

End-to-End Dense Video Captioning with Parallel Decoding

Teng Wang^{1,2}, Ruimao Zhang^{3,4}, Zhichao Lu², Feng Zheng^{2*}, Ran Cheng², Ping Luo¹

¹ The University of Hong Kong ² Southern University of Science and Technology

³ The Chinese University of Hong Kong (Shenzhen) ⁴ Shenzhen Research Institute of Big Data

tengwang@connect.hku.hk ruimao.zhang@ieee.org luzhichaocn@gmail.com

f.zheng@ieee.org ranchengcn@gmail.com pluo@cs.hku.hk

Abstract

Dense video captioning aims to generate multiple associated captions with their temporal locations from the video. Previous methods follow a sophisticated “localize-then-describe” scheme, which heavily relies on numerous hand-crafted components. In this paper, we proposed a simple yet effective framework for end-to-end dense video captioning with parallel decoding (PDVC), by formulating the dense caption generation as a set prediction task. In practice, through stacking a newly proposed event counter on the top of a transformer decoder, the PDVC precisely segments the video into a number of event pieces under the holistic understanding of the video content, which effectively increases the coherence and readability of predicted captions. Compared with prior arts, the PDVC has several appealing advantages: (1) Without relying on heuristic non-maximum suppression or a recurrent event sequence selection network to remove redundancy, PDVC directly produces an event set with an appropriate size; (2) In contrast to adopting the two-stage scheme, we feed the enhanced representations of event queries into the localization head and caption head in parallel, making these two sub-tasks deeply interrelated and mutually promoted through the optimization; (3) Without bells and whistles, extensive experiments on ActivityNet Captions and YouCook2 show that PDVC is capable of producing high-quality captioning results, surpassing the state-of-the-art two-stage methods when its localization accuracy is on par with them. Code is available at <https://github.com/ttengwang/PDVC>.

1. Introduction

As an emerging branch of video understanding, video captioning has received an increasing attention in the recent past [2, 12, 13, 21, 32, 10, 15, 19, 20], aiming to generate a natural sentence to describe one main event of

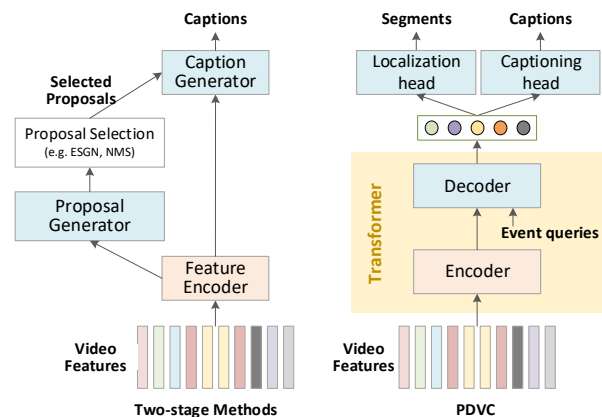


Figure 1: The defacto two-stage pipeline vs. the proposed PDVC. The two-stage “localize-then-describe” pipeline requires a dense-to-sparse proposal generation and selection process before captioning, which contains hand-crafted components and can not effectively exploit the potential mutual benefits between localization and captioning. PDVC adopts the vision transformer to learn attentive interaction of different frames, where the learnable event queries are embedded to capture the relevance between the frames and the events. Two prediction heads run in parallel over query features, leveraging the mutual benefits between two tasks and improving their performance together.

a short video. However, since realistic videos are usually long, untrimmed, and composed of a variety of events with irrelevant background contents, the above single-sentence captioning methods tend to generate sentences of blandness with less information. To circumvent the above dilemma, dense video captioning (DVC) [5, 7, 25, 29, 37] is developed for automatically localizing and captioning multiple events in the video, which could reveal detailed visual contents and generate the coherent and complete descriptions.

Intuitively, dense video captioning can be divided into two subtasks, termed event localization and event captioning. As shown in Fig. 1, the previous methods usually solve this problem by a two-stage “localize-then-describe” pipeline. It firstly predicts a set of event proposals with accurate boundaries. By extracting fine-grained semantic cues

* Corresponding author

and visual contexts of the proposal, the detailed sentence description is finally decoded by the caption generator. The above scheme is straightforward but suffers from the following issues: 1) By considering the captioning as the downstream task, the performance of such a scheme highly relies on the quality of the generated event proposals, which limits the mutual promotion of these two sub-tasks. 2) The performance of proposal generators in previous methods depends on careful anchor design [5, 31, 7, 24, 9, 34] and proposal selection post-processing (*e.g.*, non-maximum suppression [5, 31, 7, 24, 9, 34]). These hand-crafted components introduce additional hyper-parameters that highly rely on manual thresholding strategies, hindering the progress toward a fully end-to-end captioning generation.

To tackle the above issues, this paper proposes a pure end-to-end dense Video Captioning framework with Parallel Decoding termed PDVC. As shown in Fig. 1, instead of invoking the two-stage scheme, we directly feed the intermediate representation used for proposal generation into a captioning head that is parallel to the localization head. By doing so, PDVC aims to directly exploit inter-task association at the feature level. The intermediate feature vectors and the target events could be matched in a one-to-one correspondence, making the feature representations more discriminative to identify a specific event.

In practice, we consider the dense video captioning task as a set prediction problem. The proposed PDVC directly decodes the frame features, which are extracted from a Vision Transformer, into an event set with their locations and corresponding captions by applying two parallel prediction heads, *i.e.*, localization head and captioning head. Since the appropriate size of the event set is an essential indicator for dense captioning quality [9, 48], a newly proposed event counter is also stacked on the top of the Transformer decoder to further predict the number of final events. By introducing such a simple module, PDVC could precisely segment the video into a number of event pieces under the holistic understanding of the video content, avoiding the information missing as well as the replicated caption generation caused by unreliable event number estimation.

We evaluate our model on two large-scale video benchmarks, ActivityNet Captions and YouCook2. Even with a lightweight caption head (vanilla LSTM), our method can achieve comparable performance against state-of-the-art methods which adopts well-designed attention-based LSTM [24, 34] or Transformer [31]. In addition, we show quantitatively and qualitatively that the generated proposals gain benefit from the paralleling decoding design. Even with a weakly supervised setting (without location annotations), we show our model can implicitly learn the location-aware features from captions.

To summarize, the major contributions of this paper are three folds. 1) We propose a novel end-to-end dense video

captioning framework named PDVC by formulating DVC as a parallel set prediction task, significantly simplifying the traditional pipeline which highly depends on hand-crafted components. 2) We further improve PDVC with a novel event counter to estimate the number of events in the video, greatly increasing the readability of generated captions by avoiding the unrealistic event number estimation. 3) Extensive experiments on ActivityNet Captions and YouCook2 show state-of-the-art performance over existing methods.

