

Joint Event Detection and Description in Continuous Video Streams

Huijuan Xu¹ Boyang Li² Vasili Ramanishka¹ Leonid Sigal³ Kate Saenko¹

¹Boston University ²Baidu Research ³University of British Columbia

¹{hxu, vram, saenko}@bu.edu, ²boyangli@baidu.com, ³lsigal@cs.ubc.ca

Abstract

Dense video captioning is a fine-grained video understanding task that involves two sub-problems: localizing distinct events in a long video stream, and generating captions for the localized events. We propose the Joint Event Detection and Description Network (JEDDi-Net), which solves the dense video captioning task in an end-to-end fashion. Our model continuously encodes the input video stream with three-dimensional convolutional layers, proposes variable-length temporal events based on pooled features, and generates their captions. Proposal features are extracted within each proposal segment through 3D Segment-of-Interest pooling from shared video feature encoding. In order to explicitly model temporal relationships between visual events and their captions in a single video, we also propose a two-level hierarchical captioning module that keeps track of context. On the large-scale ActivityNet Captions dataset, JEDDi-Net demonstrates improved results as measured by standard metrics. We also present the first dense captioning results on the TACoS-MultiLevel dataset.

1. Introduction

The goal of automatic video description is to tell a story about events happening in a video. While early video description methods produced captions for short clips that were manually segmented to contain a single event of interest [2, 24], more recently *dense video captioning* [13] has been proposed to both segment distinct events in time and describe them in a series of coherent sentences. Figure 1 shows an example of this task for a weight-lifting video. This problem is a generalization of dense image region captioning [11, 30] and has many practical applications, such as generating textual summaries for the visually impaired, or detecting and describing important events in surveillance footage.

There are several key challenges in dense video captioning: accurately detecting the start and end of each event,

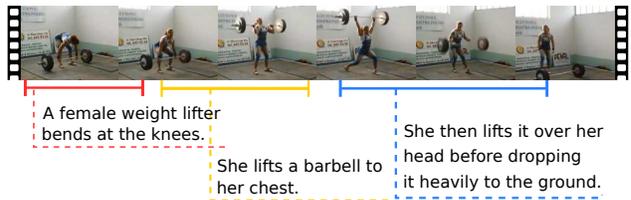


Figure 1. An example (from ActivityNet Captions [13]) of the challenges posed by the dense video captioning task. A successful model must detect the time window of each event, which significantly affects the content of predicted captions. The sequential relationship between the three activities in weight lifting suggests that visual and language contexts play a crucial role in this task.

recognizing the type of activity and objects involved, and translating this knowledge into a fluent natural language sentence. The context of the past and future sentences must also be taken into account to construct coherent stories. In [13], the authors proposed using two sets of recurrent neural networks (RNNs). The *proposal* RNN encodes convolutional features from input frames and proposes the start and end time of temporal activity segments. The separate two-layer captioning RNN receives the state vector of each activity proposal and decodes it into a sentence.

One issue with the existing approach [13] is that using the accumulated state vector of the proposal RNN to represent the visual content of a proposed segment may be inaccurate. Each state vector of the proposal RNN is used to predict a set of variable length temporal proposals, while this set of proposals use the same RNN state vector as proposal feature representation. Instead, we want to more precisely capture the activity feature by considering only the frames within that temporal segment. Another problem is that the temporal segmentation (i.e., proposal generation) stage and the caption generation stage are separately trained. As a result, errors in sentence prediction cannot be propagated back to temporal proposal generation. However, consider Figure 1: if the temporal proposal for the sentence “*she then lifts it... before dropping it...*” is shortened by a small amount, it would miss the drop part of the activity, resulting in a wrong caption.

In this work, we present a new approach to dense video captioning, the *Joint Event Detection and Description Network (JEDDi-Net)*. Our model utilizes three-dimensional convolution to extract video appearance and motion features, which are sequentially passed to the temporal event proposal network and the captioning network. Notably, the entire network is end-to-end trainable, with feature computation and temporal segmentation directly influencing captioning loss. For proposal generation and refinement, we adapt the proposal network introduced by the Region Convolutional 3D Network (R-C3D) model [29] for activity class detection. The proposal network uses 3D convolutional layers to encode the entire input video buffer and proposes variable-length temporal segments as potential activities. Spatio-temporal features are extracted for proposals using 3D Segment-of-Interest (SoI) pooling from the same convolutional feature maps shared by the proposal stage. The resulting proposal features are passed along to the captioning module. We expect to obtain more semantically accurate captioning using this proposal representation, as compared to using the accumulated RNN state representation for a set of proposals [13].

Our JEDDi-Net also uses a hierarchical recurrent caption generator: the low-level *captioner* RNN generates a sentence based on the current proposal’s features and on the context that is provided by the high-level *controller* RNN. The captioning model in [13] also provided context to its sentence generation LSTM module, in the form of visual features from the past and future weighted by their correlation with the current proposal’s features. However, the decoded sentences of preceding proposals may also provide useful context information for decoding the current one. Thus, inspired by [34, 12], our proposed hierarchical RNN captioning module incorporates both visual and linguistic context. The high-level controller RNN accumulates context from visual features and sentences generated so far, and provides it to the low-level sentence captioning module, which generates the new sentence for the target video segment.

