

Unleashing Text-to-Image Diffusion Prior for Zero-Shot Image Captioning

Jianjie Luo¹, Jingwen Chen², Yehao Li², Yingwei Pan², Jianlin Feng^{1†},
Hongyang Chao¹, and Ting Yao^{2†}

¹ Sun Yat-sen University, Guangzhou, China

² HiDream.ai Inc.

luojj26@mail3.sysu.edu.cn, {fengjlin, isschhy}@mail.sysu.edu.cn
{chenjingwen, liyehao, pandy, tiyao}@hidream.ai

Abstract. Recently, zero-shot image captioning has gained increasing attention, where only text data is available for training. The remarkable progress in text-to-image diffusion model presents the potential to resolve this task by employing synthetic image-caption pairs generated by this pre-trained prior. Nonetheless, the defective details in the salient regions of the synthetic images introduce semantic misalignment between the synthetic image and text, leading to compromised results. To address this challenge, we propose a novel **Patch-wise Cross-modal feature Mix-up (PCM)** mechanism to adaptively mitigate the unfaithful contents in a fine-grained manner during training, which can be integrated into most of encoder-decoder frameworks, introducing our PCM-Net. Specifically, for each input image, salient visual concepts in the image are first detected considering the image-text similarity in CLIP space. Next, the patch-wise visual features of the input image are selectively fused with the textual features of the salient visual concepts, leading to a mixed-up feature map with less defective content. Finally, a visual-semantic encoder is exploited to refine the derived feature map, which is further incorporated into the sentence decoder for caption generation. Additionally, to facilitate the model training with synthetic data, a novel CLIP-weighted cross-entropy loss is devised to prioritize the high-quality image-text pairs over the low-quality counterparts. Extensive experiments on MSCOCO and Flickr30k datasets demonstrate the superiority of our PCM-Net compared with state-of-the-art VLMs-based approaches. It is noteworthy that our PCM-Net ranks first in both in-domain and cross-domain zero-shot image captioning. The synthetic dataset SynthImgCap and code are available at <https://jianjieluo.github.io/SynthImgCap>.

Keywords: Zero-shot Image Captioning · CLIP · Attention Mechanism

1 Introduction

Image captioning aims to describe the content of an image with natural language. The conventional practice to resolve this task is to train an encoder-decoder model in an end-to-end manner with well-aligned pairs of images and

[†]J. Feng and T. Yao are the corresponding authors.




Fig. 1: Training paradigms for image captioning: (a) Training with well-aligned image-sentence pairs for Supervised Image Captioning (SIC); (b) Training with text-only data in VLMs-based model for Zero-shot Image Captioning (ZIC); (c) Training with synthetic image-text pairs in our PCM-Net for ZIC.

texts, widely known as Supervised Image Captioning (SIC) (Figure 1 (a)). More advanced methods are proposed [3, 13, 18–20, 25, 26, 28, 46] to better capture the correlation between visual and linguistic elements. However, the extensive labor required for the data collection makes it difficult for SIC model to scale up.

To address this challenge, Zero-shot Image Captioning (ZIC) is introduced and manages to learn the image-to-text mappings without relying on annotated data, which has attracted great interest recently. The popular recipe for ZIC is to bridge vision and language through an intermediate latent representation using unpaired images and sentences [8, 9]. Another type of VLMs-based solution (Figure 1 (b)) for ZIC involves developing an image captioning model with high-quality text-only data, where the pre-trained cross-modal prior from large-scale vision-language models (VLMs) like CLIP [32] is leveraged to align the text and image. In VLMs-based approaches, the captioning model is trained to generate caption with the corresponding textual feature in CLIP space, which will be replaced with the visual feature of the input image during inference. It is assumed that the textual and visual features sharing similar semantics should be close in CLIP space. However, this assumption may not always hold due to the inherent modality gap [21], leading to *discrepancy between training and inference*.

Motivated by the powerful capability of text-to-image diffusion model [33, 35, 36] in producing images conditioning on text prompts, we propose to generate synthetic images for the text data, and train the captioning model with the obtained synthetic image-text pair to mitigate the above issue. Nevertheless, the presence of flawed or unfaithful details in the salient regions of the synthetic images enlarges not only the distribution discrepancy between the synthetic and real images but also the semantic gap between the synthetic image and the text, thereby hindering the learning of visual-semantic alignment. Since generating high-fidelity images usually requires rigorous prompt engineering, our work primarily aims to automatically mitigate the unfaithful contents in the synthetic images during training. To achieve this goal, a novel patch-wise cross-modal feature mix-up (PCM) mechanism is devised in this paper, which can be integrated into most of encoder-decoder frameworks, introducing our PCM-Net (Figure 1 (c)). Specifically, given an input image, a set of salient visual concepts is con-

structed by performing zero-shot entity classification of images in CLIP space. Then, the fine-grained patch features of the synthetic image are mixed up with the textual features of the salient visual concepts depending on patch-wise cross-modal similarity, which effectively reduces the flawed or unfaithful details in the synthetic images. Finally, the derived features are passed through an encoder-decoder network for caption generation. Compared to the former VLMs-based approaches which rely on the global feature of CLIP, our PCM-Net mitigates the modality gap in pixel space by using synthetic images for training, where fine-grained visual features are available to boost the visual-semantic alignment. Additionally, we propose a new CLIP-weighted cross-entropy loss to facilitate the robust training of the captioning model with noisy synthetic data by adaptively re-weighting the loss for each word according to the predicted word distribution and the semantic relevance between the text and the synthetic image.

To summarize, our contributions are as follows: **(I)** We propose leveraging the powerful text-to-image diffusion model to generate synthetic images for text-only data, and further train the captioning model with these synthetic pairs, which mitigates the discrepancy between training and inference for zero-shot image captioning. **(II)** We propose a patch-wise cross-modal feature mix-up (PCM) mechanism to close the gap between synthetic and real images by automatically mitigating the defective contents in the synthetic images. **(III)** We propose a novel CLIP-weighted cross-entropy loss (CXE) to improve the training of captioning model with the noisy synthetic data. **(IV)** Extensive experiments conducted on MSCOCO and Flickr30k demonstrate our PCM-Net’s effectiveness.