2. Related Work

Temporal event proposals. Temporal event proposals (TEP), also called temporal action proposals, aims to predict temporal segments containing event instances in untrimmed videos. Mainstream approaches can be divided into two categories: anchor-based and boundary-based. Anchor-based methods [51, 49, 56, 47] pre-define a vast number of anchors at different scales with regular intervals, followed by the evaluation network to score each anchor. However, the pre-defined scales and intervals can not cover all temporal patterns, especially in videos with variable temporal scales. The boundary-based methods [59, 55, 54, 53] combine salient frames with high confidences to form proposals in a local-to-global manner. Both types of methods contain hand-crafted designs (*e.g.*, NMS and rule-based label assignment), which require careful manual threshold selection and are not a strictly end-to-end method.

Dense video captioning. Dense video captioning is a multi-task problem that combines event localization and event captioning. Krishna et al. [5] propose the first dense video captioning model, containing a multi-scale proposal module for localization and an attention-based LSTM for context-aware caption generation. Some of the following work aims to enrich the event representations by context modeling [24, 28], event-level relationships [34], or multi-modal feature fusion [35, 36], enabling more accurate and informative captioning generation.

One of the limitations of the above approaches is that the localization module can not benefit from the captioning module. Some researchers try to explore the interaction between two sub-tasks. Li et al. [7] introduce a proxy task, *i.e.*, predicting language rewards of generated sentences, as an additional optimization goal of the localization module. Zhou et al. [31] propose a differential masking mechanism to link the gradient flow from captioning loss to proposals' boundaries, enabling a joint optimization of two tasks. We argue that neither a binary mask vector [31] nor a scalar descriptiveness score [7] carries enough informative gradients of linguistic cues to guide the internal feature representation in the proposal module during the back-propagation training. Instead, the proposed PDVC exploits the inter-task interactions by enforcing the two sub-tasks share the

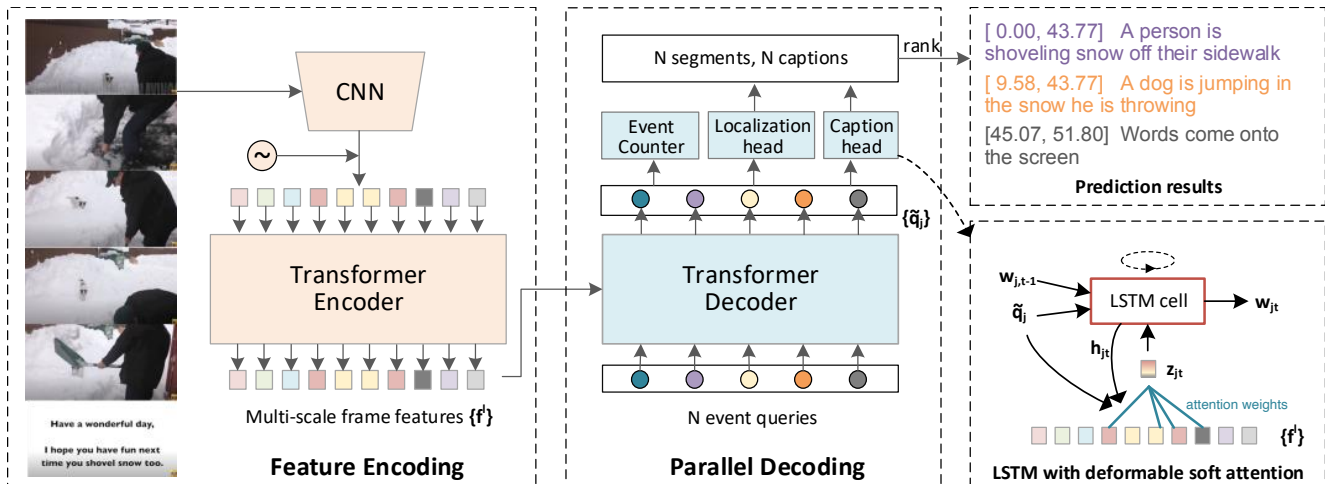


Figure 2: Overview of the proposed method. Firstly, we adopt a pre-trained video feature extractor and a transformer encoder to obtain a sequence of frame-level features. A transformer decoder and three prediction heads are then proposed to predict the locations, captions, and the number of events given learnable event queries. We provide two types of caption heads, which are based on vanilla LSTM and deformable soft-attention enhanced LSTM, respectively. In the testing stage, we select the top detected events by ranking the captioning score and localization score, with no need to remove redundancy by non-maximum suppression.

same intermediate features. Moreover, we adopt one-to-one matching between intermediate feature vectors and target event instances to obtain the discriminative features for captioning, significantly different from previous methods with a many-to-one anchor assignment strategy.

Another promising direction focuses on the coherence of generated captions. Early work [5, 7, 31, 24] usually generates a large number of proposal-caption pairs (10 times more than the number of ground-truth events) for a high recall, where massive redundancy greatly reduces the readability and coherence of generated captions. SDVC [9] is the first to tackle this problem by introducing a “localize-select-describe” pipeline. Given the output proposals produced by a TEP model, they develop an RNN-based event sequence generation network (ESGN) to select a small set of proposals, reducing the predicted proposal number from 100 to 2.85 on average. Though promising performance is achieved, SDVC is not an end-to-end model, making a multi-step training strategy necessary. The recurrent nature also restricts the application of ESGN to handle long videos with a large number of events. We parallelize localization, selection, and captioning tasks in a single end-to-end framework, largely simplifying the pipeline while enabling generating accurate and coherent captions.

Transformer-based detector. Transformer [17] is an encoder-decoder architecture based on an attention mechanism for natural language processing. Benefit from the significant ability to capture long-range relationships, Transformer has been successfully applied and shows promising performance in computer vision [46, 57, 50, 58, 45]. Detection Transformer (DETR) [39] is a newly emerging solution

to object detection, which considers object detection as a set prediction task and does not rely on any hand-crafted components. Though it offers promising performance, DETR suffers from the high training time due to the slow convergence of global attention mechanism. Deformable Transformer [38] is proposed to speed up the network training and gain better performance by attending to sparse spatial locations of images and incorporating multi-scale feature representation. Inspired by the simple design and promising performance of DETR-style detectors in the image domain, we extend Deformable Transformer into a more challenging dense video captioning task in the video domain.

