

Non-Autoregressive Coarse-to-Fine Video Captioning

Bang Yang,¹ Yuexian Zou,^{1,2*} Fenglin Liu,¹ Can Zhang¹

¹ ADSPLAB, School of ECE, Peking University, Shen Zhen, China

² Peng Cheng Laboratory

{yb.ece, zouyx, fenglinliu98, zhangcan}@pku.edu.cn

Abstract

It is encouraged to see that progress has been made to bridge videos and natural language. However, mainstream video captioning methods suffer from slow inference speed due to the sequential manner of autoregressive decoding, and prefer generating generic descriptions due to the insufficient training of visual words (*e.g.*, nouns and verbs) and inadequate decoding paradigm. In this paper, we propose a non-autoregressive decoding based model with a coarse-to-fine captioning procedure to alleviate these defects. In implementations, we employ a bi-directional self-attention based network as our language model for achieving inference speedup, based on which we decompose the captioning procedure into two stages, where the model has different focuses. Specifically, given that visual words determine the semantic correctness of captions, we design a mechanism of generating visual words to not only promote the training of scene-related words but also capture relevant details from videos to construct a *coarse-grained* sentence “template”. Thereafter, we devise dedicated decoding algorithms that fill in the “template” with suitable words and modify inappropriate phrasing via iterative refinement to obtain a *fine-grained* description. Extensive experiments on two mainstream video captioning benchmarks, *i.e.*, MSVD and MSR-VTT, demonstrate that our approach achieves state-of-the-art performance, generates diverse descriptions, and obtains high inference efficiency.

Introduction

Video captioning aims to automatically describe video contents with plausible sentences, which could be helpful for video retrieval, assisting visually-impaired people and so on. In recent years, neural captioning methods have risen to prominence and they generally adopt the encoder-decoder framework (Venugopalan et al. 2015), where videos are encoded to sequences of vectors with Convolutional Neural Networks (CNNs), and captions are often decoded from these vectors via Recurrent Neural Networks (RNNs) or Transformer (Hori et al. 2017). For real-time industrial applications, a good video captioning system may have low inference latency and describe scene-related details. However, existing methods mostly have deficiencies in both aspects.

*Corresponding author.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

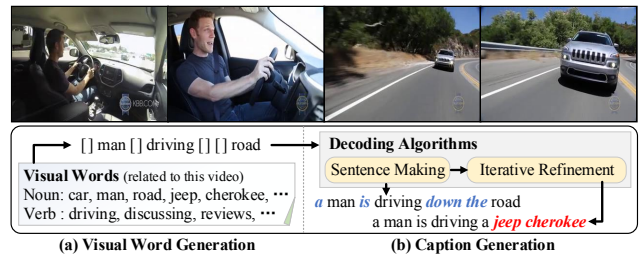


Figure 1: Illustration of the proposed coarse-to-fine captioning procedure, where (a) visual words (*i.e.*, nouns and verbs in this paper) are generated in parallel first to form a *coarse-grained* “template”, based on which (b) a *fine-grained* description is yielded via dedicated decoding algorithms.

For caption generation, current methods are stick to *autoregressive* (AR) decoding, *i.e.*, conditioning each word on the previously generated outputs. Such sequential manner results in high inference latency, which is especially amplified by the fact that fine-grained descriptions are generally long (Gella, Lewis, and Rohrbach 2018). Recently, *non-autoregressive* (NA) decoding that generates words in parallel to achieve significant inference speedup becomes an emerging focus in neural machine translation (NMT) (Gu et al. 2018; Wang et al. 2019b; Shao et al. 2020). Nevertheless, NA decoding suffers from poor approximation to the target distribution after removing the sequential dependency (Gu et al. 2018), leading to a large performance gap compared with AR counterparts.

For describing scene-related details, *visual words* (*e.g.*, nouns and verbs), which are visually-grounded and highly associated with semantic correctness (Song et al. 2017), deserve more attention than *non-visual words* (*e.g.* determiners and prepositions). By analyzing the existing caption corpora (Chen and Dolan 2011; Xu et al. 2016), it is found that visual words are much less than non-visual words in number. Therefore, treating these two kinds of words equally to train a captioning model, as most of previous works do (Afaaq et al. 2019; Huang et al. 2020), will cause insufficient training of meaningful words, which could manifest as the lack of relevant details and diversity in generated descriptions (Dai, Fidler, and Lin 2018; Sammani and Melas-Kyriazi 2020). Besides, it is challenging for the models to generate satisfy-

ing captions with the proneness of error accumulation (Bengio et al. 2015). So a flexible decoding paradigm that supports word modification is also needed.

In this paper, we propose a Non-Autoregressive Coarse-to-Fine (NACF) model to tackle the slow inference speed and unsatisfied caption quality concerns in video captioning. For achieving inference speedup, we employ a bi-directional self-attention based network (Vaswani et al. 2017) as our language model and train it with masked language modeling objective (Devlin et al. 2019) so that any subset of target words can be predicted simultaneously based on the rest ones. For improving caption quality, we propose an alternative paradigm to decompose the captioning procedure into two stages, where the model has different focuses. Specifically, we propose a mechanism of generating visual words, *i.e.*, *visual word generation*, to not only promote the training of scene-related words but also require the model to generate a *coarse-grained* sentence “template” at the first stage of inference phase. As shown in Fig. 1 (a), the generated words in the “template” (*e.g.*, “driving”) summarize a first-glance gist of the scene, which could be instructive for the subsequent generation process. Thereafter, we devise dedicated decoding algorithms to produce *fine-grained* descriptions at the second stage. As shown in Fig. 1 (b), decoding algorithms first complete the “template” by filling in suitable words, and then if necessary, iteratively mask out and reconsider some inappropriate words that the language model is least confident about to ensure sentence fluency or capture more relevant details, *e.g.*, a more precise phrase (“jeep cherokee”) is generated after deliberation.