2 Related Work

2.1 Supervised Image Captioning

Image captioning, a fundamental task in the vision-language domain, aims to describe semantic content within an image using natural language. Conventional image captioning methods train an encoder-decoder neural network on well-matched image-sentence pairs in a supervised manner. Early attempts [6, 14, 41, 43] leverage CNNs to encode visual content and RNNs to decode sentences. Later researches focus on enriching the visual feature representation by incorporating semantic attributes [49, 50] or object region features [3]. The performances of supervised image captioning are further boosted by modeling visual relationships in the image [45, 47, 48]. Recently, inspired by the success of the Transformer architecture [39], numerous Transformer-based image captioning models have emerged [5, 10, 26, 29, 44]. Despite the significant progress in supervised learning tasks, the collection of annotated training datasets is labor-intensive and time-consuming, which restricts the scalability of the captioning model.

2.2 Zero-shot Image Captioning

Zero-shot image captioning aims to learn the captioning model without human-annotated data. Recently, the vision-language pre-training models (VLMs) trained

on large-scale image-text pairs crawled from the web have demonstrated great zero-shot capability in image captioning. For example, SimVLM [42] pre-trains the VLM with a single prefix language modeling objective and infers the image caption without fine-tuning. BLIP-2 [16] only pre-trains a Querying Transformer to bridge the modality gap between the pre-trained image encoder and the large language model to further boost zero-shot image captioning. ZeroCap [38] instead steers GPT-2 to generate sentences related to the visual content under the guidance of CLIP [32] without further training. However, due to the noise in the web data, these VLMs-based approaches struggle to generate captions with correct grammar without fine-tuning on paired data.

Therefore, another approach for ZIC is to build an image captioning model based on high-quality text-only data. The prevalent methods [17, 27, 37] leverage the pre-trained cross-modal prior of the VLMs to facilitate the learning of visual-semantic alignment for zero-shot image captioning. For instance, MAGIC [37] fine-tunes GPT-2 on the training corpus, and modulates the probability of each candidate word during decoding with the cross-modal similarity in CLIP space. Furthermore, several works [7, 17, 27, 51] propose to train the captioning model with the textual features of the target caption as input at training stage, and replace it with the visual feature of the input image during inference. The pioneering CapDec [27] applies noise injection to input textual features in training to overcome the modality gap [21], while DeCap [17] projects the visual feature into the textual feature space based on the training corpus at inference for the same purpose. CgT-GAN [51] instead mitigates the modality gap by incorporating real images as input in training via additional reinforcement learning rewarded by cross-modal similarity. ViECap [7] builds a text prompt from visual entities to trigger the transferability of GPT-2. Most recently, SynTIC [24] employs the Stable Diffusion [35] model to generate synthetic images for training, and further minimize the global feature distance between synthetic and real images through contrastive learning techniques.

Summary. SynTIC [24] is most related to our work, which capitalizes on global features of the synthetic images for ZIC. Going beyond SynTIC, our PCM-Net explores spatial visual features from CLIP and leverages its capability of cross-modal alignment to mitigate flawed content in synthetic images, addressing distribution discrepancies with real images at a fine-grained level. Additionally, we introduce a novel CLIP-weighted cross-entropy loss to improve the robustness and performance of the captioning model when trained on noisy synthetic data.

3 Method

Zero-shot image captioning task aims to train an image-to-text captioning model on a text corpus only, without the associated images. Most existing methods learn to generate the caption conditioned on the corresponding textual feature during training, framing it into a task of caption reconstruction. While at inference stage, the textual feature is simply replaced with the global visual feature of the input image within the same multi-modal feature space. As a result, seman-

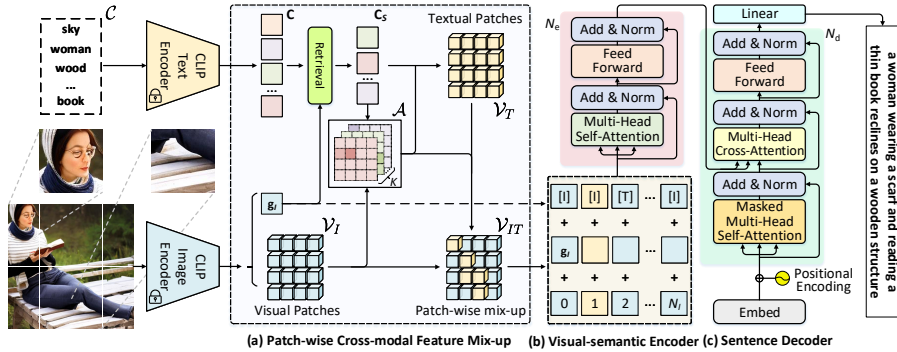


Fig. 2: An overview of our proposed PCM-Net. The flawed or unfaithful patches (e.g., poor facial details or missing limbs) in the salient regions of synthetic images would be replaced by semantically aligned textual patches in (a) Patch-wise Cross-modal Feature Mix-up. The mixed-up features are further encoded by (b) Visual-semantic Encoder, followed by (c) Sentence Decoder for caption generation.

tic misalignment arises during inference due to the inherent modality gap [21] between the caption and the image in the multi-modal space, leading to suboptimal results. To address this issue, we propose leveraging synthetic image-caption pairs generated by text-to-image diffusion models [12, 35] for training instead, unleashing the pre-trained diffusion prior for ZIC. However, the distribution discrepancy between the synthetic and real images presents a challenge for training the captioning model. In this paper, we devise a novel Patch-wise Cross-modal Feature Mix-up mechanism to bridge this gap and seamlessly incorporate it into an encoder-decoder captioning framework, introducing our proposed PCM-Net. The framework of our PCM-Net is illustrated in Figure 2. In this section, we first demonstrate the details about how to build the synthetic dataset, followed by the specifics of our PCM-Net. Lastly, the training loss is elaborated.