3. Method

To simplify the dense video captioning pipeline and explore the mutual benefits between localization task and captioning task, we directly detect a set of temporally-localized captions with an appropriate size $\{(t_j^s, t_j^e, S_j)\}_{j=1}^{N_{\text{set}}}$, where t_j^s , t_j^e , S_j represent the starting time, ending time, and the caption of an event, respectively. The set size N_{set} is also predicted by PDVC.

Specifically, a deformable transformer with an encoder-decoder structure is adopted to capture the inter-frame, inter-event, and event-frame interactions by attention mechanism and produce a set of event query features. Then, two parallel prediction heads predict the boundaries and captions for each event query simultaneously. An event counter predicts the event number N_{set} from a global view. The final results are obtained by select top N_{set} events with high confidence to ensure a complete and coherent story. Fig. 2 shows the overview of the proposed PDVC.

3.1. Preliminary: Deformable Transformer

Deformable Transformer [38] is an encoder-decoder architecture based on multi-scale deformable attention (MSDAtt). MSDAtt mitigates the slow convergency problem of the self-attention [17] in Transformer when processing image feature maps, by attending to a sparse set of sampling points around reference points. Given multi-scale feature maps $\mathbf{X} = \{\mathbf{x}^l\}_{l=1}^L$ where $\mathbf{x}^l \in \mathbb{R}^{C \times H \times W}$, a query element \mathbf{q}_j and a normalized reference point $\mathbf{p}_j \in [0, 1]^2$, MSDAtt outputs a context vector by the weighted sum of $K \times L$ sampling points across feature maps at L scales:

$$\text{MSDAtt}(\mathbf{q}_j, \mathbf{p}_j, \mathbf{X}) = \sum_{l=1}^L \sum_{k=1}^K A_{jlk} \mathbf{W} \mathbf{x}_{\tilde{\mathbf{p}}_{jlk}}^l \quad (1)$$

$$\tilde{\mathbf{p}}_{jlk} = \phi_l(\mathbf{p}_j) + \Delta \mathbf{p}_{jlk},$$

where $\tilde{\mathbf{p}}_{jlk}$ and A_{jlk} are the position and the attention weight of k -th sampled key at l -th scale for j -th query element, respectively. \mathbf{W} is the projection matrix for key elements. ϕ_l projects normalized reference points into the feature map at l -th level. $\Delta \mathbf{p}_{jlk}$ are sampling offsets w.r.t. $\phi_l(\mathbf{p}_j)$. Both A_{jlk} and $\Delta \mathbf{p}_{jlk}$ are obtained by linear projection onto the query element. Note that the original MSDAtt applies multi-head attention mechanism, while here we illustrate the single-head version for better understanding.

Deformable Transformer replaces the self-attention modules in the Transformer encoder and the cross-attention modules in the Transformer decoder with deformable attention modules, enabling a fast convergence rate and better representation ability in object detection.

3.2. Feature Encoding

To capture rich spatio-temporal features in a video, we first adopt a pre-trained action recognition network (*e.g.*, C3D [16], TSN [22]) to extract the frame-level features. We re-scale the feature map’s temporal dimension to a fixed number T by interpolation to facilitate batch processing. Then, to better utilize multi-scale features for predicting multi-scale events, we add L temporal convolutional layers (stride=2, kernel size=3) to get feature sequences across multiple resolutions, from T to $T/2^L$. The multi-scale frame features with their positional embedding [17] are fed into the deformable transformer encoder, extracting the frame-frame relationship across multiple scales. The output frame features are denoted as $\{\mathbf{f}^l\}_{l=1}^L$.

3.3. Parallel Decoding

The decoding network contains a deformable transformer decoder and three parallel heads, a captioning head for caption generation, a localization head to predict events’ boundaries with confidence scores, and an event counter to predict a suitable event number. The decoder aims to directly query the event-level features from frame features conditioned on N learnable embedding (termed event

queries) $\{\mathbf{q}_j\}_{j=1}^N$, and their corresponding scalar referent points p_j . Note that p_j is predicted by a linear projection with a sigmoid activation over \mathbf{q}_j . The event queries and reference points serve as the initial guess of the events’ features and locations (center points), and they will be refined iteratively at each decoding layer, as in [38]. The output query features and the reference point are denoted as $\tilde{\mathbf{q}}_j, \tilde{\mathbf{p}}_j$.

Localization head. Localization head performs box prediction and binary classification for each event query. Box prediction aims to predict the 2D relative offsets (center and length) of the ground-truth segment w.r.t. the reference point. Binary classification aims to generate the foreground confidence of each event query. Both box prediction and binary classification are implemented by multi-layer perceptron. After that, we obtain a set of tuples $\{t_j^s, t_j^e, c_j^{\text{loc}}\}_{j=1}^N$ to represent the detected events, where c_j^{loc} is the localization confidence of the event query $\tilde{\mathbf{q}}_j$.

Captioning head. We provide two captioning heads, a lightweight one and a standard one. The lightweight head simply feeds $\tilde{\mathbf{q}}_j$ into a vanilla LSTM at each timestamp. The word w_{jt} is predicted by an FC layer followed by a softmax activation over the hidden state \mathbf{h}_{jt} of LSTM.

However, the lightweight captioning head only receives an event-level representation $\tilde{\mathbf{q}}_j$, lacking the interactions between linguistic cues and frame features. Soft attention (SA) [26, 24, 9] is a widely-used module in video captioning, which can dynamically determine the importance of each frame when generating a word. Traditional two-stage methods [24, 9] aligns the event segments and their captions by restricting the attention area being within the event boundaries, but our captioning head can not access events’ boundaries, increasing the optimization difficulty to learn relationships between the linguistic word and frames. To alleviate this problem, we propose the deformable soft attention (DSA) to enforce the soft attention weights focus on a small area around the reference points. Specifically, when generating the t -th word w_t , we first generate K sampling points from each \mathbf{f}^l conditioned on both language query \mathbf{h}_{jt} and event query $\tilde{\mathbf{q}}_j$, following Eqn. 1, where \mathbf{h}_{jt} denotes the hidden state in LSTM. Then we consider $K \times L$ sampling points as the key/value and $[\mathbf{h}_{jt}, \tilde{\mathbf{q}}_j]$ as the query in soft attention. Since the sampling points are distributed around the reference point $\tilde{\mathbf{p}}_j$, the output features \mathbf{z}_{jt} of DSA are restricted to attend on a relatively small region. The LSTM takes as input the concatenation of the context features \mathbf{z}_{jt} , event query features $\tilde{\mathbf{q}}_j$ and previous words $w_{j,t-1}$. The probability for next word w_{jt} is obtained by an FC layer with softmax activation over \mathbf{h}_{jt} . With the evolving of LSTM, we obtain a sentence $S_j = \{w_{j1}, \dots, w_{jM_j}\}$, where M_j is the sentence length.