Contributions: JEDDi-Net can efficiently detect and describe events in long sequences of frames, including overlapping events of both long and short duration. We summarize the key contributions of our paper as follows: 1) an end-to-end model for the dense video captioning task which jointly detects events and generates their descriptions (code is released for public use¹); 2) a novel hierarchical language model that incorporates the visual and language context for each new caption and considers the relationships between events in the video; 3) a large-scale evaluation showing improved results on the ActivityNet Captions dataset [13], as well as the first dense video captioning results on the

TACoS-MultiLevel dataset [18].

2. Related Work

Activity Detection in Videos: Over the past few years, the video activity understanding task has quickly evolved from trimmed video classification [10, 17, 21, 27] to activity detection in untrimmed video, as most real-life videos are not nicely segmented and contain multiple activities. There are two types of activity detection tasks: spatio-temporal and temporal-only. Spatio-temporal activity detection [28, 33] localizes activities within spatio-temporal tubes and requires heavier annotation work to collect the training data, while temporal activity detection [3, 15, 16, 20, 22, 32] only predicts the start and end times of the activities within long untrimmed videos and classifies the overall activity without spatially localizing people and objects in the frame. Several language tasks related to activity detection have recently emerged in the literature, including the dense video captioning task, which provides detailed captions for temporally localized events [13], and the task of language-based event localization in videos [5, 9].

Our model includes a temporal activity proposal module which is inspired by the proposal network introduced by the Region Convolutional 3D Network (R-C3D) model [29] for activity class detection. Instead of employing sliding windows [20, 6] or RNN feature encodings [3, 15, 16, 22, 32, 1] to generate temporal proposals, we encode the input video segment with a fully-convolutional 3D ConvNet and use 3D SoI pooling to allow feature extraction at arbitrary proposal granularities, achieving significantly higher detection accuracy and providing better proposal features for decoding captions. Computation is saved by using 3D SoI pooling to extract proposal features from the shared convolutional feature encoding of the entire input buffer, compared to sliding window approaches which re-extract features for each window from raw input frames.

Video Captioning: Early video captioning models (e.g., [8]) generated a single caption for a trimmed video clip by first predicting the subject, verb and object in the video and then inserting them into a template sentence. More recent deep models have achieved significantly better trimmed captioning results by using RNNs/LSTMs for language modeling conditioned on CNN features [25, 24]. Attention mechanisms have also been incorporated into RNNs to choose more relevant visual features for decoding captions [31].

The video paragraph captioning task [34] has also been proposed to provide multiple detailed sentence descriptions for long video segments. In contrast to our dense captioning task, video paragraph captioning produces no temporal localization of sentences. [34] proposed a hierarchical

¹Code available: <https://github.com/VisionLearningGroup/JEDDi-Net>

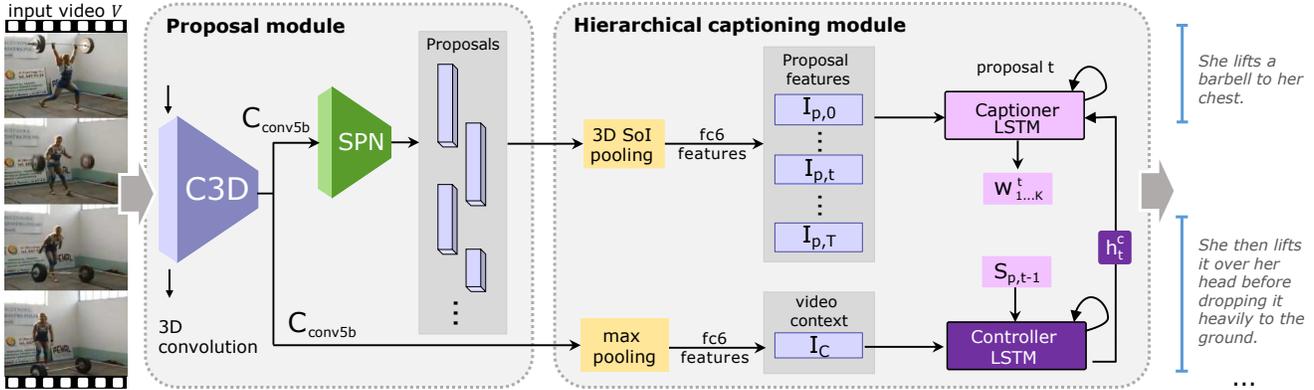


Figure 2. The overall architecture of our proposed *Joint Event Detection and Description Network (JEDDi-Net)* consists of two modules. The proposal module (Sec. 3.1) extracts features with 3D convolutional layers (C3D) and uses a Segment Proposal Network (SPN) to generate candidate segment proposals (see Fig. 3 for details). The hierarchical captioning module (Sec. 3.2) contains a controller LSTM to fuse the visual context I_c and the decoded language context $S_{p,t-1}$, and provides its hidden state h_t^c to the captioner LSTM, which decodes the next sentence. Details of LSTMs are in Fig. 4.

RNN to model the language histories when decoding multiple sentences for the video paragraph captioning task, but without explicit visual context modelling. A hierarchical RNN was also applied to image paragraph captioning [12]. However, only the visual context was recorded in the high-level controller layer, and no language history was fed into the controller. Hierarchical models have also been applied to natural language processing [19], with [14] proposing a hierarchical RNN language model that integrates sentence history to improve the coherence of documents. [13] introduced the dense-captioning task on an ActivityNet-based dataset, and modeled context using attention over past and future visual features. In this paper, we design a hierarchical captioning module which considers both the visual and language context of the video segment. Also, in contrast to [13], our proposal and captioning modules are jointly trained, with the captioning errors back-propagated to further improve the proposal features and boundaries.

