

Semantic-Conditional Diffusion Networks for Image Captioning*

Jianjie Luo, Yehao Li, Yingwei Pan, Ting Yao, Jianlin Feng, Hongyang Chao, Tao Mei
School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China
HiDream.ai Inc.

{jianjieluo.sysu,yehaoli.sysu,panyw.ustc,tingyao.ustc}@gmail.com;
{fengjlin,isschhy}@mail.sysu.edu.cn;tmei@hidream.ai

Abstract

Recent advances on text-to-image generation have witnessed the rise of diffusion models which act as powerful generative models. Nevertheless, it is not trivial to exploit such latent variable models to capture the dependency among discrete words and meanwhile pursue complex visual-language alignment in image captioning. In this paper, we break the deeply rooted conventions in learning Transformer-based encoder-decoder, and propose a new diffusion model based paradigm tailored for image captioning, namely Semantic-Conditional Diffusion Networks (SCD-Net). Technically, for each input image, we first search the semantically relevant sentences via cross-modal retrieval model to convey the comprehensive semantic information. The rich semantics are further regarded as semantic prior to trigger the learning of Diffusion Transformer, which produces the output sentence in a diffusion process. In SCD-Net, multiple Diffusion Transformer structures are stacked to progressively strengthen the output sentence with better visual-language alignment and linguistic coherence in a cascaded manner. Furthermore, to stabilize the diffusion process, a new self-critical sequence training strategy is designed to guide the learning of SCD-Net with the knowledge of a standard autoregressive Transformer model. Extensive experiments on COCO dataset demonstrate the promising potential of using diffusion models in the challenging image captioning task. Source code is available at https://github.com/YehLi/xmodaler/tree/master/configs/image_caption/scdnet.

1. Introduction

As one of the fundamental tasks in vision and language field, image captioning aims to describe the interested semantics in a natural sentence. This task naturally connects

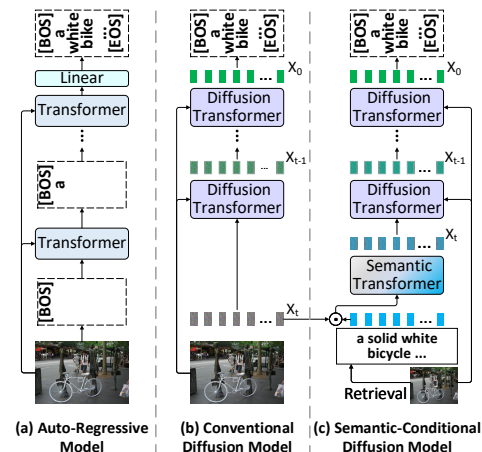


Figure 1. An illustration of (a) typical autoregressive model, (b) conventional diffusion model in [6], and (c) our semantic-conditional diffusion model in SCD-Net.

computer vision and natural language processing by perceiving visual content in a scene and interpreting them in human language, simulating a basic capability of human intelligence. The dominate force in current state-of-the-art techniques [8, 11, 15, 26, 27, 33, 34, 40, 45, 52] is to capitalize on encode-decoder structure and frame the learning process in an autoregressive manner. In particular, an image encoder is responsible for encoding visual content as high-level semantics, and a sentence decoder learns to decode the sequential sentence word-by-word (Fig.1 (a)). Nevertheless, such autoregressive progress only allows for unidirectional textual message passing, and typically relies on considerable computational resources that scale quadratically w.r.t. the sentence length.

To alleviate this limitation, recent advances [12, 13, 32, 47] start to emerge as a non-autoregressive solution that enables bidirectional textual message passing and emits all words in parallel, leading to a light-weight and scalable paradigm. However, these non-autoregressive approaches are generally inferior to the autoregressive methods. The performance degradation can be mostly attributed to the word repetition or omissions problems if leaving the sequential dependency under-exploited. In addition, the indisposed sentence quality also makes it difficult to

*The first two authors contributed equally to this work.

upgrade non-autoregressive solution with powerful self-critical sequence learning [39] widely adopted in autoregressive methods.

More recently, a superior generative module named diffusion model [16] has brought forward milestone improvements for visual content generation. This motivates a recent pioneering practice [6] to explore diffusion model for image captioning, pursuing non-autoregressive sentence generation in a high degree of parallelism. It manages to enable continuous diffusion process to produce the discrete word sequence by representing each word as binary bits. Different from the typical discrete sentence generation in a single shot, such continuous diffusion process (Fig.1 (b)) can be represented as a parameterized Markov chain that gradually adds Gaussian noise to the sentence. Each reverse state transition is thus learnt to recover the original sentence data from noise-augmented data via denoising. Despite showing comparable performances against basic autoregressive models, the severe problem of word repetition or omissions is still overlooked.

In an effort to mitigate this problem, we devise a new diffusion model based non-autoregressive paradigm, called Semantic-Conditional Diffusion Networks (SCD-Net). Our launching point is to introduce the comprehensive semantic information of the input images into the continuous diffusion process, which act as semantic prior to guide the learning of each reverse state transition (Fig.1 (c)). SCD-Net is henceforth able to encourage a better semantic alignment between visual content and the output sentence, i.e., alleviating the omission of semantic words. As a by-product, our SCD-Net enables the powerful self-critical sequence learning during the continuous diffusion process. Such sentence-level optimization of diffusion process strengthens the linguistic coherence of the output sentence, and thus alleviates word repetition issue.

Technically, our SCD-Net is composed of cascaded Diffusion Transformer structures that progressively enhance the output sentences. Each Diffusion Transformer exploits the semantic-conditional diffusion process to learn Transformer-based encoder-decoder. In particular, for each input image, Diffusion Transformer first retrieves the semantically relevant words by using an off-the-shelf cross-modal retrieval model. These semantic words are additionally incorporated into the continuous diffusion process, targeting for constraining reverse state transition with the semantic condition. More importantly, we upgrade the semantic-conditional diffusion process with a new guided self-critical sequence learning strategy. This strategy elegantly transfers the knowledge of a standard autoregressive Transformer model to our non-autoregressive Diffusion Transformer through sentence-level reward, leading to a stabilized and boosted diffusion process.