Event counter. Considering that an appropriate event number is an essential indicator for dense captioning quality:

too many events cause replicated captions and bad readability; too few detected events mean information missing and incomplete story. The event counter aims to detect the event number of the video. It contains a max-pooling layer and an FC layer with softmax activation, which first compress the most salient information of event queries $\{\tilde{q}_j\}$ to a global feature vector, and then predict a fix-size vector \mathbf{r}_{len} , where each value refers to the possibility of a specific number. During the inference stage, predicted event number is obtained by $N_{\text{set}} = \text{argmax}(\mathbf{r}_{\text{len}})$. The final outputs are obtained by selecting the top N_{set} events with accurate boundaries and good captions from N event queries. The confidence of each event query is calculated by:

$$c_j = c_j^{\text{loc}} + \frac{\mu}{M_j \gamma} \sum_{t=1}^{M_j} \log(c_{jt}^{\text{cap}}) \quad (2)$$

where c_{jt}^{cap} is the probability of the generated word. We observe that the average word confidence is not a convincing measurement of sentence-level confidence since the captioning head tends to produce overestimated confidence for short sentences. Thus, we add a modulation factor γ to rectify the influence of caption length. μ is the balance factor.

Set prediction loss. During training, PDVC produces a set of N events with their locations and captions. To match predicted events with ground truths in a global scheme, we use the Hungarian algorithm following [39] to find the best bipartite matching results. The matching cost is defined as $C = \alpha_{\text{giou}} L_{\text{giou}} + \alpha_{\text{cls}} L_{\text{cls}}$, where L_{giou} represents the generalized IOU [61] between predicted temporal segments and ground-truth segments, L_{cls} represents the focal loss [60] between the predicted classification score and the ground-truth label. The matched pairs are selected to calculate the set prediction loss, which is the weighted sum of gIOU loss, classification loss, countering loss, and caption loss:

$$L = \beta_{\text{giou}} L_{\text{giou}} + \beta_{\text{cls}} L_{\text{cls}} + \beta_{\text{ec}} L_{\text{ec}} + \beta_{\text{cap}} L_{\text{cap}} \quad (3)$$

where L_{cap} measures the cross-entropy between the predicted word probability and the ground truth normalized by the caption length, L_{ec} is also the cross-entropy loss between the predicted count distribution and the ground truth.

Note that we follow [39, 38] to add prediction heads to each layer of the transformer decoder. The final loss is the summation of the set prediction losses of all layers.

PDVC for paragraph captioning. Paragraph captioning [33, 42, 27] is a simplified version of dense video captioning, which focuses on generating a coherent paragraph and does not need to predict the temporal location of each sentence. PDVC can easily extend to paragraph captioning by removing the localization function and taking the pre-extracted proposals as input event queries. Specifically, we consider the linear embeddings of proposals' position as the event queries and use the proposals' center as reference points. Then PDVC is trained with caption loss only.

4. Experiments

4.1. Experimental Settings

Datasets. We use two large-scale benchmark datasets, ActivityNet Captions [5], and YouCook2 [30] to evaluate the effectiveness of the proposed PDVC. ActivityNet Captions contains 20k long untrimmed videos of various human activities. On average, each video lasts 120s and is annotated with 3.65 temporally-localized sentences. We follow the standard split with 10009/4925/5044 videos for training, validation, and test. YouCook2 has 2000 untrimmed videos of cooking procedures with an average duration of 320s. Each video has 7.7 annotated segments with associated sentences. We use the official split with 1333/457/210 videos for training, validation, and test.

Evaluation metrics. We evaluate our method in three aspects: 1) For localization performance, we use the average precision, average recall across IOU at $\{0.3, 0.5, 0.7, 0.9\}$ and their harmonic mean, F1 score. 2) For dense captioning performance, we follow the official evaluation tool provided by ActivityNet Challenge 2018, which calculates the average precision (measured by BLEU4 [11], METEOR [6], and CIDEr [18]) of the matched pairs between generated captions and the ground truth across IOU thresholds of $\{0.3, 0.5, 0.7, 0.9\}$. However, the official scorer does not consider the storytelling quality, *i.e.*, how well the generated captions can cover the video's whole story. We further adopt SODA_c [48] for an overall evaluation. 3) For paragraph captioning performance, we form a paragraph by sorting generated captions according to their starting time and report the paragraph-level captioning performance. Note that ActivityNet Captions has two sets of annotations for the validation set. For SODA_c, we evaluate it by two sets independently and report their average score.

Implementation details. For ActivityNet Captions, we use a C3D [16] pre-trained on Sports1M [52] to extract frame-level features. To fairly compare with state-of-the-art methods, we also test our model based on TSN [22] features provided by [31], and I3D+VGGish features provided by [36]. For YouCook2, we use the same TSN features as in [31].

We use a two-layer deformable transformer with multi-scale (4 levels) deformable attention. The deformable transformer uses a hidden size of 512 in MSDAtt layers and 2048 in feed-forward layers. The number of event queries is 10/100 for ActivityNet Captions/YouCook2. We implement a lightweight PDVC (termed PDVC.light) with the vanilla LSTM captioner and the standard PDVC with the LSTM-DSA captioner. The LSTM hidden dimension in captioning heads is 512. For the event counter, we choose the maximum count as 10/20 for ActivityNet Captions/YouCook2. In Eqn. 2, the length modulation factor γ is set to 2, and the trade-off ratio μ is set to 0.3/1.0 for PDVC.light/PDVC.

| Method | Recall | | | | | Precision | | | | | F1 |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 0.3 | 0.5 | 0.7 | 0.9 | avg | 0.3 | 0.5 | 0.7 | 0.9 | avg | |
| MFT [40] | 46.18 | 29.76 | 15.54 | 5.77 | 24.31 | 86.34 | 68.79 | 38.30 | 12.19 | 51.41 | 33.01 |
| SDVC [9] | 93.41 | 76.40 | 42.40 | 10.10 | 55.58 | 96.71 | 77.73 | 44.84 | 10.99 | 57.57 | 56.56 |
| PDVC_light | 88.78 | 71.74 | 45.70 | 17.45 | 55.92 | 96.83 | 78.01 | 41.05 | 14.69 | 57.65 | 56.77 |
| PDVC | 89.47 | 71.91 | 44.63 | 15.67 | 55.42 | 97.16 | 78.09 | 42.68 | 14.40 | 58.07 | 56.71 |

Table 1: Event localization on the ActivityNet Captions validation set.

