

Refined Semantic Enhancement towards Frequency Diffusion for Video Captioning

Xian Zhong¹, Zipeng Li¹, Shuqin Chen^{2,*}, Kui Jiang^{3,*}, Chen Chen⁴, Mang Ye³

¹ School of Computer Science and Artificial Intelligence, Wuhan University of Technology

² College of Computer, Hubei University of Education

³ School of Computer Science, Wuhan University

⁴ Center for Research in Computer Vision, University of Central Florida

{zhongx, lizipeng, csqcwx0801}@whut.edu.cn, kuijiang@whu.edu.cn, chen.chen@crcv.ucf.edu, mangye16@gmail.com

Abstract

Video captioning aims to generate natural language sentences that describe the given video accurately. Existing methods obtain favorable generation by exploring richer visual representations in encode phase or improving the decoding ability. However, the long-tailed problem hinders these attempts at low-frequency tokens, which rarely occur but carry critical semantics, playing a vital role in the detailed generation. In this paper, we introduce a novel Refined Semantic enhancement method towards Frequency Diffusion (RSFD), a captioning model that constantly perceives the linguistic representation of the infrequent tokens. Concretely, a Frequency-Aware Diffusion (FAD) module is proposed to comprehend the semantics of low-frequency tokens to break through generation limitations. In this way, the caption is refined by promoting the absorption of tokens with insufficient occurrence. Based on FAD, we design a Divergent Semantic Supervisor (DSS) module to compensate for the information loss of high-frequency tokens brought by the diffusion process, where the semantics of low-frequency tokens is further emphasized to alleviate the long-tailed problem. Extensive experiments indicate that RSFD outperforms the state-of-the-art methods on two benchmark datasets, *i.e.*, MSR-VTT and MSVD, demonstrate that the enhancement of low-frequency tokens semantics can obtain a competitive generation effect. Code is available at <https://github.com/lzp870/RSFD>.

Introduction

Video captioning is the task of understanding video content and describing it accurately. It has considerable application value in social networks, human-computer interaction, and other fields. Despite recent progress in this field, it remains a challenging task as existing models prefer to generate commonly-used words disregarding many infrequent tokens that carry critical semantics of video content, limiting the refined semantic generation in the testing phase.

The existing methods generally adopt the encoder-decoder framework (Venugopalan et al. 2015), where the encoder generates visual representation by receiving a set of consecutive frames as input, and the decoder generates captions via recurrent neural networks (RNNs) or Transformer (Hori et al. 2017). Some efforts (Hu et al. 2022a,b; Jia et al. 2022; Liao

*Corresponding authors.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

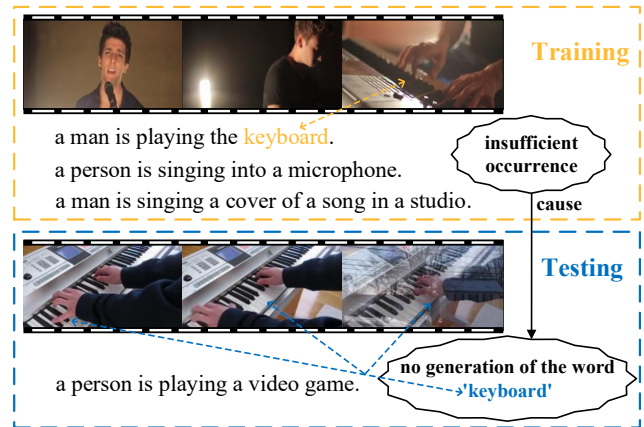


Figure 1: Illustration of existing models limitations: the insufficient occurrence of low-frequency tokens affects the information extraction during training, thus failing to generate refined low-frequency words in the testing phase.

et al. 2022b; Xu et al. 2022) are developed to explore richer visual features. Prior works (Zhang and Peng 2019; Pan et al. 2020; Zhang et al. 2020; Li et al. 2022c) enhance the spatio-temporal representations between objects, while others (Jin et al. 2020; Gao et al. 2022) tend to improve the architecture of the decoder to obtain better linguistic representations. However, they either focus on the enhancement of visual representations or the generation ability of the model, ignoring the tokens with insufficient occurrence in the training phase. Many low-frequency tokens, as a critical factor to prove caption performance, have not received adequate attention, thus challenging for the model to express similar information in other videos in the testing phase like the instance in Fig. 1. Concretely, a man is playing the keyboard and the model has partially generated “a person is playing a video game”. The result commonly occurs when the low-frequency token “keyboard” is of insufficient occurrence in the given video.

To deal with the challenges of token imbalance, the approach (Wu et al. 2016) in neural machine translation splits words into more fine-grained units. However, the token imbalance phenomenon is not fundamentally eliminated. BMI (Xu et al. 2021) has been dedicated to assigning adaptive weights to target tokens in light of the token frequency. Although

