

Enriching Local and Global Contexts for Temporal Action Localization

Zixin Zhu¹ Wei Tang² Le Wang^{1*} Nanning Zheng¹ Gang Hua³

¹Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University

²University of Illinois at Chicago

³Wormpex AI Research

Abstract

Effectively tackling the problem of temporal action localization (TAL) necessitates a visual representation that jointly pursues two confounding goals, *i.e.*, fine-grained discrimination for temporal localization and sufficient visual invariance for action classification. We address this challenge by enriching both the local and global contexts in the popular two-stage temporal localization framework, where action proposals are first generated followed by action classification and temporal boundary regression. Our proposed model, dubbed *ContextLoc*, can be divided into three sub-networks: *L-Net*, *G-Net* and *P-Net*. *L-Net* enriches the local context via fine-grained modeling of snippet-level features, which is formulated as a query-and-retrieval process. *G-Net* enriches the global context via higher-level modeling of the video-level representation. In addition, we introduce a novel context adaptation module to adapt the global context to different proposals. *P-Net* further models the context-aware inter-proposal relations. We explore two existing models to be the *P-Net* in our experiments. The efficacy of our proposed method is validated by experimental results on the THUMOS14 (54.3% at $tIoU@0.5$) and ActivityNet v1.3 (56.01% at $tIoU@0.5$) datasets, which outperforms recent states of the art. Code is available at <https://github.com/buxiangzhiren/ContextLoc>.

1. Introduction

Temporal action localization (TAL) is a fundamental task in video understanding. It aims at classifying action instances in an untrimmed video and locating their temporal boundaries. Recently, TAL has drawn increasing attention from the research community because of its wide range of applications such as action retrieval, video summarization, and intelligent security [28, 37, 27, 18].

Prior TAL methods can be divided into two categories. One-stage approaches [14, 43, 2, 26] classify and lo-

*Corresponding author.

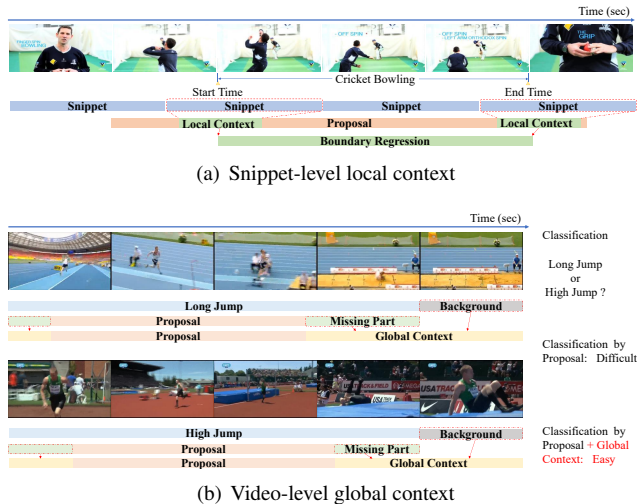


Figure 1. (a) Each proposal consists of a number of video snippets. A snippet is a small number of consecutive frames and serves as the basic unit of feature extraction. It is the snippets capturing the start and end times of the action that play an important role in action localization. (b) The video-level global context is important as it involves background and high-level activity information that can be critical to distinguish action categories of similar appearance and motion patterns.

cate action instances from an input video in a single shot. Two-stage approaches first generate category-agnostic action proposals [20, 21, 10, 25, 11, 45, 1] and then perform action classification and temporal boundary refinement [44, 39, 41, 4, 7] for each proposal. They have their respective advantages. One-stage approaches can be easily trained in an end-to-end fashion while two-stage approaches usually obtain superior performance.

Effectively tackling the task of TAL necessitates a visual representation that jointly pursues two confounding goals, *i.e.*, fine-grained discrimination for temporal localization and sufficient visual invariance for action classification. This paper addresses this challenge by leveraging rich *local* and *global contexts* in a video, in our proposed two-stage method.

The **local context** refers to snippets within a proposal. They contain fine-grained temporal information that is critical to localization. As shown in Figure 1(a), we locate the boundary of the action “cricket bowling” by the moments of bowling and catching the cricket. Therefore, it is the snippets capturing these special moments that facilitate localization in the temporal domain. However, prior approaches obtain features of a proposal by applying temporal max pooling to features of snippets within it, which unavoidably discards the fine-grained temporal information.

The **global context** refers to the entire video. It provides discriminative information complementary to features of a proposal for action classification. As shown in Figure 1(b), to distinguish “long jump” and “high jump”, we need to check not only the last a few frames of the action duration but also background frames outside the duration. In addition, the global context provides high-level activity information enforcing strong prior on the categories of actions that should appear in it. For example, it is unlikely to see sports actions in a video of household activity. Unfortunately, the video-level global context has been largely ignored by existing TAL models.

We introduce a novel network architecture, termed *ContextLoc*, to model local and global contexts in a unified framework for TAL. It consists of three sub-networks: L-Net, G-Net and P-Net. Inspired by the self-attention [36], L-Net performs a *query-and-retrieval* procedure. But different with the self-attention, the queries, keys and values in our L-Net correspond to different semantic entities, and they are specially designed to enrich the local context. Specifically, the feature vector of a proposal is taken as a *query* to match the *key* feature vectors of snippets within this proposal so that the relevant fine-grained *values* in the local context can be retrieved and aggregated to this proposal.

G-Net models the global context by integrating the video-level representation and features of each proposal. However, a naive concatenation of these two would be insufficient because the former contains not only relevant cues but also irrelevant noise. In addition, the portions of context needed to enhance different proposals are different. To effectively integrate video-level information with proposal-level features, we propose *global context adaptation*. It attends the video-level representation to the local context within each proposal so that the global context is adapted to them respectively.

