Less Is More: Picking Informative Frames for Video Captioning

Yangyu Chen[1], Shuhui Wang[2], Weigang Zhang[3] and Qingming Huang[1]

[1]University of Chinese Academy of Science, Beijing 100049, China

[2]Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China

[3]Harbin Institute of Technology, Weihai 264209, China

{yangyu.chen, shuhui.wang, weigang.zhang, qingming.huang}@vipl.ict.ac.cn

Abstract

In video captioning task, the best practice has been achieved by attention-based models which associate salient visual components with sentences in the video. However, existing study follows a common procedure which includes a frame-level appearance modeling and motion modeling on equal interval frame sampling, which may bring about redundant visual information, sensitivity to content noise and unnecessary computation cost.

We propose a plug-and-play PickNet to perform informative frame picking in video captioning. Based on a standard Encoder-Decoder framework, we develop a reinforcement-learning-based procedure to train the network sequentially, where the reward of each frame picking action is designed by maximizing visual diversity and minimizing textual discrepancy. If the candidate is rewarded, it will be selected and the corresponding latent representation of Encoder-Decoder will be updated for future trials. This procedure goes on until the end of the video sequence. Consequently, a compact frame subset can be selected to represent the visual information and perform video captioning without performance degradation. Experiment results shows that our model can use 6∼8 frames to achieve competitive performance across popular benchmarks.

1. Introduction

Human are born with the ability to identify useful information and filter redundant information. In biology, this mechanism is called sensory gating [8], which describes neurological processes of filtering out unnecessary stimuli in the brain from all possible environmental stimuli, thus prevents an overload of redundant information in the higher cortical centers of the brain. This cognitive mechanism is essentially consistent with a huge body of researches in computer vision.

As one of the strong evidences practicing on visual sensory gating, attention is introduced to identify the salient visual regions with high objectness and meaningful visual

(a) Equally sampled 30 frames from a video

(b) Informative frames

Figure 1: An illustration of the temporal redundancy in video. Video may contains many redundant information. Using only a portion of frames can clearly convey information.

patterns of an image [22, 45]. The attention has also been established on videos that contains consecutive image frames. Existing study follows a common procedure which includes a frame-level appearance modeling and motion modeling on equal interval frame sampling, say, every 3 frames or 5 frames. Visual features and motion features are extracted on the selected frame subset one by one, and they are all fed into the learning stage. Similar to image, the video attention is recognized as a spatial-temporal saliency that identifies both salient objects and their motion trajectories [28]. The video attention is also recognized as the word-frame association learned by sparse coding [39] or gaze-guided attention learning [43], which is a de-facto frame weighting mechanism. The visual attention mechanism also benefits many downstream tasks such visual captioning and visual question answering for image and video [38, 21, 41, 14].

Despite the success on bridging vision and language achieved by existing attention-based methods, there still exists critical issues to be addressed as follows.

- **Frame selection perspective.** As shown in Figure 1(a), there are many frames with duplicated and redundant visual appearance information selected with equal interval frame sampling. This will also involve remarkable computation expenditures and less performance gain as the information from the input is not appropriately sampled. For example, it takes millions of floating point calculation to extract a frame-level visual feature for a moderate-sized CNN model. Moreover, there is no guarantee that all the frames selected by equal interval sampling contain meaningful information, so it tends to be more sensitive to content noise such as motion blur, occlusion and object zoom-out.

- **Downstream video captioning task perspective.** Previous attention-based models mostly identify the spatial layout of visual saliency, but the temporal redundancy existing in neighboring frames remains unsolved as all the frames are taken into consideration. This may lead to an unexpected information overload on the visual-linguistic correlation analysis model. For example, the dense-captioning-based strategy [18, 15, 28] can potentially describe images/videos in finer levels of detail by captioning many visual regions within an image/video-clip. With an increasing number of frames, many highly similar visual regions will be generated and the problem will become prohibitive as the search space of sequence-to-sequence association becomes extremely large.

We answer the follow question: *Is there a way to use as less number of frames as possible to well approximate the performance using all the frames for video captioning?* We propose PickNet to perform informative frame picking for video captioning. Specifically, the base model for visual-linguistic association in video captioning is a standard Encoder-Decoder framework [3]. We develop a reinforcement-learning-based procedure to train the network sequentially, where the reward of each frame picking action is designed by considering both visual and textual cues. From visual perspective, we maximize the diversity between current picked frame candidate and the selected frames. From textual perspective, we minimize the discrepancy between ground truth caption and the generated one using current picked candidate. If the candidate is rewarded, it will be selected and the corresponding latent representation of Encoder-Decoder will be updated for future trials. This procedure goes on until the end of the video sequence. Consequently, a compact frame subset can be selected to represent the visual information and perform video captioning without performance degradation.

To the best of our knowledge, this is the first study on frame selection for video captioning. In fact, our framework can go beyond the Encoder-Decoder framework in

video captioning task, and serves as a complementary building block for other state-of-the-art solutions. It can also be adapted by other task-specific objectives for video analysis. In summary, the merits of our PickNet include:

- **Flexibility.** We design a plug-and-play reinforcement-learning-based PickNet to select informative frames which can pick informative frames for the next learning stage. A compact frame subset can be selected to represent the visual information and perform video captioning without performance degradation.

- **Efficiency.** The architecture can largely cut down the usage of convolution operations. It makes our method more applicable for real-world video processing.

- **Effectiveness.** Experiment shows that our model can use a small number of frames to achieve comparable or even better performance compared to state-of-the-art.

2. Related Works

2.1. Visual captioning

The visual captioning is the task that translating visual contents into natural language. Early to 2002, Kojima *et al.* [17] proposed the first video captioning system for describing human behavior. From then on, a series of image and video captioning studies have been conducted. Early approaches tackle this problem using bottom-up paradigm [11, 19, 37, 10], which first generate descriptive words of an image by attribute learning and object recognition, then combine them by language models which fit predicted words to predefined sentence templates. With the development of neural networks and deep learning, modern captioning systems are based on CNN and RNN, with the Encoder-Decoder architecture.

