mu + \sigma$	29.1	<b>26.7</b>	<b>106.6</b>	20.7	22.0	22.9	<b>63.0</b>	17.2
$\mu + 2\sigma$	<b>29.6</b>	26.1	103.5	19.6	<b>23.3</b>	21.8	58.1	15.3
<i>N</i> ( $\mu, \sigma^2$ )								
$\mu$	24.9	<b>26.7</b>	95.9	<b>21.1</b>	19.2	<b>23.2</b>	55.6	<b>17.7</b>
$\mu + \sigma$	<b>30.1</b>	26.6	<b>107.5</b>	20.4	22.3	22.5	<b>62.3</b>	16.4
$\mu + 2\sigma$	29.8	26.2	104.7	19.7	<b>23.4</b>	21.9	58.5	15.5
Best (H)	30.8	26.7	108.0	20.3	23.5	23.0	64.4	17.0

Table 11: Ablation studies of adaptive threshold  $\tau_{\text{adap}}$  of **Entity Filtering**.

video captioning, refer to Table 14.

**Image-like Retrieval:** It is crucial to identify adequate timing for injecting noise into text features for successful text-to-text retrieval that imitates image-to-text retrieval. We can split injecting timing into Pre- $\epsilon$  and Post- $\epsilon$ . We find that our setting which injects noise before performing retrieval is the best among all possible combinations. We can verify this in Table 6. The first column of the table indicates how the model performs retrieval, just for easy understanding of noise injection in retrieval.

**Fusion Module:** We utilize a cross-attention layer and transformer layer for mapping the network. In Table 8, we try multiple combinations of the number of each layer. The more layers we use, the more performance gain we can get until the number of transformer layers is 4. The performance gain is also observed when we use 8 transformer layers but it is so slight. Increasing the number of cross-attention layers is effective when the transformer layer is small, but the tendency does not last while the transformer layer grows. We conclude using 8 transformer layers and






Figure 5: Hyper-parameter search for finding best  $\sigma_r$  used in Image-like Retrieval. All experiments are conducted with the COCO test set. The X-axis denotes  $\sigma_r^2$ , and the Y-axis denotes scores of commonly used captioning metrics BLEU@4 (B@4), METEOR (M), CIDEr (C), and SPICE (S).

a single cross-attention layer shows the best performance. For a fair comparison, we detach the EF module. Also, the number of retrieved captions is crucial. We conduct ablation studies to find optimal  $k$ , which can be found in Table 7.

**Frequency-based Entity Filtering:** We need to choose 1) how many retrieved sentences  $l$ , to use and 2) the threshold  $\tau$ , for filtering nouns for EF to extract accurate and diverse entities. The former can be found in Table 9, note that in different domains, optimal  $l$  may vary. For the COCO domain, using  $l$  as 9 shows the best performance, while 7 is the best in Flickr30k.

We find the best threshold setting in a heuristic and adaptive way. In the former case Table 10, we set  $\tau$  ranging from 1 to 8, which is the minimum and maximum value of the given setting. Above 8, performance freeze due to none of the entities being retrieved. In the COCO test, we use  $l = 9$  and  $l = 7$  in the Flickr30k test split. We notice that each domain has different optimal  $\tau$ , COCO at 5 and Flickr30k at 3 for the CIDEr score. In contrast to the heuristic way, we can assume such distribution exists from frequencies  $F$ . We try Gaussian distribution and Log-normal distribution with  $\mu$ ,  $\mu + \sigma$ , and  $\mu + 2\sigma$ , capturing upper 50%, 15.8%, and 2.2% based on the frequency of entity. In Table 11, we observe  $\tau_{\text{adap}} = \mu + \sigma$  almost reproduces the performance of global optimal in the heuristic threshold. If ground truth does not exist or computing resource is limited, the adaptive threshold becomes attractive.