Our main contributions are summarized as follows¹:

- We propose a Non-Autoregressive Coarse-to-Fine model to deal with slow inference speed and unsatisfied caption quality problems for video captioning.
- We design a mechanism of generating visual words and devise dedicated decoding algorithms to achieve coarse-to-fine rather than word-by-word captioning procedure for capturing more visually-grounded details from videos.
- Extensive experiments on MSVD and MSR-VTT demonstrate the effectiveness of our approach, which achieves state-of-the-art performance, generates diverse descriptions, and obtains high inference efficiency.

Related Work

Video Captioning. With the rapid development in deep learning, the neural captioning methods that follow the encoder-decoder framework have risen to prominence. One of the first works that adopt such framework is (Venugopalan et al. 2015), where captions are generated by LSTM given the mean pooled representation over all frames. Later in (Yao et al. 2015; Song et al. 2017), temporal attention is proposed to adaptively determine which subset of frames to focus on at each decoding step. In particular, Song et al. propose to implicitly distinguish the importance of visual and non-visual words with a hierarchical attention mechanism

¹The code is available at <https://github.com/yangbang18/Non-Autoregressive-Video-Captioning>.

(Song et al. 2017). Besides exploiting the temporal structure of videos, the utilization of multi-modalities and high-level semantics also draws great attention. For instance, Xu et al. propose multimodal attention to selectively focus on content-related modalities (Xu et al. 2017). Liu et al. leverage a pre-trained network to extract attributes, which are defined as the properties observed in visual contents with rich semantic cues, and use them to align visual features (Liu et al. 2019). Most recently, Huang et al. propose to couple attribute prediction with caption generation in an end-to-end manner (Huang et al. 2020). Pan et al. propose to exploit spatio-temporal object interactions and distill such knowledge into the captioning model (Pan et al. 2020). However, all these methods adopt sequential decoding and treat visual and non-visual words equally in terms of the loss function.

Non-Autoregressive Decoding. Due to high inference efficiency, NA decoding has aroused widespread attention in the community of NMT. By removing the sequential dependency, NA decoding can generate all words in one shot to speed up decoding (Gu et al. 2018) but at the cost of inferior accuracy that manifests as token repetitions in the generated outputs (Wang et al. 2019b; Guo et al. 2019; Lee, Mansimov, and Cho 2018). To compensate for the performance degradation, Guo et al. integrate strong conditional signals into the decoder inputs to benefit the learning of internal dependencies within a sentence (Guo et al. 2019). Besides one-shot generation, some works propose to iteratively refine the sentences, so that the model can condition on parts or the whole of the previous outputs (Lee, Mansimov, and Cho 2018; Ghazvininejad et al. 2019; Gu, Wang, and Zhao 2019; Mansimov, Wang, and Cho 2019). But the downside of these methods is that from-scratch parallel generation, *i.e.*, employing the completely unknown sequences as the decoder inputs, often leads to translation errors in the early stages due to insufficient context, which could greatly influence the subsequent predictions.

Summary. Rather than designing a sophisticated architecture, our work contributes by proposing an alternative decoding paradigm, *i.e.*, generating descriptions from coarse-grained to fine-grained with the scheme of NA decoding, to generate semantically correct video captions with higher efficiency. This work pursues the iterative approaches in NMT (Ghazvininejad et al. 2019), but the core difference is that we propose to capture visual words first to formulate partially observed sequences, which can provide rich contextual information with the model to alleviate description ambiguity and thus eventually enhance the caption quality.

Approach

In this section, we first introduce the architecture of our Non-Autoregressive Coarse-to-Fine (NACF) model, followed by the proposed visual word generation. Then, we describe the coarse-to-fine captioning procedure during inference, where three dedicated decoding algorithms are presented.

Architecture

As shown in Fig. 2, our NACF comprises three modules: a CNN-based encoder, a length predictor and a bi-directional self-attention based decoder.

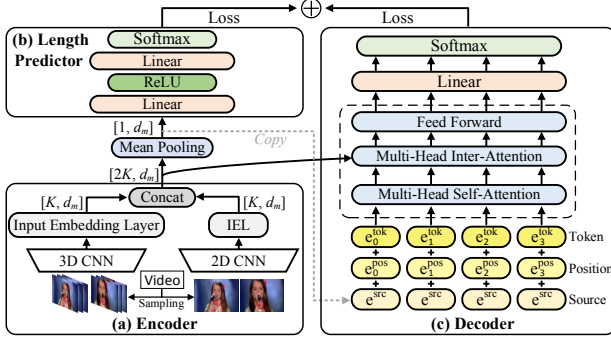


Figure 2: An overview of our proposed NACF architecture, which comprises a CNN-based encoder, a length predictor module and a bi-directional self-attention based decoder.