3.1 Synthetic Image Generation

The recent progress in text-to-image diffusion models [12, 35] has demonstrated the capability of diffusion models to generate high-fidelity images that are semantically aligned with the given texts. This highlights the promise of employing diffusion models to produce synthetic image-sentence pairs that could benefit the cross-modal vision-language learning in zero-shot image captioning. Let $\mathcal{S} = \{S_i\}_{i=1}^N$ denote a text corpus including N sentences for zero-shot image captioning, where each S_i consists of N_s words. Specifically, an off-the-shelf text-to-image generation model \mathcal{G} (i.e., Stable Diffusion [35]) is exploited to generate a synthetic image I_i for each sentence $S_i \in \mathcal{S}$, resulting in a synthetic dataset $\mathcal{D} = \{(I_i, S_i)\}_{i=1}^N$. Please note that all the experiments are conducted on \mathcal{D} .

3.2 Patch-wise Cross-modal Feature Mix-up

We have found that straightforwardly exploiting synthetic images for training yields compromised results in our experiments. The reason is that the flawed or unfaithful details in the salient areas of these images significantly enlarge not only the discrepancy between the synthetic and real images but also the semantic gap between the synthetic image and the sentence, hindering the learning of visual-semantic alignment. However, high-fidelity image generation usually requires intricate prompt engineering, and a universal solution for different diffusion models has not been established yet. Therefore, we propose a novel Patch-wise Cross-modal feature Mix-up (PCM) mechanism that automatically mitigates the possible irrelevant or defective contents in the salient areas of synthetic images during training, which avoids the complicated prompt engineering or sophisticated model design for faithful text-to-image generation. Specifically, the visual features of the holistic image I and its local patches are first extracted by the image encoder of CLIP. Let \mathbf{g}_I and $\mathcal{V}_I = \{\mathbf{v}_j\}_{j=1}^{N_I}$ denote the global image feature and the grid feature map with N_I local patch features, respectively.

Salient Visual Concepts Detection. To bypass the unfaithful or defective contents in the salient areas of the synthetic images, salient visual concepts of each image are required to be first detected, which is achieved by zero-shot entity classification in CLIP space. Specifically, a visual concept vocabulary $\mathcal{C} = \{c_j\}_{j=1}^{N_c}$ is built from the high-frequency nouns in the training corpus \mathcal{S} , where N_c is the concept vocabulary size. Then, the cosine similarity between the template sentence "A photo of $\{c_j\}$ " for each $c_j \in \mathcal{C}$ and the global visual feature \mathbf{g}_I is calculated to retrieve top- K salient visual concepts $\mathcal{C}_s = \{c_k\}_{k=1}^K$. Formally, this process is expressed as:

$$\mathcal{C}_s = \arg \max_{c_j \in \mathcal{C}} \text{CosSim}(\mathbf{c}_j, \mathbf{g}_I) = \frac{\mathbf{g}_I \cdot \mathbf{c}_j}{\|\mathbf{g}_I\| \|\mathbf{c}_j\|}, \quad (1)$$

where \mathbf{c}_j is the textual feature of the corresponding template sentence. Given \mathcal{C}_s , a sequence of textual basis vectors $\mathbf{C}_s = \{\mathbf{c}_k\}_{k=1}^K$ are naturally formed by extracting the corresponding global textual feature of each c_k .

Patch-wise Cross-modal Feature Mix-up. Our proposed PCM aims to automatically exclude the potential irrelevant and flawed parts in the salient regions of the synthetic images during training in a fine-grained manner. To achieve this goal, the original visual feature map \mathcal{V}_I is novelly transformed into a new textual feature map \mathcal{V}_T in our PCM by considering the semantic similarity between each local patch \mathbf{v}_j and all the salient concept features $\mathbf{C}_s = \{\mathbf{c}_k\}_{k=1}^K$. However, the modality gap between \mathbf{v}_j and \mathbf{c}_k inherently exists in CLIP space. Inspired by DenseCLIP [34], we employ the classification head for the [CLS] token in the image encoder of CLIP to linearly map \mathbf{v}_j to \mathbf{v}'_j for measuring patch-text similarity, resulting in $\mathcal{V}'_I = \{\mathbf{v}'_j\}_{j=1}^{N_I}$. Therefore, a patch-wise affinity mapping $\mathcal{A} \in \mathbb{R}^{N_I \times K}$ can be computed as:

$$a_{jk} = \text{CosSim}(\mathbf{v}'_j, \mathbf{c}_k) = \frac{\mathbf{v}'_j \cdot \mathbf{c}_k}{\|\mathbf{v}'_j\| \|\mathbf{c}_k\|}, \quad (2)$$

where a_{jk} is the similarity score between the j -th local patch and the k -th salient concept. Then, we softly aggregate the visual concept features in \mathbf{C}_s for each local patch \mathbf{v}'_j depending on \mathcal{A} , leading to a textual feature map $\mathcal{V}_T = \{\mathbf{v}_j^t\}_{j=1}^{N_I}$ for each synthetic image I :

$$\mathbf{v}_j^t = \sum_{k=1}^K \alpha_{jk} * \mathbf{c}_k = \sum_{k=1}^K \frac{\exp(a_{jk}/\tau)}{\sum_{m=1}^K \exp(a_{jm}/\tau)} * \mathbf{c}_k, \quad (3)$$

where τ is the temperature.