3. Approach

Overview: Figure 2 provides an overview of our proposed JEDDi-Net model. We assume training data in the form of a video V , which contains a number of ground truth segments. For each segment, we have its center position c^* and length l^* as well as the words in its caption $\{w_k\}_{k=1\dots K}$. The model consists of two main components: a segment proposal module and a captioning module.

The *Proposal Module* encodes all input frames in V using a 3D convolutional network (C3D). Based on the features obtained from the layer *conv5b*, the Segment Proposal Network (SPN) proposes temporal segments, classifies them as either potential events for captioning or back-

ground, and regresses their temporal boundaries. The C3D features C_{conv5b} for the video V are also encoded via max-pooling as video context I_c , which is utilized in captioning.

The *Hierarchical Captioning Module* generates a caption for the t^{th} proposal, $t = 1, \dots, T$. This module is composed of a caption-level controller network and a word-level sentence decoder network, both implemented with LSTMs. The controller network takes the video context vector I_c and the encoding of the previous sentence $S_{p,t-1}$ and provides a single context vector h_t^c as a summary of both visual and linguistic context. The word-level decoder network takes as input the current proposal’s features $I_{p,t}$ and the context vector h_t^c and generates the words w_k^t one by one. The entire network is trained end-to-end with three jointly optimized loss functions, including the proposal classification loss, the regression loss on the proposal’s center and length, and cross-entropy loss for word prediction. Secs. 3.1 and 3.2 introduce the segment proposal and the hierarchical captioning modules, and Sec. 3.3 explains the end-to-end optimization strategy.

3.1. Proposal Module

Video Feature Representation: The feature encoding of the input video should extract semantic appearance and dynamic features and preserve temporal sequence information. We employ the C3D architecture [23] to encode the input frames in a fully-convolutional manner. C3D consists of eight convolutional layers (from *conv1a* to *conv5b*). Convolution and pooling in spatiotemporal space allows us to retain temporal sequence information within the input video. We represent the sequence of \mathcal{L} RGB video frames of height \mathcal{H} and width \mathcal{W} as $V \in \mathbb{R}^{3 \times \mathcal{L} \times \mathcal{H} \times \mathcal{W}}$. The

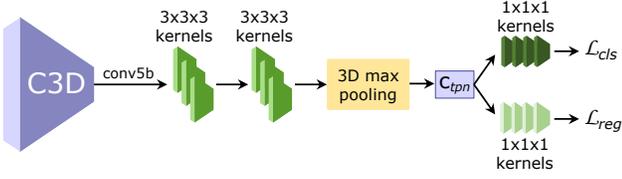


Figure 3. The structure of the Segment Proposal Network. (Sec. 3.1)

C3D convolutions encode V into feature maps $C_{conv5b} \in \mathbb{R}^{512 \times \frac{L}{8} \times \frac{H}{16} \times \frac{W}{16}}$ (512 is the channel dimension of the layer conv5b). These feature maps are used to produce the proposal features $I_{p,t}$ and video-level visual context I_C .

Segment Proposal Network (SPN): In this step, we predict the activity proposals’ start and end times. The accuracy of the proposals’ boundary will affect the proposal feature encoding, and will further affect the decoded captions, especially for short activities. To obtain feature vectors $C_{tpn} \in \mathbb{R}^{512 \times \frac{L}{8} \times 1 \times 1}$ for predicting proposals at each of $L/8$ time points, we add two 3D convolutional filters with kernel size $3 \times 3 \times 3$ on top of C_{conv5b} , followed by a 3D max-pooling filter to remove the spatial dimension. Proposed segments are predicted around a set of anchor segments [29]. Based on the 512-dimensional feature vector at each temporal location in C_{tpn} , we predict a relative offset $\{\delta c_i, \delta l_i\}$ to the center location and the length of each anchor segment $\{c_i, l_i\}_{i=1 \dots R}$, as well as a binary label indicating whether the predicted proposal contains an activity or not. This is achieved by adding two $1 \times 1 \times 1$ convolutional layers on top of C_{tpn} . A detailed diagram of the Segment Proposal Network (SPN) is shown in Figure 3.

Training: To train the binary proposal classifier in the segment proposal network, we need a training set with positive and negative examples. Only positive examples contribute to the proposal regression loss. The ground truth segments’ center location and length are transformed with respect to the positive anchor segments using Eq (1). We assign an anchor segment a positive label if it 1) overlaps with some ground-truth activity with temporal Intersection-over-Union (tIoU) higher than 0.7, or 2) has the highest tIoU overlap with some ground-truth activity. If the anchor has tIoU overlap lower than 0.3 with all ground-truth activities, then it is given a negative label. All others are held out from training. We sample balanced batches with a positive/negative ratio of 1 : 1.

