

Style-Aware Contrastive Learning for Multi-Style Image Captioning

Yucheng Zhou, Guodong Long*

Australian AI Institute, School of Computer Science, FEIT, University of Technology Sydney
yucheng.zhou-1@student.uts.edu.au, guodong.long@uts.edu.au

Abstract

Existing multi-style image captioning methods show promising results in generating a caption with accurate visual content and desired linguistic style. However, existing methods overlook the relationship between linguistic style and visual content. To overcome this drawback, we propose style-aware contrastive learning for multi-style image captioning. First, we present a style-aware visual encoder with contrastive learning to mine potential visual content relevant to style. Moreover, we propose a style-aware triplet contrast objective to distinguish whether the image, style and caption matched. To provide positive and negative samples for contrastive learning, we present three retrieval schemes: object-based retrieval, RoI-based retrieval and triplet-based retrieval, and design a dynamic trade-off function to calculate retrieval scores. Experimental results demonstrate that our approach achieves state-of-the-art performance. In addition, we conduct an extensive analysis to verify the effectiveness of our method.

1 Introduction

Stylized image captioning aims to generate a natural language description with stylized elements for a given image (Mathews et al., 2016; Gan et al., 2017). With the advance of deep learning in human-computer interaction equipment, it has been integrated into many real-world applications like education robots (Zhou, 2022), visual dialog (Das et al., 2017), and vision-language navigation (Wang et al., 2019). Therefore, it has attracted more attention from academia and industry and has become one of the essential areas in the natural language processing (NLP) and computer vision (CV) community.

However, many methods propose translating images into captions of a single caption style (Mathews et al., 2016; Gan et al., 2017). These methods need to train multiple models to handle multiple



Figure 1: (a) the caption is generated by stylized captioning model (Shuster et al., 2019); (b) and (c) show the same caption can correspond to different styles.

styles, which is very inefficient. Therefore, a rising demand for stylized image captioning is to learn an efficient model to handle multiple styles simultaneously. Recently, many efforts have been made in multi-style image captioning (Guo et al., 2019; Shuster et al., 2019; Li and Harrison, 2021).

Despite their success, existing methods suffer from a drawback: They focus on accurate visual content and desired linguistic style but overlook the relation between linguistic style and visual content. As shown in Figure 1(a), the generated caption contains accurate visual content and the desired style, but with misinformation that is easily detectable by humans, i.e., model generates “beautiful” in the caption to cater to the style of “happy”. Indeed, linguistic style (Bell, 1984) reflects personality, emotion and sentiment. In human behavior, these factors significantly influence the course of cognitive behavior (Simon, 1967). When people with different emotions see the same image are likely to describe different contents because they pay attention to different aspects. For example, people who are happy with animals may describe an image of a dog as a pretty dog with sparkling eyes and supple hair. However, one afraid of dogs may describe it as a scary dog with fierce teeth and sharp claws. The former focus on the dog’s eyes and hair, while the latter on the dog’s teeth and paws. Therefore, people with different emotions focus on different potential visual content. To

*Corresponding author.

generate human-like stylized captions, a stylized captioning model should learn to mine potential visual content relevant to different linguistic styles.

Due to the success of contrastive learning (He et al., 2020; Gao et al., 2021), some works (e.g., UNITER (Chen et al., 2020), ViLT (Kim et al., 2021)) employ contrastive learning objectives to encourage cross-modality alignment. In this work, we propose a style-aware visual encoder with contrastive learning to mine potential visual content relevant to styles. Specifically, style-aware visual encoder first mines potential visual features based on given image and style pairs. Then, we use a contrastive loss to pull potential visual features of the anchor and positive pair together while pushing those of anchor and negative pairs apart.

In addition, apart from requiring generated captions with accurate visual content and desired linguistic style, multi-style image captioning also needs to ensure that potential visual content and style are relevant. Moreover, since multi-style image captioning contains more fine-grained styles than single-stylized image captioning, it is difficult to directly distinguish the style by the caption. As shown in Figure 1(b) and (c), captions may be the same for two different styles. Therefore, multi-style image captioning is required to consider if matching among image, style and caption. Different from previous works (Guo et al., 2019) that optimize generated captions only based on style, we propose a style-aware triplet contrast loss, which can learn the triplet matching among image, style, and caption by contrasting it with the positive triplet against negative ones.

Moreover, motivated by hard negatives sampling in retriever training (Zhan et al., 2021), we present three retrieval schemes to mine positive and negative examples for contrastive learning: object-based retrieval, RoI-based retrieval and triplet-based retrieval. These three schemes calculate scores according to object overlap rate, potential visual feature similarity, and triple feature similarity, respectively. Meanwhile, we design a dynamic trade-off function to calculate retrieval scores and analyze the impact of different retrieval schemes.

We conduct an extensive evaluation on three datasets, i.e., PERSONALITY-CAPTIONS (Shuster et al., 2019), SentiCap (Mathews et al., 2016) and FlickrStyle10K (Gan et al., 2017). Our method significantly outperforms other strong competitors and achieves state-of-the-art performance.