In sum, we have made the following contributions: (I)

SCD-Net designs a novel semantic-conditional diffusion process for image captioning, pursuing a scalable non-autoregressive paradigm with better visual-language alignment. (II) SCD-Net paves a new way to couple the continuous diffusion process with a new guided self-critical sequence learning. (III) SCD-Net has been properly analyzed and verified through extensive experiments on COCO, demonstrating its encouraging potential.

2. Related Work

2.1. Autoregressive Image Captioning

RNN-based Approaches. With the advent of deep learning [3, 9, 41], modern image captioning techniques utilize CNN plus RNN to generate sentences, yielding flexible syntactical structure. For example, the earlier attempts [10, 20, 44] present an end-to-end neural network system that encodes the visual content with CNN and decodes the sentence with RNN. Attention mechanism [46] is further employed to dynamically focus on the relevant region when generating word at each decoding timestep. Meanwhile, semantic attributes [28, 50, 53, 54] are incorporated into RNN to complement visual representations for enhancing description generation. In contrast to traditional visual attention over equally-sized image regions, Up-Down [2] combines bottom-up and top-down attention mechanism to measure attention at object level. After that, GCN-LSTM [51] models the relations between objects, which are injected into the top-down attention model to boost caption generation. SGAE [49] utilizes scene graph structure to depict fine-grained semantics in the images, aiming to capture the language inductive bias.

Transformer-based Approaches. Recently, taking the inspiration from the success of Transformer [42] in NLP field, numerous Transformer-based image captioning models start to emerge. [40] directly employs the primary Transformer structure in NLP on image captioning task. [15] incorporates the spatial relations among objects into Transformer. Meshed-Memory Transformer [8] improves both encoder and decoder in Transformer by learning prior knowledge with persistent memory and exploiting low- and high-level features with a mesh-like connections across multiple layers. X-Transformer [35] further introduces higher order intra and inter-modal interactions to enhance the encoder and decoder of Transformer. Later on, APN [48] learns to parse trees unsupervisedly by imposing probabilistic graphical model on self-attention layers of Transformer for both captioning and visual question answering.

2.2. Non-Autoregressive Image Captioning

In contrast to autoregressive methods that generate sentences word-by-word, non-autoregressive methods emit all words in parallel, enabling bidirectional textual message

passing. NAT [13] first proposes non-autoregressive decoding to improve the inference speed for neural machine translation, attracting increasing attention in captioning field. MNIC [12] presents masked non-autoregressive decoding to generate captions parallelly with enhanced semantics and diversity. O2NA [32] first generates the object-oriented coarse-grained caption and then refines each object word to alleviate description ambiguity issue. SATIC [55] proposes semi-autoregressive Transformer that predicts a group of words in parallel and generates the groups from left to right, pursuing a better trade-off between speed and quality. Most recently, sparked by the success of diffusion models [16] in image generation, Bit Diffusion [6] encodes the discrete text into binary bits, and utilizes diffusion model with self-conditioning for caption generation. Nevertheless, the performance of Bit Diffusion is still inferior to state-of-the-art autoregressive Transformer. Moreover, how to optimize the diffusion model with sentence-level reward (e.g., CIDEr [43]) has not yet been explored.

Our work also falls into the latter category of non-autoregressive image captioning with diffusion models. Our SCD-Net goes beyond Bit Diffusion by strengthening the visual-language semantic alignment with semantic-conditional diffusion process. A new guided self-critical sequence learning is designed to further stabilize and boost the diffusion process with sentence-level reward.

3. Method

In this section, we present our proposed Semantic-Conditional Diffusion Networks (SCD-Net), that facilitates diffusion process based image captioning with rich semantic prior. Figure 2 depicts the cascaded framework of SCD-Net with multiple stacked Diffusion Transformers, and the detailed architecture of each Diffusion Transformer.

3.1. Problem Formulation

Notation of Diffusion Model. Suppose we have an input image with K interested objects, which will be described by a textual sentence S in N_s words. Let $\mathcal{V} = \{v_i\}_{i=1}^K$ denote the set of objects detected by Faster R-CNN [2, 38], where $v_i \in \mathbb{R}^{D_v}$ represents the D_v -dimensional feature of each object. Here we basically formulate the procedure of caption generation with diffusion model [6, 16]. Considering that the words in the textual sentence are discrete data, we follow Bit Diffusion [6] and convert each word into $n = \lceil \log_2 \mathcal{W} \rceil$ binary bits (i.e., $\{0, 1\}^n$), where \mathcal{W} is the vocabulary size. In this way, the textual sentence S is transformed into real numbers $x_0 \in \mathbb{R}^{n \times N_s}$ to trigger diffusion model. In particular, the diffusion model consists of two processes, i.e., the forward and reverse process.