We train the SPN network by jointly optimizing both the binary proposal classification and proposal boundary regression. For the i^{th} anchor segment, $\{c_i, l_i\}$ denotes the center and the length of the segment and \hat{a}_i denotes the predicted probability. The corresponding ground truth la-

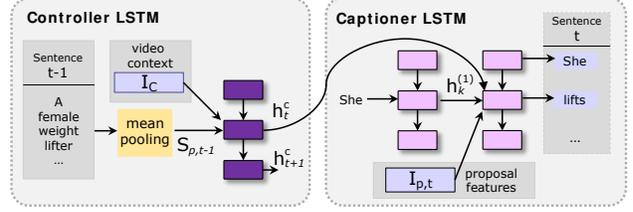


Figure 4. The structure of LSTMs in the hierarchical captioning module. (Sec. 3.2)

bels are a_i^* , c_i^* , and l_i^* . Ground truth segments are transformed with respect to positive anchor segments following the equations below:

$$\begin{aligned} \delta c_i^* &= (c_i^* - c_i) / l_i \\ \delta l_i^* &= \log(l_i^* / l_i). \end{aligned} \quad (1)$$

SPN predicts the offset $\delta \hat{c}_i$ and $\delta \hat{l}_i$. The cross-entropy loss, denoted as \mathcal{L}_{cls} , is used for binary proposal classification. The smooth L1 loss [7], \mathcal{L}_{reg} , is used for proposal boundary regression and defined as

$$\mathcal{L}_{reg}(x) = 1(|x| < 1)0.5x^2 + 1(|x| \geq 1)(|x| - 0.5) \quad (2)$$

where $1(\cdot)$ is the indicator function. The joint loss function is given by

$$\mathcal{L}_{spn} = \frac{1}{M} \sum_{i=1}^M \mathcal{L}_{cls}(\hat{a}_i, a_i^*) + a_i^* \left(\mathcal{L}_{reg}(\delta \hat{c}_i - \delta c_i^*) + \mathcal{L}_{reg}(\delta \hat{l}_i - \delta l_i^*) \right) \quad (3)$$

where M stands for the number of sampled proposals in the training batch.

At test time, we perform the inverse transformation of Eq (1) to find the center and length of predicted proposals. Then, the proposals are refined via Non-Maximum Suppression (NMS) with a tIoU threshold of 0.7.

3.2. The Hierarchical Captioning Module

Proposal Feature Encoding: To compute a visual representation of each proposed event for the captioning module, we encode predicted proposals into features $I_{p,t}$. In order to encode variable-length proposals, we adopt 3D SoI Pooling, which divides the shared feature map C_{conv5b} equally into bins, performs max-pooling within each bin, and further feeds it through the fc6 layer of the C3D network [23]. To represent visual context, we encode the entire input video segment V as a vector I_C using a max pooling layer and the shared fc6 layer.

Controller LSTM: To model context between the generated caption sentences, we adopt a hierarchical LSTM structure. The high-level Controller LSTM encodes the visual context and sentence decoding history. The low-level

Captioning LSTM decodes every proposal into a caption word by word, while being aware of visual and language context. Figure 4 illustrates this hierarchical structure.

The controller is a single layer LSTM which accepts the visual context vector I_c and the caption sentence of the previous proposal, encoded as $S_{p,t-1}$. The LSTM hidden state h_t^c of the controller encodes the visual context and the language history, and serves as a topic vector, which is fed to the sentence captioning LSTM. The recurrence equations for the controller are given as:

$$\begin{bmatrix} f_t^c \\ i_t^c \\ o_t^c \end{bmatrix} = \sigma \left(\begin{bmatrix} W_f^c \\ W_i^c \\ W_o^c \end{bmatrix} \begin{bmatrix} S_{p,t-1} \\ I_c \\ h_{t-1}^c \end{bmatrix} + \begin{bmatrix} b_f^c \\ b_i^c \\ b_o^c \end{bmatrix} \right) \quad (4)$$

$$\tilde{c}_t^c = \tanh(W_c^c [S_{p,t-1} \ I_c \ h_{t-1}^c] + b_c^c) \quad (5)$$

$$c_t^c = i_t^c \otimes \tilde{c}_t^c + f_t^c \otimes c_{t-1}^c \quad (6)$$

$$h_t^c = o_t^c \otimes \tanh(c_t^c) \quad (7)$$

where \otimes is component-wise multiplication.

The first hidden state h_0^c and the first sentence feature $S_{p,0}$ are initialized to zero. Thus, only visual features are used for decoding the first proposal. At training time, ground truth segments are sorted by ascending end time and their captions' encodings are fed to the controller LSTM in sequence. At test time, we sort the predicted proposals by their end times and decode them sequentially. For the encoding of the previous caption $S_{p,t-1}$, we experimented with two encoding methods: the mean-pooling of word vectors, or the last hidden state of the captioner LSTM. Preliminary experiments found no obvious differences in performance, so we adopt mean-pooling for simplicity.

Sentence Captioning LSTM: We design a two-layer LSTM network for decoding proposals into captions. The first layer focuses on learning the word sequence encoding and the second layer focuses on learning the fusion of visual and language information and context. Each sentence is given a maximum length K , and is padded if it is shorter than K words. As input to the first layer, each word is represented using word vectors w_k^t . The hidden state of the first layer LSTM, $h_k^{(1)}$, is fed to the second layer LSTM, along with the proposal features $I_{p,t}$ and the context vector h_t^c from the controller LSTM. The recurrence equations for the second layer LSTM are given as follows:

$$\begin{bmatrix} f_k \\ i_k \\ o_k \end{bmatrix} = \sigma \left(\begin{bmatrix} W_f \\ W_i \\ W_o \end{bmatrix} \begin{bmatrix} h_k^{(1)} \\ I_{p,t} \\ h_t^c \\ h_{k-1}^{(2)} \end{bmatrix} + \begin{bmatrix} b_f \\ b_i \\ b_o \end{bmatrix} \right) \quad (8)$$