It can be observed that \mathcal{V}_T only contains the semantics of salient concepts, which would constrain the captioning model to produce short captions. To alleviate the information loss, we further mix up \mathcal{V}_I and \mathcal{V}_T into \mathcal{V}_{IT} as per \mathcal{A} . Technically, we pinpoint the top- M visual patches in \mathcal{V}_I that are highly related to the salient concept features \mathbf{C}_s based on \mathcal{A} , and replace them with the corresponding ones from \mathcal{V}_T . In contrast to the training process with synthetic data, we concatenate the aforementioned top- M patches from \mathcal{V}_T with \mathcal{V}_I to derive \mathcal{V}_{IT} for the real images during inference. The mixed-up feature map \mathcal{V}_{IT} is then passed through an encoder-decoder framework for zero-shot image captioning.

3.3 Image Captioner

The image captioning model is framed as a typical Transformer-based encoder-decoder framework with minor adaptation for the mixed-up feature obtained in the proposed PCM. Formally, given both the global image feature \mathbf{g}_I and the mixed-up feature map \mathcal{V}_{IT} for the input synthetic image I , we first project both these features into a joint embedding space and concatenate the results into $\mathcal{V}^{(0)} = \{\mathbf{g}_I^{(0)}, \mathbf{v}_1^{(0)}, \mathbf{v}_2^{(0)}, \dots, \mathbf{v}_{N_I}^{(0)}\}$. Additionally, the position encodings and token-type embeddings are added to $\mathcal{V}^{(0)}$ following general practices, which is, in turn, fed into the encoder-decoder for caption generation.

Visual-Semantic Encoder. The visual-semantic encoder consists of N_e Transformer-based blocks, where each block includes a multi-head self-attention layer followed by a feed-forward layer. Formally, the output from a vanilla multi-head self-attention layer (MHA) with H attention heads can be computed as:

$$\begin{aligned} \text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{Concat}(\text{head}_1, \dots, \text{head}_H)W^O, \\ \text{head}_i &= \text{Attention}(\mathbf{Q}W_i^Q, \mathbf{K}W_i^K, \mathbf{V}W_i^V), \\ \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}, \end{aligned} \quad (4)$$

where $\text{Concat}(\cdot)$ is the concatenation operation. W_i^Q, W_i^K, W_i^V, W^O are the weight matrices of the i -th head, and d is a scaling factor. Therefore, the operation of the i -th block in the visual-semantic encoder can be expressed as:

$$\begin{aligned} \mathcal{V}^{(i+1)} &= \mathcal{F}(\text{LN}(\mathcal{V}^{(i)} + \text{MHA}(\mathcal{V}^{(i)}, \mathcal{V}^{(i)}, \mathcal{V}^{(i)}))), \\ \mathcal{F}(x) &= \text{LN}(x + \text{FC}(\delta(\text{FC}(x))))), \end{aligned} \quad (5)$$

where **FC**, **LN** and δ are the fully-connected layer, layer normalization, and activation function, respectively. Moreover, inter-layer global feature interaction is devised to obtain a more comprehensive global feature with the outputs from all the blocks as:

$$\tilde{\mathbf{g}}_I = W_g * \mathbf{Concat}(\mathbf{g}_I^{(0)}, \mathbf{g}_I^{(1)}, \dots, \mathbf{g}_I^{(N_e)}), \quad (6)$$

where W_g is a learnable weight matrix. Finally, \mathcal{V}_{IT} is encoded into:

$$\tilde{\mathcal{V}} = \mathbf{Concat}(\tilde{\mathbf{g}}_I, \mathbf{v}_1^{(N_e)}, \mathbf{v}_2^{(N_e)}, \dots, \mathbf{v}_{N_I}^{(N_e)}). \quad (7)$$

Sentence Decoder. Given the fine-grained visual features $\tilde{\mathcal{V}}$ from the visual-semantic encoder, a Transformer-based sentence decoder is exploited to generate sentences. Similarly, the sentence decoder is implemented as N_d Transformer blocks, where each Transformer-style block consists of a masked multi-head self-attention layer [39], a multi-head cross-attention layer, and a feed-forward layer. Let $\mathcal{H}_{0:N_s-1}^{(0)} = \{\mathbf{h}_0^{(0)}, \mathbf{h}_1^{(0)}, \dots, \mathbf{h}_{N_s-1}^{(0)}\}$ stand for the textual features of the sentence that describes the input image I , where $\mathbf{h}_t^{(0)}$ is the textual feature of the t -th word w_t in sentence S . Specifically, at the t -th decoding timestep, the i -th decoder block operates as:

$$\begin{aligned} \mathbf{h}_t^{(i+1)} &= \mathcal{F}(\mathbf{LN}(\tilde{\mathbf{h}}_t^{(i)} + \mathbf{MHA}(\tilde{\mathbf{h}}_t^{(i)}, \tilde{\mathcal{V}}, \tilde{\mathcal{V}}))), \\ \tilde{\mathbf{h}}_t^{(i)} &= \mathbf{LN}(\mathbf{h}_t^{(i)} + \mathbf{MaskedMHA}(\mathbf{h}_t^{(i)}, \mathcal{H}_{0:t}^{(i)}, \mathcal{H}_{0:t}^{(i)})), \end{aligned} \quad (8)$$

where **MaskedMHA** is the masked multi-head self-attention layer. Finally, we use the output of final block $\mathbf{h}_t^{(N_d)}$ to predict next word w_{t+1} through softmax.