P-Net models context-aware inter-proposal relations. This includes interactions between proposal-level features enhanced by the local context and interactions between global contexts adapted to different proposals. We use an existing model as P-Net and investigate two candidates: P-GCN [44] and the non-local network [40].

It is worth noting our ContextLoc is different from P-GCN [44]. P-GCN only considers inter-proposal relations,

and the features of a proposal are obtained via applying temporal max pooling to the features of snippets within it. In contrast, ContextLoc enriches the local context via fine-grained modeling of snippet-level features and enriches the global context via higher-level modeling of the video-level representation. We consider P-GCN or its relative as a useful component of our framework, *i.e.*, P-Net, and show that the local and global contexts we advocate are complementary to the inter-proposal relations.

We evaluate our proposed method on two popular benchmarks for TAL. On THUMOS14 [16], it achieves 54.3% mAP at tIoU 0.5, which outperforms the previous best method PBRNet [23]. On ActivityNet v1.3 [13], it achieves 56.01% mAP at tIoU 0.5, which outperforms the state-of-the-art method PBRNet.

The contributions of this paper are summarized below.

- To our knowledge, this is the first work attempting to exploit the snippet-level local context and the video-level global context to enhance proposal-level features within a two-stage TAL framework.
- We introduce a novel network architecture, termed *ContextLoc*, consisting of three sub-networks, *i.e.*, L-Net, G-Net, and P-Net. L-Net is the first of its kind to use a proposal to query the snippets within it and retrieve the local context to supplement it with fine-grained temporal information. G-Net augments features of each proposal by integrating the video-level representation. We introduce a novel context adaptation process to adapt the global context to different proposals. While P-Net is built on existing networks, we show that P-Net, regardless of its instantiation, is complementary to our L-Net and G-Net. Our ContextLoc unifies the respective advantages of these three sub-networks and achieves more effective TAL.
- ContextLoc is superior or comparable to the state-of-the-art performance on two popular TAL benchmarks, *i.e.*, THUMOS14 and ActivityNet v1.3.

2. Related Work

Action recognition. Action recognition [34, 35, 38, 32, 9, 15, 3] is a fundamental task in video understanding and has been widely studied in the past a few years. In this paper, we use the I3D network [3], pre-trained on Kinetics [17], to extract snippet-level features from a video. We use the two-stream strategy like Simonyan *et al.* [32] to fuse the predictions made on RGB frames and optical flows. There are a few methods exploring the global context for action recognition. He *et al.* [12] apply temporal convolutions on the local spatial-temporal feature maps to model the

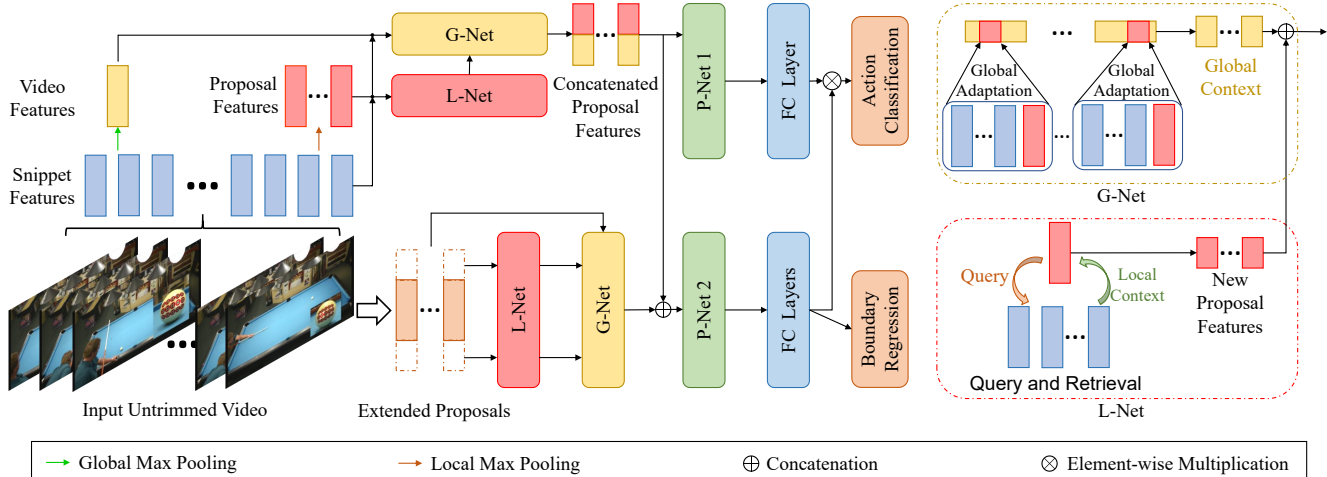


Figure 2. The pipeline of our ContextLoc. The input is an untrimmed video consisting of non-overlapping snippets. Multi-level features are extracted from snippets, proposals, and the video. L-Net, G-Net and P-Net respectively enhance the representation of each proposal via the local context, global context and context-aware inter-proposal relations. Finally, we perform action classification and temporal boundary refinement for each proposal through fully-connected layers. ContextLoc processes both original proposals and extended proposals, and fuses their predictions.

global spatial-temporal relationship. Liu *et al.* [24] establish a global temporal template by fusing local motion features over time. Different from them, our global context is adapted to each proposal by aggregating the proposal-level and snippet-level features within it for action classification and temporal boundary localization.