2 Related Work

2.1 Image Captioning

Image captioning, one of the essential tasks in multimodal research, aims to generate a description for a given image (Hodosh et al., 2015). With the progress of deep learning, many end-to-end deep neural networks for image captioning are proposed (Vinyals et al., 2015; Anderson et al., 2018; Zhou et al., 2021). Although these works have achieved excellent improvements, the caption generated by these models focuses on a single language domain, which means that their outputs can be in only one style and even dull and lack vitality sometimes.

2.2 Stylized Image Captioning

With the advance in image captioning techniques, researchers have attempted to generate an image caption with style. Mathews et al. (2016) propose a word-level regularization for captioner to model sentiment words. Gan et al. (2017) employ transforming word embeddings matrices to control style factors for generated captions. To accurately describe visual content and reflect the desired linguistic style, some methods (Mathews et al., 2018; Zhao et al., 2020) split a stylized sentence into a style-related part that reflects the linguistic style and a content-related part that contains the visual content. Zhou (2022) employ prompt-based pre-training to build a stylized captioner without any paired sketch and story. These methods alleviate the reliance on paired training data for stylized captioner training. However, these methods need to train multiple models to handle multiple styles, which is inefficient. Therefore, learning an efficient model to handle multiple styles simultaneously raises more interest. Some efforts are made on multi-style image captioning, including adversarial learning network (Guo et al., 2019) and multi-updown fusion model (Li and Harrison, 2021). To generate more diverse outputs, Shuster et al. (2019) collect PERSONALITY-CAPTIONS, a large-scale multi-style image captioning dataset.

2.3 Contrastive Learning

Recently, contrastive learning has made exciting progress in representation learning (He et al., 2020; Gao et al., 2021; Zhou et al., 2022). He et al. (2020) leverage contrastive learning to improve visual presentation learning in an unsupervised manner. Gao et al. (2021) propose a simple unsupervised contrastive learning method to perform on

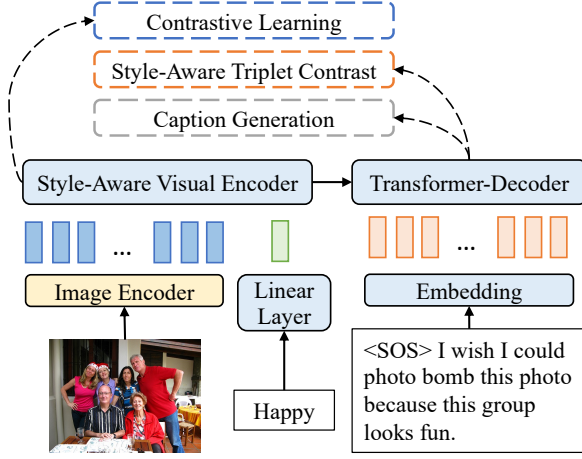


Figure 2: An overview of our proposed SACO model with three objectives. Yellow rounded rectangles denote fixed model. Blue rounded rectangles denote parameters that will be optimized.

par with previously supervised counterparts. Yang et al. (2022) propose triple contrastive learning for vision-language pre-training by leveraging both cross-modal and intra-modal self-supervision, providing complimentary benefits in representation learning. These works show the powerful ability of contrastive learning to improve representation learning.

3 Method

This section will elaborate on our Style-Aware COntrastive learning (SACO) method for multi-style image captioning, followed by our proposed novel retrieval schemes. The details of our approach are shown in Figure 2. Lastly, details about training and fine-tuning are elaborated.

3.1 Image and Style Encoding

Firstly, we pass an image I into a pre-trained convolutional neural network (CNN) to extract its visual feature V and convert the visual feature into a sequence according to the row-first direction:

$$V = \text{CNN}(I)$$

$$\text{where } V = [v_0, v_1, \dots, v_m], v_i \in \mathbb{R}^{|d'|} \quad (1)$$

In addition, the style S is represented as a one-hot vector and encoded to a style feature s by a linear layer:

$$s = \text{LinearLayer}(S), s \in \mathbb{R}^{|d|} \quad (2)$$

Since there is a different dimension between visual feature V and style feature s , a multi-layer

perceptron (MLP) built upon v_i is presented to reduce its dimension to equal the dimension of s :

$$v_i := \text{MLP}(v_i), v_i \in \mathbb{R}^{|d|}, \forall i = 1, \dots, m \quad (3)$$