Forward Process. The forward process is defined as a Markov chain that gradually adds Gaussian noise to the sentence data x_0 . For any $t \in (0, T]$, the forward state

transition [22] from x_0 to x_t is calculated as:

$$x_t = \sqrt{\text{sigmoid}(-\gamma(t'))}x_0 + \sqrt{\text{sigmoid}(\gamma(t'))}\epsilon, \quad (1)$$

where $t' = t/T$, $\epsilon \sim \mathcal{N}(0, I)$, $t \sim \mathcal{U}(0, T)$, \mathcal{N} is normal distribution, \mathcal{U} is uniform distribution, and $\gamma(t')$ is a monotonically increasing function. After that, a Diffusion Transformer $f(x_t, \gamma(t'), \mathcal{V})$ is trained to reconstruct x_0 conditioned on \mathcal{V} through denoising in ℓ_2 regression loss:

$$\mathcal{L}_{bit} = \mathbb{E}_{t \sim \mathcal{U}(0, T), \epsilon \sim \mathcal{N}(0, I)} \|f(x_t, \gamma(t'), \mathcal{V}) - x_0\|^2. \quad (2)$$

Reverse Process. In an effort to generate sentences from the learnt Diffusion Transformer based on the given image, the reverse process samples a sequence of latent states x_t from $t = T$ to $t = 0$. Specifically, given the number of steps T , we discretize time uniformly with width $1/T$ to obtain $s = t - 1 - \Delta$, $t' = t/T$, and $s' = s/T$. Next the reverse state transition x_{t-1} is measured as:

$$\begin{aligned} \alpha_s &= \sqrt{\text{sigmoid}(-\gamma(s'))}, \alpha_t = \sqrt{\text{sigmoid}(-\gamma(t'))}, \\ \sigma_s &= \sqrt{\text{sigmoid}(\gamma(s'))}, \\ c &= -\text{expm1}(\gamma(s') - \gamma(t')), \\ u(x_t; s', t') &= \alpha_s(x_t(1 - c)/\alpha_t + cf(x_t, \gamma(t'), \mathcal{V})), \\ \sigma^2(s', t') &= \sigma_s^2 c, \\ x_{t-1} &= u(x_t; s', t') + \sigma(s', t')\epsilon, \end{aligned} \quad (3)$$

where Δ denotes the time difference and $\text{expm1}(\cdot) = \text{exp}(\cdot) - 1$. After iteratively triggering Diffusion Transformer from x_T , we can obtain the estimation and a quantization operation is operated to transform it into bits x_0 .

3.2. Diffusion Transformer

The basic Diffusion Transformer is constructed as a typical Transformer-based encoder-decoder structure, including a visual encoder and a sentence decoder. Formally, given the detected objects \mathcal{V} in image, the visual encoder transforms them into visual tokens through self-attention. Then the visual tokens and the word tokens $x_t = \{w_0^t, w_1^t, \dots, w_{N_s}^t\}$ at time step t are fed into the sentence decoder for caption generation.

Visual Encoder. The visual encoder consists of N_v stacked Transformer encoder blocks. Each block is composed of a multi-head self-attention layer plus a feed-forward layer. The i -th Transformer encoder block in the visual encoder operates as:

$$\begin{aligned} \mathcal{V}^{i+1} &= \text{FFN}(\text{norm}(\mathcal{V}^i + \text{MultiHead}(\mathcal{V}^i, \mathcal{V}^i, \mathcal{V}^i))), \\ \text{FFN}(\mathbf{Z}) &= \text{norm}(\mathbf{Z} + \text{FC}(\delta(\text{FC}(\mathbf{Z})))), \\ \text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{Concat}(\text{head}_1, \dots, \text{head}_H)W^O, \\ \text{head}_i &= \text{Attention}(\mathbf{Q}W_i^Q, \mathbf{K}W_i^K, \mathbf{V}W_i^V), \\ \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}, \end{aligned} \quad (4)$$

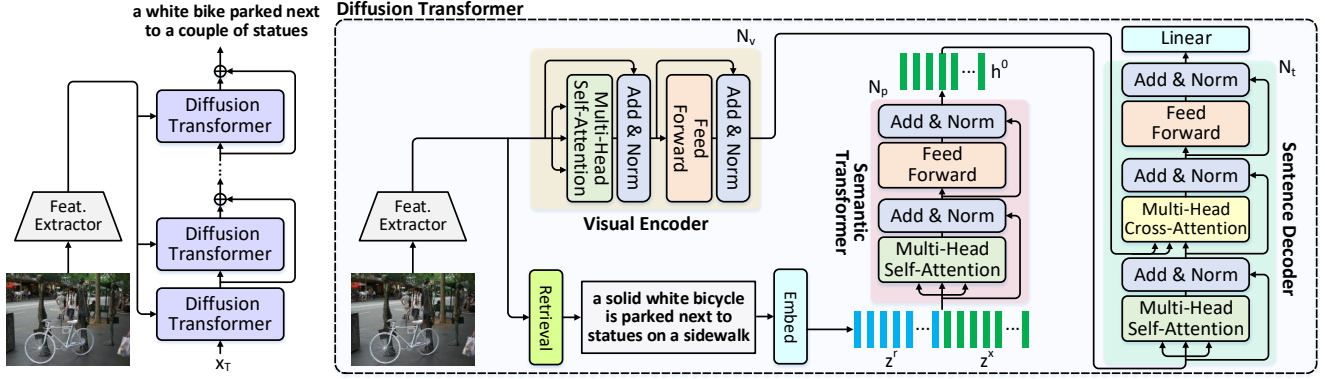


Figure 2. An overview of Semantic-Conditional Diffusion Networks (SCD-Net) consisting of multiple stacked Diffusion Transformers in a cascaded manner (left). Each Diffusion Transformer is composed of Visual Encoder, Semantic Transformer, and Sentence Decoder (right).

where **FFN** is the feed-forward layer, **MultiHead** is the multi-head self-attention layer, **norm** is layer normalization, **FC** is the fully-connected layer, **Concat**(\cdot) is the concatenation operation, δ is the activation function, W_i^Q , W_i^K , W_i^V , W^O are weight matrices of the i -th head, H is the number of head, and d is the dimension of each head. Note that the input of the first Transformer encoder block is the primary set of detected objects $\mathcal{V}^0 = \mathcal{V}$. Accordingly, after stacking N_v blocks, we achieve the contextually enhanced visual tokens $\tilde{\mathcal{V}} = \mathcal{V}^{N_v}$.