Temporal action localization. TAL is a mirror problem of image-based object detection [6, 29] in the temporal domain. Some methods tackle this task by focusing on video snippets (segments). For example, Zhao *et al.* [46] model the temporal structure of each action instance via a structured temporal pyramid. Shou *et al.* [31] exploit the effectiveness of deep networks in temporal action localization via three segment-based 3D ConvNets. Xu *et al.* [42] further consider the temporal and semantic structure by graph convolutional networks.

Some other works focus on proposals. Lin *et al.* [21] locate temporal boundaries with high probabilities, then directly combine these boundaries into proposals and evaluate whether a proposal contains an action. Chen *et al.* [5] select more reliable proposals and subsequently refine proposal boundaries by variance voting during post-processing. Different from these prior TAL methods, we refine proposal-level features with the snippet-level local context and the video-level global context, which they ignore.

Our ContextLoc is different from a recent graph-based method: G-TAD [42]. G-TAD applies a graph convolutional network (GCN) to update features of snippets, each of which is considered as a node in a snippet graph, and casts TAL as a sub-graph localization problem. By contrast, ContextLoc enhances features of each proposal by enrich-

ing the local snippet-level context and the global video-level context, which are then used for action classification and temporal boundary refinement.

Contexts in videos. Temporal contextual information is important to video understanding and has been exploited in various tasks such as visual relation detection [22], question answering [8], and representation learning [28]. Specifically, Liu *et al.* [22] proposes a novel sliding-window scheme to simultaneously predict short-term and long-term relationships. Qiu *et al.* [28] updates local and global features by modeling the diffusions between these two representations through convolutional neural networks. Different from previous work, our exploration of local and global contexts is hierarchical and rich. For proposals, snippets are local, and proposals are global. For videos, proposals are local, and videos are global. The former is modeled by L-Net and the latter is modeled by G-Net. Furthermore, the local and global contexts interact with each other through the P-Net in our proposed model.

3. The ContextLoc Model

3.1. Overview

Problem statement. The input to our ContextLoc is an untrimmed video consisting of N non-overlapping snippets. Each snippet is a small number of consecutive frames. We use the I3D network [3] to extract features from each snippet and denote the collection of all snippet-level features as $\{\mathbf{x}_j \in \mathbb{R}^D\}_{j=1}^N$, where \mathbf{x}_j is the features of the j th snippet and D is the feature dimension. Following the two-stream strategy [32], predictions made on RGB frames and opti-

cal flows are fused. The output is a set of action instances $\{\psi_i \mid \psi_i = (t_{i,s}, t_{i,e}, c_i)\}$, where $t_{i,s}, t_{i,e}$ and c_i are the starting time, ending time and action category of the i th instance, respectively.

Multi-level representations. As shown in Figure 2, ContextLoc explicitly models representations of three different levels: snippet-level representation $\mathbf{x} \in \mathbb{R}^D$ (local context), proposal-level representation $\mathbf{y} \in \mathbb{R}^D$, and video-level representation $\mathbf{z} \in \mathbb{R}^D$ (global context). The proposals are obtained by BSN [21], and each proposal i has a starting time and an ending time. The initial features of the i th proposal, denoted as \mathbf{y}_i , are obtained via temporally max-pooling the snippet-level features within its duration $\{\mathbf{x}_j \mid j \in S(i)\}$, where $S(i)$ denotes snippets whose duration is between the starting time and ending time of the i th proposal. Similarly, the initial video-level features, denoted as \mathbf{z} , are obtained by temporally max-pooling all snippet-level features in the video.

Network architecture. ContextLoc consists of three sub-networks: L-Net, G-Net and P-Net. They respectively enhance the representation of each proposal via local context, global context, and context-aware inter-proposal relations. Following previous work [30, 44, 21, 19], we also apply ContextLoc to each extended proposal with an enlarged temporal duration. The final representation of the original proposal is used for action classification (via a fully-connected layer) while that of the extended proposal is used for completeness prediction and temporal boundary refinement. These classification and completeness scores are element-wisely multiplied to make the final categorical prediction.

Below we introduce L-Net, G-Net and P-Net in Sections 3.2, 3.3 and 3.4, respectively. There is a major difference between the way we deal with the extended proposals and how previous approaches do. We detail it in Section 3.5.

3.2. L-Net (Local Context)

The initial features of a proposal obtained through temporal max-pooling are insufficient because the fine-grained temporal information critical to localization is lost. L-Net addresses this problem by finding within a proposal the snippets most relevant to it and aggregating them to retain the informative features. We call these snippets the local context because their temporal range is within a proposal and they are at a lower semantic level.

Inspired by the self-attention [36], L-Net performs a *query-and-retrieval* procedure. The *queries*, *keys*, and *values* are respectively features of each proposal, features of snippets within each proposal, and transformed features of these snippets. Specifically, a *query* proposal is matched to the *key* snippets within this proposal so that the relevant fine-grained *values* in the local context can be retrieved and aggregated to this proposal. This is achieved by building an

attention map between the query and keys and then aggregating values based on the attention weights.

Attention map. The attention weight between a proposal i and a snippet $j \in S(i)$ measures their relatedness and determines how much information will be retrieved from this snippet. It is calculated as

$$a_{i,j}^L = \frac{\sigma(s(\mathbf{y}_i, \mathbf{x}_j))}{\sum_{k \in S(i)} \sigma(s(\mathbf{y}_i, \mathbf{x}_k))}, \quad (1)$$

where σ is the ReLU function, s is the cosine similarity between two vectors. The ReLU function sets negative similarity values to 0. As a result, those irrelevant snippets will be ignored during the retrieval process. Normalization ensures the attention weights of all snippets sum to one.