## 5 Conclusion

In this paper, we propose a zero-shot captioning method, IFCap, through text-only training. IFCap performs *Image-like Retrieval* to address the gap between image-to-text retrieval and text-to-text retrieval, *Fusion Module* for interaction between existing and additional representations, and *Frequency-based Entity Filtering* during inference time to extract frequently occurring entities from the retrieved sentences. Our method can be easily applied to various tasks and provides valuable guidance for retrieval-based methods in a text-only setting. It offers clear and precise information to LLMs without relying on a limited vocabulary. The simplicity and robustness of IFCap are demonstrated through state-of-the-art performance across various datasets in image and video captioning. The future direction of our method includes the extension of our method on more complex datasets, such as region-based captioning (Kim et al., 2019a, 2021) or visual question answering (Cho et al., 2023a,b), which suffer from data issues.

## 6 Limitations

We demonstrate that IFCap exhibits superior performance across various image captioning and video captioning datasets compared to other zero-shot image captioning models with text-only training. However, the optimal value of  $\epsilon_r$  for Image-like Retrieval currently requires a heuristic approach to determine. We leave the task of finding a more convenient method for determining the optimal  $\epsilon_r$  as future work to further improve image captioning models with text-only training.

## Acknowledgements

This was partly supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government(MSIT) (No.RS-2020-II201373, Artificial Intelligence Graduate School Program(Hanyang University)) and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00245661).

## References

Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. 2019. No-caps: Novel object captioning at scale. In *Proceed-*

- ings of the *IEEE/CVF international conference on computer vision*, pages 8948–8957.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 382–398. Springer.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Steven Bird and Edward Loper. 2004. **NLTK: The natural language toolkit**. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Jae Won Cho, Dawit Mureja Argaw, Youngtaek Oh, Dong-Jin Kim, and In So Kweon. 2023a. Empirical study on using adapters for debiased visual question answering. *Computer Vision and Image Understanding*, 237:103842.
- Jae Won Cho, Dong-Jin Kim, Hyeonggon Ryu, and In So Kweon. 2023b. Generative bias for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11681–11690.
- Junjie Fei, Teng Wang, Jinrui Zhang, Zhenyu He, Chengjie Wang, and Feng Zheng. 2023. Transferable decoding with visual entities for zero-shot image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3136–3146.
- Sophia Gu, Christopher Clark, and Aniruddha Kembhavi. 2022. I can’t believe there’s no images! learning visual tasks using only language supervision. *arXiv preprint arXiv:2211.09778*.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137.
- Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. 2019a. Dense relational captioning: Triple-stream networks for relationship-based captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6271–6280.
- Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. 2019b. Image captioning with very scarce supervised data: Adversarial semi-supervised learning approach. 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)*.
- Dong-Jin Kim, Tae-Hyun Oh, Jinsoo Choi, and In So Kweon. 2021. Dense relational image captioning via multi-task triple-stream networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7348–7362.
- Dong-Jin Kim, Tae-Hyun Oh, Jinsoo Choi, and In So Kweon. 2024. Semi-supervised image captioning by adversarially propagating labeled data. *IEEE Access*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Wei Li, Linchao Zhu, Longyin Wen, and Yi Yang. 2023. Decap: Decoding clip latents for zero-shot captioning via text-only training. *arXiv preprint arXiv:2303.03032*.
- Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35:17612–17625.
- Zhiyue Liu, Jinyuan Liu, and Fanrong Ma. 2023. Improving cross-modal alignment with synthetic pairs for text-only image captioning. *Preprint*, arXiv:2312.08865.
- Ziyang Luo, Zhipeng Hu, Yadong Xi, Rongsheng Zhang, and Jing Ma. 2023. I-tuning: Tuning frozen language models with image for lightweight image captioning. In *ICASSP 2023–2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Feipeng Ma, Yizhou Zhou, Fengyun Rao, Yueyi Zhang, and Xiaoyan Sun. 2023. Image captioning with multi-context synthetic data. *Preprint*, arXiv:2305.18072.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*.