3.2 Style-Aware Visual Encoder

Originating from the observation that different styles will focus on different image regions, we propose a style-aware visual encoder with contrastive learning to mine potential visual content relevant to styles. Since self-attention (Vaswani et al., 2017) shows a strong capability to relate different positions of a sequence for feature computation, we leverage self-attention as backbone of the style-aware visual encoder. Particularly, we concatenate the visual feature V and style feature s , and apply the self-attention layers to them to derive the style-aware visual features V^s and vision-aware style feature s^v , i.e.,

$$[V^s; s^v] = \text{Self-Attention}([V; s]) \quad (4)$$

where $[\cdot; \cdot]$ denotes the operation of concatenation. Intuitively, the success of Eq. (4) requires accurately capturing the visual features relevant to styles, but a major problem arises without artifact labels: *Learning Difficulty*. Motivated by some works (Chen et al., 2020; Kim et al., 2021) implement cross-modality alignment via contrastive learning, we leverage contrastive learning for style-aware visual encoder, which drives its learning to capture potential visual content correlated to given styles without any labeled data. Specifically, we use contrastive learning to learn the representation of potential visual content by contrasting it with the positive example (\hat{V}^s, \hat{s}^v) against those of negative ones $(\bar{V}_i^s, \bar{s}_i^v)$. Particularly, we first derive features (\hat{V}^s, \hat{s}^v) and $(\bar{V}_i^s, \bar{s}_i^v)$ independently via Eq. (1-4). Then, we mine more accurate (V^s, s^v) , style-aware visual features and vision-aware style features, by contrasting it with (\hat{V}^s, \hat{s}^v) against $(\bar{V}_i^s, \bar{s}_i^v)$, i.e.,

$$\mathcal{L}^{(svc)} = -\log \frac{e^{\text{sim}(r, \hat{r})/\tau}}{e^{\text{sim}(r, \hat{r})/\tau} + \sum_{i=1}^M e^{\text{sim}(r, \bar{r}_i)/\tau}}$$

$$- \log \frac{e^{\text{sim}(s^v, \hat{s}^v)/\tau}}{e^{\text{sim}(s^v, \hat{s}^v)/\tau} + \sum_{i=1}^M e^{\text{sim}(s^v, \bar{s}_i^v)/\tau}}$$

$$\text{where } r = \text{Pooling}(V^s)$$

$$\hat{r} = \text{Pooling}(\hat{V}^s)$$

$$\bar{r} = \text{Pooling}(\bar{V}_i^s) \quad (5)$$

where $\text{sim}(\mathbf{r}, \mathbf{r}')$ is the dot product operation between ℓ_2 normalized \mathbf{r} and \mathbf{r}' (i.e. cosine similarity $\frac{\mathbf{r}^\top \mathbf{r}'}{\|\mathbf{r}\| \cdot \|\mathbf{r}'\|}$) and τ is a temperature hyperparameter to control the pull and push force; M is number of negative samples. As a result, since the style-aware visual features and vision-aware style features also offer a straightforward pathway to transmit style-aware visual information to style-aware visual encoder, it mitigates the *learning difficulty* problem.

3.3 Transformer-Decoder Based Generator

Due to the success of the pre-trained language model, there are some works (e.g., GPT (Radford et al., 2018), GPT-2 (Radford et al., 2019)) pre-train the Transformer decoder (Vaswani et al., 2017) on large-scale corpora. Recently, a fundamental paradigm of text generation tasks is to fine-tune the pre-trained model on the target data, and it can achieve exciting performance. In this work, we leverage a trained transformer decoder to initialize our generator. The generator input is adjusted to the triples $(\mathbf{V}^s, \mathbf{s}^v, \bar{\mathbf{C}})$, where $\bar{\mathbf{C}}$ refers to the segment of stylized image caption. The purpose of generator is to predict a probability distribution of the next word of the segment $\bar{\mathbf{C}}$ based on the given triple, i.e.,

$$\mathbf{h}_i = \text{Trans-Dec}(\mathbf{V}^s, \mathbf{s}^v, \bar{\mathbf{C}}) \in \mathbb{R}^{|\mathcal{V}|}$$

where $\bar{\mathbf{C}} = [c_1, \dots, c_{i-1}]$ (6)

$$\mathbf{p}_i = \text{LM-Head}(\mathbf{h}_i) \in \mathbb{R}^{|\mathcal{V}|}$$
 (7)

where \mathbf{h}_i refers to the hidden representation in i -th step; \mathcal{V} denotes token vocabulary and \mathbf{p}_i is a probability distribution over \mathcal{V} . Lastly, the caption generation objective is defined as a maximum likelihood estimation and written as:

$$\mathcal{L}^{(cap)} = -\frac{1}{|N|} \sum_{i=1}^N \log \mathbf{p}_i(c_i),$$
 (8)

where $\mathbf{p}_i(c_i)$ denotes fetching the probability of the i -th step gold token $c_i \in \mathbf{C}$ from \mathbf{p}_i . \mathbf{C} refers to the gold caption and N is its length.