Local context aggregation. We first calculate *values* by transforming features of each snippet via a fully-connected layer and then linearly combine them via the attention weights. This local context is finally aggregated with transformed features of the proposal to obtain a new representation of this proposal:

$$\mathbf{y}_i^L = \sigma(\mathbf{W}_1^L \mathbf{y}_i + \sum_{j \in S(i)} a_{i,j}^L \mathbf{W}_2^L \mathbf{x}_j), \quad (2)$$

where $\mathbf{y}_i^L \in \mathbb{R}^{D/2}$ is a new proposal-level representation carrying fine-grained temporal information, σ is the ReLU function, and $\mathbf{W}_1^L, \mathbf{W}_2^L \in \mathbb{R}^{(D/2) \times D}$ are trainable weights. In practice, we find considering the features pooled from all snippets as a special snippet and including it in $S(i)$ marginally improves the performance.

3.3. G-Net (Global Context)

The video-level global context is important as it involves background and high-level activity information that can be critical to distinguish action categories of similar appearance and motion patterns. A straightforward way to enrich the global context is to concatenate the video-level representation \mathbf{z} and features of each proposal. However, this is insufficient because the global representation contains not only relevant context but also irrelevant noise. In addition, the contexts needed to process different proposals are different. This line of analysis motivates us to adapt the global context to each proposal before integrating it.

Global context adaptation. To adapt the global context to the i th proposal, we first attend the video-level representation \mathbf{z} to the features of this proposal \mathbf{y}_i as well as the snippet-level features within it $\{\mathbf{x}_j : j \in S(i)\}$:

$$a_{i,j}^G = \frac{\sigma(s(\mathbf{z}, \mathbf{x}_j))}{\sum_{k \in S(i)} \sigma(s(\mathbf{z}, \mathbf{x}_k)) + \sigma(s(\mathbf{z}, \mathbf{y}_i))}, \quad (3)$$

$$b_i^G = \frac{\sigma(s(\mathbf{z}, \mathbf{y}_i))}{\sum_{k \in S(i)} \sigma(s(\mathbf{z}, \mathbf{x}_k)) + \sigma(s(\mathbf{z}, \mathbf{y}_i))}, \quad (4)$$

where $a_{i,j}^G$ is the attention weight between the video-level representation and the j th snippet of the i th proposal, b_i^G is the attention weight between the video-level representation and the i th proposal. With the attention weights, the adapted global context is calculated via

$$\mathbf{z}_i^G = \sigma(\mathbf{W}_1^G \mathbf{z} + \mathbf{W}_2^G (\sum_{j \in S(i)} a_{i,j}^G \mathbf{x}_j + b_i^G \mathbf{y}_i)), \quad (5)$$

where $\mathbf{z}_i^G \in \mathbb{R}^{D/2}$ is the global context adapted to the i th proposal, $\mathbf{W}_1^G, \mathbf{W}_2^G \in \mathbb{R}^{(D/2) \times D}$ are trainable weights.

Global context aggregation. Finally, G-Net concatenates the global context adapted to the i th proposal, *i.e.*, $\mathbf{z}_i^G \in \mathbb{R}^{D/2}$, and the features of this proposal obtained from L-Net, *i.e.*, $\mathbf{y}_i^L \in \mathbb{R}^{D/2}$:

$$\mathbf{y}_i^G = \mathbf{y}_i^L \oplus \mathbf{z}_i^G, \quad (6)$$

where \oplus denotes concatenation, and $\mathbf{y}_i^G \in \mathbb{R}^D$ is the output of G-Net for the i th proposal.

3.4. P-Net (Inter-proposal Relations)

P-Net takes as input $\{\mathbf{y}_i^G\}$ and outputs a new representation for each proposal. We use an existing model as P-Net and investigate two candidates: P-GCN [44] and the non-local network [40].

P-GCN constructs an action proposal graph. Each proposal is taken as a node. There are two types of relational edges. One connects overlapping proposals, and the other connects distinct but nearby proposals. Then, a GCN is applied to update the proposal-level features based on their relations. Different from P-GCN, the non-local network builds a complete graph over all proposals and dynamically calculates the edge weights based on their pairwise similarities. We will have an in-depth investigation on their effectiveness as a building block in our ContextLoc.

Since each \mathbf{y}_i^G is a concatenation of two parts as indicated in Eq. (6), P-Net models not only interactions between proposal-level features enhanced by the local context but also interactions between global contexts adapted to different proposals. We will show in the experiments the latter help improve the TAL performance. Note we do not claim any specific network architecture we adopt in P-Net as our contribution. What we would like to show is that P-Net, regardless of its instantiation, is complementary to the local and global contexts we advocate in this paper and serves as a useful component in the proposed ContextLoc.