Encoder. Given a sequence of video frames/clips of length K , we feed it into pre-trained 2D/3D CNNs to obtain visual features $V = \{v_k\}_{k=1}^K \in \mathbb{R}^{K \times d_v}$, which are further encoded to compact representations $R \in \mathbb{R}^{K \times d_m}$ via an input embedding layer (IEL), *i.e.*, $R = f_{IEL}(V)$. Here f_{IEL} adopts the shortcut connection in highway networks (Srivastava, Greff, and Schmidhuber 2015), so it can be formalized as follows (omitting biases for clarity):

$$\begin{aligned} f_{IEL}(V) &= \text{BN}(G \circ \bar{V} + (1 - G) \circ \hat{V}) \\ \bar{V} &= VW_{e1} \\ \hat{V} &= \tanh(\bar{V}W_{e2}) \\ G &= \sigma(\bar{V}W_{e3}) \end{aligned} \quad (1)$$

where BN denotes batch normalization (Ioffe and Szegedy 2015), \circ is the element-wise product, σ is sigmoid function, $W_{e1} \in \mathbb{R}^{d_v \times d_m}$, and $\{W_{e2}, W_{e3}\} \in \mathbb{R}^{d_m \times d_m}$. When considering multi-modalities, *e.g.*, image and motion modalities, we simply apply concatenation to obtain $R \in \mathbb{R}^{2K \times d_m}$.

Length Predictor. Unlike AR decoding that can automatically decide the sequence length N by predicting the end-of-sentence token, NA decoding must know N ahead. So a length predictor (LP) is introduced to predict length distribution $L \in \mathbb{R}^{N_{max}}$ given the outputs R of the encoder:

$$L = f_{LP}(R) = \text{Softmax}(\text{ReLU}(\text{MP}(R)W_{l1})W_{l2}) \quad (2)$$

where MP denotes mean pooling, $W_{l1} \in \mathbb{R}^{d_m \times d_m}$, $W_{l2} \in \mathbb{R}^{d_m \times N_{max}}$, and N_{max} is the predefined maximum sequence length. Given the ground-truth length distribution L^* , whose j -th element denotes the percentage of sentences of length j in the training corpus for a specific video, we minimize the Kullback-Leibler (KL) divergence between L and L^* :

$$\mathcal{L}_{len} = D_{KL}(L^* || L) = - \sum_{j=1}^{N_{max}} l_j^* \log \frac{l_j}{l_j^*} \quad (3)$$

During training, we directly use the sequence length of ground-truth sentences. As for inference, we will depict the utilization of the predicted length distribution L in experimental settings.

Decoder. To obtain a non-autoregressive decoder, we adopt one-layer decoder of Transformer (Vaswani et al. 2017) with two modifications. One is that we remove the causal mask in the self-attention layer. By doing so, our decoder becomes bi-directional, thus the prediction of each token can use both left and right contexts. Another is that we pursue the works in NMT (Lee, Mansimov, and Cho 2018; Guo et al. 2019) to enhance decoder inputs by integrating the copied source information (the dashed line in Fig. 2(c)). To train the model, we use the masked language modeling objective (*a.k.a.* the cloze task) in BERT (Devlin et al. 2019). Specifically, some tokens in a ground-truth sentence Y^* are randomly masked out to obtain a partially-observed sequence Y_{obs} and a masked (unobserved) sequence $Y_{mask} = Y^* \setminus Y_{obs}$. Then the decoder takes Y_{obs} and representations R as inputs to predict the probability distribution over words:

$$p_\theta(y|Y_{obs}, R) = f_{dec}(Y_{obs}, R) \quad (4)$$

where f_{dec} denotes the transformation within the decoder². We only minimize the negative log likelihood of Y_{mask} :

$$\mathcal{L}_{mlm} = - \sum_{y \in Y_{mask}} \log p_\theta(y|Y_{obs}, R) \quad (5)$$

Unlike BERT that uses a small masking ratio (*e.g.*, 15%), we use a uniformly distributed ratio ranging from β_l to β_h so that the model can be trained with examples of different difficulties. Next, we will elaborate on how to generate meaningful visual words with the proposed decoder.

Visual Word Generation

For non-autoregressive video captioning, visual word generation is proposed for two purposes: promoting the training of scene-related words and generating coarse-grained “templates” that serve as starting points at the inference phase. To achieve that, we directly use the proposed decoder without introducing extra parameters. Formally, given a ground-truth sentence Y^* of length N , we first construct a corresponding target sequence $Y^{vis} = \{y_n^{vis}\}_{n=1}^N$ as follows:

$$y_n^{vis} = \begin{cases} y_n^* & \text{if POS}(y_n^*) \in \{\text{noun, verb}\} \\ [mask] & \text{otherwise} \end{cases} \quad (6)$$

where POS(\cdot) denotes the part-of-speech of a word. Then the decoder is forced to predict Y^{vis} without available word information, *i.e.*, $Y_{obs}^{vis} = \emptyset_{[vis]}$ (a sequence of the same special token $[vis]$ in practice). Hence the loss function for visual word generation is defined as:

$$\mathcal{L}_{vis} = - \sum_{y \in Y^{vis}} \log p_\theta(y|\emptyset_{[vis]}, R) \quad (7)$$

As the generation process solely depends on R , the generated visual words may not be comprehensive. But as we will verify later, they are instructive for the follow-up caption generation process. Finally, the overall loss function of our approach is formulated as:

$$\mathcal{L}_{NACF} = \mathcal{L}_{len} + \mathcal{L}_{mlm} + \lambda \mathcal{L}_{vis} \quad (8)$$

where λ is set to 0.8 empirically.

²Detailed formulation of f_{dec} is left to the technical appendix.