| Method | Features | Ground-truth proposals | | | Predicted proposals | | | |
|-------------------|------------|------------------------|--------------|--------------|---------------------|-------------|--------------|-------------|
| | | B4 | M | C | B4 | M | C | SODA_c |
| DCE [5] | C3D | 1.60 | 8.88 | 25.12 | 0.17 | 5.69 | 12.43 | - |
| TDA-CG [24]* | C3D | - | 9.69 | - | 1.31 | 5.86 | 7.99 | - |
| DVC [7] | C3D | 1.62 | 10.33 | 25.24 | 0.73 | 6.93 | 12.61 | - |
| SDVC [9] | C3D | - | - | - | - | 6.92 | - | - |
| Efficient [37] | C3D | - | - | - | 1.35 | 6.21 | 13.82 | - |
| ECHR [34] | C3D | 1.96 | 10.58 | 39.73 | 1.29 | 7.19 | 14.71 | 3.22 |
| PDVC_light | C3D | 2.61 | 10.48 | 47.83 | 1.51 | 7.11 | 26.21 | 5.17 |
| PDVC | C3D | 2.64 | 10.54 | 47.26 | 1.65 | 7.50 | 25.87 | 5.26 |
| MT [31]* | TSN | 2.71 | 11.16 | 47.71 | 1.15 | 4.98 | 9.25 | - |
| PDVC | TSN | 3.07 | 11.27 | 52.53 | 1.78 | 7.96 | 28.96 | 5.44 |
| MDVC [35]*† | I3D+VGGish | 1.98 | 11.07 | 42.67 | 1.01 | 6.86 | 7.77 | - |
| BMT [36]*† | I3D+VGGish | 1.99 | 10.90 | 41.85 | 1.88 | 7.43 | 11.94 | - |
| PDVC † | I3D+VGGish | 3.12 | 11.26 | 53.65 | 1.96 | 8.08 | 28.59 | 5.42 |

Table 3: Dense captioning on the ActivityNet Captions validation set. B4/M/C is short for BLEU4/METEOR/CIDEr. * indicates results re-evaluated by the same evaluation toolkit. † means results with part of the dataset (9% videos missing).

The cost ratios in the bipartite matching are $\alpha_{\text{giou}}:\alpha_{\text{cls}}=2:1$ and the loss ratios are $\beta_{\text{giou}}:\beta_{\text{cls}}:\beta_{\text{ec}}:\beta_{\text{cap}}=2:1:1:1$. We use Adam [4] optimizer with an initial learning rate of $5e-5$ and the mini-batch size of 1 video.

4.2. Comparison with State-of-the-art Methods

Localization performance. The comparison of event localization quality is shown in Table 1. SDVC and MFT generate event proposals by a sophisticated “localize-select-describe” workflow. In contrast, PDVC removes the hand-crafted designs in the traditional proposal modules and directly outputs the proposals in a parallel manner, which is more efficient to deal with long sequences than recurrent counterparts. We surpass MFT by a large margin and achieve similar (slightly better) performance to SDVC, which shows the effectiveness of parallel set prediction in our method. Besides, the choice of the captioning head can slightly influence the balance of precision and recall.

Dense captioning performance. In Table 3, we list the performance of state-of-the-art models with cross-entropy training¹ on ActivityNet Captions. With ground-

¹A few methods [9, 34] incorporates Reinforcement Learning (RL) [14] after the cross-entropy training to further boost the performance. Note that we do not compare with these methods since RL training requires a

| Method | Predicted proposals | | | |
|-------------------|---------------------|-------------|--------------|-------------|
| | B4 | M | C | SODA_c |
| MT [31] | 0.30 | 3.18 | 6.10 | - |
| ECHR [34] | - | 3.82 | - | - |
| PDVC_light | 0.89 | 4.56 | 23.07 | 4.34 |
| PDVC | 0.80 | 4.74 | 22.71 | 4.42 |

Table 2: Dense captioning on YouCook2.

| Method | Features | B4 | M | C |
|-------------------------------|----------|--------------|--------------|--------------|
| Ground-truth proposals | | | | |
| HSE [41] | V | 9.84 | 13.78 | 18.78 |
| MART [42] | V+F | 10.33 | 15.68 | 23.42 |
| VTrans [31] | V+F | 9.75 | 15.64 | 22.16 |
| Trans-XL [43] | V+F | 10.39 | 15.09 | 21.67 |
| GVD [44] | V+F+O | 11.04 | 15.71 | 21.95 |
| GVDsup [44] | V+F+O | 11.30 | 16.41 | 22.94 |
| AdvInf [33] | V+F+O | 10.04 | 16.60 | 20.97 |
| PDVC | V+F | 11.80 | 15.93 | 27.27 |
| Predicted proposals | | | | |
| MFT [40] | V+F | 10.29 | 14.73 | 19.12 |
| PDVC | V+F | 10.24 | 15.80 | 20.45 |

Table 4: Paragraph captioning on the ActivityNet Captions *ae-val* set [44]. V/F/O refers to visual/flow/object features.

truth proposals, PDVC achieves considerable improvement over the state-of-the-art on BLEU4 and CIDEr, which shows a deformable transformer plus an LSTM captioner can give good caption quality. With predicted proposals, PDVC with C3D features achieves the best performance on BLEU4/METEOR/CIDEr/SODA_c, giving a 22.22%/4.31%/75.87%/63.35% relative improvement over state-of-the-art scores. We find that PDVC with ground-truth proposals does not show much superiority over ECHR on METEOR but surpasses ECHR with predicted proposals, indicating the generated proposals of PDVC are much better. Even with a lightweight LSTM as a captioner, PDVC_light can surpass most two-stage approaches on BLEU4/CIDEr/SODA_c. The reason mainly comes from the parallel decoding of the captioning head and localization head, which helps to generate proposals with high descriptiveness and discriminative internal representation.

Table 2 shows the dense captioning performance on the YouCook2 validation set. Our method achieve state-of-the-art performance with a considerable performance gain over other methods on all metrics.

more complex captioning network (e.g., Hierarchical RNN [27]) and extra-long training time, which is opposite to the design philosophy of PDVC. Moreover, RL training tends to produce longer sentences with repeated phrase [23], lowering the coherence and readability of generated captions.

Paragraph captioning performance. Table 4 shows the comparison between PDVC and state-of-the-art paragraph captioning methods. With ground-truth proposals, PDVC with a deformable transformer plus an attention-based LSTM can surpass several transformer-based captioning models, like MART, VTrans, and Trans-XL, indicating the strong representation ability of deformable attention in the encoder-decoder and the LSTM-DSA. It is promising for PDVC to get a further performance boost by incorporating a transformer captioner. We leave this for future work.