3.5. Extended Proposals

A common practice in TAL [30, 44, 21, 19] is to extend each proposal on both ends (*e.g.*, by 50% of the temporal duration). The predictions obtained from these extended proposals and the original ones are fused as described in Section 3.1. Prior methods like P-GCN treat an extended

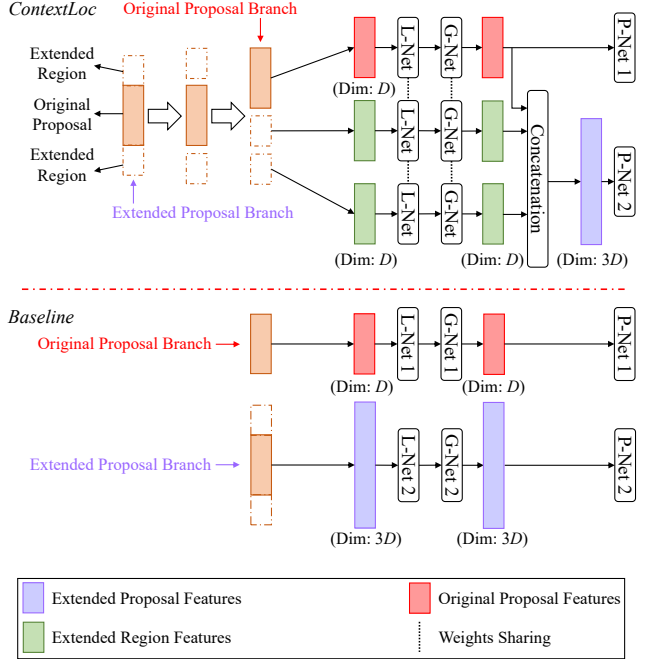


Figure 3. Different methods of processing the extended proposals. To better highlight the original and extended proposals, we do not include other entities (*e.g.*, snippets and the video) in this figure. Bottom: prior approaches treat an extended proposal as a single proposal and process it in a separate branch. Top: our new method treats the extended proposal as three proposals in L-Net and G-Net. P-Net, *i.e.*, “P-Net 2”, treats an extended proposal as a single proposal and processes the concatenated features.

proposal as a single proposal independent of the original proposal. As we will show in the experiments, this strategy does not work well for L-Net and G-Net. On the one hand, setting the feature dimension of an extended proposal much larger than that of the original proposal can account for the enlarged temporal duration but significantly increases the model complexity. On the other hand, setting the feature dimension of an extended proposal the same as that of the original proposal leads to inferior performance. In addition, the extended proposal and the original one are processed separately. Their internal connection (*i.e.*, the original proposal is a part of the extended proposal) is ignored.

To address this problem, we treat the extended proposal as three proposals in L-Net and G-Net, *i.e.*, the original proposal and the extended regions on the two sides. The duration of each extended region is 50% as long as that of the original proposal. As shown in Figure 3, L-Net and G-Net process these three proposals separately but use shared weights. Then, we concatenate their new representations. Finally, P-Net, *i.e.*, “P-Net 2” in Figure 3 Top, treats an extended proposal as a single proposal and processes these concatenated features. Note the processing of the original

proposal in L-Net and G-Net is now a part of the processing of the extended proposal. This not only reduces the model size and computational complexity but also reflects the connection between the original proposal and the extended one.

4. Experiments

Datasets. The THUMOS14 [16] dataset contains a large number of open source videos of human actions in real environments. It includes four parts: the training set, verification set, background set, and test set. The training set is based on the UCF-101 dataset [33], and the videos are edited (each video usually contains an instance). For the remaining three parts (verification, background and test sets), the videos are untrimmed. Following the common setting in THUMOS14, we apply 200 videos in the validation set for training and conduct evaluation on the 213 annotated videos from the test set. ActivityNet v1.3 [13] is currently the largest dataset of action analysis in videos, including 20,000 Youtube videos with 200 action categories. The training set contains about 10,000 videos. Both the validation set and the test set contain about 5,000 videos. On average, each video has 1.5 action instances. Following the standard practice, we train our method on the training videos and test it on the validation videos.

Evaluation metric. We evaluate the performance of ContextLoc using mean average precision (mAP) values at different tIoU thresholds. On THUMOS14, the tIoU thresholds are chosen from {0.3, 0.4, 0.5, 0.6, 0.7}. On ActivityNet v1.3, the tIoU thresholds are from {0.5, 0.75, 0.95}, and we also report the average mAP of the tIoU thresholds between 0.5 and 0.95 with a step of 0.05.

Implementation details. We divide each input video into 64-frame snippets and obtain snippet-level features (1024-dimensional feature vectors) via the I3D network [3]. We implement ContextLoc using PyTorch. It is trained on three NVIDIA GeForce GTX TITAN XP GPUs with a batch size of 32. The stochastic gradient descent (SGD) solver is adopted for optimization. On THUMOS14, the initial learning rate is 0.01 and it is divided by 10 after 15 epochs. On ActivityNet v1.3, the initial learning rate is 0.01 for the first 15 epochs and it is divided by 10 after 10 epochs. The final score for calculating mAP is obtained by multiplying the scores of action classification and completeness prediction. The predictions from the RGB and optical flow streams are fused with the ratio of 5:6. Non-Maximum Suppression (NMS) removes duplicated action proposals and obtains the final predicted proposals. Unless otherwise stated, P-GCN is taken as the P-Net.

Loss functions. The action classification of original proposals is trained with the cross-entropy loss, the completeness prediction is trained with the hinge loss, and the boundary regression is trained with the smooth L_1 loss.

Method	RGB	Flow	Fusion	tIoU@0.5
P-Net (non-local [40])	32.16	35.13	39.76	45.81
+L-Net	34.48	38.13	42.64	50.46
+G-Net	34.31	36.97	41.33	47.74
+L-Net + G-Net	35.77	40.07	44.03	51.00
P-Net (P-GCN [44])	34.93	39.26	42.43	49.10
+L-Net	36.40	40.48	45.12	53.17
+G-Net	36.15	40.42	43.51	50.11
+L-Net + G-Net	37.23	42.52	45.70	54.30

Table 1. Ablation study on the effectiveness of L-Net and G-Net. P-Net is taken as a strong baseline. The average mAP (%) of tIoU thresholds from 0.1 to 0.9 and the mAP of tIoU@0.5 are reported on the THUMOS14 test set.