Coarse-to-Fine Captioning

To yield plausible descriptions during inference, our captioning procedure is decomposed into two stages. At the first stage, we generate a coarse-grained “template” $Y^{(0)}$ and collect its confidence $C^{(0)}$ given the predicted sequence length N and representations R :

$$y_n^{(0)}, c_n^{(0)} = (\arg) \max_w p_\theta(y = w | \mathcal{O}_{[vis]}, R) \quad (9)$$

where $y_n^{(0)}$ is either the $[mask]$ token or a visual word (see Eq. 6). In the subsequent generation process, four variables are introduced: the number of iterations T , the observed sequence at t -th ($t \in [1, T]$) iteration $Y_{obs}^{(t)}$, the prediction result $Y^{(t)}$ and its confidence $C^{(t)}$. Specifically, $Y_{obs}^{(1)}$ is initialized by the visual words in the coarse-grained “template” $Y^{(0)}$:

$$Y_{obs}^{(1)} = \{y_n^{(0)} | y_n^{(0)} \neq [mask]\} \quad (10)$$

Then at the second stage, three dedicated decoding algorithms, *i.e.*, Mask-Predict (MP) (Ghazvininejad et al. 2019), Easy-First (EF) and Left-to-Right (L2R) are introduced to generate fine-grained descriptions.

Mask-Predict (MP). This algorithm iterates over two steps at t -th iteration: *Mask* where m_t tokens with lowest confidence are masked out and *Predict* where those masked tokens are reconsidered based on the rest $N - m_t$ tokens. The order of these two steps are switched in this paper because $Y_{obs}^{(1)}$ is already obtained. We only update the prediction results $Y^{(t)}$ and confidence $C^{(t)}$ for those unobserved tokens given $Y_{obs}^{(t)}$ and R :

$$y_n^{(t)}, c_n^{(t)} = \begin{cases} (\arg) \max_w p_\theta(y = w | Y_{obs}^{(t)}, R) & \text{if } n \in I_t \\ y_n^{(t-1)}, c_n^{(t-1)} & \text{otherwise} \end{cases} \quad (11)$$

where I_t denotes the index set of unobserved tokens:

$$I_t = \{n | y_n^{(t-1)} \notin Y_{obs}^{(t)}\} \quad (12)$$

As low confidence $c_n^{(t)}$ means the token $y_n^{(t)}$ is incompatible with others, reconsidering such token could benefit caption quality. So for the next iteration, $Y_{obs}^{(t+1)}$ is defined as:

$$Y_{obs}^{(t+1)} = \{y_j^{(t)} | j \in \text{topk}(C^{(t)}, k = N - m_{t+1})\} \quad (13)$$

where we use a linear decay ratio r to decide m_t and make sure there is at least one token to be reconsidered:

$$r = \frac{T - t + 1}{T}, \quad m_t = \max(\lfloor N \cdot r \rfloor, 1) \quad (14)$$

Easy-First (EF). This algorithm generates q tokens with highest confidence among the unobserved tokens at each iteration. Given N and u (the cardinality of $Y_{obs}^{(1)}$), EF algorithm needs $T = \lceil (N - u)/q \rceil$ iterations and can be briefly formulated as:

$$Y_{obs}^{(t+1)} = Y_{obs}^{(t)} \cup \{y_{I_t, j}^{(t)} | j \in \text{topk}(C_{I_t}^{(t)}, k = q)\} \quad (15)$$

where $Y_{I_t}^{(t)}$ and $C_{I_t}^{(t)}$ denote the prediction and confidence of unobserved tokens at t -th iteration respectively. $Y^{(t)}$ and

	$Y_{obs}^1 = \emptyset$	$Y_{obs}^1 = \{\text{girl, field}\}$
	a girl playing playing a a ball a girl is playing on a ball a girl is jumping on a trampoline MP	a girl is playing in a field a girl is jumping in a field a girl is jumping in a field CT-MP
	a girl _ _ _ _ a girl is _ _ a _ a girl is jumping on a _ a girl is jumping on a trampoline EF	a girl _ _ _ a field a girl is _ in a field a girl is jumping in a field a girl is jumping in a field CT-EF
	a girl _ _ _ _ a girl is playing _ _ _ _ a girl is playing with a _ a girl is playing with a ball L2R	a girl is _ _ field a girl is playing in _ field a girl is playing in a field a girl is playing in a yard CT-L2R

Figure 3: Illustration of how to generate captions using different decoding algorithms (best viewed in color). The newly generated words (in color) are predicted based on the observed words (in black). We set $T = 3$ for MP algorithm, $q = 2$ for L2R and EF algorithms. The prefix “CT-” mean using the coarse-grained “templates” (Eq. 10).

$C^{(t)}$ are calculated by Eq. 11 while I_t is computed by Eq. 12. Since the visual words in $Y_{obs}^{(1)}$ are not modified during generation, we can reconsider them based on $Y = Y^{(T)} \setminus Y_{obs}^{(1)}$ and R :

$$y_n^{(T)}, c_n^{(T)} = (\arg) \max_w p_\theta(y = w | Y, R) \quad \text{s.t. } y_n^{(0)} \neq [mask] \quad (16)$$

Left-to-Right (L2R). In contrast to EF, this algorithm is monotonous, *i.e.*, it generates q tokens among the unobserved tokens from left to right at each iteration. L2R algorithm also needs $T = \lceil (N - u)/q \rceil$ iterations and can be briefly defined as:

$$Y_{obs}^{(t+1)} = Y_{obs}^{(t)} \cup \{y_{I_t, 1}^{(t)}, \dots, y_{I_t, q}^{(t)}\} \quad (17)$$

where $Y_{I_t}^{(t)}$ denote the prediction of unobserved tokens at t -th iteration. $Y^{(t)}$ (also $C^{(t)}$) is computed by Eq. 11 while I_t is calculated by Eq. 12. Similar to EF, one more iteration can be added to reconsider the visual words (Eq. 16). Although we make both L2R and EF to generate q words at each iteration, a fixed number of iterations T can be set for both of them to produce captions in constant time. Besides, both L2R and EF algorithms can cooperate with the MP algorithm to iteratively refine sentences if necessary.