- David Nukrai, Ron Mokady, and Amir Globerson. 2022. Text-only training for image captioning using noise-injected clip. *arXiv preprint arXiv:2211.00575*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Rita Ramos, Bruno Martins, Desmond Elliott, and Yova Kementchedjheva. 2023. Smallcap: lightweight image captioning prompted with retrieval augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2840–2849.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Yixuan Su, Tian Lan, Yahui Liu, Fangyu Liu, Dani Yogatama, Yan Wang, Lingpeng Kong, and Nigel Collier. 2022. Language models can see: Plugging visual controls in text generation. *arXiv preprint arXiv:2205.02655*.
- Yoad Tewel, Yoav Shalev, Roy Nadler, Idan Schwartz, and Lior Wolf. 2022a. Zero-shot video captioning with evolving pseudo-tokens. *arXiv preprint arXiv:2207.11100*.
- Yoad Tewel, Yoav Shalev, Idan Schwartz, and Lior Wolf. 2022b. Zerocap: Zero-shot image-to-text generation for visual-semantic arithmetic. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17918–17928.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Junyang Wang, Ming Yan, Yi Zhang, and Jitao Sang. 2023. From association to generation: Text-only captioning by unsupervised cross-modal mapping. *arXiv preprint arXiv:2304.13273*.
- Junyang Wang, Yi Zhang, Ming Yan, Ji Zhang, and Jitao Sang. 2022. Zero-shot image captioning by anchor-augmented vision-language space alignment. *arXiv preprint arXiv:2211.07275*.
- Zuxuan Wu, Ting Yao, Yanwei Fu, and Yu-Gang Jiang. 2017. *Deep learning for video classification and captioning*, page 3–29. Association for Computing Machinery and Morgan & Claypool.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Zequan Zeng, Yan Xie, Hao Zhang, Chiyu Chen, Zhengjue Wang, and Bo Chen. 2024. Meacap: Memory-augmented zero-shot image captioning. *arXiv preprint arXiv:2403.03715*.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. *Vinvl: Revisiting visual representations in vision-language models*. *Preprint*, arXiv:2101.00529.

## A Image-like Retrieval

Method	COCO			
	B@4	M	C	S
Knight	27.8	<b>26.4</b>	98.9	19.6
Knight + ILR	<b>29.8</b>	25.6	<b>102.7</b>	<b>19.7</b>

Table 12: Effect of **Image-like Retrieval** on Knight.

HyperParameters	COCO	Flickr30k	NoCaps	MSVD	MSR-VTT
Epochs	5	30	-	10	10
$l$	9	7	7	7	7
$\tau$	5	3	3	5	6

Table 13: Hyperparameter table.

Method	In-domain				Cross-domain										Video Captioning							
	COCO		Flickr		COCO $\Rightarrow$ NoCaps Val										COCO $\Rightarrow$ Flickr		Flickr $\Rightarrow$ COCO		MSR-VTT		MSVD	
	C	S	C	S	In		Near		Out		Entire		C	S	C	S	C	S	C	S		
ViECap	92.9	18.2	47.9	13.6	61.1	10.4	64.3	9.9	65.0	8.6	66.2	9.5	38.4	11.2	54.2	12.5	-	-	-	-		
Knight	98.9	19.6	56.3	16.3	-	-	-	-	-	-	-	-	48.9	14.2	64.4	15.1	31.9	<b>8.5</b>	63.8	5.0		
IFCap*	102.0	20.0	59.8	15.8	70.1	11.2	<b>72.5</b>	10.9	<b>72.1</b>	<b>9.6</b>	<b>74.0</b>	10.5	47.5	12.7	60.7	13.6	20.8	4.1	40.2	3.4		
IFCap	<b>108.0</b>	<b>20.3</b>	<b>64.4</b>	<b>17.0</b>	<b>75.8</b>	<b>12.4</b>	72.3	<b>11.6</b>	60.2	8.9	70.5	<b>10.8</b>	<b>59.2</b>	<b>15.6</b>	<b>76.3</b>	<b>17.3</b>	<b>38.9</b>	6.7	<b>83.9</b>	<b>6.3</b>		

Table 14: Overall comparison among baselines and IFCap.  $\star$ : without Entity Filtering module in the inference time.