$$\tilde{c}_k = \tanh(W_c [h_k^{(1)} \ I_{p,t} \ h_t^c \ h_{k-1}^{(2)}] + b_c) \quad (9)$$

$$c_k = i_k \otimes \tilde{c}_k + f_k \otimes c_{k-1} \quad (10)$$

$$h_k = o_k \otimes \tanh(c_k) \quad (11)$$

The hidden state $h_k^{(2)}$ goes through a softmax and is used to predict the word at the k^{th} position in the caption. We optimize the normalized log likelihood over all T ground truth proposals and all K unrolled timesteps in the sentence captioning module:

$$\mathcal{L}_{caption} = -\frac{1}{KT} \sum_{t,k} \log P(w_k^t | I_{p,t}, h_t^c, w_1^t, \dots, w_{k-1}^t). \quad (12)$$

3.3. End-to-End Optimization

JEDDi-Net can be trained end-to-end with the proposal and hierarchical captioning modules optimized jointly. The overall loss is as follows; we set $\lambda = 1$.

$$\mathcal{L}_{total} = \mathcal{L}_{spn} + \lambda \mathcal{L}_{caption} \quad (13)$$

Our end-to-end training allows us to propagate gradient information back to the underlying C3D network and optimize the convolutional filters for better proposal features and visual context encoding. In activity detection, multiple positive and negative proposals are generated according to tIoU thresholds with ground truth segments in a single video and selectively form balanced training mini-batches. In dense captioning, however, a video contains only a few ground-truth captions. Further, the same captions always appear together in the same mini-batch with one video as input during end-to-end training. We find the lack of diversity to severely disrupt proper optimization.

We propose a more effective training strategy. We first extract intermediate ground truth segment features from the pretrained SPN and C3D classification networks. We then shuffle these and form a relatively large training batch with diverse captions to pre-train the captioning module. After pretraining, the entire network is trained end-to-end following the conventional strategy with a reduced learning rate. In the experimental section, we show substantial performance improvements after end-to-end training compared to the separately trained models. In the next section, we present experimental results illustrating the benefits of end-to-end training on both proposal prediction and caption generation.

4. Experiments

We evaluate JEDDi-Net on the large-scale ActivityNet Captions dataset proposed by [13]. For proposal evaluation, we use the conventional Area Under the AR vs AN curve (AUC) with tIoU threshold 0.8. When evaluating captions, we follow [13] by computing the average precision (BLEU, METEOR, CIDEr and ROUGE-L) across tIoU thresholds of 0.3, 0.5, 0.7, 0.9 for the top 1000 proposals. In addition, we report results on the TACoS-MultiLevel dataset [18].

Table 1. Proposal evaluation results on ActivityNet Captions dataset (in percentage). AUC at IoU threshold 0.8 and average AUC at tIoU thresholds $\alpha \in (0.5, 0.95)$ with step 0.05 are reported.

	$\alpha = 0.8$	$\alpha \in (0.5, 0.95)$
DAP [13]	30	-
multi-scale DAP [13]	38	-
pretrain SPN	57.75	57.12
JEDDi-Net(joint training)	59.13	58.70
JEDDi-Net(joint training w/ context)	58.21	58.24

4.1. Experiments on the ActivityNet Captions

Dataset and Setup: The ActivityNet Captions dataset [13] with around 20k videos are split into training, validation and testing with a 50%/25%/25% ratio. Each video contains at least two ground truth segments and each segment is paired with one ground truth caption. We keep all the words that appear at least 5 times. The height \mathcal{H} and width \mathcal{W} of all input frames are set to 112 each following [23]. We set the number of frames \mathcal{L} to 768, breaking the arbitrary length input video into 768 frame chunks and zero-padding if necessary. The maximum caption length is set to 30, which covers over 97% of captions in the training set. We sample frames at 3 fps and set the number of anchor segment scales to be 36 to generate candidate proposals². In the hierarchical captioning module, we set the hidden state dimension to 20 in the controller LSTM and 512 in the captioner LSTMs.

We train the SPN using the temporal annotation of ground truth segments in the ActivityNet Captions dataset with Sports-1M pretrained C3D weight initialized [23]. We also extract fc6 features for ground truth proposals from pretrained SPN, shuffle the proposal features and paired ground truth captions, and form batches of size 32 to train the captioner LSTM from scratch. The pretrained SPN and captioner LSTM will serve as initialization weights for our end-to-end model. We refer our full JEDDi-Net which is jointly trained for SPN and hierarchical captioning modules as ‘JEDDi-Net(joint training with context)’. After removing the controller LSTM of the hierarchical captioning module in ‘JEDDi-Net(joint training with context)’, we refer this ablation model as ‘JEDDi-Net(joint training)’. To show the effectiveness of end-to-end training in our model, we extract proposal features from the separately trained SPN and decode captions using the separately trained captioner LSTM, and refer to this model as ‘JEDDi-Net(separate training)’.