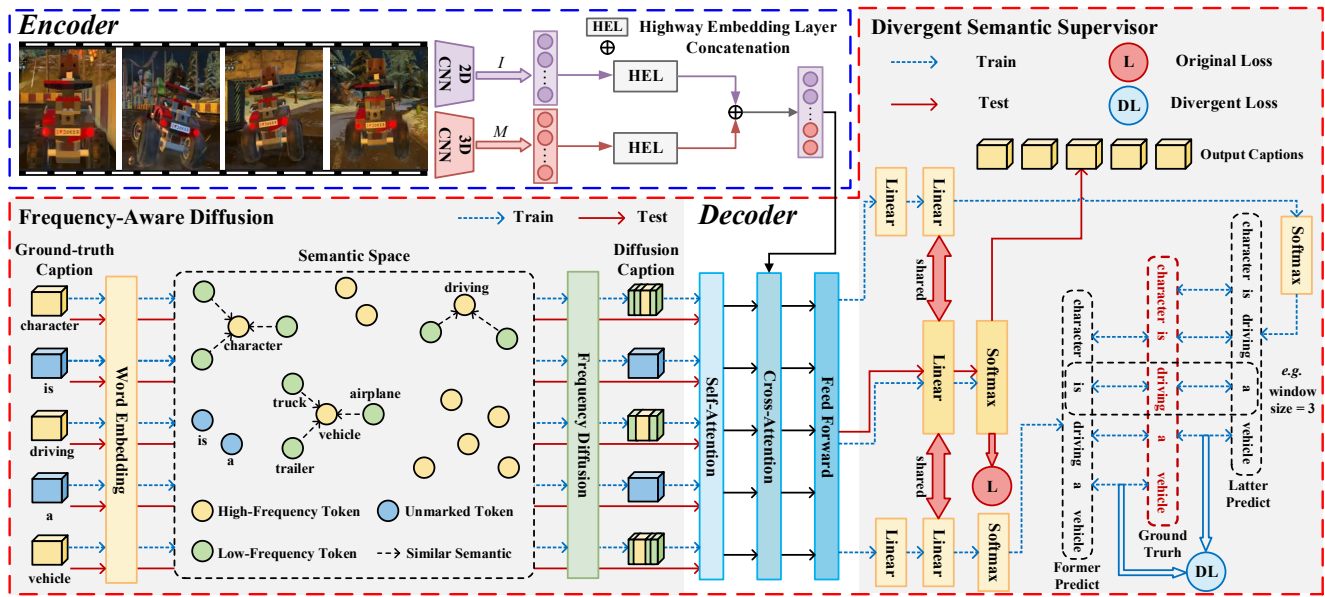


Figure 2: Overview of the proposed RSFD architecture. It mainly consists of the encoder in the top-left box and the decoder with FAD and DSS modules in another box. In the training phase, FAD promotes the model comprehending the refined information by mapping the ground-truth caption to the semantic space and fusing it in frequency diffusion. DSS supervises the central word to obtain its distinctive semantics. In the testing phase, only the transformer-based parts are implemented for sentence generation.

they increase the exposure of low-frequency tokens during training, they potentially damage the information of high-frequency ones. In video captioning, few researchers focus on the learning effect of low-frequency tokens. ORG-TRL (Zhang et al. 2020) introduces an external language model to extend the low-frequency tokens, which depends heavily on external knowledge. Taking the above attempts as a reference, the motivation of our design can be depicted as follows: 1) With the help of high-frequency tokens more times trained by the model, low-frequency tokens entrust their semantics to high-frequency ones obtaining adequate exposure to alleviate the effect of token imbalance. 2) Instruction from multiple teachers to high-frequency tokens boosts specific semantics suitable for them as compensation due to information loss by the previous stage, and low-frequency tokens further enhance their generation.

In this work, we propose a novel method named Refined Semantic enhancement towards Frequency Diffusion (RSFD) for video captioning to address the mentioned issue. In RSFD, unlike previous effort (Zhang et al. 2020) that with the help of external modules, we propose a module named Frequency-Aware Diffusion (FAD) to explore the semantics of low-frequency tokens inside our model. Firstly, in light of the occurrence numbers of each token in the whole corpus and in each video they appear, we split them into three categories (high-frequency, low-frequency, and unmarked). Secondly, we exploit the diffusion process to comprehend low-frequency tokens by integrating them with high-frequency ones. With the assistance of the diffusion of low-frequency tokens, the model develops the ability to absorb the semantics of low-frequency tokens. Afterward, to alleviate the semantic

loss of high-frequency tokens (Pure high-frequency words can be entirely trained, whereas they carry low-frequency semantics that disturbs the model’s learning of them) by the diffusion process and further excavate the low-frequency information, we further design a Divergent Semantic Supervisor (DSS) module to exploit adjacency words to supervise the central word. With the adjacent supervisory information, the central high-frequency token can obtain the semantics from multiple teachers adapted to itself for compensation, and the central low-frequency one can further promote its generation. We concentrate on the representation of low-frequency tokens, which helps the model capture information about relatively rare but vital words in the testing phase. The improvement of low-frequency tokens enables the generation caption to have a refined description.

Our main contributions are summarized threefold:

- We put the emphasis on token imbalance that has not been focused seriously in video captioning and dedicate to concentrating on infrequent but vital tokens to enhance the refined semantics of generated captions.
- We devise a unique module, termed Frequency-Aware Diffusion (FAD), in the form of diffusion to learn tokens with insufficient occurrence that are most likely to carry critical semantics. To the best of our knowledge, this is the first time that the frequency of tokens is attached to great importance in video captioning.
- We further design a Divergent Semantic Supervisor (DSS) module as a constraint for balancing different frequent tokens by leveraging distinctive semantics adapted to the token itself, and quantitatively validating its effectiveness.

Related Works

RNN- and Transformer-Based Methods

With the development of deep learning, researchers extract video features through deep convolution neural networks and combine them with RNNs or Transformer (Hori et al. 2017) to make the generated words more accurate. Early in (Venugopalan et al. 2015), the encoder-decoder framework is applied first to the captioning model, where visual features are leveraged by performing mean-pooling over each frame, and long short-term memory (LSTM) (Hochreiter and Schmidhuber 1997) is adopted to generate sentences. Benefiting from parallel computing ability and scalability of Transformer (Vaswani et al. 2017), it is applied to many multi-modal tasks. Inspired by it, Transformer has been introduced for video captioning (Chen et al. 2018a). SBAT (Jin et al. 2020) is intended to reduce the redundancy of consecutive frames. D² (Gao et al. 2022) utilizes the syntactic prior to measuring the contribution of visual features and syntactic prior information for each word.

Diffusion Model

Some efforts (Li et al. 2022a; Ye et al. 2022b) have been conducted on data noise. In recent years, Diffusion model (Sohl-Dickstein et al. 2015) has made remarkable progress. DDPM (Ho, Jain, and Abbeel 2020) develops the denoising ability by artificially adding noise in the diffusion process and matching each small step of the inverse process corresponding to the forward process. TimeGrad (Rasul et al. 2021) integrates the diffusion model with the autoregressive model by sampling from the data distribution at each time step by estimating its gradient. The diffusion approaches (Liao et al. 2022a; Zhu et al. 2022) have obtained significant progress in computer vision. Inspired by the potential of learning noise by noise addition that is just aiming at learning the noise purely, RSFD regards the information on low-frequency tokens as noise and adds it to high-frequency tokens, developing the model’s ability to comprehend low-frequency tokens.