We observe that Image-like Retrieval is also applicable to other models that employ text-to-text retrieval (Wang et al., 2023). Based on Fig. 5, we perform **ILR** with  $\epsilon_r = 0.04$  in the training time of Knight. In the COCO test set, every metric except METEOR is improved compared to vanilla Knight (Wang et al., 2023), verifying the effectiveness of our **ILR**.

## B Hyperparameter

We include the details about our experiments in each dataset in Table 13.

## C Comparison with Baselines

We compare baselines (Fei et al., 2023; Wang et al., 2023) with IFCap and IFCap\* in every domain, including in-domain captioning, cross-domain captioning, and video captioning. Results can be found in Table 14.

## D Qualitative Results

We show additional qualitative results in Fig. 6.





**Knight** : A silver passenger train traveling down a track next to an elevated walkway.  
**ViECap** : A car is shown in front of a large billboard.  
**ViECap entity** : []  
**IFCap** : A monorail train traveling down tracks next to a building.  
**IFCap entity** : [monorail, train]  
**GT** : A monorail making it's way down the track above a bunch of cars.



**Knight** : A man with a beard and a dog on a couch.  
**ViECap** : A man standing next to a brown and white dog.  
**ViECap entity** : [dog]  
**IFCap** : A man and a dog are smiling in front of a Christmas tree.  
**IFCap entity** : [man, dog]  
**GT** : A man in front of a Christmas tree with his dog.



**Knight** : A view of a mountain range with an airplane in the background.  
**ViECap** : A large airplane flying through a blue sky.  
**ViECap entity** : [airplane]  
**IFCap** : The wing of an airplane with mountains in the background.  
**IFCap entity** : [mountain, wing, airplane]  
**GT** : The view out of an airplane with part of the wing.




**Knight** : A giraffe standing next to a large tree.  
**ViECap** : Two giraffes standing next to each other in a grassy area.  
**ViECap entity** : [giraffe]  
**IFCap** : A giraffe standing next to a tree in the water.  
**IFCap entity** : [giraffe, tree, water]  
**GT** : A giraffe in a field next to tree and body of water.




**Knight** : A group of men racing each other on a course.  
**ViECap** : A skier in a red jacket is skiing down a hill.  
**ViECap entity** : [skis]  
**IFCap** : Two cross country skiers racing down a hill.  
**IFCap entity** : [country]  
**GT** : Two cross country skiers heading onto the trail.



**Knight** : A motorcycle is parked on the side of the road next to a tree.  
**ViECap** : A basket full of bananas hanging from a tree.  
**ViECap entity** : []  
**IFCap** : A motorcycle that is sitting on top of a fence.  
**IFCap entity** : [motorcycle]  
**GT** : A motorcycle sitting on top of a fence as décor.



**Knight** : A group of traffic lights sitting on top of a road.  
**ViECap** : A street filled with traffic lights next to a tall building.  
**ViECap entity** : [traffic light]  
**IFCap** : A bunch of traffic lights at an intersection.  
**IFCap entity** : [light, intersection]  
**GT** : A photo taken from one vehicle of another at an intersection.



**Knight** : A bowl of fruit sitting on top of a counter.  
**ViECap** : A close up of fruits and vegetables on a table.  
**ViECap entity** : []  
**IFCap** : A close up of a bowl of oranges and apples on a counter.  
**IFCap entity** : [bowl, apple, orange, counter]  
**GT** : a bowl of apples and a bowl of oranges.



**Knight** : A black vase with a white flower in it.  
**ViECap** : A black and silver spoon with a tooth brush in it.  
**ViECap entity** : [spoon]  
**IFCap** : A black and white vase with a flower in it.  
**IFCap entity** : [vase, flower]  
**GT** : Thin black and white vase with black flowers.

Figure 6: Qualitative result on the COCO test set. We highlight the retrieved entities and their appearance in the generated captions with **IFCap** , **ViECap** and **Intersection** .