Proposal Evaluation: The proposal evaluation result is shown in Table 1. The dense video captioning model in [13] uses DAP [3] as its proposal network, extends DAP to a

²Specifically, we chose the following anchor scales based on cross-validation - [1,2,3,4,5,6,7,8,10,12,14,16,20,24,28,32,40,48,56,64,66,68,70,72,74,76,78,80,82,84,86,88,90,92,94,96].

multi-scale version and shows improved proposal results in AUC at tIoU 0.8. Our pretrained SPN model achieves 57.75% at tIoU 0.8 in AUC, 19.75% higher than [13], indicating our superior ability to segment events of interest. Following the traditional evaluation of the temporal localization task in ActivityNet, we also report the average AUC result across ten different tIoU thresholds uniformly distributed between 0.5 and 0.95 with 1000 proposals per video. The average AUC for our pretrained SPN is 57.12%, which is on par with tIoU at 0.8, indicating robust performance of SPN across different tIoUs.

Dense Captioning Evaluation: The average dense video captioning results across four tIoUs using the evaluation code released by [13] are shown in Table 2. We list the two baseline results from [13], the model without visual context and the one with visual context. Our first ‘JEDDi-Net(separate training)’ model without end-to-end training already achieves reasonable results with a METEOR score 2.55% higher than the best context model in [13]. This indicates that our decoded captions are more semantically meaningful and closer to human descriptions. These results further motivate our proposal feature encoding method, which employs 3D SoI pooling directly on the conv features of the input video segment, rather than using the LSTM hidden state for a set of proposals. Our ‘JEDDi-Net(separate training)’ and ‘JEDDi-Net(joint training)’ models without context do better than [13]’s ‘no context’ model on all evaluation metrics. After end-to-end training, both ‘JEDDi-Net(joint training)’ and ‘JEDDi-Net(joint training with context)’ improve on all evaluation metrics compared to ‘JEDDi-Net(separate training)’. This shows the benefits of joint parameter training for dense video captioning. Our ‘JEDDi-Net(joint training with context)’ model that incorporates visual and language context further improves all the language evaluation metrics compared to the no context version.

Our full model outperforms the context model in [13] on all evaluation metrics except for Bleu₄. In particular, we achieve a 78% relative improvement on METEOR, the only metric used by the test server. The reason for lower BLEU₄ might be that we did not leverage the power of beam search due to limited computational resources. We decoded the captions with greedy search (in Table 2), selecting the most probable word at each timestep. Experiments in several papers [26, 4] show that beam search can improve some evaluation metrics, especially Bleu₃ and Bleu₄.

Applying the same JEDDi-Net(joint training with context) on the test server yields an average METEOR score of 8.81%, which is on the same level as the average METEOR score 8.58% on the validation set. This demonstrates that our model generalizes well to unseen data.

Table 3 shows all the evaluation metrics for all the four tIoUs in details for our ‘JEDDi-Net(joint training with con-

Table 2. Dense video captioning results on ActivityNet Captions dataset (in percentage). The average Bleu_1-4 (B1-B4), METEOR (M), CIDEr (C) and ROUGE.L (R) across tIoU thresholds of 0.3, 0.5, 0.7, 0.9 are reported.

Model	B1	B2	B3	B4	M	C	R
R. Krishna et al. [13] (no context)	12.23	3.48	2.1	0.88	3.76	12.34	-
R. Krishna et al. [13] (with context)	17.95	7.69	3.86	2.20	4.82	17.29	-
JEDDi-Net(separate training)	16.72	6.65	2.65	1.07	7.37	14.65	16.47
JEDDi-Net(joint training)	19.27	8.69	3.78	1.54	8.30	19.81	18.86
JEDDi-Net(joint training w/ context)	19.97	9.10	4.06	1.63	8.58	19.88	19.63
JEDDi-Net(joint training w/ context) on test server	-	-	-	-	8.81	-	-

Table 3. Dense video captioning results at different tIoU thresholds α on ActivityNet Captions dataset (in percentage). The Bleu_1-4 (B1-B4), METEOR (M), CIDEr (C), and ROUGE.L (R) at different tIoU thresholds α are reported for our JEDDi-Net(joint training with context) with greedy search decoding.

α	B1	B2	B3	B4	M	C	R
0.3	19.72	8.84	4.04	1.65	8.44	13.40	19.80
0.5	20.31	9.26	4.22	1.71	8.75	16.53	20.41
0.7	20.86	9.70	4.37	1.76	8.97	21.52	20.74
0.9	19.00	8.60	3.61	1.39	8.17	28.09	17.57
avg $\alpha \in (0.3, 0.5, 0.7, 0.9)$	19.97	9.10	4.06	1.63	8.58	19.88	19.63

text)’. As tIoU α increases from 0.3 to 0.7, Bleu_1-4, METEOR and ROUGE.L increase steadily, with the highest scores at $\alpha = 0.7$. The reason might be that our SPN network is trained with tIoU greater than 0.7 as positive examples, and tested with post processed NMS at 0.7. However, Bleu_1-4, METEOR and ROUGE.L decrease significantly at tIoU 0.9, possibly because much fewer proposals have been left for evaluation using the tIoU 0.9 criterion. The CIDEr metric is consistently improved across tIoU α values from 0.3 to 0.9, which indicates the sensitivity of the CIDEr score to the number of evaluation proposals. CIDEr measures the diversity of the captions. When a small subset of proposals is kept with higher tIoU, the captions are more diverse and the CIDEr score is higher, and vice versa.