Example. Each of the decoding algorithms mentioned above can start with either the generated coarse-grained “template” (Eq. 10) or a completed unknown sequence by setting $Y_{obs}^{(1)} = \emptyset$ and $I_1 = \{1, 2, \dots, N\}$ (Eq. 12). To differentiate these two versions, we name the algorithms utilizing coarse-grained “templates” with the prefix “CT-”. As shown in Fig. 3, the original version of algorithms, *i.e.*, MP, EF and L2R, hallucinate some concepts (*e.g.*, “trampoline” and “ball”) that does not exist in visual contents due to the limited contextual information of completed unknown sequences. But under the guidance of some generated visual words (*i.e.*, “girl” and “field”), the algorithms can produce more precise descriptions, which will be verified later.

Model	MSVD			MSR-VTT		
	BLEU@4	METEOR	CIDEr-D	BLEU@4	METEOR	CIDEr-D
STAT (Yan et al. 2019)	52.0	33.3	73.8	39.3	27.1	43.8
GRU-EVE (Aafaq et al. 2019)	47.9	35.0	78.1	38.3	28.4	48.1
POS-CG (Wang et al. 2019a)	52.5	34.1	88.7	42.0	28.2	48.7
MARN (Pei et al. 2019)	48.6	35.1	92.2	40.4	28.1	47.1
MAD-SAP (Huang et al. 2020)	53.3	35.4	90.8	41.3	28.3	48.5
STG-KD (Pan et al. 2020)	52.2	36.9	93.0	40.5	28.3	47.1
AR-B	48.7	35.3	91.8	40.5	28.7	49.1
NA-B	53.7	35.5	92.8	40.4	28.0	47.6
Our NACF	55.6	36.2	96.3	42.0	28.7	51.4

Table 1: Comparison with the state-of-the-art methods on MSVD and MSR-VTT, where AR-B and NA-B are our baselines.

Experiments

In this section, we evaluate our NACF on two datasets: Microsoft Video Description (MSVD) (Chen and Dolan 2011) and MSR-Video To Text (MSR-VTT) (Xu et al. 2016).

Experimental Settings

Datasets. MSVD contains 1,970 video clips and roughly 80,000 English sentences. We follow the split settings in prior works (Pei et al. 2019; Pan et al. 2020), *i.e.*, 1,200, 100 and 670 videos for training, validation and testing, respectively. MSR-VTT consists of 10,000 video clips, each of which has 20 captions and a category tag. Following the official split, we use 6,513, 497 and 2,990 videos for training, validation and testing, respectively. The vocabulary size of MSVD is 9,468, whereas that of MSR-VTT is 10,547.

Feature Extraction. We follow (Pei et al. 2019) and opt for the same type of features, *i.e.*, 2048-D image features from ResNet-101 (He et al. 2016) pre-trained on the ImageNet dataset (Deng et al. 2009), 2048-D motion features from ResNeXt-101 with 3D convolutions (Hara, Kataoka, and Satoh 2018) pre-trained on the Kinetics dataset (Kay et al. 2017), and all category tags included in MSR-VTT.

Length Beam and Teacher Rescoring. Following the common practice of noisy parallel decoding during inference (Gu et al. 2018; Wang et al. 2019b), we select top B length candidates from the predicted length distribution L , and decode the same example with different lengths in parallel. An autoregressive counterpart (*i.e.*, the AR-B introduced later) is then used to re-score these B candidates. We finally select captions with the highest confidence as hypotheses.

Parameter Settings. The maximum sequence length N_{max} is set to 20 for MSVD, whereas $N_{max} = 30$ for MSR-VTT. We empirically set $K = 8$ for each modality. For the decoder, we adopt 1 decoder layer, 512 model dimensions, 2,048 hidden dimensions and 8 attention heads per layer. Both word and position embeddings are implemented by trainable 512-D embedding layers. For regularization, we use 0.5 dropout and 5×10^{-4} L_2 weight decay. We train batches of 64 video-sentence pairs using ADAM (Kingma and Ba 2015) with an initial learning rate of 5×10^{-3} . We stop training our model until 50 epochs are reached. We use NLTK toolkit (Bird, Klein, and Loper 2009) for part-

of-speech tagging. In the following experiments, our NACF uses CT-MP algorithm with the number of iterations T of 5 and beam size B of 6 unless otherwise specified.

Evaluation Metrics. We report three common metrics including BLEU (Papineni et al. 2002), METEOR (Banerjee and Lavie 2005) and CIDEr-D (Vedantam, Lawrence Zitnick, and Parikh 2015). All metrics are computed by Microsoft COCO Evaluation Server (Chen et al. 2015).