Translation Methods with Frequency Modeling

In neural machine translation, the token imbalance data drive the model preferentially generates high-frequency words while ignoring the rarely-used words. GNMT (Wu et al. 2016) proposes to divide words into a limited set of common sub-word units to handle the rare words naturally. BMI (Xu et al. 2021) assigns an adaptive weight to promote token-level adaptive training. In video captioning, the large-scale training corpus demonstrates the phenomenon of the imbalance occurrence of tokens. External language model (Zhang et al. 2020) is introduced to extend the ground-truth tokens to deal with the token imbalance problem, which heavily depends on external knowledge. In contrast, RSFD highlight relatively rare words that are most likely refined semantics in the sentence, alleviating token imbalance inside our model.

Proposed Method

We devise our Refined Semantic enhancement towards Frequency Diffusion (RSFD) for video captioning. As shown in

Fig. 2, our overall framework follows the encoder-decoder structure. During the training process, Frequency-Aware Diffusion (FAD) encourages the model to add low-frequency token noise to learn its semantics. Then the diffusion features of tokens are fused with the corresponding visual features according to the cross-attention mechanism. At the head of the decoder, Divergent Semantic Supervisor (DSS) obtains distinctive semantic features by updating the gradient that adapts to the token itself. In the testing phase, only the original Transformer architecture is retained to generate captions.

Encoder-Decoder Framework

Encoder. Image and motion features are processed by two separate encoders, which give a set of consecutive video clips of length K . Here we use $V \in \mathbb{R}^{K \times d_v}$ to represent the features of these two types, where d_v denotes dimension of each feature. We feed them into highway embedding layer (HEL) to obtain representations $R \in \mathbb{R}^{K \times d_h}$, i.e., $R = f_{\text{HEL}}(V)$, where d_h is output dimension of HEL, and f_{HEL} can be trained directly through simple gradient descent (Srivastava, Greff, and Schmidhuber 2015), thus it can be formulated as:

$$f_{\text{HEL}}(V) = \text{BN}(T(V) \circ H(V) + (1 - T(V)) \circ P(V)), \quad (1)$$

where $H(V) = VW_{\text{eh}}$, $P(V) = \tanh(H(V)W_{\text{ep}})$, and $T(V) = \phi(H(V)W_{\text{et}})$. BN denotes batch normalization operation (Ioffe and Szegedy 2015), \circ is the element-wise product, and ϕ is the sigmoid function. $W_{\text{eh}} \in \mathbb{R}^{d_v \times d_h}$, and $\{W_{\text{ep}}, W_{\text{et}}\} \in \mathbb{R}^{d_h \times d_h}$. The multi-modalities, e.g., image and motion features are applied concatenation to obtain $R \in \mathbb{R}^{2K \times d_h}$ to feed into the decoder.

Decoder. The decoder block consists of self-attention, cross-attention, and feed-forward layer. The self-attention layer is formulated as:

$$\begin{aligned} \text{SelfAtt}(E_{<t}) &= \text{MultiHead}(E_{<t}, E_{<t}, E_{<t}) \\ &= \text{Concat}(\text{head}_{1:h})W_h, \end{aligned} \quad (2)$$

where $E_{<t} \in \mathbb{R}^{T \times d_h}$ denotes the feature vector of embedded captions, T represents the length of caption. Concat denotes concatenation operation, and h denotes the number of multiple attention heads. $W_h \in \mathbb{R}^{d_h \times d_h}$ is a trainable parameter. Residual connection and layer normalization are adopted after the self-attention layer:

$$E'_{<t} = \text{LayerNorm}(E_{<t} + \text{SelfAtt}(E_{<t})), \quad (3)$$

where LayerNorm smooths the size relationship between different samples.

So far, the self-attention process has been completed, and the result $E'_{<t}$ is employed for cross-attention as a query:

$$D_{<t} = \text{LayerNorm}\left(E'_{<t} + \text{MultiHead}\left(E'_{<t}, R, R\right)\right), \quad (4)$$

where $D_{<t}$ denotes the sentence feature of cross-attention output. We adopt a feed-forward network (FFN) that employs non-linear transformation, and apply LayerNorm to calculate the probability distribution of words by Softmax:

$$D_t = \text{LayerNorm}(D_{<t} + \text{FFN}(D_{<t})), \quad (5)$$

$$P_t(y_t|y_{<t}, R) = \text{Softmax}(D_t W_p), \quad (6)$$

where R denotes the encoded video features, $W_p \in \mathbb{R}^{d_n \times d_e}$ is a trainable variable. The objective is to minimize the cross-entropy loss:

$$\mathcal{L}_t = - \sum_{t=1}^T \log P_t(y_t^* | y_{<t}^*, R), \quad (7)$$

where y_t^* denotes the ground-truth token at time step t .

Frequency-Aware Diffusion

ORG-TRL (Zhang et al. 2020) introduces an external language model to extend the ground-truth tokens to cope with the token imbalance problem. However, it needs to be trained in large-scale corpus in advance, divorced from the captioning model, and the external model consumes additional computing resources and memory. Instead, we propose an FAD module inside captioning model to sufficiently grasp low-frequency semantics. We first introduce our split method of high-frequency and low-frequency tokens and present a reasonable explanation for it. Then we elaborate on the diffusion process to learn the semantics of low-frequency tokens.

Split of Distinct Frequency Words. Due to the severe imbalance of tokens in the corpus, a small number of high-frequency tokens occupy many occurrences while a large number of low-frequency tokens appear rarely. Here we indicate HFT for high-frequency tokens, LFT for low-frequency ones, and UMT for unmarked tokens. Then we define a token belongs to different categories of frequency as:

$$\begin{cases} \text{HFT, } \frac{|\text{vid}(\text{tok})|}{|\text{vid}_{\text{all}}(\text{tok})|} \geq \gamma, & \frac{|\text{tok}|}{|\text{cap}|} \geq \delta, \\ \text{LFT, } \frac{|\text{vid}(\text{tok})|}{|\text{vid}_{\text{all}}(\text{tok})|} < \gamma, & \end{cases} \quad (8)$$

$$\begin{cases} \text{UMT, } \frac{|\text{vid}(\text{tok})|}{|\text{vid}_{\text{all}}(\text{tok})|} \geq \gamma, & \frac{|\text{tok}|}{|\text{cap}|} < \delta, \\ \text{LFT, } \frac{|\text{vid}(\text{tok})|}{|\text{vid}_{\text{all}}(\text{tok})|} < \gamma, & \end{cases} \quad (9)$$

where δ and γ represent hyper-parameter to decide which category of the frequency the token belongs to. In video captioning, a video has multiple ground-truth captions. $|\text{vid}(\cdot)|$ and $|\text{vid}_{\text{all}}(\cdot)|$ respectively represent the occurrence numbers of the token in the corresponding video and the sum number of tokens in all captions of its video. $|\text{tok}|$ denotes the number of occurrences of this token in all ground-truth captions of the whole corpus, and $|\text{cap}|$ indicates the number of all captions in the dataset.