Sentence Decoder. The sentence decoder contains N_t stacked Transformer decoder blocks. Each block consists of one multi-head self-attention layer, one multi-head cross-attention layer, and one feed-forward layer. In contrast to conventional Transformers [40, 42] that utilize mask to prevent positions from attending to subsequent positions, the multi-head self-attention layer in Diffusion Transformer is bi-directional without masks. In this way, the i -th Transformer decoder block operates as:

$$\begin{aligned} h^{i+1} &= \text{FFN}(\text{norm}(\tilde{h}^i + \text{MultiHead}(\tilde{h}^i, \tilde{\mathcal{V}}, \tilde{\mathcal{V}}))), \\ \tilde{h}^i &= \text{norm}(h^i + \text{MultiHead}(h^i, h^i, h^i)). \end{aligned} \quad (5)$$

The inputs of the first Transformer decoder block at time step t are the word tokens $h^0 = \{w_0^t, w_1^t, \dots, w_{N_s}^t\}$. After stacking N_t blocks, the hidden state outputted by the last block h^{N_t} is utilized to predict the probability distribution of each output word, which is calculated by:

$$p_i = \text{softmax}(W^T h_i^{N_t}), \quad (6)$$

where W^T is the weight matrix, $h_i^{N_t}$ and $p_i \in \mathbb{R}^{\mathcal{W}}$ is the hidden state vector and probability distribution corresponding to the i -th word, respectively. After that, we map the probability distribution of p_i into bits b_i by taking weighted average over all the \mathcal{W} bits in the vocabulary:

$$b_i = \sum_{c=1}^{\mathcal{W}} p_i^c B^c, \quad (7)$$

where p_i^c is the c -th probability of p_i and B^c is the bit representation of the c -th word in the vocabulary. To speedup convergence of Diffusion Transformer, we also integrate the diffusion process objective \mathcal{L}_{bit} in Eq. (2) with a typical cross entropy loss \mathcal{L}_{XE} over the probability distribution p during training. Therefore, the final objective is calculated as: $\mathcal{L} = \mathcal{L}_{XE} + \mathcal{L}_{bit}$.

3.3. Semantic Condition

Recent practice attempts to employ diffusion model for image captioning by directly estimating x_0 from the latent state x_T sampled from noise distribution $\mathcal{N}(0, I)$ in the reverse process. Nevertheless, this way overlooks the complex visual-language alignment and inherent sequential dependency among words during diffusion process, thereby resulting in the word repetition or omissions problems. To mitigate this limitation, we upgrade the diffusion model with additional semantic prior, yielding a semantic-conditioned diffusion process for Diffusion Transformer.

Technically, given an image, we first search the semantically relevant sentence from training sentence pool by using an off-the-shelf cross-modal retrieval model. The relevant sentence is further represented as a sequence of word tokens s_r , which reflect the comprehensive semantic information. After that, we take s_r as semantic condition to constrain the diffusion process of Diffusion Transformer. In particular, at each time step t , we capitalize on an additional semantic Transformer with N_p semantic Transformer blocks to contextually encode current latent state x_t with semantic prior s_r . Here we concatenate x_t with the previous timestep prediction \tilde{x}_0 along channel dimension as in [6]. The textual features of latent state x_t and semantic prior s_r are thus calculated as:

$$\begin{aligned} z^x &= \text{FC}(\text{Concat}(x_t, \tilde{x}_0)) + \varphi(\gamma(t')), \\ z^r &= \text{FC}(s_r). \end{aligned} \quad (8)$$

where φ is a multi-layer perception. The positional encodings are omitted for simplicity. Next, we concatenate the

textual features z^x and z^r as ($\mathcal{W}^0 = [z^x, z^r]$), which is further fed into semantic Transformer to achieve the semantic-conditional latent state. In this way, the i -th semantic Transformer block is calculated as:

$$\mathcal{W}^{i+1} = \mathbf{FFN}(\mathbf{norm}(\mathcal{W}^i + \mathbf{MultiHead}(\mathcal{W}^i, \mathcal{W}^i, \mathcal{W}^i))) \quad (9)$$

Finally, by denoting the outputs of semantic Transformer as $\mathcal{W}^{N_p} = [\mathcal{W}_x^{N_p}, \mathcal{W}_r^{N_p}]$, we take $\mathcal{W}_x^{N_p}$ as the strengthened semantic-conditional latent state $\mathcal{W}_x^{N_p} = h^0 = \{w_0^t, w_1^t, \dots, w_{N_s}^t\}$, which will be fed into sentence decoder for caption generation in diffusion model.

3.4. Cascaded Diffusion Transformer

Inspired by the success of cascaded diffusion models for image generation [17], our SCD-Net stacks multiple Diffusion Transformers in a cascaded fashion. Such cascaded structure aims to progressively strengthen the output sentence with better visual-language alignment and linguistical coherence. Formally, this cascaded diffusion process can be represented as:

$$F(x_t, \gamma(t'), \mathcal{V}) = f_M \circ f_{M-1} \circ \dots \circ f_1(x_t, \gamma(t'), \mathcal{V}), \quad (10)$$

where M is the total number of stacked Diffusion Transformer, and f_1 is the first Diffusion Transformer equipped with aforementioned semantic condition. In this way, each Diffusion Transformer f_i ($i \geq 2$) operates diffusion process conditioned on the sentence prediction x_0^{i-1} of previous Diffusion Transformer f_{i-1} . Accordingly, for each Diffusion Transformer f_i ($i \geq 2$), we slightly modify its structure to take the additional semantic cues of x_0^{i-1} into account. Specifically, given the latent state x_t , previous timestep prediction \tilde{x}_0 , and the previous Diffusion Transformer prediction x_0^{i-1} , we remould Eq. (8) by measuring textual features of latent state z^x as:

$$z^x = \mathbf{FC}(\mathbf{Concat}(x_t, \tilde{x}_0, x_0^{i-1})) + \varphi(\gamma(t')). \quad (11)$$

Then the textual features z^x are concatenated with semantic prior z^r , and we feed them into the semantic Transformer. At inference, the output of each Diffusion Transformer f_i is directly fused with the output of previous Diffusion Transformer f_{i-1} at each timestep.