Compared Approaches. We compare our NACF with the following state-of-the-art methods: **STAT** (Yan et al. 2019) and **GRU-EVE** (Aafaq et al. 2019) which capture spatio-temporal dynamics by attention mechanism and Short Fourier Transform respectively, **POS-CG** (Wang et al. 2019a) which uses the global syntactic part-of-speech information, **MARN** (Pei et al. 2019) which leverages memory network to capture cross-video contents, **MAD-SAP** (Huang et al. 2020) which predicts a concise set of attributes at each decoding step, and **STG-KD** (Pan et al. 2020) which distills the knowledge of spatio-temporal object interactions. Additionally, we consider two baselines, namely autoregressive baseline (**AR-B**) and non-autoregressive baseline (**NA-B**). Specifically, AR-B has similar architecture as our NACF, but it excludes the length predictor and includes causal mask in the self-attention layer. NA-B is the same with our NACF but it excludes visual word generation, *i.e.*, $\lambda = 0$ (Eq. 8).

Performance Comparison

The quantitative results in Table 1 illustrate that our NACF achieves state-of-the-art performance on both MSVD and MSR-VTT datasets, which is mainly beneficial from the proposed coarse-to-fine captioning procedure. Since we opt for the same type of features as MARN and similar features as MAD-SAP and STG-KD, fair comparisons between these methods and our approach can be guaranteed. Notably, our NACF achieves significant improvement on CIDEr-D, *e.g.*, a relative improvement of 3.5% on MSVD compared with STG-KD while 6.0% on MSR-VTT compared with MAD-SAP. As the CIDEr-D metric is to punish the often-seen but uninformative n-grams in the dataset, the superior performance on CIDEr-D indicates that our NACF can capture more scene-related keywords from videos. It is noteworthy that our NACF is slightly worse than STG-KD on the ME-

exp	Model	B4	M	C_D
1	NA-B	40.4	28.0	47.6
2	NA-B w/ \mathcal{L}_{vis}	40.8	28.2	49.4
3	NA-B w/ \mathcal{L}_{vis} , CT (NACF)	42.0	28.7	51.4
4	AR-B	40.5	28.7	49.1
5	AR-B w/ \mathcal{L}_{vis}	41.4	29.0	50.8

Table 2: Effect of the visual word generation loss (\mathcal{L}_{vis} , Eq. 7) and the generated coarse-grained templates (CT) on the performance on MSR-VTT in terms of BLEU@4 (B4), METEOR (M) and CIDEr-D (C_D).




Figure 4: The average relative growth rate of word frequency after training with visual word generation (\mathcal{L}_{vis}). Here we compare exp2 with exp1 (listed in Table 2), and detail the results in all 20 categories of videos from MSR-VTT. “doc.” and “ads.” are short for documentary and advertisement.

TEOR metric in MSVD, which is because the latter learn spatio-temporal object interactions well in a small dataset with few portions of animations (Pan et al. 2020).

We also present the performance of baselines in Table 1, and obtain two observations. (1) Compared with AR-B, NA-B obtains superior performance on MSVD while performs poorly on MSR-VTT, showing that the removal of sequential dependency makes NA decoding face more a severe “multi-modality” problem³ (Gu et al. 2018) on a larger dataset. (2) Our NACF surpasses NA-B by a large margin on both datasets. As we will show in the next subsection, this superior performance is attributed to the proposal of visual word generation, which not only generates informative gradients to promote the training of visual words but also alleviates the “multi-modality” problem by providing a warm start with the model during inference.

Ablation Studies and Analyses

Visual Word Generation. Our proposed visual word generation task plays a critical role in both training and inference phases. During training, this task generates auxiliary gradients, thus the effect of \mathcal{L}_{vis} on performance is worth exploring. As shown in Table 2, \mathcal{L}_{vis} brings promising performance gains for both NA-B (exp2 vs. exp1) and AR-B (exp5 vs. exp4), especially on the CIDEr-D metric. To figure out the word-level effect of \mathcal{L}_{vis} , first we collect the captions

³For example, a NA model considers two possible captions C_1 and C_2 , it could predict one token from C_1 while another token from C_2 due to the conditional independence.

Model	Novel	Unique	Vocab Usage
AR-B	17.19	25.79	3.36
NA-B	23.88	31.24	3.24
NACF	34.35	42.47	3.83

Table 3: Diversity of generated captions at various aspects (%) on MSR-VTT.

Algorithm	B1	B2	B3	B4	M	C_D
MP	81.0	67.5	53.5	40.8	28.2	49.4
EF	81.5	67.9	54.1	41.4	28.7	50.6
L2R	81.0	67.3	53.4	40.8	28.4	48.7
CT-MP	82.2	68.7	54.7	42.0	28.7	51.4
CT-EF	82.1	68.4	54.4	41.7	28.8	51.8
CT-L2R	81.7	68.2	54.3	41.7	28.7	50.6

Table 4: Performance of our NACF using different decoding algorithms on MSR-VTT. Here $T = 5$ and $q = 1$.

generated by the model with or without \mathcal{L}_{vis} , then we measure the relative growth rates of word frequency, and finally we average the results of words of the same type. As shown in Fig. 4, visual words get an overall boost in various categories of videos. All these results demonstrate that training with \mathcal{L}_{vis} can improve the caption quality by addressing the insufficient training of visual words.

During inference, visual word generation can produce coarse-grained “templates” (CT). An obvious improvement of taking CT as starting points can be observed in Table 2 (exp3 vs. exp2) and Table 4, which indicates that decoding with some known visual words can generate more semantically correct captions than the from-scratch generation that starts with completely unknown sequences. In summary, our proposed visual word generation is versatile and it benefits a lot for non-autoregressive video captioning.