We use γ to assess the frequency of the token in intra-video, for it only indicates the frequency inside its located video (termed intra-frequency), and δ evaluates the frequency of the token in inter-video for it expresses the frequency in the whole dataset (termed inter-frequency). We split tokens into three kinds of frequency tokens in four cases: When inter-frequency exceeds the δ , if the token has higher intra-frequency, it is divided into HFT owing to its universality; Otherwise, proving that it cannot integrate sufficiently with the corresponding video and be divided into LFT, even if it is a common token; When inter-frequency is below the δ while the intra-frequency is greater than the γ , it shows that even though it is not a frequent token, it can be adequately

comprehended by the video to which it belongs. Thus it is divided into UMT for it can be self-sufficient without disturbing other high-frequency words; Instead, when both of them are below their respective hyper-parameters, the token is obviously classified as LFT, for it requires more trains.

Noising Frequency Diffusion. Our diffusion is an intuitive diffusion exploiting one noise addition step rather than the vanilla Markov chain with multiple adding noise steps. Inspired by denoising after adding noise (Ho, Jain, and Abbeel 2020), we regard low-frequency tokens as noise and add them to high-frequency tokens so that our diffusion is a process from low-frequency to high-frequency. We respectively define $\text{Token}^L = \{\text{token}_1^L, \text{token}_2^L, \dots, \text{token}_m^L\}$, $\text{Token}^H = \{\text{token}_1^H, \text{token}_2^H, \dots, \text{token}_n^H\}$ as low-frequency and high-frequency token features list. A similarity matrix S is constructed to measure the similarity semantics between i -th LFT and j -th HFT:

$$S_{ij} = \frac{\exp s_{ij}}{\sum_{j=1}^n \exp s_{ij}}, \quad (10)$$

where S_{ij} is the (i, j) -th element of $S \in \mathbb{R}^{m \times n}$. We adopt softmax function as the normalization function, s_{ij} denotes the cosine similarity score between two tokens. We choose j -th HFT corresponding to max value in each line as the diffusion object of i -th LFT:

$$\text{LoH} = \arg \max (S, \dim = 1), \quad (11)$$

where LoH denotes the list of high-frequency tokens corresponding to low-frequency tokens one by one, and the value of LoH represents the index of Token^H . In this way, we have established the corresponding relationship between low-frequency tokens and their semantically most similar high-frequency token. Afterward, we obtain the parameter of the noise according to the similarity:

$$\alpha_i = \frac{\exp \sigma(\text{token}_i^L, \text{token}_{\text{loh}}^H)}{\sum_{i \in N} \exp \sigma(\text{token}_i^L, \text{token}_{\text{loh}}^H)}, \quad (12)$$

where loh denotes i -th value of LoH, α_i represents the significance of i -th noise, N denotes the indexes of low-frequency words corresponding to one high-frequency word, σ denotes the cosine similarity of semantic between these two tokens, then we obtain the diffusion $\text{token}_{\text{loh}}^H$ of high-frequency:

$$\text{token}_{\text{loh}}^H = \text{token}_{\text{loh}}^H + \sum_{i \in N} \alpha_i \circ \text{token}_i^L. \quad (13)$$

Note that unlike DDPM (Ho, Jain, and Abbeel 2020) and TimeGrad (Rasul et al. 2021), which learns the noise first and then denoising, our model only learns low-frequency tokens regardless of how to remove them. By adding appropriate low-frequency tokens as noise to high-frequency ones, RSFD can generate relatively rarely used but more refined words.

Divergent Semantic Supervisor

Due to the noise semantics of low-frequency in FAD, the semantics of high-frequency is constrained to some extent. To complement the semantics of the central high-frequency tokens and to further encourage the formation of the central low-frequency tokens, DSS offers the contextual cue of adjacent tokens to the central tokens.

Skip-gram (Mikolov et al. 2013) designs a window centered on the central word and supervises the generation of the central word using the contextual words. Motivated by it, we design a DSS obtaining the gradient that adapts to the token itself to understand the word adequately.

As shown in Fig. 2, we project the hidden features to the new feature space before the last linear layer and then project new features to the space of the dimension of corpus size:

$$D'_t = D_t W_a W_p, \quad (14)$$

$$P'_t(y_t | y_{<t}, R) = \text{Softmax}(D'_t), \quad (15)$$

where D_t is the result of Eq. (5), $W_a \in \mathbb{R}^{d_h \times d_h}$ is a trainable variable, $W_p \in \mathbb{R}^{d_h \times d_c}$ is also a trainable variable shared with Eq. (6). We use D'_{if} and D'_{il} to respectively represent the former and the latter linear layer of the same layer as the original last union linear projection:

$$P'_{\text{if}}(y_t | y_{<t-1}, R) = \text{Softmax}(D'_{\text{if}}), \quad (16)$$

$$P'_{\text{il}}(y_t | y_{<t+1}, R) = \text{Softmax}(D'_{\text{il}}). \quad (17)$$

The divergent loss is calculated as:

$$\mathcal{L}_{\text{div}} = - \sum_{t=1}^T \log P'_{\text{if}}(y_t^* | y_{<t-1}^*, R) - \sum_{t=1}^T \log P'_{\text{il}}(y_t^* | y_{<t+1}^*, R), \quad (18)$$

where two items denote the former and the latter divergent losses, respectively. The final loss is calculated as:

$$\mathcal{L} = \mathcal{L}_t + \lambda \mathcal{L}_{\text{div}}, \quad (19)$$

where λ is a variable parameter to decide the significance of \mathcal{L}_{div} . In the above example, we only use one word before and after to supervise the central word. In the experimental results section, we display the number of words selected as the supervisor to achieve the best generation effect.