3.4 Style-Aware Triplet Contrast

Different from stylized image captioning for a single style, multi-style image captioning contains more fine-grained styles, and it is difficult to distinguish the style by the caption directly. So multi-style image captioning is required to consider if matching among image, style and caption. Recently, contrastive learning has shown its power capability in alignment between positive pairs and

dispersion between negative ones. Inspired by advances in contrastive learning, we propose a style-aware triplet contrast loss to learn the triplet matching among image, style, and caption by contrasting it with the positive triplet against negative ones. Given an image \mathbf{I} , style \mathbf{S} and caption \mathbf{C} , we first obtain the representation for this triplet. Since visual and style features have fed into the decoder, the triplet representation can be extracted by the hidden representation of the decoder, i.e.,

$$\mathbf{h} = -\frac{1}{|N|} \sum_{i=1}^N \text{MLP}_{\text{triplet}}(\mathbf{h}_i)$$
 (9)

where N is the same as that in Eq.8 and denotes the length of the caption; $\text{MLP}_{\text{triplet}}$ refers to a Multilayer Perceptron; and \mathbf{h}_i can be derived from Eq.6. Then, we enhance the triplet representation \mathbf{h} by contrasting it with a positive triplet $\hat{\mathbf{h}}$ against negative ones $\bar{\mathbf{h}}$, i.e.

$$\mathcal{L}^{(stc)} = -\log \frac{e^{\text{sim}(\mathbf{h}, \hat{\mathbf{h}})/\tau}}{e^{\text{sim}(\mathbf{h}, \hat{\mathbf{h}})/\tau} + \sum_{i=1}^M e^{\text{sim}(\mathbf{h}, \bar{\mathbf{h}})/\tau}}$$
 (10)

where $\text{sim}(\cdot, \cdot)$ is the dot product operation as same as that in Eq.5.

3.5 Retrieval Schemes

In our proposed contrastive learning objective, positive and negative samples are important elements. Since caption datasets are not designed with positive and negative samples, we propose three heuristics to derive positive and negative samples for a triplet of image \mathbf{I} , style \mathbf{S} and caption \mathbf{C} :

Object-based Retrieval. We first leverage a well-trained object detection model (Ren et al., 2015) to obtain object classes in images. Then, we retrieve image $\hat{\mathbf{I}}$ from image set \mathbb{I} according to object overlap with \mathbf{I} , and derive a probability for sampled examples:

$$P_{obj} = \frac{N_{overlap}}{N_{\mathbf{I}}}$$
 (11)

where $N_{\mathbf{I}}$ denotes number of objects in the image \mathbf{I} , and $N_{overlap}$ refers to the number of overlapped objects both in \mathbf{I} and $\hat{\mathbf{I}}$.

RoI-based Retrieval. In this retrieval scheme, the region of interest (RoI) refers to potential visual content relevant to style. We retrieve image $\hat{\mathbf{I}}$ based on the similarity between its representation of potential visual content and that of image \mathbf{I} :

$$P_{roi} = \text{sim}(\mathbf{V}, \hat{\mathbf{V}})$$
 (12)

where V and \hat{V} are the style-aware visual features of I and \hat{I} , and details of their calculation can refer to Eq.5.

Triplet-based Retrieval. Since triplet matching is essential for multi-style image captioning, we retrieve image \hat{I} according to the similarity between triplets:

$$P_{tri} = \text{sim}(\mathbf{h}, \hat{\mathbf{h}}) \quad (13)$$

where \mathbf{h} and $\hat{\mathbf{h}}$ are the triplet representation for the triplet (I, S, C) and $(\hat{I}, \hat{S}, \hat{C})$, and details of their calculation can refer to Eq.10.

Dynamic Trade-off Function. We combine the above three novel retrieval schemes to rank samples, and score of each sample can be defined as:

$$P = \theta^\mu P_{obj} + (1 - \theta^\mu)(\phi P_{roi} + (1 - \phi)P_{tri}) \quad (14)$$

where P denotes the score of samples and ϕ is a trade-off parameter; θ denotes a decay factor and μ is the current training epoch. During training, we randomly select a sample among the top-10 samples as positive one, and the negative samples are randomly sampled based on $P < \max(0.1, P_{max} - \omega^\mu)$.

3.6 Training and Fine-tuning

For model training, we adopt the same 2-stage training scheme (training and fine-tuning) as in (Anderson et al., 2018). In the training stage, we optimize the model according to the three objectives proposed above, and the loss function of our model can be integrated into the following:

$$\mathcal{L} = \mathcal{L}^{(cap)} + \alpha \mathcal{L}^{(svc)} + \beta \mathcal{L}^{(stc)} \quad (15)$$

where α and β are trade-off parameters.

In the fine-tuning stage, we employ the CIDEr score to optimize our model as same as (Rennie et al., 2017), i.e., returning a reward for generated caption \hat{C} .

$$\mathcal{R}^{(CIDEr)} = R_{CIDEr}(\hat{C}) - b \quad (16)$$

where b is a baseline, i.e., the reward $R_{CIDEr}(C^*)$ for the generated caption C^* with greedy search.

4 Experiments

4.1 Dataset and Evaluation Metrics

We evaluate our proposed approach on three datasets, PERSONALITY-CAPTIONS (Shuster et al., 2019), SentiCap (Mathews et al., 2016) and FlickrStyle10K (Gan et al., 2017).

PERSONALITY-CAPTIONS. Shuster et al. (2019) collect a large-scale multi-style image captioning dataset, PERSONALITY-CAPTIONS, which includes 201,858 images, 215 personality traits, and 241,858 stylized captions. We divide the dataset following (Shuster et al., 2019), and the size of the training set, validation set, and test set are 186,858, 5,000, and 10,000, respectively. In the test set, each image contains five reference captions. Following Shuster et al. (2019), we use the same metrics to report our results, and the evaluation is based on the coco-caption code¹. The evaluation metrics include: BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004), CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016).