3.5. Guided Self-Critical Sequence Training

Conventional autoregressive image captioning techniques [2, 39, 40] commonly utilize the Self-Critical Sequence Training [39] (SCST) to boost up performances with sentence-level optimization (e.g., CIDEr metric [43]):

$$\mathcal{L}_R(\theta) = -\mathbb{E}_{y_{1:N_s} \sim p_\theta} [R(y_{1:N_s})], \quad (12)$$

where R denotes the CIDEr score function. The gradient of the loss can be approximated as:

$$\nabla_\theta \mathcal{L}_R(\theta) \approx -(R(y_{1:N_s}^s) - R(\hat{y}_{1:N_s})) \nabla_\theta \log p_\theta(y_{1:N_s}^s), \quad (13)$$

where $y_{1:N_s}^s$ is the sampled caption and $R(\hat{y}_{1:N_s})$ denotes the sentence-level reward of baseline in greedily decoding inference. However, it is not trivial to directly apply SCST to the diffusion process in Diffusion Transformer. The difficulty mainly originates from two aspects. First, the non-autoregressive inference procedure of Diffusion Transformer contains multiple steps (e.g., more than 20), and thus it is impractical to sample sentences directly from the noise x_T . Furthermore, since each output word in the Diffusion Transformer is sampled independently, simply assigning the same reward to each word would result in typical word repetition problem [14]. To address these limitations, we propose a new Guided Self-Critical Sequence Training, which nicely guides the learning of SCD-Net with the knowledge derived from a standard autoregressive Transformer model.

Technically, we first train a standard autoregressive Transformer teacher model that shares same architecture of Diffusion Transformer. Next, for each training image, this Transformer teacher model predicts high-quality sentence S^{tea} as additional semantic guidance. Then we feed the predicted high-quality sentence S^{tea} into the cascaded Diffusion Transformer, instead of the ground-truth sentences. Rather than randomly sampling several captions as in conventional SCST, we enforce one of the sampled sentences to be the same as the predicted sentence S^{tea} . Let $y_{1:N_s}^{s_j} \Big|_{j=0}^{N_y}$ denote the sampled sentences containing S^{tea} , where N_y is the number of sampled sentences. The gradient of $\mathcal{L}_R(\theta)$ is thus measured as:

$$\nabla_\theta \mathcal{L}_R(\theta) \approx -\frac{1}{N_y} \sum_{j=0}^{N_y} (R(y_{1:N_s}^{s_j}) - R(\hat{y}_{1:N_s})) \nabla_\theta \log p_\theta(y_{1:N_s}^{s_j}). \quad (14)$$

In this way, the sampling of the high-quality sentence S^{tea} tends to receive a positive reward, that encourages Diffusion Transformer to produce high-quality sentence. The possibility of other low-quality sentences (e.g., sentences with word repetition) is thus suppressed. It is worthy to note that when the diffusion model becomes saturated, we replace the sentence S^{tea} derived from autoregressive Transformer model with the estimated S' by the diffusion model, if the quality of S' (measured in CIDEr) is higher than S^{tea} .

4. Experiments

In this section, we empirically verify the effectiveness of our SCD-Net by conducting experiments over the widely adopted COCO benchmark [30] for image captioning.

4.1. Datasets and Experimental Settings

Dataset. COCO is the most popular dataset in image captioning field, which comprises 82,783 training images, 40,504 validation images, and 40,775 testing images. Each image is annotated with five descriptions. Considering that the annotations of the official testing set are not provided,

Table 1. Comparison results of SCD-Net with other state-of-the-art autoregressive and non-autoregressive approaches on COCO Karpathy test split for image captioning. † denotes the use of a superior object detector Pix2seq [5].

	Cross-Entropy Loss								CIDEr Score Optimization							
	B@1	B@2	B@3	B@4	M	R	C	S	B@1	B@2	B@3	B@4	M	R	C	S
Autoregressive																
SCST [39]	-	-	-	30.0	25.9	53.4	99.4	-	-	-	-	34.2	26.7	55.7	114.0	-
RFNet [19]	76.4	60.4	46.6	35.8	27.4	56.5	112.5	20.5	79.1	63.1	48.4	36.5	27.7	57.3	121.9	21.2
Up-Down [2]	77.2	-	-	36.2	27.0	56.4	113.5	20.3	79.8	-	-	36.3	27.7	56.9	120.1	21.4
GCN-LSTM [51]	77.3	-	-	36.8	27.9	57.0	116.3	20.9	80.5	-	-	38.2	28.5	58.3	127.6	22.0
LBPf [37]	77.8	-	-	37.4	28.1	57.5	116.4	21.2	80.5	-	-	38.3	28.5	58.4	127.6	22.0
SGAE [49]	77.6	-	-	36.9	27.7	57.2	116.7	20.9	80.8	-	-	38.4	28.4	58.6	127.8	22.1
ORT [15]	-	-	-	-	-	-	-	-	80.5	-	-	38.6	28.7	58.4	128.3	22.6
Transformer [40]	76.1	59.9	45.2	34.0	27.6	56.2	113.3	21.0	80.2	64.8	50.5	38.6	28.8	58.5	128.3	22.6
AoANet [18]	77.4	-	-	37.2	28.4	57.5	119.8	21.3	80.2	-	-	38.9	29.2	58.8	129.8	22.4
M^2 Transformer [8]	-	-	-	-	-	-	-	-	80.8	-	-	39.1	29.2	58.6	131.2	22.6
Non-Autoregressive																
MNIC [12]	75.4	57.7	42.6	30.9	27.5	55.6	108.1	21.0	-	-	-	-	-	-	-	-
MIR [24]	-	-	-	32.5	27.2	-	109.5	20.6	-	-	-	-	-	-	-	-
CMAL [14]	78.5	-	-	35.3	27.3	56.9	115.5	20.8	80.3	-	-	37.3	28.1	58.0	124.0	21.8
SATIC [55]	77.3	-	-	32.9	27.0	-	111.0	20.5	80.6	-	-	37.9	28.6	-	127.2	22.3
Bit Diffusion † [6]	-	-	-	34.7	-	58.0	115.0	-	-	-	-	-	-	-	-	-
SCD-Net	79.0	63.4	49.1	37.3	28.1	58.0	118.0	21.6	81.3	66.1	51.5	39.4	29.2	59.1	131.6	23.0