Even with predicted proposals, we observe PDVC has a comparable performance with previous methods with ground-truth proposals, indicating that query features contain rich information covering main segments in videos. While most previous paragraph captioning methods require ground-truth annotation at testing, our model reduces the captioning module’s dependence on accurate proposals by parallel decoding, raising the possibility of generating good paragraphs without human annotations of the test video.

Efficiency. We compare the inference time of PDVC against two-stage methods TDA-CG [24], MT [31] under the same hardware environment in Table 5. Our methods are more efficient since: 1) Only a few event proposals with their captions are predicted in parallel; 2) We do not require a dense-to-sparse proposal selection like NMS; 3) MSDAtt is an efficient operator due to the sparse sampling.

4.3. Interaction between Localization & Captioning

In this part, we go deeper into the mutual effect between two subtasks. It is straightforward that localization can aid captioning since the localization supervision guides the query features to specific ground-truth regions, which contains rich semantics matching the target captions. Therefore, we focus on how captioning task affects proposals’ quality, which is less explored in the previous literature.

Captioning supervision helps generate proposals with descriptiveness. To better study the quality of proposals generated by PDVC, we use the same pre-trained event captioning model [26] to evaluate the descriptiveness of generated proposals of different models. We also reimplemented two mainstream proposal generation modules SST and SST+ESGN for comparison. Both SST and SST+ESGN are trained with localization loss only, while PDVC is trained with both localization and captioning loss. As shown in Table 6, PDVC achieves a slightly lower F1 score but the best descriptiveness score among the four models.

We match each generated proposal with one ground-truth segment with the highest overlap. Fig. 3 demonstrates the statistics of matching results. Surprisingly, incorporating caption supervision yields a considerable boost in caption quality of high-precision proposals (*i.e.*, $IOU > 0.9$). The reason may be that the captioning head is trained based

| Method | BAF-CG [24] | MT [31] | PDVC _{light} | PDVC |
|------------|-------------|---------|-----------------------|-------------|
| Time(secs) | 2.39 | 2.05 | 0.09 | 0.16 |

Table 5: Inference speed. We report average inference time (secs/video) of 100 sampled videos with a single Tesla V100 GPU.

| Method | Loss | #p | Rec. | Pre. | F1 | B4 | M | C |
|---------------------------|-----------|------|-------|-------|-------|-------------|-------------|--------------|
| SST [1] | loc. | 3.00 | 42.00 | 60.99 | 49.74 | 0.98 | 6.70 | 17.34 |
| SST+ESGN [9] | loc. | 2.79 | 53.80 | 61.37 | 57.33 | 1.09 | 6.80 | 19.67 |
| Ours _{loc} _only | loc. | 3.26 | 56.35 | 58.69 | 57.49 | 0.98 | 6.71 | 19.36 |
| Ours (PDVC) | loc.+cap. | 3.03 | 55.42 | 58.07 | 56.71 | 1.24 | 7.03 | 21.91 |

Table 6: Proposal quality with different loss types. Rec./Pre./F1 measures the localization performance, while B4/M/C measures dense captioning performance. #p is the number of proposals.

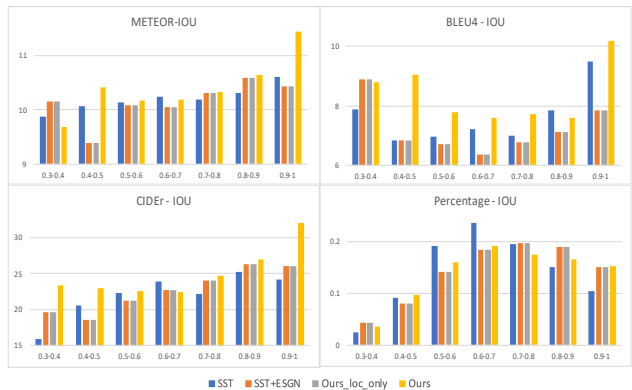


Figure 3: Caption quality vs. IOU. We omit the pairs with $IOU < 0.3$ (less than 2% of all pairs).

on event query features corresponding to accurate proposals, so PDVC learns to enhance the descriptiveness of the high-precision proposals. The last subfigure shows the IOU distribution of matched pairs. Most proposals produced by SST are not very accurate (mainly with $0.5 < IOU < 0.8$). When further incorporating ESGN for adaptively proposal selection, the majority of proposals are with $0.6 < IOU < 0.9$. Ours and Ours_{loc}_only achieve a similar IOU distribution to SST+ESGN, but do not introduce any hand-crafted components like anchor generation and NMS.

Generally speaking, descriptiveness is positively correlated to the precision of proposals with an ideal captioner. However, the performance of existing captioners is still far from satisfactory, which means they generate wrong or boring captions for some proposals. To reduce improper captions of the final results, it is essential to generate not only location-accurate but caption-aware proposals. Our model provides an effective solution to explore the mutual benefits between localization and captioning by parallel decoding.

Captioning supervision helps learn location-aware features. Another advantage of parallel decoding is that we can directly remove the localization head to study the behavior of captioning head. We train an event proposal generation module based on merely captioning supervision, by




Figure 4: The distribution of predicted proposals without localization supervision. We plot the predicted proposals of 200 randomly sampled videos in the YouCook2 validation set. Horizontal and vertical axes represent the re-scaled center position and re-scaled length of proposals, respectively. Each sub-figure contains 30 clusters with different colors corresponds to 30 input event queries. R and P refer to recall and precision of the 30 generated proposals, respectively.

| Transformer | | Captioning head | | M | SODA_c | #q | counter | Rec. | Pre. | M | SODA_c | γ | Rec. | Pre. | M | SODA_c |
|-------------|------------|-----------------|--------|------|--------|-----|---------|-------|-------|------|--------|----------|-------|-------|------|--------|
| Vanilla | Deformable | LSTM | SA DSA | | | | | | | | | | | | | |
| | | | | 6.10 | 3.06 | 5 | ✓ | 57.46 | 57.10 | 6.96 | 5.02 | 0.0 | 48.67 | 47.35 | 5.78 | 5.23 |
| ✓ | | ✓ | | 7.11 | 5.17 | 10 | ✓ | 55.92 | 57.65 | 7.11 | 5.17 | 0.5 | 50.63 | 50.08 | 6.59 | 5.25 |
| | ✓ | ✓ | | 6.15 | 3.40 | 30 | ✓ | 53.35 | 59.08 | 7.18 | 4.90 | 1.0 | 51.52 | 53.31 | 7.35 | 5.02 |
| | ✓ | ✓ | ✓ | 7.50 | 5.26 | 100 | ✓ | 51.88 | 59.27 | 7.33 | 4.59 | 2.0 | 55.92 | 57.65 | 7.11 | 5.17 |
| | ✓ | ✓ | ✓ | | | 10 | × | 77.67 | 44.88 | 6.62 | 4.30 | 3.0 | 55.72 | 57.87 | 6.90 | 5.19 |

(a) Ablating the deformable operations

(b) Varying query number & event counter

(c) Varying γ

Table 7: Ablation studies on the ActivityNet Captions validation set. Subfigure (b) and (c) are based on PDVC.light.

making some modifications to the original PDVC to stabilize training, such as fixing the sampling offsets in the decoder and using the captioning cost in bipartite matching. More details can be found in the supplementary material. After iterative refinement in the decoder, we directly regard the reference points corresponding to event queries in the last decoder layer as event proposals. Fig. 4 shows the position distribution of predicted proposals on YouCook2. We also report quantitative results such as recall and precision.