We show two videos with predicted dense captions from JEDDi-Net as qualitative examples from the ActivityNet Captions dataset in Figure 5(a). Our model generates continuous and fluent descriptions of the activities of jumping over the mat and making a cocktail, taking context into account. We note that the ground truth caption for segment A in the first video is “A man is seen...”, while our prediction is “A person is seen...”. Though these two 4-grams have the same meaning in this video, such predictions will not be counted as positive in the Bleu_4 score, indicating a potential reason for the lower value.

4.2. Experiments on the TACoS-MultiLevel Dataset

Dataset and Setup: The TACoS-MultiLevel dataset [18] contains cooking videos with the start and end time for cap-

Table 4. Proposal evaluation results on TACoS-MultiLevel dataset, showing AUC at tIoU threshold 0.8 and average AUC at tIoU $\alpha \in (0.5, 0.95)$ with step 0.05.

	$\alpha = 0.8$	$\alpha \in (0.5, 0.95)$
pretrain SPN	36.88	41.90
JEDDi-Net(joint training)	36.85	43.30
JEDDi-Net(joint training w/ context)	36.31	43.23

tions and activity labels, which can be used for dense video captioning. Compared to ActivityNet Captions, TACoS has more ground truth annotations per video with an average of 284 sentences per video. We use the same 143/42 video split for training and testing as in [34]. All words are kept in the vocabulary and the maximum caption length is set to 15. Frames are sampled at 5 fps. Other settings are identical with the ActivityNet Captions experiments. We evaluate three ablated models on proposal detection and caption generation.

Results: Table 4 shows results of proposal evaluation. We report the average AUC result across ten tIoU thresholds uniformly distributed between 0.5 and 0.95 for top 1000 proposals per video. We also measure the improvement of proposal detection after end-to-end training. The average AUC for both ‘JEDDi-Net(joint training)’ and ‘JEDDi-Net(joint training with context)’ improve compared with the pretrained SPN, while AUC at tIoU 0.8 stays almost the same.

Table 5 shows results for caption generation averaged across four tIoUs. No dense captioning results have been

Table 5. Dense video captioning results on TACoS-MultiLevel dataset (in percentage). The average Bleu_1-4 (B1-B4), METEOR (M), CIDEr (C) and ROUGE.L (R) across tIoU thresholds of 0.3, 0.5, 0.7, 0.9 are reported.

models	B1	B2	B3	B4	M	C	R
JEDDi-Net(separate training)	45.2	32.3	19.7	13.1	20.7	65.4	46.2
JEDDi-Net(joint training)	48.7	36.4	24.6	17.4	23.3	99.7	50.0
JEDDi-Net(joint training w/ context)	49.2	37.1	25.2	18.1	23.9	104.0	50.9

previously reported on this dataset, so ours is the first set of such results. The previously reported trimmed video captioning results can be considered as the upper bound for our task on the same annotations, as it is noted in [13]. TACoS-MultiLevel dataset [18] reports a Bleu_4 value of 27.5% for trimmed video captioning, which can be seen as the upper bound of our reported Bleu_4 value with the consideration of tIoU overlaps. Compared to ‘JEDDi-Net(separate training)’, all evaluation metrics for both ‘JEDDi-Net(joint training)’ and ‘JEDDi-Net(joint training with context)’ improve after end-to-end training, which indicates the benefits of our approach. Also ‘JEDDi-Net(joint training with context)’ further improves all evaluation metrics through modelling of visual and language context in the hierarchical captioning module, compared to ‘JEDDi-Net(joint training)’ without explicitly modeling context.

Figure 5(b) provides two examples of video predictions from TACoS-MultiLevel dataset. Though JEDDi-Net missed some objects in the generated captions like “a measuring cup”, JEDDi-Net could still provide fine-grained descriptions of certain activities involving small objects such as the orange and the egg. The network likely benefited from learning object representations from the captions in end-to-end training.

5. Conclusion

In this paper, we proposed JEDDi-Net, an end-to-end deep neural network designed to perform the dense video captioning task, and introduced an optimization strategy for training it end-to-end. The visual and language context is incorporated by the controller in the hierarchical captioning module, to provide context for decoding each proposal rather than training and decoding each proposal independently. Our end-to-end framework can be further extended to solve other vision and language tasks, such as natural language localization in videos.

Acknowledgements: Supported in part by IARPA (contract number D17PC00341) and DARPA’s XAI program.³

³The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.



JEDDi-Net (ours)
A: A person is seen walking down a gym with a large group of people on the side.
B: The man then jumps into the air and begins to jump around the mat.
C: The man then jumps off the mat and walks away.

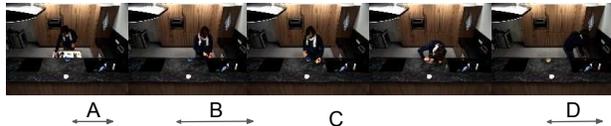
Ground Truth
A: A man is seen sitting before a large mat with another person laying in the background.
B+C: The man then runs down the mats on stilts and jumps over a large mat.



JEDDi-Net (ours)
A: A man is standing in a kitchen talking to a camera.
B: The man pours a glass of ice and pours it into a glass.
C: The man then puts the drink into the glass and pours it into the glass.

Ground Truth
A: A man is standing behind a bar with several bottles of drinks.
B: He fills a cup with ice and pours all the ingredients into the glass.
C: He takes out another glass and pours the drink into the new glass and places a straw into it.