3.4 CLIP-weighted Cross-Entropy Loss

Following the conventional training strategy in image captioning, the captioning model parametrized as θ is optimized by maximizing the likelihood of the ground-truth sentence (i.e., standard cross-entropy loss):

$$\theta^* = \arg \max_{\theta} \sum_{(I,S)} \log p(S|I; \theta). \quad (9)$$

By applying the chain rule, the log probability of the sentence can be decomposed into the sum of the log probabilities over the ground-truth words:

$$\log p(S|I; \theta) = \sum_{t=1}^{N_s} \log p(w_t|I, \{w_i\}_{i=0}^{t-1}; \theta). \quad (10)$$

However, due to the existing distribution discrepancy between synthetic and real images, the standard cross-entropy loss is not optimal in this scenario, where it treats low-quality synthetic images as important as high-quality ones, thereby

impeding cross-modal learning of the captioning model. To address this challenge, we propose a novel CLIP-weighted cross-entropy loss. Unlike the standard cross-entropy loss, which treats all synthetic image-text pairs equally, our proposed CLIP-weighted cross-entropy loss dynamically adjusts the training loss for each sample based on the global vision-language similarity score [11]:

$$\begin{aligned} \log p(S|I; \theta) &= c_I \sum_{t=1}^{N_s} \log p(w_t|I, \{w_i\}_{i=0}^{t-1}; \theta), \\ c_I = CLIPScore(I, S) &= \min(1.0, w * \frac{\mathbf{g}_I \cdot \mathbf{g}_S}{\|\mathbf{g}_I\| \|\mathbf{g}_S\|}), \end{aligned} \quad (11)$$

where w is the scaling factor, \mathbf{g}_I and \mathbf{g}_S are the global CLIP embeddings of the synthetic image and the ground-truth sentence, respectively. With this additional weighting coefficient, CLIP-weighted cross-entropy loss penalizes the low-quality synthetic pairs for improved learning of vision-language alignment.

4 Experiments

4.1 Datasets and Experimental Settings

Datasets. We empirically verify and analyze the effectiveness of our PCM-Net on two widely adopted image captioning benchmarks: MSCOCO [23] and Flickr30k [31]. There are five human-annotated sentences per image in these datasets. For fair comparisons, we follow the Karpathy split [14] and take 113,287 images for training, 5,000 images for validation, and 5,000 images for testing on MSCOCO. For Flickr30k which consists of 31,783 images, we use 1,000 images for validation, 1,000 for testing and the rest for training. It is worth noting that we only use the text corpus in training split following the zero-shot image captioning settings. We use Stable Diffusion [35] to synthesize an image for each caption in training corpus, resulting in SynthImgCap dataset consisting of 542,401 and 144,541 synthetic image-text pairs for MSCOCO and Flickr30k, respectively.

Implementation Details. In PCM-Net, the visual-semantic encoder and sentence decoder are built with $N_e = N_d = 3$ Transformer blocks. The size of hidden state in each Transformer block is 512. We utilize CLIP-ViT-L/14 to extract visual and textual features as in CgT-GAN [51]. As a result, each input image is represented as a 768-dimensional global feature vector plus a 1,024-dimensional grid feature map for local patches. Gaussian noise is injected into global feature vector in training following CapDec [27]. We build the visual concept vocabulary \mathcal{C} from the high-frequency nouns in MSCOCO training corpus, and each concept is represented as a 768-dimensional textual feature in CLIP space. The visual and textual patch features extracted by CLIP are further mapped into the common space with 512 dimensions through a fully connected layer. During the training stage, PCM-Net is optimized with Adam [15] optimizer on CXE loss. The whole optimization process takes 12 epochs with the learning rate scheduling strategy in [39]. The warmup steps and batch size are set as 10,000 and 32.

Table 1: Performance of our PCM-Net and other state-of-the-art approaches on the test split of the MSCOCO and Flickr30k datasets under the in-domain zero-shot setting. † denotes the use of real images in training.

Methods	Backbone	MSCOCO					Flickr30k				
		B4	M	R	C	S	B4	M	R	C	S
ZeroCap [38]	ViT-B/32	7.0	15.4	31.8	34.5	9.2	5.4	11.8	27.3	16.8	6.2
MAGIC [37]	ViT-B/32	12.9	17.4	39.9	49.3	11.3	6.4	13.1	31.6	20.4	7.1
CapDec [27]	RN50x4	26.4	25.1	51.8	91.8	-	17.7	20.0	43.9	39.1	-
DeCap [17]	ViT-B/32	24.7	25.0	-	91.2	18.7	21.2	21.8	-	56.7	15.2
ViECap [7]	ViT-B/32	27.2	24.8	-	92.9	18.2	21.4	20.1	-	47.9	13.6
SynTIC [24]	ViT-B/32	29.9	25.8	53.2	101.1	19.3	22.3	22.4	47.3	56.6	16.6
PCM-Net	ViT-B/32	31.5	25.9	53.9	103.8	19.7	26.9	23.0	50.1	61.3	16.8
CgT-GAN † [51]	ViT-L/14	30.3	26.9	54.5	108.1	20.5	24.1	22.6	48.2	64.9	16.1
PCM-Net	ViT-L/14	33.6	26.9	55.4	113.6	20.8	28.5	24.3	51.4	69.5	18.2

At inference, we adopt beam search strategy with the beam size as 3. We report the performances of PCM-Net over five evaluation metrics: BLEU-4 [30] (B4), METEOR [4] (M), ROUGE [22] (R), CIDEr [40] (C), and SPICE [2] (S).

4.2 Performance Comparison

In-domain Zero-shot Image Captioning. We conduct experiments on MSCOCO and Flickr30k datasets under the in-domain settings, where the model is trained on the training corpus and evaluated on the test split of the same dataset. Table 1 summarizes the performance comparisons between the state-of-the-art models and our PCM-Net. All runs are briefly grouped into three directions: (1) zero-shot methods without further training (e.g., ZeroCap [38]) that repurpose the text-to-image matching models to generate captions; (2) traditional methods (e.g., MAGIC [37], CapDec [27], DeCap [17], ViECap [7]) that exploit the potent cross-modal alignment capabilities of CLIP to bridge the modality gap; (3) the approaches (e.g., SynTIC [24]) that utilize synthetic image-text pairs created by off-the-shelf text-to-image generation models. As shown in the table, our PCM-Net consistently exhibits better performances than the state-of-the-art methods across all the metrics on both MSCOCO and Flickr30k datasets. In particular, the CIDEr score of PCM-Net can achieve 113.6%, which leads to an absolute improvement of 5.5% over CgT-GAN (CIDEr: 108.1%). This demonstrates the effectiveness of leveraging synthetic image-text data for mitigating the modality gap in pixel space. Compared to the method that does not involve further training (e.g., ZeroCap), traditional methods (e.g., ViECap [7]) improve the performances by exploiting the cross-modal alignment of CLIP to bridge the modality gap. Though SynTIC [24] manages to enhance performances by training on synthetic image-text pairs in a similar spirit, our PCM-Net substantially outperforms it with a significant margin. This demonstrates the merits of exploring fine-grained spatial visual features extracted from CLIP and applying patch-wise cross-modal feature mix-up to mitigate unfaithful content in synthetic

Table 2: Performance of our PCM-Net and other state-of-the-art approaches on the test split of the MSCOCO and Flickr30k datasets under the cross-domain zero-shot setting. † denotes the use of real images in training.