Diversity. To quantify caption diversity, we compute three metrics following (Dai, Fidler, and Lin 2018), namely *Novel* (the percentage of captions that have not been seen in the training data), *Unique* (the percentage of captions that are unique among the generated captions) and *Vocab Usage* (the percentage of words in the vocabulary that are adopted to generate captions). As shown in Table 3, NACF achieves the best performance across all metrics, indicating the advantage of our approach in terms of caption diversity.

Decoding Algorithm. In Table 4, we present the performance of all aforementioned decoding algorithms. As we can observe, (CT-)EF and (CT-)MP always outperform (CT-)L2R. Therefore, we can conclude that adaptive generation rather than monotonic generation is requisite for our NACF to generate plausible descriptions.

Inference Efficiency. We measure latency⁴ following (Guo et al. 2019; Wang et al. 2019b), and conduct experiments in PyTorch on a single NVIDIA Titan X. Fig. 5 shows

⁴Latency is computed as the time to decode a single sentence without minibatching, averaged over the whole test set.




Figure 5: Relative decoding speed-up versus CIDEr-D on (a) MSVD and (b) MSR-VTT. The circles in red denote NACF (CT-MP) while in purple denote NACF (CT-EF). The squares in green denote NA-B (MP). The triangles in blue denote AR-B. We take the latency of AR-B ($B = 5$) as a benchmark, which costs 37.4 ms and 43.8 ms on MSVD and MSR-VTT, respectively.

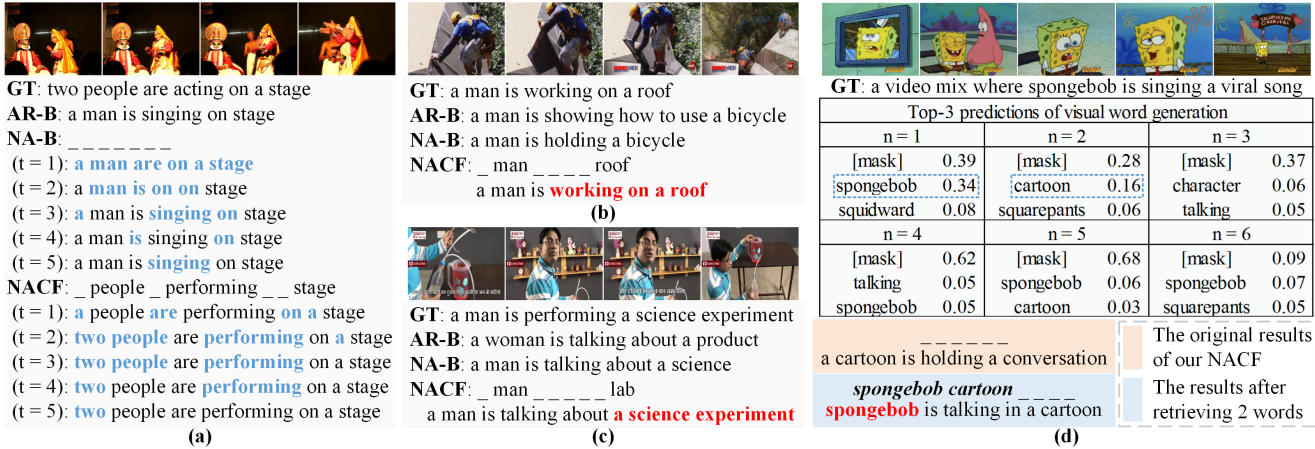


Figure 6: Qualitative results on MSVD and MSR-VTT. The incomplete sentences of our NACF are generated coarse-grained “templates” (CT). In (a), words in blue denote the update during iterative refinement, while for the rest examples, accurate keywords are highlighted in red. In (d), our NACF generates an unsatisfying caption due to the inadequate generation of visual words, which can be alleviated by retrieving two visual words (“spongebob” and “cartoon”) that are potential to be predicted.

the speed-performance trade-off on MSVD and MSR-VTT, where we take the latency of AR-B ($B = 5$) as a benchmark. Notably, with CT-MP algorithm, our NACF can generate captions over 3.2 times ($B = 4, T = 1$) faster than AR-B ($B = 5$) on MSVD while 2.2 times ($B = 5, T = 3$) on MSR-VTT without performance degradation. These results demonstrate the high inference speed of our NACF.

Qualitative Analysis

Fig. 6 shows a few visualized examples of generated captions for different models. As we can see in (a), (b), and (c), while baselines AR-B and NA-B mistaken the video contents, our NACF can capture relevant visual words (e.g., “roof” in (b) and “lab” in (c)) and thus generate more precise captions. Specifically in (a), where the intermediate process of iterative refinement is presented, we can observe that the generated visual words of our NACF provide rich contextual information to predict the quantifier “two”. However in (d), our NACF suffers from inadequate visual word generation. To figure out which visual words are potential to be

predicted, we visualize top-3 predictions in (d) and find that most of them are somewhat relevant. According to the result that the description gets improved after retrieving “spongebob” and “cartoon”, a more robust mechanism of generating visual words deserves further study.