we utilize the widely adopted Karpathy split [2, 39] and take the 113,287 images for training, 5,000 images for validation, and 5,000 image for testing. Moreover, we also report the performances on the official testing set through online testing server. We convert all the training sentences into lower case and filter out rare words which occur less than four times following [2]. In this way, the final vocabulary consists of 10,199 unique words.

Implementation Details. The whole SCD-Net is implemented in X-modaler codebase [25]. In SCD-Net, we employ the Faster-RCNN pre-trained on ImageNet [9] and Visual Genome [23] to extract image region features. The dimension of the original region feature is 2,048 and we transform it into a 512-dimensional feature by a fully-connected layer. Each word is converted into 14-bits as in [6]. The visual encoder, sentence decoder, and the semantic Transformer contains $N_v = 3$, $N_t = 3$, and $N_p = 3$ Transformer blocks. The hidden size of each Transformer block is set as 512. The training of SCD-Net includes two stages. In the first stage, the whole architecture of SCD-Net is optimized via Adam [21] optimizer on four V100 GPUs with ℓ_2 regression loss and labels smoothing. The optimization process includes 60 epochs with a batch size of 16 and the same learning rate scheduling strategy as in [42] (warmup: 20,000 iterations). In the second stage, we select the initialization model trained by the first stage that achieves the best CIDEr score on validation set. Then SCD-Net is optimized with CIDEr score via our guided self-critical sequence training for 60 epochs. We fix the batch size and

learning rate as 16 and 0.00001, respectively. At inference, the time step number and the time difference is set as 50 and 0, respectively. Five types of evaluation metrics are adopted: BLEU@N [36] (B@1-4), METEOR [4] (M), ROUGE [29] (R), CIDEr [43] (C), and SPICE [1] (S). All the metrics are measured through the source codes released by COCO Evaluation Server [7].

4.2. Performance Comparison

Offline Evaluation. Table 1 summaries the performances of our SCD-Net on the offline COCO Karpathy test split under two different optimization strategies, i.e., the standard cross-entropy loss and the sentence-level optimization in CIDEr score. Considering that our SCD-Net belongs to non-autoregressive image captioning techniques, here we include two groups of baselines (autoregressive and non-autoregressive methods).

In general, there is a clear performance gap between the most state-of-the-art autoregressive and non-autoregressive baselines under both optimization strategies. The results basically verify the weakness of non-autoregressive techniques that incurs word repetition or omissions problems. In between, even a basic autoregressive Transformer encoder-decoder model (Transformer [40]) can achieve comparable performances with the best competitors in non-autoregressive baselines (Bit Diffusion and SATIC). In contrast, by constraining the typical diffusion model with additional semantic prior, our SCD-Net consistently exhibits better performances than autoregressive Transformer

Table 2. Comparison results of SCD-Net with other state-of-the-art approaches on the official test split in online test server.

Model	B@1		B@2		B@3		B@4		M		R		C	
	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40
Ensemble Model														
SCST [39]	78.1	93.7	61.9	86.0	47.0	75.9	35.2	64.5	27.0	35.5	56.3	70.7	114.7	116.7
Up-Down [2]	80.2	95.2	64.1	88.8	49.1	79.4	36.9	68.5	27.6	36.7	57.1	72.4	117.9	120.5
RFNet [19]	80.4	95.0	64.9	89.3	50.1	80.1	38.0	69.2	28.2	37.2	58.2	73.1	122.9	125.1
GCN-LSTM [51]	80.8	95.2	65.5	89.3	50.8	80.3	38.7	69.7	28.5	37.6	58.5	73.4	125.3	126.5
APN [48]	-	-	-	-	-	-	38.9	70.2	28.8	38.0	58.7	73.7	126.3	127.6
AoANet [18]	81.0	95.0	65.8	89.6	51.4	81.3	39.4	71.2	29.1	38.5	58.9	74.5	126.9	129.6
Single Model														
CMAL [14]	79.8	94.3	63.8	87.2	48.8	77.2	36.8	66.1	27.9	36.4	57.6	72.0	119.3	121.2
CAVP [31]	80.1	94.9	64.7	88.8	50.0	79.7	37.9	69.0	28.1	37.0	58.2	73.1	121.6	123.8
SGAE [49]	80.6	95.0	65.0	88.9	50.1	79.6	37.8	68.7	28.1	37.0	58.2	73.1	122.7	125.5
SCD-Net	80.2	95.1	64.9	89.3	50.1	80.1	38.1	69.4	29.0	38.2	58.5	73.5	126.2	129.2

across all metrics. In particular, the CIDEr and SPICE score of SCD-Net under cross-entropy loss is 118.0% and 21.6%, which manifest the absolute improvement of 4.7% and 0.6% against autoregressive Transformer, respectively. This demonstrates the key advantage of our semantic-conditional diffusion model for image captioning in a non-autoregressive fashion. Nevertheless, the performances of our SCD-Net under cross-entropy loss are still inferior to state-of-the-art autoregressive methods with deliberate designs of upgraded Transformer structures (e.g., attention on attention in AoANet and meshed memory in M^2 Transformer). When further equipped with guided self-critical sequence training, our SCD-Net manages to outperform the most competitive autoregressive baseline with an upgraded Transformer structure (M^2 Transformer), leading to slight improvements in CIDEr (0.4%) and SPICE (0.4%). Such encouraging performances confirm the effectiveness of guided self-critical sequence training tailored for diffusion model in SCD-Net.