An active branch of captioning is utilizing the attention mechanism to weight the input features. For image captioning, the mechanism is typically in the form of spatial attention. Xu *et al.* [36] first introduced an attention based model that automatically learn to fix its gaze on salient objects while generating the corresponding words in the output sequence. For video captioning, the temporal attention is added. Yao *et al.* [39] took into account both the local and global temporal structure of videos to produce descriptions, and their model learned to automatically select the most relevant temporal segments given the text-generating RNN. However, the attention based methods, especially temporal attention, are operated on full observed condition, which is not suitable in some real world applications, such as blind navigation. Our method do not require the global information of videos, which is more effective in these applications.

2.2. Frame selection

The main battle of studying how to select informative frames of videos is in the video summarization field. This problem may be formulated as image searching. For example, Song *et al.* [31] considered images related to the video title that can serve as a proxy for important visual concepts, so they developed a co-archetypal analysis technique that learns canonical visual concepts shared between video and images, and used it to summarize videos. Other people use sparse learning to deal with this problem. Zhao *et al.* [44] proposed to learn a dictionary from given video using group sparse coding, and the summary video was then generated by combining segments that cannot be sparsely reconstructed using the learned dictionary.

Some video analysis task cooperates with frame selection mechanism. For example, in action detection, Yeung *et al.* [40] designed a policy network to directly predict the temporal bounds of actions, which decreased the need for processing the whole video, and improved the detection performance. However, the prediction made by this method is in the form of normalized global position, which requires the knowledge of the video length, and this makes it unable to deal with real video streams. Different from the above methods, our model select frames based on both semantic and visual information, and do not need to know the global length of videos.

3. Method

Our method can be viewed as the combination of two parts: the Encoder-Decoder based sentence generator and the PickNet.

3.1. Preliminary

Like most of video captioning methods, our model is built on the Encoder-Decoder based sentence generator. In this subsection, we briefly introduce this building block.


Figure 2: The encode-decode procedure for video captioning.

Encoder. Given an input video, we use a recurrent video encoder which takes a sequence of visual features $(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ as input and outputs a fixed size vector $\mathbf{v}$ as the representation of this video. The encoder is built on top of a Long Short-Term Memory (LSTM) [13] unit, which has been widely used for video encoding, since it is

known to properly deal with long range temporal dependencies. Different from vanilla recurrent neural network unit, LSTM introduces a memory cell $\mathbf{c}$ which maintains the history of the inputs observed up to a time-step. Specifically, we use the following equations:

$$\mathbf{i}_t = \sigma(W_{ix}\mathbf{x}_t + W_{ih}\mathbf{h}_{t-1} + \mathbf{b}_i) \qquad (1)$$

$$\mathbf{f}_t = \sigma(W_{fx}\mathbf{x}_t + W_{fh}\mathbf{h}_{t-1} + \mathbf{b}_f) \qquad (2)$$

$$\tilde{\mathbf{c}}_t = \phi(W_{gx}\mathbf{x}_t + W_{gh}\mathbf{h}_{t-1} + \mathbf{b}_g) \qquad (3)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \qquad (4)$$

$$\mathbf{o}_t = \phi(W_{ox}\mathbf{x}_t + W_{oh}\mathbf{h}_{t-1} + \mathbf{b}_o) \qquad (5)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \phi(\mathbf{c}_t), \qquad (6)$$

where $\odot$ denotes the element-wise Hadamard product, σ is the sigmoid function, ϕ is the hyperbolic tangent $\mathtt{tanh}$, W_* are learned weight matrices and $\mathbf{b}_*$ are learned biases vectors. The hidden state $\mathbf{h}$ and memory cell $\mathbf{c}$ are initialized to zero. And the last hidden state $\mathbf{h}_T$ is used as the final encoded video representation $\mathbf{v}$.

Decoder and sentence generation. Once the representation of the video has been generated, the recurrent decoder can employ it to generate the corresponding description. At every time-step of the decoding phase, the decoder unit uses the encoded vector $\mathbf{v}$, previous generated one-hot representation word $\mathbf{w}_{t-1}$ and previous internal state $\mathbf{p}_{t-1}$ as input, and outputs a new internal state $\mathbf{p_t}$. Like [3], our decoder unit is the Gated Recurrent Unit (GRU) [7], a simplified version of LSTM, which is good at language decoding. The output of GRU is modulated via two sigmoid gates: a reset gate $\mathbf{r}_t$ and an update gate $\mathbf{z}_t$. The operation detail is as the following:

$$\mathbf{z}_t = \sigma(W_{zw}W_w\mathbf{w}_{t-1} + W_{zv}\mathbf{v} + W_{zp}\mathbf{p}_{t-1} + \mathbf{b}_z) \quad (7)$$

$$\mathbf{r}_t = \sigma(W_{rw}W_w\mathbf{w}_{t-1} + W_{rv}\mathbf{v} + W_{rp}\mathbf{p}_{t-1} + \mathbf{b}_r). \quad (8)$$

Exploiting the values of the above gates, the output of the decoder at timestep t is computed as:

$$\tilde{\mathbf{p}}_t = \phi(W_{pw}W_w\mathbf{w}_{t-1} + W_{pv}\mathbf{v} + W_{pp}(\mathbf{r}_t \odot \mathbf{p}_{t-1}) + \mathbf{b}_p) \qquad (9)$$

$$\mathbf{p}_t = (1 - \mathbf{z}_t) \odot \mathbf{p}_{t-1} + \mathbf{z}_t \odot \tilde{\mathbf{p}}_t, \qquad (10)$$

where W_* and $\mathbf{b}_*$ are learned weights and biases and W_w transforms the one-hot encoding of words to a dense lower dimensional embedding. Again, $\odot$ denotes the element-wise product, σ is the sigmoid function and ϕ is the hyperbolic tangent. A softmax function is applied on $\mathbf{p}_t$ to compute the probability of producing certain word at current time-step:

$$p_\omega(\mathbf{w}_t|\mathbf{w}_{t-1}, \mathbf{w}_{t-2}, ..., \mathbf{w}_1, \mathbf{v}) = \mathbf{w}_t^T \mathrm{softmax}(W_p\mathbf{p}_t), \qquad (11)$$

where W_p is used to project the output of the decoder to the dictionary space and ω denotes all parameters of the

Encoder-Decoder. Also, the internal state $\mathbf{p}$ is initialized to zero. We use the greedy decode routine to generate every word. It means that at every time-step, we choose the word that has the maximal $p_\omega(\mathbf{w}_t | \mathbf{w}_{t-1}, \mathbf{w}_{t-2}, ..., \mathbf{w}_1, \mathbf{v})$ as the current output word. Specifically, we use a special token <BOS> as $\mathbf{w}_0$ to start the decoding, and when the decoder generates another special token <EOS>, the decoding procedure is terminated.