(a) ActivityNet Captions dataset



JEDDi-Net (ours)
A: The person took a knife from the drawer.
B: The person cut the orange in half.
C: The person juiced the orange half.
D: The person put the peel in the trash.

Ground Truth
A: The person took a knife from the drawer.
B: The person placed one half of the orange on the top of the juicer.
C: The person juiced an orange half using the juicer.
D: The person threw the orange rind into the trash.



JEDDi-Net (ours)
A: The person entered the kitchen.
B: The person took out a drawer.
C: The person took an egg from the refrigerator.
D: The person put the egg in the bowl.
E: The person placed the egg back on the counter.

Ground Truth
A: The person entered the kitchen.
B: The person took out a measuring cup.
C: The person took an egg from the refrigerator.
D: The person put the egg-white in one glass.
E: The person threw the shell in the garbage.

(b) TACoS-MultiLevel dataset

Figure 5. Qualitative visualization of the predicted dense captions by JEDDi-Net (best viewed in color). Figure (a) and (b) show results for two videos each in ActivityNet captions dataset and TACoS-MultiLevel dataset.

References

- [1] S. Buch, V. Escorcia, C. Shen, B. Ghanem, and J. C. Niebles. Sst: Single-stream temporal action proposals. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6373–6382. IEEE, 2017.
- [2] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [3] V. Escorcia, F. C. Heilbron, J. C. Niebles, and B. Ghanem. DAPs: Deep Action Proposals for Action Understanding. In *European Conference on Computer Vision*, pages 768–784, 2016.
- [4] H. Fang, S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, et al. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482, 2015.
- [5] J. Gao, C. Sun, Z. Yang, and R. Nevatia. Tall: Temporal activity localization via language query. *arXiv preprint arXiv:1705.02101*, 2017.
- [6] J. Gao, Z. Yang, K. Chen, C. Sun, and R. Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3628–3636, 2017.
- [7] R. Girshick. Fast R-CNN. In *IEEE International Conference on Computer Vision*, pages 1440–1448, 2015.
- [8] S. Guadarrama, N. Krishnamoorthy, G. Malkarnenkar, S. Venugopalan, R. Mooney, T. Darrell, and K. Saenko. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 2712–2719, 2013.
- [9] L. A. Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell. Localizing moments in video with natural language. *arXiv preprint arXiv:1708.01641*, 2017.
- [10] S. Ji, W. Xu, M. Yang, and K. Yu. 3D Convolutional Neural Networks for Human Action Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:221–231, 2013.
- [11] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [12] J. Krause, J. Johnson, R. Krishna, and L. Fei-Fei. A hierarchical approach for generating descriptive image paragraphs. *arXiv preprint arXiv:1611.06607*, 2016.
- [13] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles. Dense-captioning events in videos. In *International Conference on Computer Vision (ICCV)*, 2017.
- [14] R. Lin, S. Liu, M. Yang, M. Li, M. Zhou, and S. Li. Hierarchical recurrent neural network for document modeling. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 899–907, 01 2015.
- [15] S. Ma, L. Sigal, and S. Sclaroff. Learning Activity Progression in LSTMs for Activity Detection and Early Detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1942–1950, 2016.
- [16] A. Montes, A. Salvador, and X. G. i Nieto. Temporal Activity Detection in Untrimmed Videos with Recurrent Neural Networks. *arXiv preprint arXiv:1608.08128*, 2016.
- [17] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond Short Snippets: Deep Networks for Video Classification. In *IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015.
- [18] A. Rohrbach, M. Rohrbach, W. Qiu, A. Friedrich, M. Pinkal, and B. Schiele. Coherent multi-sentence video description with variable level of detail. In *German conference on pattern recognition*, pages 184–195. Springer, 2014.
- [19] I. V. Serban, A. Sordoni, Y. Bengio, A. C. Courville, and J. Pineau. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, pages 3776–3784, 2016.
- [20] Z. Shou, D. Wang, and S.-F. Chang. Temporal Action Localization in Untrimmed Videos via Multi-stage CNNs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [21] K. Simonyan and A. Zisserman. Two-stream Convolutional Networks for Action Recognition in Videos. In *Neural Information Processing Systems*, pages 568–576, 2014.
- [22] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A Multi-Stream Bi-Directional Recurrent Neural Network for Fine-Grained Action Detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [23] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [24] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko. Sequence to sequence – video to text. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [25] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics – Human Language Technologies (NAACL HLT)*, 2015.
- [26] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [27] H. Wang and C. Schmid. Action Recognition with Improved Trajectories. In *IEEE International Conference on Computer Vision*, pages 3551–3558, 2013.
- [28] P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Learning to Track for Spatio-Temporal Action Localization. In *IEEE International Conference on Computer Vision*, 2015.
- [29] H. Xu, A. Das, and K. Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. *arXiv preprint arXiv:1703.07814*, 2017.

- [30] L. Yang, K. Tang, J. Yang, and L.-J. Li. Dense captioning with joint inference and visual context. *arXiv preprint arXiv:1611.06949*, 2016.
- [31] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville. Describing videos by exploiting temporal structure. In *Proceedings of the IEEE international conference on computer vision*, pages 4507–4515, 2015.
- [32] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end Learning of Action Detection from Frame Glimpses in Videos. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2678–2687, 2016.
- [33] G. Yu and J. Yuan. Fast Action Proposals for Human Action Detection and Search. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [34] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu. Video paragraph captioning using hierarchical recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4584–4593, 2016.