SentiCap & FlickrStyle10K. SentiCap (Mathews et al., 2016) and FlickrStyle10K (Gan et al., 2017) are two publicly stylized image caption datasets. According to (Li et al., 2021), we process SentiCap and FlickrStyle10K datasets, and use samples from MSCOCO (Lin et al., 2014) and Flickr30K (Hodosh et al., 2015) as large-scale paired factual data. Following Li et al. (2021), we use BLEU, METEOR (Banerjee and Lavie, 2005), CIDEr, style classification accuracy (cls.) and the average perplexity (ppl.) as evaluation metrics. Style classification accuracy is measured by a well-trained BERT, which achieves accuracies of 95.9%, 98.0%, 98.1%, and 99.5% on the test sets of humorous, romantic, negative, and positive styles, respectively. The average perplexity is measured by a well-trained trigram-based statistical language model using SRILM toolkit. A lower score denotes that the generated caption is more fluent and reflects the desired linguistic style better.

4.2 Implementation Details

Our approach adopts the same 2-stage training scheme (training and fine-tuning). For training stage, input images are resized to the size of 256×256 , and then we randomly crop the image size as 224×224 as model input. To ensure a fair comparison, we use the ResNeXt-IG-3.5B (Mahajan et al., 2018) as the pre-trained image encoder, as same as (Shuster et al., 2019), and the size of output features is $7 \times 7 \times 2048$. Then, the visual features are reshaped as 49×2048 according to the row-first direction. The style-aware visual encoder is constructed with three self-attention layers. We use a distilled version of pre-trained GPT-2 (Sanh

¹<https://github.com/tylin/coco-caption>

et al., 2019) as our transformer-decoder based generator, as same as (Nguyen et al., 2020). The layers and attention heads of the decoder are 6 and 8. The dimension of embedding vectors and hidden states in the decoder are 768 and 1024. Special tokens for the beginning and end of sentences are <SOS> and <EOS>. In addition, the size of the style linear layer are 215×768 and 5×768 for PERSONALITY-CAPTIONS dataset and SentiCap and FlickrStyle10K datasets. For model training, we utilize the Adam optimizer (Kingma and Ba, 2015) with learning rate of $1e-4$. The batch size, warm-up proportion, weight decay, maximum training epoch and temperature hyperparameter τ are 128, 0.1, 0.01, 10 and 0.08. The trade-off parameter ϕ and decay factor θ in Eq.14 are 0.5 and 0.9. ω for negatives sampling is 0.8. Trade-off parameters α and β in Eq.15 are 0.5 and 0.7. For fine-tuning stage, the maximum training epoch and learning rate are 3 and $1e-5$. Other experimental details are the same as that of training stage. For testing, we use beam search with beam size of 3 to generate captions with maximum sentence length of 30. Our model is trained on one V100 GPU.

4.3 Baselines

We compare our model with the following state-of-the-art baselines: (1) **ShowTell** propose by (Shuster et al., 2019) and is a variant of (Vinyals et al., 2015) that concatenates the style features with the input word vectors at each decoding step. (2) **ShowAtt-Tell**, proposed by (Xu et al., 2015), aims to enhance the correlation between the text and image by an attention mechanism, and the used model is a variant proposed by (Shuster et al., 2019). (3) **UpDown** with a decoder of two LSTMs can adapt to generate attention weights and use it to generate captions (Anderson et al., 2018), and the used model is a variant proposed by (Shuster et al., 2019). (4) **GPT** with an image encoder is fine-tuned on the captioning dataset (Radford et al., 2019). (5) **GPT-Speaker** (Nguyen et al., 2020) employs the language model GPT2 as a language prior for both the speaker and listener in the multi-agent communication framework. (6) **3M** (Li and Harrison, 2021) is a multi-style image captioner that is a multi-UpDown encoder-decoder model integrated with multi-modal features.

4.4 Main Results

Comparison results on PERSONALITY-CAPTIONS test set are shown in Table 1.

Method	B@1	B@4	R	C	S
ShowTell	38.4	7.3	24.3	9.6	1.6
ShowAttTell	43.3	7.1	27.0	12.6	3.6
UpDown	44.4	8.0	27.4	16.5	5.2
GPT	49.2	9.1	29.0	19.0	6.3
GPT-Speaker	52.1	8.4	30.2	19.9	7.3
GPT-Speaker*	52.3	8.2	30.1	20.0	7.4
3M	43.0	8.0	27.6	18.6	4.8
SACO (Ours)	54.8	9.7	32.6	21.0	8.1

Table 1: Comparison results on PERSONALITY-CAPTIONS test set. B@1, B@4, R, C and S denote BLEU@1, BLEU@4, ROUGE-L, CIDEr and SPICE, respectively. * refers to the baseline of our reproduction.