Adapt.	Inter.	RGB	Flow	Fusion	Class%
		34.98	39.45	42.27	76.2
✓		37.15	40.62	44.14	77.8
	✓	35.26	41.05	43.23	76.8
✓	✓	37.23	42.52	45.70	78.5

Table 2. Ablation study on the individual and combinatorial effects of global context adaptation in G-Net (denoted as “Adapt.”) and interactions between adapted global contexts achieved by P-Net (denoted as “Inter.”). The average mAP (%) of tIoU thresholds from 0.1 to 0.9 and the action classification accuracy (Class%) are reported on the THUMOS14 test set.

4.1. Ablation Study

To explore the effectiveness of each component under a controlled setting, we remove or change the corresponding part of ContextLoc in each ablation experiment while keeping the others (not mentioned) the same. All results are reported on THUMOS14.

ContextLoc sub-networks. Since prior work has already proved that modeling inter-proposal relations helps TAL, we take P-Net as a strong baseline to validate the effectiveness of L-Net (local context) and G-Net (global context), which are the main contribution of this paper. We consider two instantiations of P-Net, *i.e.*, P-GCN and the non-local network. The results are shown in Table 1.

When the non-local network is taken as P-Net, adding L-Net alone before it improves 2.88% on average mAP and 4.65% on tIoU 0.5. Adding G-Net alone before P-Net improves 1.57% on average mAP and 1.93% on tIoU 0.5. Our complete ContextLoc (L-Net+G-Net+P-Net) further improves the performance.

When P-GCN is taken as P-Net, adding L-Net alone before it improves 2.69% on average mAP and 4.07% on tIoU 0.5. Adding G-Net alone before P-Net improves 1.08% on average mAP and 1.01% on tIoU 0.5. The complete ContextLoc (L-Net+G-Net+P-Net) further improves the performance.

Comparing the two sets of results achieved via different

Method	Dim.	Fusion	tIoU@0.5	#Params	Flops
Baseline	1024	40.04	45.87	6.5M	3.6G
Baseline	3072	44.57	52.68	25.6M	16.2G
Ours	3072	45.70	54.30	6.7M	3.1G

Table 3. Ablation study on different methods of processing extended proposals (as illustrated in Figure 3). “Dim.” denotes the dimension of features of an extended proposal. The average mAP (%) of tIoU thresholds from 0.1 to 0.9 and the mAP of tIoU@0.5 are reported on the THUMOS14 test set.

Method	Dim.	Fusion	tIoU@0.5	#Params	Flops
P-GCN	1024	42.43	49.10	4.6M	1.7G
Deep P-GCN	1024	42.20	48.76	6.4M	2.6G
Deeper P-GCN	1024	42.03	48.86	8.8M	3.3G
ContextLoc	1024	45.70	54.30	6.7M	3.1G
ContextLoc	512	44.27	52.32	3.4M	1.6G

Table 4. Comparison of model sizes and complexities. Deep P-GCN and Deeper P-GCN respectively add one and two graph convolutional layers to P-GCN. “Dim.” denotes the dimension of features of a proposal, *i.e.*, D . The average mAP (%) is calculated on tIoU thresholds from 0.1 to 0.9. Results are reported on the THUMOS14 test set.

Method	Class%	tIoU%
P-Net (P-GCN)	74.3	78.2
+ L-Net	75.9	81.4
+ G-Net	77.4	79.1
+ L-Net + G-Net	78.5	82.0

Table 5. Ablation study on classification and localization. The action classification accuracy (Class%) and tIoU w.r.t. ground truth (tIoU%) after regression are reported on the THUMOS14 test set.

instantiations of P-Net, we can observe that P-GCN outperforms the non-local network regardless of whether P-Net works alone or is integrated with other sub-networks. The reason is that the non-local network is not specially designed for TAL. Therefore, in all other experiments, we will use P-GCN as the P-Net.

Several conclusions can be drawn from this experiment. (1) Either the local context modeled via L-Net or the global context modeled via G-Net benefits TAL. (2) The local context, global context and inter-proposal relations are complementary. Our ContextLoc unifies their advantages and achieves the best performance. (3) The effectiveness of L-Net, G-Net and ContextLoc does not depend on any specific instantiation of P-Net. They consistently improve the performance over the strong baselines.

Global context adaptation & interactions between adapted global contexts. This ablation study investigates the individual and combinatorial effects of global context adaption in G-Net and interactions between adapted global contexts achieved by P-Net. To deactivate the latter, we

let global context aggregation occur after the proposal-level features pass through P-Net. Specifically, we first send proposal-level features obtained from L-Net, *i.e.*, $\{\mathbf{y}_i^L\}$, to P-Net and then concatenate the new representations of each proposal with their corresponding adapted global contexts, *i.e.*, $\{\mathbf{z}_i^G\}$.

The results are shown in Table 2. We can see that both global context adaptation and interactions between adapted global contexts improve the performance. The former benefits the RGB stream more than the flow stream while the latter play a more important role in the flow stream. They together help ContextLoc achieve the best performance.

Extended proposals. Prior approaches process the extended proposal as a single proposal. We instead process it as three proposals in L-Net and G-Net and process the concatenation of updated features in P-Net. Results are shown in Table 3. When the feature dimension of an extended proposal in the baseline method is the same as that of the original proposal, *i.e.*, 1024, its performance is much worse than that of our new method. When the feature dimension of an extended proposal in the baseline method is much larger than that of the original proposal, *i.e.*, 3072, its performance gets better (with significantly more parameters and flops) but is still inferior to ours.