Conclusion

In this paper, we propose a novel Non-Autoregressive Coarse-to-Fine (NACF) model for video captioning, which is based on a masked language model for parallelization and equipped with visual word generation and dedicated decoding algorithms to generate accurate and diverse captions in a coarse-to-fine manner. Extensive experiments on two video captioning benchmarks, MSVD and MSR-VTT, demonstrate that our proposed NACF achieves state-of-the-art performances with two unique advantages, i.e., generating more vivid descriptions and decoding faster (e.g., 3.2 times speed-up without performance degradation) than autoregressive captioning models. In our future study, we will focus on interpretable and controllable caption generation.

Acknowledgements

Special acknowledgements are given to AOTO-PKUSZ Joint Research Center of Artificial Intelligence on Scene Cognition technology Innovation for its support.

References

- Aafaq, N.; Akhtar, N.; Liu, W.; Gilani, S. Z.; and Mian, A. 2019. Spatio-temporal dynamics and semantic attribute enriched visual encoding for video captioning. In *CVPR*.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *ACL workshop*.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in NIPS*.
- Bird, S.; Klein, E.; and Loper, E. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Chen, D.; and Dolan, W. B. 2011. Collecting highly parallel data for paraphrase evaluation. In *ACL-HLT*, 190–200.
- Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Dai, B.; Fidler, S.; and Lin, D. 2018. A neural compositional paradigm for image captioning. In *Advances in NIPS*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Gella, S.; Lewis, M.; and Rohrbach, M. 2018. A Dataset for Telling the Stories of Social Media Videos. In *EMNLP*.
- Ghazvininejad, M.; Levy, O.; Liu, Y.; and Zettlemoyer, L. 2019. Mask-predict: Parallel decoding of conditional masked language models. In *EMNLP-IJCNLP*.
- Gu, J.; Bradbury, J.; Xiong, C.; Li, V. O.; and Socher, R. 2018. Non-autoregressive neural machine translation. In *ICLR*.
- Gu, J.; Wang, C.; and Zhao, J. 2019. Levenshtein transformer. In *Advances in NIPS*.
- Guo, J.; Tan, X.; He, D.; Qin, T.; Xu, L.; and Liu, T.-Y. 2019. Non-autoregressive neural machine translation with enhanced decoder input. In *AAAI*.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hori, C.; Hori, T.; Lee, T.-Y.; Zhang, Z.; Harsham, B.; Hershey, J. R.; Marks, T. K.; and Sumi, K. 2017. Attention-based multimodal fusion for video description. In *ICCV*.
- Huang, Y.; Chen, J.; Ouyang, W.; Wan, W.; and Xue, Y. 2020. Image captioning with end-to-end attribute detection and subsequent attributes prediction. *IEEE Transactions on Image Processing* 29.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 448–456.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Lee, J.; Mansimov, E.; and Cho, K. 2018. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *EMNLP*.
- Liu, F.; Liu, Y.; Ren, X.; He, X.; and Sun, X. 2019. Aligning visual regions and textual concepts for semantic-grounded image representations. In *Advances in NIPS*.
- Mansimov, E.; Wang, A.; and Cho, K. 2019. A generalized framework of sequence generation with application to undirected sequence models. *arXiv preprint arXiv:1905.12790*.
- Pan, B.; Cai, H.; Huang, D.-A.; Lee, K.-H.; Gaidon, A.; Adeli, E.; and Niebles, J. C. 2020. Spatio-Temporal Graph for Video Captioning with Knowledge Distillation. In *CVPR*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*.
- Pei, W.; Zhang, J.; Wang, X.; Ke, L.; Shen, X.; and Tai, Y.-W. 2019. Memory-attended recurrent network for video captioning. In *CVPR*.
- Sammani, F.; and Melas-Kyriazi, L. 2020. Show, Edit and Tell: A Framework for Editing Image Captions. In *CVPR*.
- Shao, C.; Zhang, J.; Feng, Y.; Meng, F.; and Zhou, J. 2020. Minimizing the bag-of-ngrams difference for non-autoregressive neural machine translation. In *AAAI*.
- Song, J.; Guo, Z.; Gao, L.; Liu, W.; Zhang, D.; and Shen, H. T. 2017. Hierarchical lstm with adjusted temporal attention for video captioning. In *IJCAI*.
- Srivastava, R. K.; Greff, K.; and Schmidhuber, J. 2015. Training very deep networks. In *Advances in NIPS*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in NIPS*.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *CVPR*.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R.; and Saenko, K. 2015. Translating Videos to Natural Language Using Deep Recurrent Neural Networks. In *NAACL*.
- Wang, B.; Ma, L.; Zhang, W.; Jiang, W.; Wang, J.; and Liu, W. 2019a. Controllable video captioning with POS sequence guidance based on gated fusion network. In *ICCV*.

Wang, Y.; Tian, F.; He, D.; Qin, T.; Zhai, C.; and Liu, T.-Y. 2019b. Non-autoregressive machine translation with auxiliary regularization. In *AAAI*.

Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.

Xu, J.; Yao, T.; Zhang, Y.; and Mei, T. 2017. Learning multimodal attention LSTM networks for video captioning. In *ACM Multimedia*.

Yan, C.; Tu, Y.; Wang, X.; Zhang, Y.; Hao, X.; Zhang, Y.; and Dai, Q. 2019. STAT: Spatial-Temporal Attention Mechanism for Video Captioning. *IEEE Transactions on Multimedia*.

Yao, L.; Torabi, A.; Cho, K.; Ballas, N.; Pal, C.; Larochelle, H.; and Courville, A. 2015. Describing videos by exploiting temporal structure. In *ICCV*.