Experimental Results

Datasets and Evaluation Metrics

MSR-VTT (Xu et al. 2016) consists of 10,000 video clips, each annotated with 20 English captions and a category. MSVD (Chen and Dolan 2011) is a widely-used benchmark video captioning dataset collected from YouTube, composed of 1,970 video clips and roughly 80,000 English sentences. Following the standard split, we use the same setup as previous works (Pan et al. 2020; Ye et al. 2022a), which takes 6,513 video clips for training, 497 video clips for validation, and 2,990 video clips for testing on MSR-VTT, as well as 1,200, 100, and 670 videos for training, validation, and testing on MSVD. The vocabulary size of MSR-VTT is 10,547, whereas that of MSVD is 9,468.

We present the results on four commonly used metrics to evaluate caption generation quality: BLEU-4 (B-4) (Papineni et al. 2002), METEOR (M) (Banerjee and Lavie 2005), ROUGE.L (R) (Lin 2004), and CIDEr (C) (Vedantam, Zitnick, and Parikh 2015). BLEU-4, METEOR, and ROUGE.L are generally used in machine translation. CIDEr is proposed for captioning tasks specifically and is considered more consistent with human judgment.

Implementation Details

In our experiments, we follow (Pei et al. 2019) to extract image and motion features to encode video information. Specifically, we use ImageNet (Deng et al. 2009) pre-trained ResNet-101 (He et al. 2016) to extract 2D scene features for each frame. We also utilize Kinetics (Kay et al. 2017) pre-trained ResNeXt-101 with 3D convolutions (Hara, Kataoka, and Satoh 2018).

In our implementation, size of video feature d_v and the hidden size d_h are set to 2,048 and 512. Empirically, we set the sampled frames $K = 8$ for each video clip. We set maximum sequence length T to 30 on MSR-VTT, whereas $T = 20$ on MSVD. Transformer decoder has a decoder layer, 8 attention heads, 0.5 dropout ratio, and 0.0005 ℓ_2 weight decay. We implement word embeddings by trainable 512 dimensions of embedding layers. In the training phase, we adopt Adam (Kingma and Ba 2015) with an initial learning rate of 0.005 to optimize our model. The batch size is set to 64, and the training epoch is set to 50. During testing, we use the beam-search method with size 5 to generate the predicted sentences. γ and δ for deciding the category of the token are respectively set to 0.015 and 0.0015. We set $\lambda = 0.07$ on MSR-VTT and $\lambda = 0.4$ on MSVD to demonstrate the significance of the divergent loss. All our experiments are conducted on two NVIDIA Tesla PH402 SKU 200.

Comparison with the State-of-the-Arts

The comparison results with the state-of-the-arts in Table 1 illustrate that RSFD achieves a significant performance improvement. On MSR-VTT, RSFD achieves the best results under 3 out of 4 metrics. Specifically, RSFD outperforms existing methods except for ORG-TRL (Zhang et al. 2020) under all the metrics on MSR-VTT, proving that our focus on solving the insufficient occurrence of low-frequency words is practical. In particular, RSFD achieves 53.1% and 96.7% under CIDEr on MSR-VTT and MSVD, respectively, making an improvement of 2.0% and 5.3% over AR-B (Yang et al. 2021). As CIDEr captures human judgment of consensus better than other metrics, superior performance under CIDEr indicates that RSFD can generate semantically more similar to human understanding. The boost in performance demonstrates the advantages of RSFD, which exploits FAD and DSS, and models the low-frequency token for detailed and comprehensive textual representations.

RSFD is slightly lower than ORG-TRL under BLEU-4 on MSR-VTT, which uses an external language model to integrate linguistic knowledge. But it is undeniable that RSFD is second only to it, demonstrating the superiority of RSFD. OABTG (Zhang and Peng 2019), STG-KD (Pan et al. 2020), and ORG-TRL perform better under METEOR and ROUGE.L on MSVD. They either capture the temporal interaction of objects well in a small dataset or bring complementary representation by considering video semantics. Despite not exploring the above information, RSFD focuses on generating refined descriptions and outperforms other methods under CIDEr, benefiting from its emphasis on refined low-frequency tokens. Indicated that CIDEr best matches human consensus judgment.

Method	Venue	MSR-VTT				MSVD			
		B-4	M	R	C	B-4	M	R	C
M3-VC (Wang et al. 2018b)	CVPR '18	38.1	26.6	-	-	51.8	32.5	-	-
RecNet (Wang et al. 2018a)	CVPR '18	39.1	26.6	59.3	42.7	52.3	34.1	69.8	80.3
PickNet (Chen et al. 2018b)	ECCV '18	41.3	27.7	59.8	44.1	52.3	33.3	69.6	76.5
OA-BTG (Zhang and Peng 2019)	CVPR '19	41.4	28.2	-	46.9	56.9	36.2	-	90.6
MARN (Pei et al. 2019)	CVPR '19	40.4	28.1	60.7	47.1	48.6	35.1	71.9	92.2
MGSA (Chen and Jiang 2019)	AAAI '19	42.4	27.6	-	47.5	53.4	35.0	-	86.7
GRU-EVE (Aafaq et al. 2019)	CVPR '19	38.3	28.4	60.7	48.1	47.9	35.0	71.5	78.1
POS-CG (Wang et al. 2019)	ICCV '19	42.0	28.2	61.6	48.7	52.5	34.1	71.3	88.7
STG-KD (Pan et al. 2020)	CVPR '20	40.5	28.3	60.9	47.1	52.2	36.9	73.9	93.0
SAAT (Zheng, Wang, and Tao 2020)	CVPR '20	40.5	28.2	60.9	49.1	46.5	33.5	69.4	81.0
ORG-TRL (Zhang et al. 2020)	CVPR '20	43.6	28.8	62.1	50.9	54.3	36.4	73.9	95.2
SBAT (Jin et al. 2020)	IJCAI '20	42.9	28.9	61.5	51.6	53.1	35.3	72.3	89.5
TTA (Tu et al. 2021)	PR '21	41.4	27.7	61.1	46.7	51.8	35.5	72.4	87.7
SibNet (Liu, Ren, and Yuan 2021)	TPAMI '21	41.2	27.8	60.8	48.6	55.7	35.5	72.6	88.8
SGN (Ryu et al. 2021)	AAAI '21	40.8	28.3	60.8	49.5	52.8	35.5	72.9	94.3
AR-B (Yang et al. 2021) (Baseline) †	AAAI '21	42.1	29.1	61.2	51.1	49.2	35.3	72.1	91.4
FrameSel (Li et al. 2022b)	TCSVT '22	38.4	27.2	59.7	44.1	50.4	34.2	70.4	73.7
TVRD (Wu et al. 2022)	TCSVT '22	43.0	28.7	62.2	51.8	50.5	34.5	71.7	84.3
RSFD (Ours)		43.4	29.3	62.3	53.1	51.2	35.7	72.9	96.7