Style	Method	B@1	B@3	M	C	cls.	ppl.
Humor	SF-LSTM	27.4	8.5	11.0	39.5	-	-
	MSCap	16.3	1.9	5.3	15.2	91.3	22.7
	SAN	29.5	9.9	12.5	47.2	99.4	13.7
	SACO	35.2	10.3	13.4	51.2	99.7	12.9
Roman	SF-LSTM	27.8	8.2	11.2	37.5	-	-
	MSCap	17.0	2.0	5.4	10.1	88.7	20.4
	SAN	30.9	10.9	13.0	53.3	99.6	13.1
	SACO	35.8	13.5	14.1	55.8	99.9	12.4
Pos	SF-LSTM	50.5	19.1	16.6	60.0	-	-
	MSCap	46.9	16.2	16.8	55.3	92.5	19.6
	SAN	53.0	23.4	18.1	72.0	100.0	11.7
	SACO	56.3	24.7	19.5	72.2	99.8	11.5
Neg	SF-LSTM	50.3	20.1	16.2	59.7	-	-
	MSCap	45.5	15.4	16.2	51.6	93.4	19.2
	SAN	51.2	20.5	17.6	67.0	100.0	14.8
	SACO	53.2	23.6	19.1	70.5	100.0	13.3

Table 2: Comparison results on FlickrStyle10K and SentiCap test set. M denotes METEOR.

From the table, we can make three observations. First, we can observe that our methods achieve state-of-the-art performance on the PERSONALITY-CAPTIONS dataset. Second, models based on GPT reach a better performance than LSTM-based models, which shows the powerful generation capability of transformer-based decoder in stylized caption generation. Lastly, our model significantly outperforms the GPT-2 model by a large margin (i.e., improving 2.0 in CIDEr). That is, style-aware contrastive learning can improve the model by strengthening different style understanding.

In addition, we also show comparison results on FlickrStyle10K and SentiCap test set in Table 2. Competitors include SF-LSTM (Chen et al., 2018), MSCap (Guo et al., 2019) and SAN (Li et al., 2021). Results show that transformer-based methods (i.e., SAN and SACO) outperform LSTM-based methods (i.e., SF-LSTM and MSCap). More-

Method	B@1	B@4	R	C	S
SACO	54.8	9.7	32.6	21.0	8.1
◇ w/o CIDEr	53.0	9.6	31.9	19.1	7.2
◇ w/o STC	52.9	9.0	31.2	20.0	7.5
◇ w/o SVC	52.4	8.7	30.8	19.8	7.3
◇ w/o SVC, STC	48.9	9.0	29.1	18.7	6.0
◇ $\mathcal{L}^{(cap)}$ Only	47.3	8.7	29.2	16.0	5.4

Table 3: Ablation study. “w/o CIDEr” denotes removing fine-tuning stage for model training; “w/o STC” denotes removing the style-aware triplet contrast loss; “w/o SVC” denotes removing the style-aware visual contrast encoder; “w/o SVC, STC” denotes removing both contrastive learning methods in our model; and “ $\mathcal{L}^{(cap)}$ Only” denotes our model only train with the caption generation objective and without fine-tuning.

Method	B@1	B@4	R	C	S
SACO	54.8	9.7	32.6	21.0	8.1
SACO (Dec w/o style)	53.8	9.4	31.9	20.5	7.8
◇ w/o SVC	52.9	9.0	31.2	20.0	7.5

Table 4: Impact of style-aware visual contrast encoder.

over, compared with strong competitors, our approach achieves state-of-the-art performance on FlickrStyle10K and SentiCap.

4.5 Ablation Study

To demonstrate the effectiveness of our method, we conduct an ablation study and results are shown in Table 3. We first investigate the impact of the style-aware triplet contrast objective by removing it and find that model performance is decreased. Next, we investigate our method without style-aware visual contrast encoder and observe that the performance drops, which is worse than that without SVC. Moreover, we remove both contrastive learning objectives, and results show the performance is degraded again. The observations above demonstrate the effectiveness of style-aware contrastive learning.

4.6 Analysis

4.6.1 Impact of Contrastive Learning

Impact of Style-Aware Visual Contrast Encoder.

To further investigate the impact of the style-aware visual contrast encoder, we remove the input style for our decoder, i.e., the decoder input only consists of V^s, \bar{C} in Eq.6. Results are shown in Table 4. Results show that our decoder without input style outperforms the model without SVC, which demonstrates that style-aware visual contrast encoder can capture potential visual content relevant to the given style.

Method	B@1	B@4	R	C	S
SACO	54.8	9.7	32.6	21.0	8.1
SACO (STC-pointwise)	53.2	9.0	31.4	20.1	7.3
SACO (comp.)	53.0	9.1	31.2	20.2	7.4
◇ w/o STC	52.4	8.7	30.8	19.8	7.3

Table 5: Impact of style-aware triplet contrast.




Figure 3: Retrieval schemes ablation.