3.2. Our approach

3.2.1 Architecture

The PickNet aims to select informative video content without knowing the global information. It means that the pick decision can only be based on the current observation and the history, which makes it more difficult than video summarization tasks. The more challenging thing is, we do not have supervised information to guide the learning of PickNet in video captioning tasks. Therefore, we formulate the problem as a reinforcement learning task, *i.e.*, given an input image sequence sampled from a video, the agent should select a subset of them under certain policy to retain video content as much as possible. Here, we use PickNet to produce the picking policy. Figure 4 shows the architecture of PickNet.


Figure 4: The PickNet uses the flattened difference grayscale image as input and produces a binomial distribution to indicate picking the current frame or not.

For the consideration of computation efficiency, we use a simple two-layer feedforward neural network as the prototype of PickNet. The network has two outputs, which indicate the probabilities to pick or drop the current observed frame. We model the frame picking process as the glance-and-compare fashion. For each input frame $\mathbf{z}_t$, we first convert the colored image into grayscale image, and then resize it into a smaller image $\mathbf{g}_t$, which can be viewed as a "glance" of current frame. Then we subtract the current glance $\mathbf{g}_t$ by the glance of the last picked frame $\tilde{\mathbf{g}}$, to get a grayscale difference image $\mathbf{d}_t$; this can be seen as the "compare". Finally we flatten the 2D grayscale difference image into a 1D fixed size vector, and feed it to PickNet to produce a binomial distribution that the pick decision is

sampled from:

$$\mathbf{s}_t = W_2 \cdot (\max(W_1 \cdot \text{vec}(\mathbf{d}_t) + \mathbf{b}_1, \mathbf{0})) + \mathbf{b}_2 \quad (12)$$

$$p_\theta(a_t | \mathbf{z}_t, \tilde{\mathbf{g}}) \sim \text{softmax}(\mathbf{s}_t), \quad (13)$$

where W_* are learned weight matrices and $\mathbf{b}_*$ are learned biases vectors. During training, we use stochastic policy, i.e., the action is sampled according to Equation (13). When testing, the policy becomes determined, hence the action with higher probability is chosen. If the policy decides to pick the current frame, the frame feature will be extracted by a pretrained CNN and embedded into a lower dimension, then passed to the encoder unit, and the template will be updated:

$$\tilde{\mathbf{g}} \leftarrow \mathbf{g}_t. \quad (14)$$

We force PickNet to pick the first frame, thus the encoder will always process at least one frame, which makes the training procedure more robust. Figure 3 shows how PickNet works with the encoder. It is worth noting that the input of PickNet can be of any other forms, such as the difference between optical flow maps, which may handle the motion information more properly.

3.2.2 Rewards

The design of rewards is very essential to reinforcement learning. For the purpose of picking informative video frames, we consider two parts of reward: the language reward and visual diversity reward.

Language reward. First of all, the picked frames should contain rich semantic information, which can be used to effectively generate language description. In the video captioning task, it is natural to use the evaluated language metrics as the language reward. Here, we choose CIDEr [32] score. Given a video v_i and a collection of human generated reference sentences $S_i = \{s_{ij}\}$, the goal of CIDEr is to measure the similarity of the machine generated sentence c_i to a majority of how most people describe the video. So the language reward r_l is defined as:

$$r_l(v_i, S_i) = \text{CIDEr}(c_i, S_i) \quad (15)$$

Visual diversity reward. Also, we want the picked frames that have good diversity in visual features. Using only language reward may miss some important visual information, so we introduce the visual diversity reward r_v. For all the selected frame features $\{\mathbf{x}_i \in \mathbb{R}^D\}$, we use the standard deviation of feature vectors as the visual diversity reward:

$$r_v(v_i) = \sum_{j=1}^{D} \sqrt{\frac{1}{N_p} \sum_{i=1}^{N_p} (\mathbf{x}_i^{(j)} - \mu^{(j)})^2}, \quad (16)$$

Figure 3: A typical frame picking and encoding procedure of our framework. $\mathcal{F}$ denotes PickNet. E is the encoder unit and $\mathbf{v}$ is the encoded video representation. The design choice is the balance between processing time and computation cost. The system can simultaneously extract convolutional features and decide whether to pick the frame or not at each time-step. If it decides not to pick the frame at certain time-step, the convolutional neural network can stop early to save computation cost.

where N_p is the number of picked frames, $\mathbf{x}_i^{(j)}$ is the j-th value of the i-th visual feature, and $\mu^{(j)} = \frac{1}{N_p} \sum_{i=1}^{N_p} \mathbf{x}_i^{(j)}$ is the mean of all the j-th value of visual features.

Picks limitation. If the number of picked frames is too large or too small, it may lead to poor performances in either efficiency or effectiveness, so we assign a negative reward to discourage this situations. Empirically, we set the minimum picked number $N_{\min}$ as 3, which stands for beginning, highlight and ending. The maximum picked number $N_{\max}$ is initially set as the $\frac{1}{3}$ of total frame number, and will be shrunk down along with the training process, until decreased to a minimum value τ.

In summary, we merge the two parts of reward, and the final reward can be written as

$$r(v_i) = \begin{cases} \lambda_l r_l(v_i, S_i) + \lambda_v r_v(v_i) & \text{if} \quad N_{\min} \leq N_p \leq N_{\max} \\ R^- & \text{otherwise}, \end{cases} \tag{17}$$

where λ_* is the weighting hyper-parameters and R^- is the penalty reward.

3.3. Training

The training procedure is splitted into three stages. The first stage is to pretrain the Encoder-Decoder. We call it *supervision* stage. In the second stage, we fix the Encoder-Decoder and train PickNet by reinforcement learning. It is called *reinforcement* stage. And the final stage is the joint training of PickNet and the Encoder-Decoder. We call it *adaptation* stage. We use standard back-propagation to train the Encoder-Decoder, and REINFORCE [34] to train Pick-Net.