Table 1: Performance (%) comparison with the state-of-the-arts on MSR-VTT and MSVD. † indicates the reproduced method. The best results are shown in bold.

Dataset	size	B-4	M	R	C
MSR-VTT	3	42.5	29.2	61.6	52.2
	5	43.4	29.3	62.3	53.1
	7	42.9	29.2	61.5	51.9
	9	42.7	29.1	61.4	51.6
MSVD	3	50.9	35.3	72.5	94.1
	5	51.2	35.7	72.9	96.7
	7	50.6	35.4	72.6	93.8
	9	50.7	35.3	72.4	93.0

Table 2: Performance (%) comparison of different configurations of window size in DSS module on MSR-VTT and MSVD. The best results are shown in bold.

Dataset	Method	B-4	M	R	C
MSR-VTT	AR-B (Baseline)	42.1	29.1	61.2	51.1
	RSFD w FAD	43.0	29.2	62.0	52.9
	RSFD w DSS	42.5	29.2	61.6	52.2
	RSFD (Ours)	43.4	29.3	62.3	53.1
MSVD	AR-B (Baseline)	49.2	35.3	72.1	91.4
	RSFD w FAD	50.9	35.5	72.7	94.1
	RSFD w DSS	50.6	35.5	72.3	93.1
	RSFD (Ours)	51.2	35.7	72.9	96.7

Table 3: Performance (%) comparison of different components in RSFD on MSR-VTT and MSVD. The best results are shown in bold.

Study on Trade-off Parameter

Trade-off Parameter of Supervision Window. To explore the impact of the number of words, we design a supervision window of different sizes to extract adjacent supervision information shown in Table 2. It shows that when the window size increases, the overall performance shows an upward trend, but when the size exceeds 5, the performance begins to decline. When the supervision window is too small, insufficient adjacent information cannot play a supervisory role. And in an oversized supervision window, not only low-frequency tokens, all tokens receive unique gradient updates, causing overfitting of other tokens that have been trained enough.

Trade-off Hyper-Parameter λ . To evaluate the effectiveness of λ and find an appropriate value for λ , we adjust the value of Eq. (19) under all metrics on MSR-VTT are given in Fig. 3(a). Note that only using original loss ($\lambda = 0$) shows the worst performances. Thus it can be concluded again that

adding divergent loss boosts the performance of the captioning model, which benefited from gradient correction of adjacent words to central words. Considering all the metrics comprehensively, we empirically set $\lambda = 0.07$ on MSR-VTT.

Trade-off Hyper-Parameter γ and δ . We conduct experiments on MSR-VTT to explore the effect of hyper-parameter γ for the intra-frequency and hyper-parameter δ for the inter-frequency. All metrics scores of different settings are illustrated in Figs. 3(b) and 3(c). If δ is too high, the reduction of the number of high-frequency tokens makes the information of low-frequency tokens cannot be well absorbed and fully integrated with visual features. Similarly, if γ is too high, some words that should not be split into low-frequency tokens are divided into them, causing the overfitting of these words. In contrast, if γ is set too low, low-frequency tokens that need to be sufficiently trained are ignored. Considering the above discussion, we set γ and δ as 0.015 and 0.0015, respectively.




Figure 3: Analysis of (a) different parameter λ for the significance of divergent semantic supervisor, (b) different parameter γ deciding intra-frequency, and (c) different parameter δ determining inter-frequency on MSR-VTT. We show the relative results on all metrics.

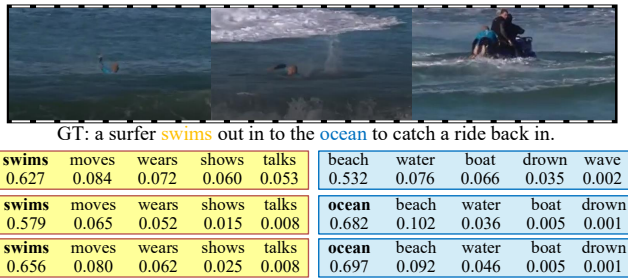


Figure 4: Instances of AR-B (baseline), RSFD with FAD, and RSFD model in the testing phase. The yellow and blue words indicate the high-frequency and low-frequency words to be predicted. The yellow box on the left corresponds to the three methods predicting the yellow word, while the blue box on the right generates the blue word.

Ablation Study

Effectiveness of Different Components. The results are presented in Table 3. The first line method AR-B (Yang et al. 2021) (baseline) only applies a Transformer-based model to the encoder-decoder framework. The methods in the second and third rows refer to methods with FAD and DSS, respectively. It can be observed that FAD, through sufficiently exploiting the semantics of low-frequency tokens, outperforms AR-B under popular metrics, particularly bringing significant CIDEr improvements. Although FAD slightly affects the semantics of high-frequency tokens, we integrate it with DSS, which makes up the information of high-frequency tokens to some extent. Our complete method RSFD results demonstrate that combining the above two modules can achieve the best performance.