As the training epoch increases, proposals’ center tends to spread uniformly, and the proposals’ length tends to focus on a relatively small value. Though a noticeable gap exists between the distributions of predicted proposals and ground-truth proposals, we see that the predicted proposals are gradually approaching the ground truth during training. The recall/precision at epoch 13 is 30.32/14.24, which is better than that at initialization (24.14/13.56). Based on the above findings, we argue that our method can implicitly capture the location-aware features from caption supervision, helping the optimization of the event localization.

4.4. Ablation Studies

Deformable components. As shown in Table 7a, when removing deformable operations from deformable transformer or LSTM-DSA, the performance degrades considerably. We conclude that: 1) Adding locality into transformer helps to extract temporally-sensitive features for localization-aware tasks; 2) Focusing on a small segment around the proposals rather than the whole video helps the optimization of the event captioning.

Query number & event counter. As shown in Table 7b,

only a few queries are sufficient for good performance. Too many queries lead to high precision and METEOR, but low Recall and SODA_c. We choose an appropriate query number for striking a balance of recall and precision. The final event number also controls the balance of precision and recall. The event counter can predict a reasonable number of event instances, making the generated captions reveal a whole story in the video.

Length modulation. Table 7c shows that modulating the caption length ($\gamma > 1$) obtains a better trade-off between METEOR & SODA_c and Precision & Recall than averaging ($\gamma = 1$) or summing ($\gamma = 0$) the word scores.

5. Conclusion

This paper presents PDVC, an end-to-end dense video captioning framework with parallel decoding, which formulates dense video captioning as a set prediction task. PDVC directly produces a set of temporally-localized sentences without a dense-to-sparse proposal generation and selection process, significantly simplifying the traditional “localize-then-describe” pipeline. Prediction heads for event localization and event captioning run in parallel to exploit the inter-task mutual benefits. Experiments on two benchmark datasets show that PDVC can generate high-quality captions and surpass state-of-the-art methods.

Acknowledgements. This work was supported by the National Natural Science Foundation of China No. 61972188, 61903178, 61906081, and U20A20306, the General Research Fund of Hong Kong No. 27208720, and the Program for Guangdong Introducing Innovative and Entrepreneurial Teams No. 2017ZT07X386.

References

- [1] S. Buch, V. Escorcia, C. Shen, B. Ghanem, and J. C. Niebles, "SST: Single-stream temporal action proposals," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6373-6382. [7](#)
- [2] L. Gao, Z. Guo, H. Zhang, X. Xu, and H. T. Shen, "Video captioning with attention-based LSTM and semantic consistency," *IEEE Trans. Multimedia*, vol. 19, no. 9, pp. 2045-2055, Sep. 2017. [1](#)
- [3] R. Girdhar, J. Carreira, C. Doersch, and A. Zisserman, "Video action transformer network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 244-253. [3](#)
- [4] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1-13. [6](#)
- [5] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, "Dense-captioning events in videos," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 706-715. [1](#), [2](#), [3](#), [5](#), [6](#)
- [6] A. Lavie and A. Agarwal, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *Proc. ACL Workshop Intrinsic Extrinsic Eval. Measures Mach. Transl. Summarization*, 2005, pp. 65-72. [5](#)
- [7] Y. Li, T. Yao, Y. Pan, H. Chao, and T. Mei, "Jointly localizing and describing events for dense video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7492-7500. [1](#), [2](#), [3](#), [6](#)
- [8] M. Marszalek, I. Laptev, and C. Schmid, "Actions in context," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 2929-2936.
- [9] J. Mun, L. Yang, Z. Ren, N. Xu, and B. Han, "Streamlined dense video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 6588-6597. [2](#), [3](#), [4](#), [6](#), [7](#)
- [10] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui, "Jointly modeling embedding and translation to bridge video and language," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4594-4602. [1](#)
- [11] K. Papineni, S. Roukos, T. Ward, and W. Zhu, "BLEU: a method for automatic evaluation of machine translation," in *Proc. Annu. Meet. Assoc. Comput. Linguistics*, 2002, pp. 311-318. [5](#)
- [12] W. Pei, J. Zhang, X. Wang, L. Ke, X. Shen and Y. W. Tai, "Memory-attended recurrent network for video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8347-8356. [1](#)
- [13] M. Qi, Y. Wang, A. Li, and J. Luo, "Sports video captioning via attentive motion representation and group relationship modeling," *IEEE Trans. Circuits Syst. Video Technol.*, early access, doi: 10.1109/TCSVT.2019.2921655. [1](#)
- [14] S. J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, and V. Goel, "Self-critical sequence training for image captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 7008-7024. [6](#)
- [15] M. Rohrbach, W. Qiu, I. Titov, S. Thater, M. Pinkal, and B. Schiele, "Translating video content to natural language descriptions," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 433-440. [1](#)
- [16] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3D convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 4489-4497. [4](#), [5](#)
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Proc. Sys.*, 2017, pp. 5998-6008. [3](#), [4](#)
- [18] R. Vedantam, C. L. Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 4566-4575. [5](#)
- [19] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, "Translating videos to natural language using deep recurrent neural networks," in *Proc. North Amer. Chapter Assoc. Comput. Linguistics-Human Language Technol.*, 2015, pp. 1494-1504. [1](#)
- [20] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko, "Sequence to sequence-video to text," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 4534-4542. [1](#)
- [21] B. Wang, L. Ma, W. Zhang, and W. Liu, "Reconstruction network for video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7622-7631. [1](#)
- [22] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. V. Gool, "Temporal Segment Networks for Action Recognition in Videos," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 11, pp. 2740-2755, Sep. 2019. [4](#), [5](#)
- [23] Q. Wang and A. B. Chan, "Describing like humans: On diversity in image captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4195-4203. [6](#)
- [24] J. Wang, W. Jiang, L. Ma, W. Liu, and Y. Xu, "Bidirectional attentive fusion with context gating for dense video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7190-7198. [2](#), [3](#), [4](#), [6](#), [7](#)
- [25] H. Xu, B. Li, V. Ramanishka, L. Sigal, and K. Saenko, "Joint event detection and description in continuous video streams," in *Proc. IEEE Winter Conf. Applicat. Comput. Vis.*, 2019, pp. 396-405. [1](#)
- [26] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville, "Describing videos by exploiting temporal structure," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 4507-4515. [4](#), [7](#)
- [27] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu, "Video paragraph captioning using hierarchical recurrent neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4584-4593. [5](#), [6](#)
- [28] D. Yang and C. Yuan, "Hierarchical context encoding for events captioning in videos," in *Proc. IEEE Int. Conf. Image Process.*, 2018, pp. 1288-1292. [2](#)