Supervision stage. When training the Encoder-Decoder, traditional method maximizes the likelihood of the next ground-truth word given previous ground-truth words using back-propagation. However, this approach causes the *exposure bias* [26], which results in error accumulation during generation at test time, since the model has never been exposed to its own predictions. In order to alleviate this phenomenon, the schedule sampling [4] procedure is used, which feeds back the model's own predictions and slowly increases the feedback probability during training. We use SGD with cross entropy loss to train the Encoder-Decoder. Given the ground-truth sentences $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m)$, the loss is defined as:

$$L_X(\omega) = -\sum_{t=1}^{m} \log(p_\omega(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots \mathbf{y}_1, \mathbf{v})), \tag{18}$$

where $p_\omega(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots \mathbf{y}_1, \mathbf{v})$ is given by the parametric model in Equation (11).

Reinforcement stage. In this stage, we fix the Encoder-Decoder and treat it as the *environment*, which can produce language reward to reinforce PickNet. The goal of training is to minimize the negative expected reward:

$$L_R(\theta) = -\mathbb{E}_{\mathbf{a}^s \sim p_\theta} \left[r(\mathbf{a}^s) \right], \tag{19}$$

where θ denotes all parameters of PickNet, p_θ is the learned policy parameterized by Equation (13), and $\mathbf{a}^s = (a_1^s, a_2^s, \dots, a_n^s)$ while a_t^s is the action sampled from the learned policy at the time step t.

We train PickNet by using REINFORCE algorithm, which is based on the observation that the gradient of a non-differentiable expected reward can be computed as follows:

$$\nabla_\theta L_R(\theta) = -\mathbb{E}_{\mathbf{a}^s \sim p_\theta} \left[r(\mathbf{a}^s) \nabla_\theta \log p_\theta(\mathbf{a}^s) \right]. \tag{20}$$

Using the chain rule, the gradient can be rewritten as:

$$\nabla_\theta L_R(\theta) = \sum_{t=1}^{n} \frac{\partial L_R(\theta)}{\partial \mathbf{s}_t} \frac{\partial \mathbf{s}_t}{\partial \theta} \tag{21}$$

$$= \sum_{t=1}^{n} -\mathbb{E}_{\mathbf{a}^s \sim p_\theta} r(\mathbf{a}^s)(p_\theta(a_t^s) - \mathbf{1}_{a_t^s}) \frac{\partial \mathbf{s}_t}{\partial \theta}, \tag{22}$$

where $\mathbf{s}_t$ is the input to the softmax function. In practice, the gradient can be approximated using a single Monte-Carlo sample $\mathbf{a}^s = (a_1^s, a_2^s, \ldots, a_n^s)$ from p_θ:

$$\nabla_\theta L_R(\theta) \approx -\sum_{t=1}^{n} r(\mathbf{a}^s)(p_\theta(a_t^s) - \mathbf{1}_{a_t^s}) \frac{\partial \mathbf{s}_t}{\partial \theta}. \tag{23}$$

When using REINFORCE to train the policy network, we need to estimate a baseline reward b to diminish the variance of gradients. Here, the *self-critical* [27] strategy is used to estimate b. In brief, the reward obtained by current model under inferencing used at test stage, denoted as $r(\hat{\mathbf{a}})$, is treated as the baseline reward. Therefore, the final gradient expression is:

$$\nabla_\theta L_R(\theta) \approx -(r(\mathbf{a}^s) - r(\hat{\mathbf{a}})) \sum_{t=1}^{n} (p_\theta(a_t^s) - \mathbf{1}_{a_t^s}) \frac{\partial \mathbf{s}_t}{\partial \theta}. \tag{24}$$

Adaptation stage. After the first two stages, the Encoder-Decoder and PickNet are well pretrained, but there exists a gap between them because the Encoder-Decoder use the full video frames as input while PickNet just selects a portion of frames. So we need a joint training stage to integrate this two parts together. However, the pick action is not differentiable, so the gradients introduced by cross-entropy loss can not flow into PickNet. Hence, we follow the approximate joint training scheme. In each iteration, the forward pass generates frame picks which are treated just like fixed picks when training the Encoder-Decoder, and the backward propagation and REINFORCE updates are performed as usual. It acts like performing dropout in time sequence, which can improve the versatility of the Encoder-Decoder.

4. Experimental Setup
4.1. Datasets

We evaluate our model on two widely used video captioning benchmark datasets: the Microsoft Video Description (MSVD) [5] and the MSR Video-to-Text (MSR-VTT) [35].

Microsoft Video Description (MSVD). The Microsoft Video Description is also known as YoutubeClips. It contains 1,970 Youtube video clips, each labeled with around 40 English descriptions collected by Amazon Mechanical Turkers. As done in previous works [33], we split the dataset into three parts: the first 1,200 videos for training, then the followed 100 videos for validation and the reset 670 videos for test. This dataset mainly contains short video clips with a single action, and the average duration is about 9 seconds. So it is very suitable to use only a portion of frames to represent the full video.

MSR Video-to-Text (MSR-VTT). The MSR Video-to-Text is a large-scale benchmark for video captioning. It provides 10,000 video clips, and each video is annotated with 20 English descriptions. Thus, there are 200,000 video-caption pairs in total. This dataset is collected from a commercial video search enginging and so far it covers the most comprehensive categories and diverse visual contents. Following the original paper, we split the dataset in contiguous groups of videos by index number: 6,513 for training, 497 for validation and 2,990 for test.

4.2. Metrics

We employ four popular metrics for evaluation: BLEU [25], ROUGE$_L$ [20], METEOR [2] and CIDEr. As done in previous video captioning works, we use METEOR and CIDEr as the main comparison metrics. In addition, Microsoft COCO evaluation server [6] has implemented these metrics and release evaluation functions[1], so we directly call such evaluation functions to test the performance of video captioning. Also, the CIDEr reward is computed by these functions.

4.3. Video preprocessing

First, we sample equally-spaced 30 frames for every video, and resize them into 224×224 resolution. Then the images are encoded with the final convolutional layer of ResNet152 [12], which results in a 2,048-dimensional vector. Most video captioning models use motion features to improve performance. However, we only use the appearance features in our model, because extracting motion features is very time-consuming, which deviates from our purpose that cutting down the computation cost for video captioning.

4.4. Text preprocessing

We tokenize the labeled sentences by converting all words to lowercases and then utilizing the word_tokenize function from NLTK[2] toolbox to split sentences into words and remove punctuation. Then, the word with frequency less than 3 is removed. As a result, we obtain the vocabulary with 5,491 words from MSVD and 13,065 words from MSR-VTT. For each dataset, we use the one-hot vector (1-of-N encoding, where N is the size of vocabulary) to represent each word.