Model size & complexity. The model size and computational complexity of our ContextLoc are 6.7M and 3.1G flops, respectively. In contrast, P-GCN has 4.6M parameters and 1.7G flops. This is expected because ContextLoc includes P-Net as a building component. To have a fair comparison, we construct a Deep P-GCN and a Deeper P-GCN by adding one and two graph convolutional layers to P-GCN, respectively. Table 4 shows that naively increasing the model complexity and network depth can not improve the performance. This further proves that the local and global contexts we advocate is complementary to the inter-proposal relations.

We also experiment on a light version of ContextLoc by halving the feature dimensions. With fewer parameters (3.4M versus 4.6M) and flops (1.6G versus 1.7G) than P-GCN, it outperforms P-GCN by 3.22% at tIoU@0.5. This again verifies that the effectiveness of our ContextLoc is not caused by a larger model size or more flops.

Concrete analysis. Table 5 reports the classification and localization accuracy of high-quality proposals, whose tIoU w.r.t. ground truth are greater than 0.7. The results indicate L-Net contributes more to localization (tIoU%) and G-Net more to classification, which function as expected.

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

THUMOS14. As summarized in Table 6, our ContextLoc outperforms all the compared methods on the THUMOS14 test set. At tIoU 0.5, ContextLoc outperforms the previously best method PBRNet [23] by 3.0% absolute im-

Method	0.3	0.4	0.5	0.6	0.7
CDC [30]	40.1	29.4	23.3	13.1	7.9
TCN [7]	-	33.3	25.6	15.9	9.0
TURN-TAP [11]	44.1	34.9	25.6	-	-
R-C3D [41]	44.8	35.6	28.9	-	-
SS-TAD [2]	45.7	-	29.2	-	9.6
SSN [46]	51.9	41.0	29.8	-	-
BSN [21]	53.5	45.0	36.9	28.4	20.0
MGG [25]	53.9	46.8	37.4	29.5	21.3
BMN [20]	56.0	47.4	38.8	29.7	20.5
G-TAD [42]	54.5	47.6	40.2	30.8	23.4
TAL-Net [4]	53.2	48.5	42.8	33.8	20.8
Huang <i>et al.</i> [14]	60.2	54.1	44.2	32.3	19.1
Zhao <i>et al.</i> [45]	53.9	50.7	45.4	38.0	28.5
P-GCN [44]	63.6	57.8	49.1	-	-
PBRNet [23]	58.5	54.6	51.3	41.8	29.5
ContextLoc (ours)	68.3	63.8	54.3	41.8	26.2


Table 6. Results on the THUMOS14 test set. The mAP (%) at different tIoU thresholds are reported. **Bold** fonts indicate the best performance.

Method	0.5	0.75	0.95	Average
R-C3D [41]	26.80	-	-	-
TAL-Net [4]	38.23	18.30	1.30	20.22
SSN [46]	43.26	28.70	5.63	28.28
BSN [21]	46.45	29.96	8.02	30.03
BMN [20]	50.07	34.78	8.29	33.85
G-TAD [42]	50.36	34.60	9.02	34.09
Zhao <i>et al.</i> [45]	43.47	33.91	9.21	30.12
PBRNet [23]	53.96	34.97	8.98	35.01
P-GCN [44]	42.90	28.14	2.47	26.99
ContextLoc (ours)	51.24	31.40	2.83	30.59
P-GCN* [44]	48.26	33.16	3.27	31.11
ContextLoc* (ours)	56.01	35.19	3.55	34.23

Table 7. Results on the ActivityNet v1.3 validation set. The mAP (%) at different tIoU thresholds and the average mAP of IoU thresholds from 0.5 to 0.95 are reported. Symbol (*) denotes using external video-level classes predictions. **Bold** fonts indicate the best performance.

provement. This demonstrates the benefit of enriching local and global contexts for TAL.

ActivityNet v1.3. The results on ActivityNet v1.3 are shown in Table 7. Zeng *et al.* [44] found taking into account video-level classes predicted by UntrimmedNet [39] helps improve the performance of P-GCN. We also adopt the same strategy to obtain ContextLoc*. Specifically, we add a branch of video-level representation z to predict the video-level classes. More details about ContextLoc* can be found in the supplementary material. At tIoU 0.5, ContextLoc* reaches an mAP of 56.01% which is 2.05% higher than the



(a) Qualitative results on an example from the THUMOS14 test set (top) and an example from the ActivityNet v1.3 validation set (bottom) show that our ContextLoc locates the temporal boundaries more accurately than P-GCN.



(b) Qualitative results on an example from the THUMOS14 test set (top) and an example from the ActivityNet v1.3 validation set (bottom) show that our ContextLoc correctly classifies action instances which P-GCN fails on.

Figure 4. Results from the ground truth, ContextLoc and P-GCN are respectively illustrated using blue, orange and green bars.

current best 53.96% achieved by PBRNet [23]. At tIoU 0.75, ContextLoc* outperforms the previously best method PBRNet by 0.22%. At tIoU 0.95, ContextLoc* does not perform as well as previous methods and only marginally outperforms P-GCN. In the supplementary material, we visualized some results to explain this gap.

4.3. Qualitative Results

To validate the effectiveness of temporal boundary location and action classification, we visualize qualitative results in Figure 4(a) and Figure 4(b), respectively. Compared with P-GCN, ContextLoc predicts a more precise temporal boundary and classification for these examples.