Methods	Backbone	MSCOCO \Rightarrow Flickr30k					Flickr30k \Rightarrow MSCOCO				
		B4	M	R	C	S	B4	M	R	C	S
MAGIC [37]	ViT-B/32	6.2	12.2	31.3	17.5	-	5.2	12.5	30.7	18.3	-
CapDec [27]	RN50x4	17.3	18.6	42.7	35.7	-	9.2	16.3	36.7	27.3	-
DeCap [17]	ViT-B/32	16.3	17.9	-	35.7	11.1	12.1	18.0	-	44.4	10.9
ViECap [7]	ViT-B/32	17.4	18.0	-	38.4	11.2	12.6	19.3	-	54.2	12.5
SynTIC [24]	ViT-B/32	17.9	18.6	42.7	38.4	11.9	14.6	19.4	40.9	47.0	11.9
PCM-Net	ViT-B/32	20.8	19.2	45.2	45.5	12.9	17.1	19.6	43.2	54.9	12.8
CgT-GAN † [51]	ViT-L/14	17.3	19.6	43.9	47.5	12.9	15.2	19.4	40.9	58.7	13.4
PCM-Net	ViT-L/14	23.9	21.2	47.8	55.9	14.2	17.9	20.3	44.0	61.3	13.5

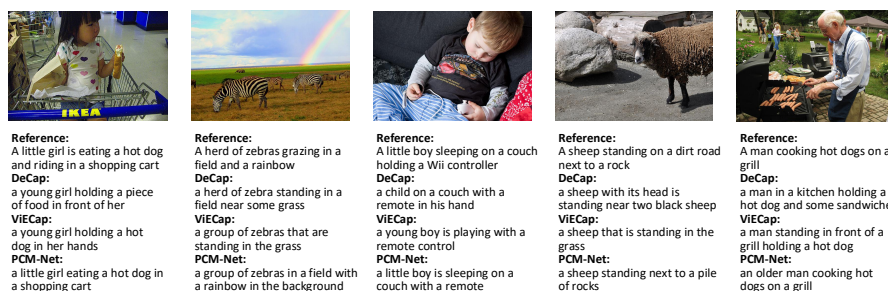


Fig. 3: Examples of image captioning results generated by DeCap [17], ViECap [7] and our PCM-Net, coupled with the corresponding ground-truth sentences(Reference).

images. Furthermore, PCM-Net leverages CLIP to prioritize high-quality synthetic image-text pairs. Similar trends are also observed in the Flickr30k dataset. This again confirms the advantage of our proposal.

Cross-domain Zero-shot Image Captioning. Next, we evaluate our PCM-Net under the cross-domain settings, where the model is evaluated on the test split of a different dataset. As shown in Table 2, the performance trends in cross-domain settings are similar to those in in-domain settings. Concretely, our PCM-Net surpasses the state-of-the-art techniques (CgT-GAN) by an absolute improvement of 8.4% in CIDEr score on the MSCOCO \Rightarrow Flickr30k task. The results again demonstrate the effectiveness of unfaithful content mitigation and emphasis on high-quality synthetic pairs for zero-shot image captioning.

Qualitative Analysis. Figure 3 shows several qualitative results of our PCM-Net and two approaches(i.e., DeCap and ViECap) on MSCOCO dataset, coupled with one human-annotated ground-truth sentence (Reference). From these exemplar results, it is easy to see that our PCM-Net can predict more semantically relevant and logically correct sentences. For instance, in the first example, both DeCap and ViECap are only aware of the major objects (girl

Table 3: Ablation study on each design in PCM-Net on MSCOCO under the in-domain zero-shot setting. **Base** denotes the vanilla Transformer-based encoder-decoder model. **Mix-up** represents the use of patch-wise cross-modal feature mix-up mechanism. **CXE** refers to the CLIP-weighted cross-entropy loss.

Base	Mix-up	CXE	MSCOCO				
			B4	M	R	C	S
✓			33.0	26.3	54.8	108.9	20.2
✓	✓		33.3	26.8	55.0	111.6	20.7
✓		✓	33.2	26.4	55.1	112.0	20.5
✓	✓	✓	33.6	26.9	55.4	113.6	20.8




Fig. 4: Ablation study on hyperparameters in PCM-Net on MSCOCO.

and hot dog), while ignoring the salient object of shopping cart. Instead, by employing a text-to-image diffusion prior for zero-shot image captioning and applying PCM plus CXE loss to boost the training, our PCM-Net effectively captures all significant objects in the image (girl, hot dog, and shopping cart), yielding more accurate and descriptive sentences.