---

[1] https://github.com/tylin/coco-caption
[2] http://www.nltk.org/

Ours: a cat is playing with a dog
GT: a dog is playing with a cat

Ours: a person is solving a rubik's cube
GT: person playing with toy

Figure 5: Example results on MSVD (left) and MSR-VTT (right). The green boxes indicate picked frames. (Best viewed in color and zoom-in. Frames are organized from left to right, then top to bottom in temporal order.)

4.5. Implementation details

We use the validation set to tune some hyperparameters of our framework. The learning rates for three training stages are set to 3×10^{-4}, 3×10^{-4} and 1×10^{-4}, respectively. The training batchsize is 128 for MSVD and 256 for MSR-VTT, while each stage is trained up to 100 epoches and the best model is used to initialize the next stage. The minimum value of maximum picked frames τ is set to 7, and the penalty reward R^- is -1. To regularize the training and avoid over-fitting, we apply the well known regularization technique Dropout with retain probability 0.5 on the input and output of the encoding LSTMs and decoding GRUs. Embeddings for video features and words have size 512, while the sizes of all recurrent hidden states are empirically set to 1,024. For PickNet, the size of glance is 56×56, and the size of hidden layer is 1,024. The Adam [16] optimizer is used to update all the parameters.

5. Results and Discussion

Figure 5 gives some example results on the test sets of two datasets. As it can be seen, our PickNet can select informative frames, so the rest of our model can use these selected frames to generate reasonable descriptions. More results are offered in supplemental materials. In order to demonstrate the effectiveness of our framework, we compare our approach with some state-of-the-art methods on the two datasets, and analyze the learned picks of PickNet.

5.1. Comparison with the state-of-the-arts

We compare our approach on MSVD with six state-of-the-art approaches for video captioning: TA [39], S2VT [33], LSTM-E [24], p-RNN [42] HRNE [23] and BA [3]. LSTM-E uses a visual-semantic embedding to

generate better captions. TA uses temporal attention while p-RNN use both temporal and spatial attention. BA uses a hierarchical encoder while HRNE use a hierarchical decoder to describe videos. S2VT uses stack LSTMs both for the encode and decode stage. All of these methods use motion features (C3D or optical flow) and extract visual features frame by frame. Besides, we report the performance of our baseline model, which encodes all the sampled frames. In order to compare our PickNet with other picking policies, we conduct two other trials that pick frames by randomly selecting and k-means clustering, respectively. Also, for analyzing the effect of different rewards, we conduct some ablation studies on them. As it can be noticed in Table 3, our method improves plain techniques and achieves the state-of-the-art performance on MSVD. This result outperforms the most recent state-of-the-art method by a considerable margin of $\frac{76.0-65.8}{65.8} \approx 15.5\%$ on the CIDEr metric. Further, we try to compare the time efficiency among these approaches. However, most of state-of-the-art methods do not release executable codes, so the accurate performance may not be available. Instead, we estimate the running time by the complexity of visual feature extractors and the number of processed frames. The details of running time estimation are listed in supplemental materials. Thanks to the PickNet, our captioning model is 4~33 times faster than other methods.

On MSR-VTT, we compare four state-of-the-art approaches: ruc-uva [9], Aalto [29], DenseCap [28] and MS-RNN [30]. ruc-uva incorporates the Encoder-Decoder with two new stages called early embedding which enriches input with tag embeddings, and late reranking which rescore generated sentences in terms of their relevance to a specific video. Aalto first trains two models which are separately based on attribute and motion features, and then trains

Model	BLEU@4	ROUGE-L	METEOR	CIDEr	Time
Previous Work					
TA[39]	41.9	-	29.6	51.7	4x
S2VT[33]	-	-	29.8	-	13x
LSTM-E[24]	45.3	-	31.0	-	5x
p-RNN[42]	**49.9**	-	32.6	65.8	5x
HRNE[23]	43.8	-	**33.1**	-	33x
BA[3]	42.5	-	32.4	63.5	12x
Our Models					
Baseline	44.8	68.5	31.6	69.4	5x
Random	35.6	64.5	28.4	49.2	2.5x
k-means (k=6)	45.2	68.5	32.4	70.9	1x
PickNet (V)	43.6	68.4	32.4	75.6	1x
PickNet (L)	**49.9**	**69.3**	32.9	74.7	1x
PickNet (V+L)	46.1	69.2	**33.1**	**76.0**	1x

Table 1: Experiment results on MSVD. L denotes using language reward and V denotes using visual diversity reward. k is set to the average number of picks $\bar{N}_p$ on MSVD. ($\bar{N}_p \approx 6$)

Model	BLEU@4	ROUGE-L	METEOR	CIDEr	Time
Previous Work					
ruc-uva [9]	38.7	58.7	26.9	45.9	4.5x
Aalto [29]	39.8	59.8	26.9	45.7	4.5x
DenseCap [28]	41.4	61.1	28.3	48.9	3.5x
MS-RNN [30]	39.8	59.3	26.1	40.9	10x
Our Models					
Baseline	36.8	59.0	26.7	41.2	3.8x
Random	31.3	55.7	25.2	32.6	1.9x
k-means (k=8)	37.8	59.1	26.9	41.4	1x
PickNet (V)	35.6	58.2	26.8	41.0	1x
PickNet (L)	37.3	58.9	27.0	41.9	1x
PickNet (V+L)	38.9	59.5	27.2	42.1	1x

Table 2: Experiment results on MSR-VTT. k is set to the average number of picks $\bar{N}_p$ on MSR-VTT. ($\bar{N}_p \approx 8$)

a evaluator to choose the best candidate generated by the two captioning model. DenseCap generates multiple sentences with regard to video segments and uses a winner-take-all scheme to produce the final description. MS-RNN uses a multi-modal LSTM to model the uncertainty in videos to generate diverse captions. Compared with these methods, our method can be simply trained in end-to-end fashion, and does not rely upon any attribute information. The performance of these approaches and that of our solution is reported in Table 4. We observe that our approach is able to achieve competitive result even without utilizing attribute information, while other methods take advantage of attributes and auxiliary information sources. Also, our model is fastest. It is also worth noting that the PickNet can be easily integrated with the compared methods, since none of them incorporated with frame selection algorithm. For example, DenseCap generates region-sequence candidates based on equally sampled frames. It can alternatively utilize PickNet to reduce the time for generating candidates by cutting down the number of selected frames.