5. Conclusion

This paper introduces a novel network architecture, termed *ContextLoc*, for TAL. It models the local context, global context and context-aware inter-proposal relations in a unified framework. Ablation experiments under controlled settings indicate the effectiveness of each component of ContextLoc. Results on two datasets demonstrate that ContextLoc outperforms state-of-the-art TAL methods.

Acknowledgment. This work was supported partly by National Key R&D Program of China Grant 2018AAA0101400, NSFC Grants 62088102, 61976171, and XXXXXXXX, and Young Elite Scientists Sponsorship Program by CAST Grant 2018QNRC001.

References

- [1] Yueran Bai, Yingying Wang, Yunhai Tong, Yang Yang, Qiyue Liu, and Junhui Liu. Boundary content graph neural network for temporal action proposal generation. In *ECCV*, pages 121–137, 2020.
- [2] Shyamal Buch, Victor Escorcia, Bernard Ghanem, Li Fei-Fei, and Juan Carlos Niebles. End-to-end, single-stream temporal action detection in untrimmed videos. In *BMVC*, 2017.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017.
- [4] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018.
- [5] Yunze Chen, Mengjuan Chen, Rui Wu, Jiagang Zhu, Zheng Zhu, and Qingyi Gu. Refinement of boundary regression using uncertainty in temporal action localization. In *BMVC*, 2020.
- [6] Zhao-Min Chen, Xin Jin, Borui Zhao, Xiu-Shen Wei, and Yanwen Guo. Hierarchical context embedding for region-based object detection. In *ECCV*, pages 633–648, 2020.
- [7] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S. Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *ICCV*, pages 5727–5736, 2017.
- [8] Chenyou Fan, Xiaofan Zhang, Shu Zhang, Wensheng Wang, Chi Zhang, and Heng Huang. Heterogeneous memory enhanced multimodal attention model for video question answering. In *CVPR*, pages 1999–2007, 2019.
- [9] Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti. Action recognition in the premotor cortex. *Brain : a journal of neurology*, 119 (Pt 2):593–609, 1996.
- [10] Jiyang Gao, Kan Chen, and Ram Nevatia. Ctap: Complementary temporal action proposal generation. In *ECCV*, pages 68–83, 2018.
- [11] Jiyang Gao, Zhenheng Yang, Chen Sun, Kan Chen, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. In *ICCV*, pages 3648–3656, 2017.
- [12] Dongliang He, Zhichao Zhou, Chuang Gan, Fu Li, Xiao Liu, Yandong Li, Limin Wang, and Shilei Wen. Stnet: Local and global spatial-temporal modeling for action recognition. In *AAAI*, pages 8401–8408, 2019.
- [13] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- [14] Yupan Huang, Qi Dai, and Yutong Lu. Decoupling localization and classification in single shot temporal action detection. In *ICME*, pages 1288–1293, 2019.
- [15] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:221–231, 2013.
- [16] Yu-Gang Jiang, Haroon Idrees, Amir Zamir, Alex Gorban, Ivan Laptev, Mubarak Shah, and Rahul Sukthankar. Thumos challenge: Action recognition with a large number of classes. 2014.
- [17] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *ArXiv*, abs/1705.06950, 2017.
- [18] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, pages 8545–8552, 2019.
- [19] Chuming Lin, Jian Li, Yabiao Wang, Ying Tai, Donghao Luo, Zhipeng Cui, Chengjie Wang, Jilin Li, Feiyue Huang, and Rongrong Ji. Fast learning of temporal action proposal via dense boundary generator. In *AAAI*, 2020.
- [20] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3888–3897, 2019.
- [21] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, pages 3–19, 2018.
- [22] Chenchen Liu, Yang Jin, Kehan Xu, Guoqiang Gong, and Yadong Mu. Beyond short-term snippet: Video relation detection with spatio-temporal global context. In *CVPR*, pages 10837–10846, 2020.
- [23] Qinying Liu and Zilei Wang. Progressive boundary refinement network for temporal action detection. In *AAAI*, 2020.
- [24] Yang Liu, Zhaoyang Lu, Jing Li, Tao Yang, and Chao Yao. Global temporal representation based cnns for infrared action recognition. *IEEE Signal Processing Letters*, 25(6):848–852, 2018.
- [25] Yuan Liu, Lin Ma, Yifeng Zhang, Wei Liu, and Shih-Fu Chang. Multi-granularity generator for temporal action proposal. In *CVPR*, pages 3604–3613, 2019.
- [26] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *CVPR*, pages 344–353, 2019.
- [27] Thien Nguyen and Ralph Grishman. Graph convolutional networks with argument-aware pooling for event detection. In *AAAI*, pages 5900–5907, 2018.
- [28] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, pages 12056–12065, 2019.
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, 2016.
- [30] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 1417–1426, 2017.
- [31] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, pages 1049–1058, 2016.
- [32] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, pages 568–576, 2014.
- [33] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *ArXiv*, abs/1212.0402, 2012.

- [34] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [35] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5999–6009, 2017.
- [37] Heng Wang, Dan Oneata, Jakob Verbeek, and Cordelia Schmid. A robust and efficient video representation for action recognition. *International Journal of Computer Vision*, 119(3):219–238, 2016.
- [38] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, pages 3551–3558, 2013.
- [39] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, pages 4325–4334, 2017.
- [40] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018.
- [41] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *ICCV*, pages 5794–5803, 2017.
- [42] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, pages 10153–10162, 2020.
- [43] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016.
- [44] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *CVPR*, pages 7094–7103, 2019.
- [45] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *ECCV*, pages 539–555, 2020.
- [46] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, pages 2914–2923, 2017.