Online Evaluation. In addition, we also evaluate our SCD-Net on the official testing set by submitting the results to online testing server¹. Table 2 shows the performance comparisons. It is worthy to note that here we only report the performances of our SCD-Net with single model, without using any model ensemble as in some baselines (e.g., APN and AoANet). Similar to the observations in offline evaluation, our SCD-Net again surpasses all the single-model baselines across most metrics, including both autoregressive methods (e.g., SGAE) and non-autoregressive approach (CMAL). Moreover, the single model of our SCD-Net even attains comparable performances with the ensemble version of some state-of-the-art autoregressive methods (e.g., APN). The results basically prove the advantage of exploiting semantic condition in diffusion model for image captioning.

¹<https://competitions.codalab.org/competitions/3221>



Figure 3. Examples of image captioning results generated by Transformer [40], Bit Diffusion [6] and our SCD-Net, as well as the corresponding ground-truths (GT).

Qualitative Analysis. Figure 3 further illustrates the image captioning ground-truths and qualitative results of different approaches for five different images. In general, all the three methods are able to produce somewhat semantically relevant and linguistically coherent captions. However, the explicit modeling of key semantics in an image is overlooked in Transformer, which might result in word omissions (e.g., the missing of “blanket” in the second image). Moreover, the non-autoregressive Bit Diffusion sometimes produces low-quality caption with word repetition problem (e.g., the repetition of “blanket” in the second im-

Table 3. Ablation study on each design in SCD-Net on COCO Karpathy test split. **Base** denotes the conventional diffusion model. **Semantic** represents the use of semantic condition in diffusion model. **RL** refers to the sentence-level optimization in reinforcement learning, which can be set as the typical Self-Critical Sequence Training (SCST) [39] or our proposed Guided SCST (GSCST). **Cascade** denotes the use of the cascaded structure.

Base	Semantic	RL	Cascade	B@4	M	R	C	S	
✓				35.9	27.3	57.2	114.5	20.7	
✓	✓			36.4	27.8	57.4	116.2	21.2	
✓	✓		SCST	34.6	27.7	57.1	120.8	21.5	
✓	✓		GSCST	38.5	29.1	58.6	128.6	22.9	
✓	✓		GSCST	✓	39.4	29.2	59.1	131.6	23.0

age). To address these issues, our SCD-Net novelly capitalizes on semantic-conditioned diffusion model and guided self-critical sequence training to facilitate image captioning, leading to more accurate and descriptive captions.

4.3. Experimental Analysis

Ablation Study. Here we investigate how each design in our SCD-Net influences the overall image captioning performance on COCO Karpathy test split. Table 3 details the performances across different ablated runs of SCD-Net. We start from a base model (**Base**) that directly leverages the typical diffusion process as in [6] to train single Diffusion Transformer, which achieves 114.5% in CIDEr under cross-entropy loss. Note that our implemented Base model is slightly inferior to primary Bit Diffusion (CIDEr: 115.0%) which is equipped with a stronger object detector. Next, by upgrading conventional diffusion model with the additional semantic prior (**Semantic**), we observe clear performance boosts for such semantic-conditional diffusion model. When further training our Diffusion Transformer with sentence-level optimization (**RL**) via **SCST**, only the CIDEr and SPICE scores are slightly improved, while the other scores even decrease. The results basically reveal the difficulty of applying SCST to non-autoregressive solution, due to the indisposed sentence quality of randomly sampled captions in SCST. As an alternative, our Guided SCST (**GSCST**) guides the reinforcement learning of Diffusion Transformer with high-quality sentence derived from an autoregressive Transformer teacher model, thereby leading to significant improvements in terms of all metrics. Finally, after stacking multiple Diffusion Transformers in a cascaded manner (**Cascade**), our SCD-Net obtains the best image captioning performance.

Ablation on Different Transformer Block Number. We further examine how image captioning performances are affected when capitalizing on different numbers of Transformer blocks in single Diffusion Transformer. As shown in Table 4, the performances are relatively smooth when the Transformer block number varies between 3 and

Table 4. Ablation on different Transformer block number in single Diffusion Transformer (without cascaded structure) on COCO Karpathy test split.

# Transformer Block	B@4	M	R	C	S
3	38.5	29.1	58.6	128.6	22.9
4	38.3	29.1	58.5	128.6	23.0
5	38.4	29.0	58.5	128.1	22.9
6	38.2	29.0	58.4	128.1	22.9

Table 5. Ablation on different Diffusion Transformer number in SCD-Net (with cascaded structure) on COCO Karpathy test split.

# Diffusion Transformer	B@4	M	R	C	S
1	38.5	29.1	58.6	128.6	22.9
2	39.4	29.2	59.1	131.6	23.0
3	39.5	29.2	59.1	131.7	23.0

4. When enlarging the block number to 5, the performances slightly decreases. We speculate that this may be the results of unnecessary context information mining among the input tokens with more Transformer blocks in visual encoder and sentence decoder. Accordingly, the Transformer block number in Diffusion Transformer is set as 3, which achieves the best performances with less computational cost.

Ablation on Different Diffusion Transformer Number in SCD-Net. To explore the relationship between the performance and the Diffusion Transformer number in SCD-Net, we show the performances by varying this number from 1 to 3. As listed in Table 5, more stacked Diffusion Transformer in a cascaded manner generally leads to performance improvements. This finding basically highlights the merit of progressively strengthening the output sentence in a cascaded fashion. In particular, once the number of Diffusion Transformer is larger than 1, the performances are less affected, easing the selection of Diffusion Transformer number in SCD-Net practically. Finally, we empirically set the Diffusion Transformer number as 2.