4.3 Experimental Analysis

Ablation Study. We conduct ablation study to investigate how each component in our PCM-Net influences the overall performances on MSCOCO Karpathy test split. Table 3 presents the performance comparisons among different ablated runs of our PCM-Net. We start with a Transformer-based encoder-decoder model (**Base**), which is trained on synthetic pairs in a supervised manner. Next, by incorporating the patch-wise cross-modal feature mix-up mechanism (**Base+Mix-up**) during training, we observe clear performance boosts. This implies that the mixed cross-modal features can mitigate unfaithful content in synthetic images, and thus improve visual-semantic learning. In addition, we also apply the CLIP-weighted cross-entropy loss to the base model. In this way, **Base+CXE** exhibits better performance, verifying the merit of concentration on high-quality pseudo pairs. Finally, we integrate both patch-wise cross-modal feature mix-up mechanism and the CLIP-weighted cross-entropy loss, **Base+Mix-up+CXE** (i.e., our PCM-Net) achieves the best performances across all the evaluation metrics.

Table 4: Performance of ViECap and two variants equipped with our proposals on the validation split of NoCaps [1] benchmark under the cross-domain zero-shot setting.

Methods	Backbone	MSCOCO \Rightarrow NoCaps val							
		In-domain		Near-domain		Out-of-domain		Overall	
		C	S	C	S	C	S	C	S
ViECap [7]	ViT-B/32+GPT-2 _{Base}	61.1	10.4	64.3	9.9	65.0	8.6	66.2	9.5
SYN-ViECap	ViT-B/32+GPT-2 _{Base}	61.7	10.5	68.5	10.3	71.4	9.4	70.5	10.0
PCM-ViECap	ViT-B/32+GPT-2 _{Base}	66.0	10.7	72.7	10.6	75.7	9.7	74.7	10.3

Ablation on Hyperparameters in PCM and CXE. We perform ablation analyses on three crucial hyperparameters in our PCM-Net. Firstly, we investigate the impact of varying the number (K) of salient visual concepts in PCM. In Figure 4 (a), a consistent performance improvement is observed as K increases from 1 to 5, with a slight decline beyond 5, likely due to the introduction of noisy concepts with an excessively large K . Secondly, we explore the influence of replacing different numbers (M) of visual patches in PCM. Figure 4 (b) exhibits similar trends, indicating that excessive replacement of meaningful visual patches by text-based patches leads to a dominance of text features and subsequent degradation in representation capacity. Lastly, we scrutinize the hyperparameter in CXE loss by varying the scaling factor (w) from 1 to 6. Figure 4 (c) shows the best results when w is set to 4, but performance degrades when w exceeds 4 due to higher CLIPScore weights being assigned to low-quality pairs, reducing CXE loss to vanilla cross-entropy loss. Based on these empirical analyses, we set K , M , and w as 5, 5, and 4, respectively.

4.4 Framework compatibility

To verify the generalizability of our methods, we follow the training mechanism of ViECap [7] and evaluate the models on another widely used benchmark, NoCaps [1]. We first upgrade ViECap to SYN-ViECap by replacing the CLIP textual features of the input sentence with the CLIP visual features of the generated image by diffusion model for training. As shown in Table 4, SYN-ViECap outperforms ViECap across all metrics on the NoCaps validation split, demonstrating the advantage of text-to-image diffusion priors for ZIC. By further integrating PCM and CXE into SYN-ViECap, PCM-ViECap achieves a substantial performance boost across all metrics, particularly excelling in the Out-of-domain category. This indicates its effectiveness in narrowing the distribution discrepancy between synthetic and real images. In summary, our proposed PCM mechanism and CXE loss can generalize well on various model structures and achieve remarkable zero-shot transferability.

4.5 Effect of the Training Data Size

We also explore the effect of training data size on PCM-Net following [7, 24]. This is done by randomly sampling various proportions of data from MSCOCO

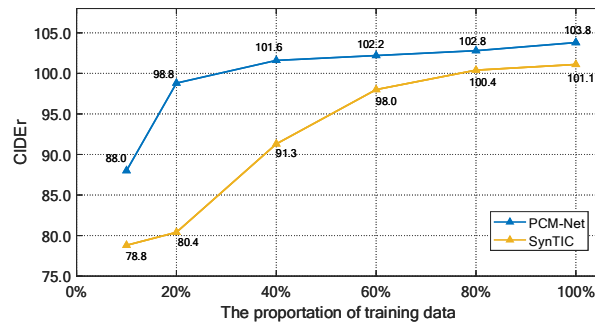


Fig. 5: Performance comparison of PCM-Net and SynTIC on MSCOCO with various proportions of training data.

to train PCM-Net. As depicted in Figure 5, PCM-Net consistently outperforms SynTIC across all data scales, especially in low-data scenarios. Notably, PCM-Net trained with only 40% data surpasses SynTIC that uses the entire training dataset (101.6 vs. 101.1). In summary, PCM-Net performs robustly with limited training data and can be further improved by increasing the training data size.

5 Conclusion

In this paper, we propose PCM-Net for zero-shot image captioning, which aims to mitigate the modality gap of CLIP for visual-semantic learning by leveraging synthetic image-text data for training. Particularly, we employ an off-the-shelf text-to-image diffusion model to build a synthetic dataset, i.e., SynthImgCap. For each input image, salient visual concepts are first constructed by performing zero-shot entity classification of images in CLIP space. After that, the patch-wise cross-modal feature mix-up mechanism is proposed to mix the fine-grained patch features of the synthetic image with the textual features, reducing the flawed or unfaithful details in the synthetic images. Finally, a visual-semantic encoder-decoder is exploited to refine the derived features and generate a caption. To improve the training of captioning model with noisy synthetic data, we propose a novel CLIP-weighted cross-entropy loss to prioritize the high-quality image-text pairs over the low-quality counterparts. Extensive experiments conducted on MSCOCO and Flickr30k datasets demonstrate the superiority of our PCM-Net.

Broader Impact. Our PCM-Net is trained to produce image captions based on learnt statistics of the training corpus and synthetic images, and thus the biases rooted in those data will be reflected in the outputs, resulting in negative societal impacts. Hence more future research is necessary for addressing the issue.

Acknowledgments

This work is partially supported by China NSFC under Grant No. 61772563.