(a) Dist. of the number of picks.   (b) Dist. of the position of picks.

Figure 6: Statistics on the behavior of our PickNet.

5.2. Analysis of learned picks

We collect statistics on the properties of our PickNet. Figure 6 shows the distributions of the number and position of picked frames on the test sets of MSVD and MSR-VTT. As observed in Figure 6(a), in the vast majority of the videos, less than 10 frames are picked. It implies that in most case only $\frac{10}{30} \approx 33.3\%$ frames are necessary to be encoded for captioning videos, which can largely reduce the computation cost. Specifically, the average number of picks is around 6 for MSVD and 8 for MSR-VTT. Looking at the distributions of position of picks in Figure 6(b), we observe a pattern of *power law distribution*, *i.e.*, the probability of picking a frame is reduced as time goes by. It is reasonable since most videos are single-shot and the anterior frames are sufficient to represent the whole video.

5.3. Captioning for streaming video

a cat is playing → a rabbit is playing → a rabbit is being petted → **a person is petting a rabbit** ×3

Figure 7: An example of online video captioning.

One of the advantage of our method is that it can be applied to streaming video. Different from offline video captioning, captioning for streaming video requires the model to tackle with unbounded video and generate descriptions immediately when the visual information has changed, which meets the demand of practical applications. For this online setting, we first sample frames at 1fps, and then sequentially feed the sampled frames to PickNet. If certain frame is picked, the pretrained CNN will be used to extract visual features of this frame. After that, the encoder will receive this feature, and produce a new encoded representation of the video stream up to current time. Finally, the decoder will generate a description based on the encoded representation. Figure 7 demonstrates an example of online video captioning with the picked frames and corresponding descriptions. As it is shown, the descriptions will be more appropriate and more determined as the informative frames are picked. More results can be seen in supplemental materials.

6. Conclusion

In this work, we design a plug-and-play reinforcement-learning-based PickNet to select informative frames for the task of video captioning, which achieves promising performance on effectiveness, efficiency and flexibility on popular benchmarks. This architecture can largely cut down the usage of convolution operations by picking only 6∼8 frames for a video clip, while other video analysis methods usually require more than 40 frames. This property makes our method more applicable for real-world video processing. The proposed PickNet has a good flexibility and could be potentially employed to other video-related applications, such as video classification and action detection, which will be further addressed in our future work.

References

[1] CNN-Benchmarks: benchmarks for popular CNN models. https://github.com/jcjohnson/cnn-benchmarks. 11

[2] S. Banerjee and A. Lavie. Meteor: an automatic metric for mt evaluation with improved correlation with human judgments. In *ACL*, pages 65–72, Ann Arbor, 2005. Association for Computational Linguistics. 6

[3] L. Baraldi, C. Grana, and R. Cucchiara. Hierarchical boundary-aware neural encoder for video captioning. In *CVPR*, 2017. 2, 3, 7, 8, 11

[4] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*, pages 1171–1179, Montreal, 2015. Neural Information Processing Systems Foundation. 5

[5] D. L. Chen and W. B. Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL HLT*, pages 190–200, 2011. 6

[6] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, C. L. Zitnick, and P. Dollár. Microsoft coco captions: Data collection and evaluation server. *arXiv*, 2015. 6

[7] K. Cho, B. van Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014. 3

[8] H. C. Cromwell, R. P. Mears, L. Wan, and N. N. Boutros. Sensory gating: A translational effort from basic to clinical science. *Clinical EEG and Neuroscience*, 39(2):69–72, 2008. 1

[9] J. Dong, X. Li, W. Lan, Y. Huo, and C. G. M. Snoek. Early embedding and late reranking for video captioning. In *ACM MM*, pages 1082–1086, Amsterdam, 2016. ACM Press. 7, 8, 11

[10] H. Fang, S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollar, J. Gao, X. He, M. Mitchell, J. C. Platt, C. L. Zitnick, and G. Zweig. From captions to visual concepts and back. In *CVPR*, pages 1473–1482. IEEE, jun 2015. 2

[11] A. Farhadi, M. Hejrati, M. A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth. Every picture tells a story: Generating sentences from images. In *ECCV*, pages 15–29, Crete, Greece, 2010. Springer. 2

[12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 171–180, Las Vega, 2016. IEEE Computer Society Press. 6

[13] S. Hochreiter and J. J. J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 3

[14] C. Hori, T. Hori, T.-Y. Lee, K. Sumi, J. R. Hershey, and T. K. Marks. Attention-based multimodal fusion for video description. In *ICCV*, 2017. 1

[15] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *CVPR*, pages 4565–4574, 2016. 2

[16] D. P. Kingma and J. L. Ba. Adam: a method for stochastic optimization. In *ICLR*, pages 1–15, San Diego, 2015. 7

[17] A. Kojima, T. Tamura, and K. Fukunaga. Natural language description of human activities from video images based on concept hierarchy of actions. *IJCV*, 50(2):171–184, 2002. 2

[18] J. Krause, J. Johnson, R. Krishna, and L. Fei-Fei. A hierarchical approach for generating descriptive image paragraphs. In *CVPR*, 2017. 2

[19] G. Kulkarni, V. Premraj, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Baby talk: Understanding and generating image descriptions. In *CVPR*. Citeseer, 2011. 2

[20] C. Y. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL workshop*, number 1, pages 25–26, 2004. 6

[21] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical co-attention for visual question answering. In *NIPS*, pages 1–11, 2016. 1

[22] V. Mnih, N. Heess, A. Graves, and K. Kavukcuoglu. Recurrent models of visual attention. In *NIPS*, pages 1–9, 2014. 1

[23] P. Pan, Z. Xu, Y. Yang, F. Wu, and Y. Zhuang. Hierarchical recurrent neural encoder for video representation with application to captioning. In *CVPR*, pages 1029–1038, Las Vega, 2016. IEEE Computer Society Press. 7, 8, 11

[24] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. In *CVPR*, pages 4594–4602, Las Vega, may 2016. IEEE Computer Society Press. 7, 8, 11

[25] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318, Philadelphia, 2002. Association for Computational Linguistics. 6

[26] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba. Sequence level training with recurrent neural networks. In *ICLR*, pages 1–16, 2016. 5

[27] S. J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, and V. Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. 6