- [29] Z. Zhang, D. Xu, W. Ouyang, and C. Tan, "Show, tell and summarize: Dense video captioning using visual cue aided sentence summarization," *IEEE Trans. Circuits Syst. Video Technol.*, 2019, early access, doi: 10.1109/TCSVT.2019.2936526. **1**
- [30] L. Zhou, C. Xu, and J. J. Corso, "Towards automatic learning of procedures from web instructional videos," in *Proc. AAAI Conf. Artificial Intell.*, 2018, pp. 7590-7598. **5**
- [31] L. Zhou, Y. Zhou, J. J. Corso, R. Socher, and C. Xiong, "End-to-end dense video captioning with masked transformer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8739-8748. **2, 3, 5, 6, 7**
- [32] S. Chen, J. Chen, Q. Jin, and A. Hauptmann, "Video captioning with guidance of multimodal latent topics," in *Proc. ACM Multimedia*, 2017, pp. 1838-1846. **1**
- [33] J. S. Park, M. Rohrbach, T. Darrell, and A. Rohrbach, "Adversarial inference for multi-sentence video description," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 6598-6608. **5, 6**
- [34] T. Wang and H. Zheng and M. Yu and Q. Tian and H. Hu, "Event-centric hierarchical representation for dense video captioning," in *IEEE Trans. Circuits Syst. Video Technol.*, vol. 21, no. 5, pp. 1890-1900, 2020. **2, 6**
- [35] V. Iashin and E. Rahtu, "Multi-modal dense video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 4117-4126. **2, 6**
- [36] V. Iashin and E. Rahtu, "A better use of audio-visual cues: Dense video captioning with bi-modal transformer," in *Proc. British Mach. Vis. Conf.*, 2020. **2, 5, 6**
- [37] M. Suin and A. N. Rajagopalan, "An efficient framework for dense video captioning," in *Proc. AAAI Conf. Artificial Intell.*, 2020, pp. 12039-12046. **1, 6**
- [38] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable detr: Deformable transformers for end-to-end object detection," in *Proc. Int. Conf. Learn. Represent.*, 2021. **3, 4, 5**
- [39] N. Carion, M. Francisco, S. Gabriel, U. Nicolas, K. Alexander and Z. Sergey, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 213-229. **3, 5**
- [40] Y. Xiong, B. Dai, and D. Lin, "Move forward and tell: A progressive generator of video descriptions," in *Proc. Eur. Conf. Comput. Vis.*, 2018. **6**
- [41] B. Zhang, H. Hu, and F. Sha, "Crossmodal and hierarchical modeling of video and text," in *Proc. Eur. Conf. Comput. Vis.*, 2018. **6**
- [42] J. Lei, L. Wang, Y. Shen, D. Yu, T. L. Berg and M. Bansal, "Mart: Memory-augmented recurrent transformer for coherent video paragraph captioning," in *Proc. ACL*, 2020. **5, 6**
- [43] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," in *Proc. ACL*, 2019. **6**
- [44] L. Zhou, Y. Kalantidis, X. Chen, J. J. Corso, and M. Rohrbach, "Grounded video description," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019. **6**
- [45] H. Chen, Y. Wang, T. Guo, C. Xu, Y. Deng, Z. Liu, S. Ma, C. Xu, C. Xu, and W. Gao, "Pre-trained image processing transformer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 12299-12310. **3**
- [46] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent.*, 2021. **3**
- [47] V. Escorcia, F. C. Heilbron, J. C. Niebles, and B. Ghanem, "Daps: Deep action proposals for action understanding," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 768-784. **2**
- [48] S. Fujita, T. Hirao, H. Kamigaito, M. Okumura, and M. Nagata, "Soda: Story oriented dense video captioning evaluation framework," in *Proc. Eur. Conf. Comput. Vis.*, 2020. **2, 5**
- [49] J. Gao, Z. Yang, K. Chen, C. Sun, and R. Nevatia, "Turn tap: Temporal unit regression network for temporal action proposals," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3628-3636. **2**
- [50] R. Girdhar, J. Carreira, C. Doersch, and A. Zisserman, "Video action transformer network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 244-253. **3**
- [51] F. C. Heilbron, J. C. Niebles, and B. Ghanem, "Fast temporal activity proposals for efficient detection of human actions in untrimmed videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1914-1923. **2**
- [52] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1725-1732. **5**
- [53] C. Lin, J. Li, Y. Wang, Y. Tai, D. Luo, Z. Cui, C. Wang, J. Li, F. Huang, and R. Ji, "Fast learning of temporal action proposal via dense boundary generator," in *Proc. AAAI Conf. Artificial Intell.*, 2020, pp. 11499-11506. **2**
- [54] T. Lin, X. Liu, X. Li, E. Ding, and S. Wen, "Bmn: Boundary-matching network for temporal action proposal generation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 3889-3898. **2**
- [55] T. Lin, X. Zhao, H. Su, C. Wang, and M. Yang, "Bsn: Boundary sensitive network for temporal action proposal generation," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 3-19. **2**
- [56] Z. Shou, D. Wang, and S. Chang, "Temporal action localization in untrimmed videos via multi-stage cnns," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1049-1058. **2**
- [57] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 10347-10357. **3**
- [58] F. Yang, H. Yang, J. Fu, H. Lu, and B. Guo, "Learning texture transformer network for image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 5791-5800. **3**

- [59] Y. Zhao, Y. Xiong, L. Wang, Z. Wu, X. Tang, and D. Lin, "Temporal action detection with structured segment networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2914-2923. [2](#)
- [60] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017. [5](#)
- [61] H. Rezaatofghi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, and S. Savarese, "Generalized intersection over union: A metric and a loss for bounding box regression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 658-666. [5](#)