Evaluation of FAD and DSS. We analyze the effect of each component in the testing phase shown in Fig. 4. It can be seen that FAD generates “ocean” shown in the second line, which as an infrequent token, carries critical semantics but is challenging to predict by AR-B (Yang et al. 2021). FAD helps the model comprehend low-frequency tokens relative to video content. Although FAD slightly narrows the high-frequency token “swims” probability from 0.627 to 0.579, DSS compensates for it to 0.656 and further enhances the semantics of the low-frequency token “ocean” from 0.682 to 0.697. In general, our method promotes the generation of refined captions owing to our focus on low-frequency tokens

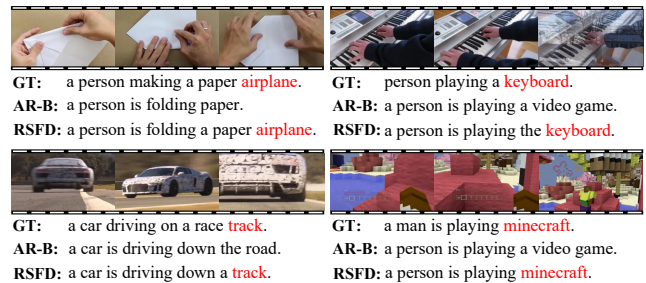


Figure 5: Qualitative results generated on the testing sets of MSR-VTT with AR-B (baseline) and RSFD. The red words indicate the refined information in the ground truth as well as the generated sentences of RSFD.

and effectively alleviates the long-tailed problem in video captioning, demonstrating a competitive generation effect.

Qualitative Results

We present some results generated on MSR-VTT in Fig. 5. It can be observed that the content of captions generated by RSFD is more refined than AR-B (Yang et al. 2021) (baseline). Taking the first result as an example, AR-B only understands the rough meaning of the video, *e.g.*, “folding”, missing the abundant description, *e.g.*, “airplane”. By contrast, our model captures more details and generates them accurately, demonstrating that RSFD produces relatively less common but critical words. The rest of the results have similar characteristics.

Conclusion

In this paper, we present a Refined Semantic enhancement towards Frequency Diffusion (RSFD) for video captioning. RSFD addresses the problem of the Transformer-based or RNN-based architecture on the insufficient occurrence of low-frequency tokens that limits the capability of comprehensive semantic generation. Our method incorporates the Frequency-Aware Diffusion (FAD) and Divergent Semantic Supervisor (DSS), which capture more detailed semantics of relatively less common but critical tokens to provide refined semantic enhancement with the help of low-frequency tokens. RSFD achieves competitive performances by outperforming baselines by large margins of 2.0% and 5.3% in terms of CIDEr on MSR-VTT and MSVD, respectively.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants 62271361 and 62066021, the Department of Science and Technology, Hubei Provincial People's Government under Grant 2021CFB513, and the Hubei Province Educational Science Planning Special Funding Project under Grant 2022ZA41.

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 *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 12487–12496.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proc. ACL Annu. Meeting Assoc. Comput. Linguistics Workshop*, 65–72.
- Chen, D. L.; and Dolan, W. B. 2011. Collecting Highly Parallel Data for Paraphrase Evaluation. In *Proc. ACL Annu. Meeting Assoc. Comput. Linguistics*, 190–200.
- Chen, M.; Li, Y.; Zhang, Z.; and Huang, S. 2018a. TVT: Two-View Transformer Network for Video Captioning. In *Proc. PMLR Asian Conf. Mach. Learn.*, 847–862.
- Chen, S.; and Jiang, Y. 2019. Motion Guided Spatial Attention for Video Captioning. In *Proc. AAAI Conf. Artif. Intell.*, 8191–8198.
- Chen, Y.; Wang, S.; Zhang, W.; and Huang, Q. 2018b. Less Is More: Picking Informative Frames for Video Captioning. In *Proc. Springer Eur. Conf. Comput. Vis.*, 367–384.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Li, F. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 248–255.
- Gao, Y.; Hou, X.; Suo, W.; Sun, M.; Ge, T.; Jiang, Y.; and Wang, P. 2022. Dual-Level Decoupled Transformer for Video Captioning. In *Proc. ACM Int. Conf. Multimedia Retr.*, 219–228.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet? In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 6546–6555.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 770–778.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Adv. Neural Inf. Process. Syst.*
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 9(8): 1735–1780.
- Hori, C.; Hori, T.; Lee, T.; Zhang, Z.; Harsham, B.; Hershey, J. R.; Marks, T. K.; and Sumi, K. 2017. Attention-Based Multimodal Fusion for Video Description. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 4203–4212.
- Hu, M.; Jiang, K.; Liao, L.; Nie, Z.; Xiao, J.; and Wang, Z. 2022a. Progressive Spatial-temporal Collaborative Network for Video Frame Interpolation. In *Proc. ACM Int. Conf. Multimedia*, 2145–2153.
- Hu, M.; Jiang, K.; Liao, L.; Xiao, J.; Jiang, J.; and Wang, Z. 2022b. Spatial-Temporal Space Hand-in-Hand: Spatial-Temporal Video Super-Resolution via Cycle-Projected Mutual Learning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 3564–3573.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proc. JMLR Int. Conf. Mach. Learn.*, 448–456.
- Jia, X.; Zhong, X.; Ye, M.; Liu, W.; and Huang, W. 2022. Complementary Data Augmentation for Cloth-Changing Person Re-Identification. *IEEE Trans. Image Process.*, 31: 4227–4239.
- Jin, T.; Huang, S.; Chen, M.; Li, Y.; and Zhang, Z. 2020. SBAT: Video Captioning with Sparse Boundary-Aware Transformer. In *Proc. IJCAI Int. Joint Conf. Artif. Intell.*, 630–636.
- 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:1705.06950*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *Proc. Int. Conf. Learn. Represent. Workshop*.
- Li, J.; Zha, S.; Chen, C.; Ding, M.; Zhang, T.; and Yu, H. 2022a. Attention Guided Global Enhancement and Local Refinement Network for Semantic Segmentation. *IEEE Trans. Image Process.*, 31: 3211–3223.
- Li, L.; Zhang, Y.; Tang, S.; Xie, L.; Li, X.; and Tian, Q. 2022b. Adaptive Spatial Location with Balanced Loss for Video Captioning. *IEEE Trans. Circuits Syst. Video Technol.*, 17–30.
- Li, Z.; Zhong, X.; Chen, S.; Liu, W.; Huang, W.; and Li, L. 2022c. Background Disturbance Mitigation for Video Captioning via Entity-Action Relocation. In *Proc. IEEE Int. Conf. Acoustics Speech Signal Process.*
- Liao, L.; Chen, W.; Xiao, J.; Wang, Z.; Lin, C.; and Satoh, S. 2022a. Unsupervised Foggy Scene Understanding via Self Spatial-Temporal Label Diffusion. *IEEE Trans. Image Process.*, 31: 3525–3540.
- Liao, L.; Xu, K.; Wu, H.; Chen, C.; Sun, W.; Yan, Q.; and Lin, W. 2022b. Exploring the Effectiveness of Video Perceptual Representation in Blind Video Quality Assessment. In *Proc. ACM Int. Conf. Multimedia*, 837–846.
- Lin, C. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Proc. ACL Annu. Meeting Assoc. Comput. Linguistics Workshop*, 74–81.
- Liu, S.; Ren, Z.; and Yuan, J. 2021. SibNet: Sibling Convolutional Encoder for Video Captioning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 3259–3272.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. In *Proc. Int. Conf. Learn. Represent. Workshop*.
- Pan, B.; Cai, H.; Huang, D.; Lee, K.; Gaidon, A.; Adeli, E.; and Niebles, J. C. 2020. Spatio-Temporal Graph for Video