[28] Z. Shen, J. Li, Z. Su, M. Li, Y. Chen, Y.-G. Jiang, and X. Xue. Weakly supervised dense video captioning. In *CVPR*, 2017. 1, 2, 7, 8, 11

[29] R. Shetty and J. Laaksonen. Frame-and segment-level features and candidate pool evaluation for video caption generation. In *ACM MM*, pages 1073–1076, Amsterdam, 2016. ACM Press. 7, 8, 11

[30] J. Song, Y. Guo, L. Gao, X. Li, A. Hanjalic, and H. T. Shen. From deterministic to generative: Multi-modal stochastic rnns for video captioning. *arXiv*, 2017. 7, 8, 11

[31] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes. Tvsum: Summarizing web videos using titles. In *CVPR*, pages 5179–5187, 2015. 3

[32] R. Vedantam, C. L. Zitnick, and D. Parikh. Cider: consensus-based image description evaluation. In *CVPR*, pages 4566–4575, Boston, 2015. IEEE Computer Society Press. 4

[33] S. Venugopalan, M. Rohrbach, T. Darrell, J. Donahue, K. Saenko, and R. Mooney. Sequence to sequence - video to text. In *ICCV*, pages 4534–4542, Santiago, 2015. IEEE Computer Society Press. 6, 7, 8, 11

[34] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992. 5

[35] J. Xu, T. Mei, T. Yao, and Y. Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 6

[36] K. Xu, J. L. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, Y. Bengio, J. L. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: neural image caption generation with visual attention. In *ICML*, pages 2048–2057, Lille, 2015. 2

[37] Y. Yang, C. L. Teo, H. Daumé III, and Y. Aloimonos. Corpus-guided sentence generation of natural images. In *EMNLP*, pages 444–454. Association for Computational Linguistics, 2011. 2

[38] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *CVPR*, number 1, pages 21–29, 2015. 1

[39] L. Yao, K. Cho, N. Ballas, C. Paí, and A. Courville. Describing videos by exploiting temporal structure. In *ICCV*, pages 4507–4515, Santiago, 2015. IEEE Computer Society Press. 1, 2, 7, 8, 11

[40] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016. 3

[41] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo. Image captioning with semantic attention. In *CVPR*, page 10, 2016. 1

[42] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu. Video paragraph captioning using hierarchical recurrent neural networks. In *CVPR*, pages 4584–4593, Las Vega, oct 2016. IEEE Computer Society Press. 7, 8, 11

[43] Y. Yu, J. Choi, Y. Kim, K. Yoo, S.-H. Lee, G. Kim, and G. Kim. Supervising neural attention models for video captioning by human gaze data. In *CVPR*, Honolulu, 2017. 1

[44] B. Zhao and E. P. Xing. Quasi real-time summarization for consumer videos. *CVPR*, pages 2513–2520, 2014. 3

[45] H. Zheng, J. Fu, and T. Mei. Look closer to see better : Recurrent attention convolutional neural network for fine-grained image recognition. *CVPR*, pages 4438–4446, 2017. 1

Supplemental Materials

7. Details on Time Estimation

We estimate the running time by the complexity of visual feature extractors and the number of processed frames, and our PickNet (V+L) is treated as the baseline. For visual features, both appearance and motion are taken into consideration. The relative running time of different CNNs is based on public reports [1]. If a specific model uses motion features, the total computation time will be doubled, since extracting motion features is very time-consuming. For each model, the number of processed frames is set to the expected number of sampled frames under the sampling method. However, some model sample every 5 or 10 frames from input video, so the expected number depends on video length. In order to compare these model

Model	Appearance	Motion	Sampling method	Frame num.	Time
Previous Work					
TA[39]	GoogleNet (0.5x)	C3D (2x)	uniform sampling 26 frames	26 (4x)	$0.5 \times 2 \times 4 = 4x$
S2VT[33]	VGG (0.5x)	OF (2x)	uniform sampling 80 frames	80 (13x)	$0.5 \times 2 \times 13 = 13x$
LSTM-E[24]	VGG (0.5x)	C3D (2x)	uniform sampling 30 frames	30 (5x)	$0.5 \times 2 \times 5 = 5x$
p-RNN[42]	VGG (0.5x)	C3D (2x)	uniform sampling 30 frames	30 (5x)	$0.5 \times 2 \times 5 = 5x$
HRNE[23]	GoogleNet (0.5x)	C3D (2x)	first 200 frames	200 (33x)	$0.5 \times 2 \times 33 = 33x$
BA[3]	ResNet (0.5x)	C3D (2x)	every 5 frames	72 (12x)	$0.5 \times 2 \times 12 = 12x$
Our Models					
Baseline	ResNet (1x)	$\times$	uniform sampling 30 frames	30 (5x)	$1 \times 5 = 5x$
Random	ResNet (1x)	$\times$	randomly sampling	15 (2.5x)	$1 \times 2.5 = 2.5x$
k-means (k=6)	ResNet (1x)	$\times$	k-means clustering	6 (1x)	$1 \times 1 = 1x$
PickNet (V)	ResNet (1x)	$\times$	picking	6 (1x)	$1 \times 1 = 1x$
PickNet (L)	ResNet (1x)	$\times$	picking	6 (1x)	$1 \times 1 = 1x$
PickNet (V+L)	ResNet (1x)	$\times$	picking	6 (1x)	$1 \times 1 = 1x$

Table 3: Running time estimation on MSVD. OF means optical flow. BA uses ResNet50 while our models use ResNet152. k is set to the average number of picks $\bar{N}_p$ on MSVD. ($\bar{N}_p \approx 6$)

Model	Appearance	Motion	Sampling method	Frame num.	Time
Previous Work					
ruc-uva [9]	GoogleNet (0.5x)	C3D (2x)	every 10 frames	36 (4.5x)	$0.5 \times 2 \times 4.5 = 4.5x$
Aalto [29]	GoogleNet (0.5x)	C3D+IDT (2x)	one frame every second	36 (4.5x)	$0.5 \times 2 \times 4.5 = 4.5x$
DenseCap [28]	ResNet (0.5x)	C3D (2x)	uniform sampling 30 frames	30 (3.5x)	$0.5 \times 2 \times 3.5 = 3.5x$
MS-RNN [30]	ResNet (1x)	C3D (2x)	uniform sampling 40 frames	40 (5x)	$1 \times 2 \times 5 = 10x$
Our Models					
Baseline	ResNet (1x)	$\times$	uniform sampling 30 frames	30 (3.8x)	$1 \times 3.8 = 3.8x$
Random	ResNet (1x)	$\times$	randomly sampling	15 (1.9x)	$1 \times 1.9 = 1.9x$
k-means (k=8)	ResNet (1x)	$\times$	k-means clustering	8 (1x)	$1 \times 1 = 1x$
PickNet (V)	ResNet (1x)	$\times$	picking	8 (1x)	$1 \times 1 = 1x$
PickNet (L)	ResNet (1x)	$\times$	picking	8 (1x)	$1 \times 1 = 1x$
PickNet (V+L)	ResNet (1x)	$\times$	picking	8 (1x)	$1 \times 1 = 1x$