References

1. Agrawal, H., Desai, K., Wang, Y., Chen, X., Jain, R., Johnson, M., Batra, D., Parikh, D., Lee, S., Anderson, P.: Nocaps: Novel object captioning at scale. In: ICCV. pp. 8948–8957 (2019)
2. Anderson, P., Fernando, B., Johnson, M., Gould, S.: Spice: Semantic propositional image caption evaluation. In: ECCV (2016)
3. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: CVPR (2018)
4. Banerjee, S., Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: ACL Workshop (2005)
5. Cornia, M., Stefanini, M., Baraldi, L., Cucchiara, R.: Meshed-memory transformer for image captioning. In: CVPR (2020)
6. Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
7. Fei, J., Wang, T., Zhang, J., He, Z., Wang, C., Zheng, F.: Transferable decoding with visual entities for zero-shot image captioning. In: ICCV. pp. 3136–3146 (2023)
8. Feng, Y., Ma, L., Liu, W., Luo, J.: Unsupervised image captioning. In: CVPR. pp. 4125–4134 (2019)
9. Gu, J., Joty, S., Cai, J., Zhao, H., Yang, X., Wang, G.: Unpaired image captioning via scene graph alignments. In: ICCV. pp. 10323–10332 (2019)
10. Herdade, S., Kappeler, A., Boakye, K., Soares, J.: Image captioning: Transforming objects into words. NeurIPS (2019)
11. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. In: EMNLP (2021)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS (2020)
13. Huang, L., Wang, W., Chen, J., Wei, X.Y.: Attention on attention for image captioning. In: ICCV (2019)
14. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: CVPR (2015)
15. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
16. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: ICML (2023)
17. Li, W., Zhu, L., Wen, L., Yang, Y.: Decap: Decoding clip latents for zero-shot captioning via text-only training. In: ICLR (2023)
18. Li, Y., Pan, Y., Yao, T., Chen, J., Mei, T.: Scheduled sampling in vision-language pretraining with decoupled encoder-decoder network. In: AAAI (2021)
19. Li, Y., Pan, Y., Yao, T., Mei, T.: Comprehending and ordering semantics for image captioning. In: CVPR. pp. 17990–17999 (2022)
20. Li, Y., Yao, T., Pan, Y., Chao, H., Mei, T.: Pointing novel objects in image captioning. In: CVPR. pp. 12497–12506 (2019)
21. Liang, V.W., Zhang, Y., Kwon, Y., Yeung, S., Zou, J.Y.: Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. NeurIPS (2022)
22. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: ACL Workshop (2004)

23. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
24. Liu, Z., Liu, J., Ma, F.: Improving cross-modal alignment with synthetic pairs for text-only image captioning. In: AAAI (2024)
25. Luo, J., Li, Y., Pan, Y., Yao, T., Chao, H., Mei, T.: Coco-bert: Improving video-language pre-training with contrastive cross-modal matching and denoising. In: ACM MM. pp. 5600–5608 (2021)
26. Luo, J., Li, Y., Pan, Y., Yao, T., Feng, J., Chao, H., Mei, T.: Semantic-conditional diffusion networks for image captioning. In: CVPR. pp. 23359–23368 (2023)
27. Nukrai, D., Mokady, R., Globerson, A.: Text-only training for image captioning using noise-injected clip. In: EMNLP Findings. pp. 4055–4063 (2022)
28. Pan, Y., Mei, T., Yao, T., Li, H., Rui, Y.: Jointly modeling embedding and translation to bridge video and language. In: CVPR. pp. 4594–4602 (2016)
29. Pan, Y., Yao, T., Li, Y., Mei, T.: X-linear attention networks for image captioning. In: CVPR (2020)
30. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: ACL (2002)
31. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: ICCV. pp. 2641–2649 (2015)
32. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763. PMLR (2021)
33. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
34. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Denseclip: Language-guided dense prediction with context-aware prompting. In: CVPR. pp. 18082–18091 (2022)
35. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
36. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. NeurIPS 35, 36479–36494 (2022)
37. Su, Y., Lan, T., Liu, Y., Liu, F., Yogatama, D., Wang, Y., Kong, L., Collier, N.: Language models can see: Plugging visual controls in text generation. arXiv preprint arXiv:2205.02655 (2022)
38. Tewel, Y., Shalev, Y., Schwartz, I., Wolf, L.: Zerocap: Zero-shot image-to-text generation for visual-semantic arithmetic. In: CVPR. pp. 17918–17928 (2022)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
40. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: CVPR (2015)
41. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: CVPR (2015)
42. Wang, Z., Yu, J., Yu, A.W., Dai, Z., Tsvetkov, Y., Cao, Y.: Simvlm: Simple visual language model pretraining with weak supervision. In: ICLR (2022)
43. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: ICML (2015)

44. Yang, X., Gao, C., Zhang, H., Cai, J.: Auto-parsing network for image captioning and visual question answering. In: ICCV (2021)
45. Yang, X., Tang, K., Zhang, H., Cai, J.: Auto-encoding scene graphs for image captioning. In: CVPR (2019)
46. Yao, T., Pan, Y., Li, Y., Mei, T.: Incorporating copying mechanism in image captioning for learning novel objects. In: CVPR. pp. 6580–6588 (2017)
47. Yao, T., Pan, Y., Li, Y., Mei, T.: Exploring visual relationship for image captioning. In: ECCV (2018)
48. Yao, T., Pan, Y., Li, Y., Mei, T.: Hierarchy parsing for image captioning. In: ICCV (2019)
49. Yao, T., Pan, Y., Li, Y., Qiu, Z., Mei, T.: Boosting image captioning with attributes. In: ICCV (2017)
50. You, Q., Jin, H., Wang, Z., Fang, C., Luo, J.: Image captioning with semantic attention. In: CVPR (2016)
51. Yu, J., Li, H., Hao, Y., Zhu, B., Xu, T., He, X.: Cgt-gan: Clip-guided text gan for image captioning. In: ACM MM. pp. 2252–2263 (2023)