- Captioning with Knowledge Distillation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 10867–10876.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proc. ACL Annu. Meeting Assoc. Comput. Linguistics*, 311–318.
- Pei, W.; Zhang, J.; Wang, X.; Ke, L.; Shen, X.; and Tai, Y. 2019. Memory-Attended Recurrent Network for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 8347–8356.
- Rasul, K.; Seward, C.; Schuster, I.; and Vollgraf, R. 2021. Autoregressive Denoising Diffusion Models for Multivariate Probabilistic Time Series Forecasting. In *Proc. JMLR Int. Conf. Mach. Learn.*, 8857–8868.
- Ryu, H.; Kang, S.; Kang, H.; and Yoo, C. D. 2021. Semantic Grouping Network for Video Captioning. In *Proc. AAAI Conf. Artif. Intell.*, 2514–2522.
- Sohl-Dickstein, J.; Weiss, E. A.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *Proc. JMLR Int. Conf. Mach. Learn.*, 2256–2265.
- Srivastava, R. K.; Greff, K.; and Schmidhuber, J. 2015. Training Very Deep Networks. In *Adv. Neural Inf. Process. Syst.*, 2377–2385.
- Tu, Y.; Zhou, C.; Guo, J.; Gao, S.; and Yu, Z. 2021. Enhancing the Alignment between Target Words and Corresponding Frames for Video Captioning. *Pattern Recognit.*, 107702.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. In *Adv. Neural Inf. Process. Syst.*, 5998–6008.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDEr: Consensus-Based Image Description Evaluation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 4566–4575.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R. J.; and Saenko, K. 2015. Translating Videos to Natural Language Using Deep Recurrent Neural Networks. In *Proc. ACL Conf. North Am. Chapter Assoc. Comput. Linguistics*, 1494–1504.
- Wang, B.; Ma, L.; Zhang, W.; Jiang, W.; Wang, J.; and Liu, W. 2019. Controllable Video Captioning with POS Sequence Guidance Based on Gated Fusion Network. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2641–2650.
- Wang, B.; Ma, L.; Zhang, W.; and Liu, W. 2018a. Reconstruction Network for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 7622–7631.
- Wang, J.; Wang, W.; Huang, Y.; Wang, L.; and Tan, T. 2018b. M3: Multimodal Memory Modelling for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 7512–7520.
- Wu, B.; Niu, G.; Yu, J.; Xiao, X.; Zhang, J.; and Wu, H. 2022. Towards Knowledge-Aware Video Captioning via Transitive Visual Relationship Detection. *IEEE Trans. Circuits Syst. Video Technol.*, 32: 6753–6765.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; Klingner, J.; Shah, A.; Johnson, M.; Liu, X.; Kaiser, L.; Gouws, S.; Kato, Y.; Kudo, T.; Kazawa, H.; Stevens, K.; Kurian, G.; Patil, N.; Wang, W.; Young, C.; Smith, J.; Riesa, J.; Rudnick, A.; Vinyals, O.; Corrado, G.; Hughes, M.; and Dean, J. 2016. Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *arXiv:1609.08144*.
- Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 5288–5296.
- Xu, X.; Liu, W.; Wang, Z.; Hu, R.; and Tian, Q. 2022. Towards Generalizable Person Re-Identification with a Bi-Stream Generative Model. *Pattern Recognit.*, 132: 108954.
- Xu, Y.; Liu, Y.; Meng, F.; Zhang, J.; Xu, J.; and Zhou, J. 2021. Bilingual Mutual Information Based Adaptive Training for Neural Machine Translation. In *Proc. ACL Annu. Meeting Assoc. Comput. Linguistics*, 511–516.
- Yang, B.; Zou, Y.; Liu, F.; and Zhang, C. 2021. Non-Autoregressive Coarse-to-Fine Video Captioning. In *Proc. AAAI Conf. Artif. Intell.*, 3119–3127.
- Ye, H.; Li, G.; Qi, Y.; Wang, S.; Huang, Q.; and Yang, M.-H. 2022a. Hierarchical Modular Network for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 17939–17948.
- Ye, M.; Li, H.; Du, B.; Shen, J.; Shao, L.; and Hoi, S. C. H. 2022b. Collaborative Refining for Person Re-Identification with Label Noise. *IEEE Trans. Image Process.*, 31: 379–391.
- Zhang, J.; and Peng, Y. 2019. Object-Aware Aggregation with Bidirectional Temporal Graph for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 8327–8336.
- Zhang, Z.; Shi, Y.; Yuan, C.; Li, B.; Wang, P.; Hu, W.; and Zha, Z. 2020. Object Relational Graph with Teacher-Recommended Learning for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 13275–13285.
- Zheng, Q.; Wang, C.; and Tao, D. 2020. Syntax-Aware Action Targeting for Video Captioning. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 13093–13102.
- Zhu, H.; Yuan, J.; Yang, Z.; Zhong, X.; and Wang, Z. 2022. Fine-Grained Fragment Diffusion for Cross Domain Crowd Counting. In *Proc. ACM Int. Conf. Multimedia*, 5659–5668.