Impact of Style-Aware Triplet Contrast. To investigate the impact of the style-aware triplet contrast objective, we replace the objective with two binary classification based objectives: point-wise classification and the comp. objective in (Nguyen et al., 2020). As shown in Table 5, results show that binary classification based objectives underperform our contrastive objective, which demonstrates that contrasting positive and negative triplets is helpful for the model to understand triplet matching.

4.6.2 Retrieval Schemes Ablation

Due to contrastive learning depending on positive and negative sampling, we conduct an ablative analysis on retrieval schemes. Results are shown in Figure 3. From left figure, we can make two observations: (1) The decay factor θ set of 0.9 can perform better than that of 0.8. It demonstrates that object-based retrieval is essential in the early training phase. (2) The trade-off parameter ϕ set of 0.5 can achieve the best performance, which shows the vital role of both style-aware contrastive objectives in positive and negative sampling. From right figure, based on the sampling function in §3.5, as ω becomes greater, negative sampling range during training will become greater. We can see that a greater negative sampling range is beneficial for contrasting learning. The reason is that the larger sampling range allows model to access harder negatives. In addition, we find that the performance gradually drops when ω is greater than 0.8, which shows too hard negatives hinder model learning.

4.6.3 Qualitative Comparison

To extensively evaluate our model, we conduct a qualitative comparison of our model and GPT-

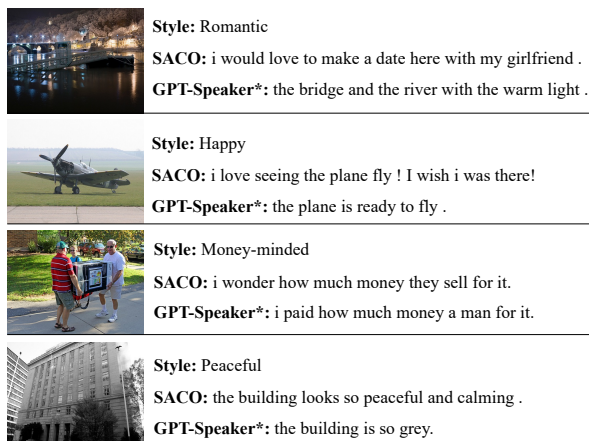


Figure 4: Random sampling examples generated by SACO and GPT-Speaker*.



Figure 5: Interpretable visualization analysis of random sampling examples. The brighter image areas mean more relevant to style.

Speaker, and some random sampling examples are shown in Figure 4. For example, in the first line, we can observe that GPT-Speaker can capture objects in the image and describe them in a caption, but it does not entail style. Therefore, the caption generated by our model is shown to more natural and with desired style. For instance, in the second and last lines, we can find that the caption generated by GPT-Speaker is more like a factual description. In contrast, the caption generated by our model is more expressive of style.

4.6.4 Interpretable Visualization Analysis

To investigate the effectiveness of style-aware visual contrast encoder, we conduct an interpretable visualization analysis, as shown in Figure 5. In the attention map of self-attention layers, the brighter image areas mean greater attention weights, i.e., these areas are more relevant to the given style. As shown in the first example, under the style "money-minded", the object is more paid attention to than the human, which is very intuitive. Next, in the second example, under the style "ex-

Type of evaluation	Win Percentage	
	SACO	GPT-Speaker*
Engagingness	62.9	37.1
Visual Relevance	64.9	35.1
Personalized Relevance	63.4	36.6

Table 6: Human Evaluation.

citing", "a dog of running" are paid more attention and described with "love" and "cute", which are reasonable. Therefore, the results show that the style-aware visual contrast encoder can effectively capture potential visual content related to style.

4.6.5 Human Evaluation

To comprehensively evaluate our method, we conducted a human evaluation to compare our model and GPT-Speaker. Following Nguyen et al. (2020), we considered the engagingness and relevance of captions. Engagingness evaluation considers human preference for the naturalness and appropriateness of the captions, while relevance evaluation involves visual and stylized relevance. Therefore, there are three types of evaluation. We randomly sampled 50 samples from the test set for each evaluation type above. Each sample includes an image and a style. Then, we use our model and GPT-Speaker to generate captions for these samples. We displayed the selected image-style pairs and their caption generated from our model and GPT-Speaker to 7 recruited annotators. They need to judge which captions are better quality based on the type of evaluation. As shown in Table 6, results show that the performance of our model is significantly better than GPT-Speaker, i.e., our model can generate fascinating captions.

5 Conclusion

In this work, we dive into the relationship between linguistic style and visual content. The first is potential visual content has varied for different styles. We propose a style-aware visual encoder that learns to capture the representation of potential visual content. Second, since the model is required to distinguish whether the image, style and caption are matched, we present a style-aware triplet contrast objective to improve the model's capability to discriminate triplet matching. In addition, we propose three novel retrieval schemes to sample positive and negative examples for contrastive learning. Results show that our method delivers new state-of-the-art performance.

Limitations

Although our proposed method can effectively mine latent visual content related to style, it still suffers from weaknesses in generating multiple stylized captions for the same image and style pair. Specifically, our method relies on beam search to generate diverse stylized captions for the same image and style pair and lacks the capability to control content for caption generation interactively. In further work, we will study how to generate stylized captions by interactively selecting specified regions in latent visual content.