5. Conclusion

In this work, we delve into the idea of strengthening visual-language alignment and linguistical coherence in diffusion model for image captioning. To verify our claim, we shape a new semantic-conditional diffusion process that upgrades diffusion model with additional semantic prior. A guided self-critical sequence training strategy is further devised to stabilize and boost the diffusion process. We empirically validate the superiority of our proposal against state-of-the-art non-autoregressive approaches. At the same time, we are happy to see that our new diffusion model based paradigm manages to outperform the competitive autoregressive method sharing the same Transformer encoder-decoder structure. The results basically demonstrate the encouraging potential of diffusion model in image captioning.

References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016. 6
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 2, 3, 5, 6, 7
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015. 2
- [4] Satyanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop*, 2005. 6
- [5] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. In *ICLR*, 2021. 6
- [6] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022. 1, 2, 3, 4, 6, 7, 8
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6
- [8] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *CVPR*, 2020. 1, 2, 6
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 6
- [10] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 2
- [11] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lin Liang, Zhe Gan, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Injecting semantic concepts into end-to-end image captioning. In *CVPR*, 2022. 1
- [12] Junlong Gao, Xi Meng, Shiqi Wang, Xia Li, Shanshe Wang, Siwei Ma, and Wen Gao. Masked non-autoregressive image captioning. *arXiv preprint arXiv:1906.00717*, 2019. 1, 3, 6
- [13] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. In *ICLR*, 2018. 1, 3
- [14] Longteng Guo, Jing Liu, Xinxin Zhu, Xingjian He, Jie Jiang, and Hanqing Lu. Non-autoregressive image captioning with counterfactuals-critical multi-agent learning. In *IJCAI*, 2021. 5, 6, 7
- [15] Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. Image captioning: Transforming objects into words. *NeurIPS*, 2019. 1, 2, 6
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3
- [17] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 2022. 5
- [18] Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on attention for image captioning. In *ICCV*, 2019. 6, 7
- [19] Wenhao Jiang, Lin Ma, Yu-Gang Jiang, Wei Liu, and Tong Zhang. Recurrent fusion network for image captioning. In *ECCV*, 2018. 6, 7
- [20] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 2
- [21] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [22] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *NeurIPS*, 2021. 3
- [23] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 6
- [24] Jason Lee, Elman Mansimov, and Kyunghyun Cho. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *EMNLP*, 2020. 6
- [25] Yehao Li, Yingwei Pan, Jingwen Chen, Ting Yao, and Tao Mei. X-modaler: A versatile and high-performance codebase for cross-modal analytics. In *ACM MM*, 2021. 6
- [26] Yehao Li, Yingwei Pan, Ting Yao, Jingwen Chen, and Tao Mei. Scheduled sampling in vision-language pretraining with decoupled encoder-decoder network. In *AAAI*, 2021. 1
- [27] Yehao Li, Yingwei Pan, Ting Yao, and Tao Mei. Comprehending and ordering semantics for image captioning. In *CVPR*, 2022. 1
- [28] Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, and Tao Mei. Pointing novel objects in image captioning. In *CVPR*, 2019. 2
- [29] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *ACL Workshop*, 2004. 6
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [31] Daqing Liu, Zheng-Jun Zha, Hanwang Zhang, Yongdong Zhang, and Feng Wu. Context-aware visual policy network for sequence-level image captioning. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1416–1424, 2018. 7
- [32] Fenglin Liu, Xuancheng Ren, Xian Wu, Bang Yang, Shen Ge, and Xu Sun. O2na: An object-oriented non-autoregressive approach for controllable video captioning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021. 1, 3
- [33] Jianjie Luo, Yehao Li, Yingwei Pan, Ting Yao, Hongyang Chao, and Tao Mei. Coco-bert: Improving video-language pre-training with contrastive cross-modal matching and denoising. In *ACM Multimedia*, 2021. 1

- [34] Yingwei Pan, Yehao Li, Jianjie Luo, Jun Xu, Ting Yao, and Tao Mei. Auto-captions on gif: A large-scale video-sentence dataset for vision-language pre-training. In *ACM Multimedia*, 2022. 1
- [35] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. In *CVPR*, 2020. 2
- [36] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 6
- [37] Yu Qin, Jiajun Du, Yonghua Zhang, and Hongtao Lu. Look back and predict forward in image captioning. In *CVPR*, 2019. 6
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 3
- [39] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. 2, 5, 6, 7, 8
- [40] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypertexted, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 1, 2, 4, 5, 6, 7
- [41] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NeurIPS*, 2014. 2
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 4, 6
- [43] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 3, 5, 6
- [44] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015. 2
- [45] Lingxiang Wu, Min Xu, Lei Sang, Ting Yao, and Tao Mei. Noise augmented double-stream graph convolutional networks for image captioning. *IEEE TCSVT*, 2020. 1
- [46] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 2
- [47] Xu Yan, Zhengcong Fei, Zekang Li, Shuhui Wang, Qingming Huang, and Qi Tian. Semi-autoregressive image captioning. In *ACM Multimedia*, 2021. 1
- [48] Xu Yang, Chongyang Gao, Hanwang Zhang, and Jianfei Cai. Auto-parsing network for image captioning and visual question answering. In *ICCV*, 2021. 2, 7
- [49] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *CVPR*, 2019. 2, 6, 7
- [50] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Incorporating copying mechanism in image captioning for learning novel objects. In *CVPR*, 2017. 2
- [51] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018. 2, 6, 7
- [52] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Hierarchy parsing for image captioning. In *ICCV*, 2019. 1
- [53] Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. Boosting image captioning with attributes. In *ICCV*, 2017. 2
- [54] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016. 2
- [55] Yuanen Zhou, Yong Zhang, Zhenzhen Hu, and Meng Wang. Semi-autoregressive transformer for image captioning. In *ICCV*, 2021. 3, 6