Table 4: Running time estimation on MSR-VTT. IDT means improved dense trajectory. DenseCap uses ResNet50 while our models use ResNet152. k is set to the average number of picks $\bar{N}_p$ on MSR-VTT. ($\bar{N}_p \approx 8$)

with others, we consider the input video with a duration of 10 seconds and a frame rate of 36fps, thus the total number of frames for input video is fixed to 360. Table 3 and Table 4 show the detail of time estimation on MSVD and MSR-VTT.

It is worth mentioning that the estimated time on other compared approaches is just rough. In fact, the actual speedup of our model will be higher than we estimated, because most of the compared models use complex pipeline. For instance, in the p-RNN, the temporal- and spatial-attention mechanisms are exploited to selectively focus on visual elements during sentence generation. Also, in the DenseCap, a lexical FCN is utilized to extract region features of every frames, then the cost-effective lazy forward (CELF) method is employed to generate region-sequence candidates, and finally a bi-directional LSTM is used to encode each region-sequence candidate. All the above procedures are far more complex than our Encoder-Decoder pipeline, therefore will consume much more processing time.

## 8. More Result Examples

8.1. Offline video captioning

Figure 8 shows more results for offline video captioning on MSVD and MSR-VTT. As mentioned before, our PickNet can select a minor portion of frames to represent the whole video and describe the video with regard to picked frames. Moreover, in the left column, only 3 or 4 frames are picked, and these picked frames are all in the front part of the video. We suppose that it is because these videos contain univocal content, and in each video most frames are similar. Under these circumstances, it is enough to pick a few of frames at the beginning to describe these video. Meanwhile, in the right column, more frames are picked, and the picked frames are scattered. We suggest that it is because the content of these videos is diverse. In this situation, the PickNet will traverse the whole video and select more frames.

Altogether, two characteristics of picked frames can be found. **The first characteristic is that the picked frames are concise and highly related to the generated descriptions.** For example, in Figure 8(a), our model only selects four frames, which correspond to *holding the gun*, *porting arms*, *aiming*, and *shooting*, respectively. All other frames are more or less duplicated visually or semantically, so those redundant frames are ignored. In Figure 8(c), our model selects the 5th frame instead of the 6th frame. Although the 6th frame is more diverse than the 5th frame, it is not related to the description, so our model does not select it but pick the 5th frame to confirm that the clip is about *playing a guitar*. In Figure 8(f), the picked frames appropriately describe the shot change, therefore the model can focus on the two women and understand they are talking to each other. **The second one is that the adjacent frames may be picked to represent action.** For example, in Figure 8(b), our model selects a pair of adjacent frames, i.e., the 6th and the 7th frames, which can properly represent the *seasoning* action. In Figure 8(d), the first two frames are picked to represent the *chopping* action. In Figure 8(e), the 5th to 7th frames are selected to represent the *playing* action. And in Figure 8(f), the 15th and the 16th frames are chosen to represent the *talking* action.

With these characteristics, our model may generate more accurate descriptions than ground-truths. For example, in Figure 8(b), our model explicitly indicates there is a *woman*, while the ground-truth only use *someone* to refer to it. In Figure 8(f), the generated sentence correctly describes the content of the video, while the ground-truth just tells it is a movie clip.

8.2. Online video captioning

For online video captioning, we first sample frames at 1fps, and then sequentially feed the sampled frames to PickNet. If a certain frame is picked, the pre-trained CNN will be used to extract visual features of this frame. After that, the encoder will receive this feature, and produce a new encoded representation of the video stream up to current time. Finally, the decoder will generate a description based on the encoded representation. Figure 9 demonstrates some examples of online video captioning with the picked frames and the corresponding descriptions. As we discussed before, the descriptions will be more appropriate and determined as the informative frames are picked. In example (a), the most salient object in the first picked frame is the *cat*, so the generated description is just about the cat. After observing enough frames, the model knows this video is about *a woman is playing with a kitten*, and produces the correct description. In example (b), the model first generates a description that *a boy is running*, since the man in a blue shirt is more prominent and the motion pattern seems like *running*. Along with picking the following frames, the other two persons are noticed and their actions are recognized as *dancing*, hence the model produces a more accurate description that *three persons are dancing*. And at the beginning of example (c), the model is only aware of that there is a *man* with a *sword*, and do not know what the man is doing. After picking the third frame, it is clear that the man is *stabbing* a target, then the word *target* is substituted by a more precise word *silhouette* when more frames are picked. The video version of online captioning results can be seen in the uploaded videos.

(a) **Ours: a man is shooting a gun**
GT: a man is shooting

(b) **Ours: a woman is seasoning meat**
GT: someone is seasoning meat

(c) **Ours: a man is playing a guitar**
GT: a man is playing a guitar

(d) **Ours: a woman is chopping some leaves**
GT: a woman is slicing some leaves

(e) **Ours: a person is playing a video game**
GT: a game is being played

(f) **Ours: there is a woman is talking with a woman**
GT: it is a movie

Figure 8: Example results on the test set of MSVD and MSR-VTT. The green boxes indicate picked frames. (Best viewed in color and zoom-in. Frames are organized from left to right, then top to bottom in temporal order.)

(a) **a cat is licking its lips** → **a woman is a baby** → **a woman is a baby** → **a woman is feeding a baby** → a woman is playing with a kitten

(b) **a boy is running** → **a boy is running** → **a boy is running** → **the boys are dancing** → three persons are dancing

(c) **a man is a sword** → **a boy is doing a** → **a man with a sword stabs a target** → a man is stabbing a silhouette with a sword ×2

Figure 9: Example results of online video captioning.