References

- 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*, volume 9909 of *Lecture Notes in Computer Science*, pages 382–398. Springer.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. **Bottom-up and top-down attention for image captioning and visual question answering**. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. IEEE Computer Society.
- 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.
- Allan Bell. 1984. Language style as audience design. *Language in society*, 13(2):145–204.
- Tianlang Chen, Zhongping Zhang, Quanzeng You, Chen Fang, Zhaowen Wang, Hailin Jin, and Jiebo Luo. 2018. **"factual" or "emotional": Stylized image captioning with adaptive learning and attention**. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part X*, volume 11214 of *Lecture Notes in Computer Science*, pages 527–543. Springer.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. **UNITER: universal image-text representation learning**. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 104–120. Springer.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M. F. Moura, Devi Parikh, and Dhruv Batra. 2017. **Visual dialog**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1080–1089. IEEE Computer Society.
- Chuang Gan, Zhe Gan, Xiaodong He, Jianfeng Gao, and Li Deng. 2017. **Stylenet: Generating attractive visual captions with styles**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 955–964. IEEE Computer Society.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. **Simcse: Simple contrastive learning of sentence embeddings**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6894–6910. Association for Computational Linguistics.
- Longteng Guo, Jing Liu, Peng Yao, Jiangwei Li, and Hanqing Lu. 2019. **Mscap: Multi-style image captioning with unpaired stylized text**. In *CVPR 2019*, pages 4204–4213. Computer Vision Foundation / IEEE.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. **Momentum contrast for unsupervised visual representation learning**. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. Computer Vision Foundation / IEEE.
- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2015. **Framing image description as a ranking task: Data, models and evaluation metrics (extended abstract)**. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 4188–4192. AAAI Press.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. **Vilt: Vision-and-language transformer without convolution or region supervision**. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Chengxi Li and Brent Harrison. 2021. **3m: Multi-style image caption generation using multi-modality features under multi-updown model**. In *Proceedings of the Thirty-Fourth International Florida Artificial Intelligence Research Society Conference, North Miami Beach, Florida, USA, May 17-19, 2021*.

- Guodun Li, Yuchen Zhai, Zehao Lin, and Yin Zhang. 2021. [Similar scenes arouse similar emotions: Parallel data augmentation for stylized image captioning](#). In *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 5363–5372. ACM.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft COCO: common objects in context](#). In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer.
- Dhruv Mahajan, Ross B. Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. 2018. [Exploring the limits of weakly supervised pre-training](#). In *ECCV 2018*, volume 11206 of *Lecture Notes in Computer Science*, pages 185–201. Springer.
- Alexander Patrick Mathews, Lexing Xie, and Xuming He. 2016. [Senticap: Generating image descriptions with sentiments](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3574–3580. AAAI Press.
- Alexander Patrick Mathews, Lexing Xie, and Xuming He. 2018. [Semstyle: Learning to generate stylised image captions using unaligned text](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 8591–8600. IEEE Computer Society.
- Minh Thu Nguyen, Duy Phung, Minh Hoai, and Thien Huu Nguyen. 2020. [Structural and functional decomposition for personality image captioning in a communication game](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4587–4593. Association for Computational Linguistics.
- 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, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. [Faster R-CNN: towards real-time object detection with region proposal networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. [Self-critical sequence training for image captioning](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1179–1195. IEEE Computer Society.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Kurt Shuster, Samuel Humeau, Hexiang Hu, Antoine Bordes, and Jason Weston. 2019. [Engaging image captioning via personality](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12516–12526. Computer Vision Foundation / IEEE.
- Herbert A. Simon. 1967. Motivational and emotional controls of cognition. In *Psychological Review*, volume 74, pages 29–39.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. [Show and tell: A neural image caption generator](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3156–3164. IEEE Computer Society.
- Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. 2019. [Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation](#). In *CVPR 2019*, pages 6629–6638. Computer Vision Foundation / IEEE.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *Proceedings of the 32nd International*

Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2048–2057. JMLR.org.

Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. 2022. [Vision-language pre-training with triple contrastive learning](#). *CoRR*, abs/2202.10401.

Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. [Optimizing dense retrieval model training with hard negatives](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1503–1512. ACM.

Wentian Zhao, Xinxiao Wu, and Xiaoxun Zhang. 2020. [Memcap: Memorizing style knowledge for image captioning](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 12984–12992. AAAI Press.

Yucheng Zhou. 2022. [Sketch storytelling](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*, pages 4748–4752. IEEE.

Yucheng Zhou, Tao Shen, Xiubo Geng, Guodong Long, and Daxin Jiang. 2022. [ClarET: Pre-training a correlation-aware context-to-event transformer for event-centric generation and classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2559–2575, Dublin, Ireland. Association for Computational Linguistics.

Yucheng Zhou, Wei Tao, and Wenqiang Zhang. 2021. [Triple sequence generative adversarial nets for unsupervised image captioning](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pages 7598–7602. IEEE.