

Video Activity Localisation with Uncertainties in Temporal Boundary

Jiabo Huang^{1,4}, Hailin Jin², Shaogang Gong¹, and Yang Liu^{*3,5}

¹ Queen Mary University of London
{jiabo.huang, s.gong}@qmul.ac.uk

² Adobe Research
hljin@adobe.com

³ Wangxuan Institute of Computer Technology, Peking University
yangliu@pku.edu.cn

⁴ Vision Semantics Limited

⁵ Beijing Institute for General Artificial Intelligence

Abstract. Current methods for video activity localisation over time assume implicitly that activity temporal boundaries labelled for model training are determined and precise. However, in unscripted natural videos, different activities mostly transit smoothly, so that it is intrinsically ambiguous to determine in labelling precisely when an activity starts and ends over time. Such uncertainties in temporal labelling are currently ignored in model training, resulting in learning mis-matched video-text correlation with poor generalisation in test. In this work, we solve this problem by introducing Elastic Moment Bounding (EMB) to accommodate flexible and adaptive activity temporal boundaries towards modelling universally interpretable video-text correlation with tolerance to underlying temporal uncertainties in pre-fixed annotations. Specifically, we construct elastic boundaries adaptively by mining and discovering frame-wise temporal endpoints that can maximise the alignment between video segments and query sentences. To enable both more accurate matching (segment content attention) and more robust localisation (segment elastic boundaries), we optimise the selection of frame-wise endpoints subject to segment-wise contents by a novel Guided Attention mechanism. Extensive experiments on three video activity localisation benchmarks demonstrate compellingly the EMB’s advantages over existing methods without modelling uncertainty.

1 Introduction

The goal of video activity localisation is to locate temporally video moments-of-interest (MoIs) of a specific activity described by a natural language query of an untrimmed continuous long video (often unscripted and unstructured) that contains many different activities [28,29,21].

One straightforward solution to the task, denoted as proposal-free methods (Fig. 1 (b)), is to predict directly the start and end frames of a target moment

* Corresponding author.




Fig. 1. An illustration of different activity localisation methods. (a) Activity’s temporal boundaries are intrinsically uncertain in manual labelling (break-down and highlighted in ‘red’). (b) Proposal-free methods learn to identify the frame-wise temporal endpoints. (c) Proposal-based methods learn the holistic alignment of video segments and query sentences between their feature spaces. (d) *Elastic Moment Bounding* (EMB) optimises simultaneously endpoints selection with maximisation of segment content agreement between visual and textual representations.

that align to the given query [27,7,28,14]. Such paradigm deploys directly the fixed manual activity endpoints labels for model training, implicitly assuming these labels are well-defined and ignoring uncertainties in the labels. However, unlike labelling object spatial bounding-boxes, there is a considerable variation in how activities occur in unconstrained scenarios. There may not even be a precise definition of the exact temporal extent of an activity. Fitting such uncertain temporal endpoints will inevitably lead to semantically mis-matched visual-textual correlations which are not universally interpretable and result in poor generalisation in test. For example, the two queries Q_1 and Q_2 in Fig. 1 (a) are semantically similar in describing ‘putting on shoes’. Nonetheless, the annotated activity (gray bars on top) for Q_1 starts from putting down a box while Q_2 begins with sitting down on a sofa. By training a model with such uncertain temporal endpoints (Fig. 1 (b)), the model is trained to match notably different visual features of ‘putting down a box’ and ‘sitting down on a sofa’ with the same query on ‘putting on shoes’. Clearly, the model suffers from poor learning due to uncertainty in visual cues. Moreover, as observed in [12], *the annotation bias can be inconsistent from different annotators*. Giving the same videos and query sentences to 5 different annotators, only 42% and 35% of their annotated activity boundary are mutually agreed (with at least 50% IoU) on Charades-STA [5] and ActivityNet-Captions [10], respectively. This highlights the extent of activity label uncertainties in model training inherent to the current proposal-free methods, and the potential significant misinformation in training such models. Another solution (Fig. 1 (c)) is to generate many candidate proposals for a target moment and aligns segment-level video features with the query sentences [1,5,29]. By formulating the localisation task as a matching problem, the proposal-based methods consider alignment by the whole moment with less focus on the exact boundary matching [22,21]. By doing so, it can be less sensitive to the boundary

labels but more reliance on salient content (attention) between proposals and the target segment. This can make them more tolerant to the uncertainties in temporal annotations. However, the problem of detecting accurately the start- and end-point of a target activity moment remains unsolved especially when learning without exhaustive moment proposals for efficiency concerns.

In this work, we introduce *Elastic Moment Bounding* (EMB) to address the limitation of proposal-free paradigm by modelling explicitly the label uncertainty in the temporal boundaries of an activity moment. Instead of forcing a model to fit manually labelled *rigid* activity endpoints, each MoIs are modelled by an elastic boundary with a set of candidate endpoints. The model then learns to select optimally from consistent visual-textual correlations among semantically similar activities. This introduces model robustness to label uncertainty. Specifically, we conduct a proposal-based segment-wise content alignment in addition to learning of frame-wise boundary identification. As the predicted segment is required to be highly aligned with the query textual description, we represent the gap between the predicted endpoints and the manual labelled endpoints as an elastic boundary (Fig. 1 (d)). This process imposes explicitly label uncertainties to model training. To enable activity localisation to be both more attention driven (accurate) and sensitive to an elastic boundary (robust), we introduce an interaction between the segment-wise content representations and frame-wise boundary features by assembling representations through a *Guided Attention* mechanism. The segment-wise boundary-guided attention helps minimise redundant frames in each elastic boundary whilst the frame-wise content-guided attention highlighting transitional frames with apparent visual changes indicating the potential start and end points of an activity.

We make three contributions in this work: (1) We introduce a model to explore collaboratively both proposal-free and proposal-based mechanisms for learning to detect more accurate activity temporal boundary localisation *when training labels are inherently uncertain*. We formulate a new *Elastic Moment Bounding* (EMB) method to expand a manually annotated single pair of fixed activity endpoints to an elastic set. (2) To reinforce directly robust content matching (the spirit of proposal-based) as a condition to accurate endpoints localisation (the spirit of proposal-free) of activities in videos, we introduce a Guided Attention mechanism to explicitly optimise frame-wise boundary visual features subject to segment-wise content representations and vice versa, so to minimise redundant frames in each elastic boundary whilst highlighting frames signalling activity transitions subject to segment content holistically. (3) Our EMB model provides a state-of-the-art performance on three video activity localisation benchmark datasets, improving existing models that suffer from sensitivity to uncertainties in activity training labels.

2 Related work

Proposal-based content alignment. By aggregating all the frames within a video segment and aligning them holistically with the query sentences, the

segment-wise approaches [1,5,29,30,6] are insensitive to the boundary as its most salient and semantically aligned parts are not necessarily at its two ends. The endpoint frames play a significant role to help differentiate video moments from their overlapping counterparts containing redundant frames, hence, critical for video activity localisation. Therefore, we explicitly associate the segment-wise content information with the frame-wise boundary information and complement them by each other through a novel guided attention mechanism.

Proposal-free boundary identification. In contrast, the proposal-free methods learn to directly regress the start and end timestamps of the target moments or predict the per-frame probabilities of being the endpoints [4,27,7,28,14,33,31,11]. In either case, they take the temporal boundaries provided manually as the oracles for learning exactly the same predictions. However, this is prone to be misled by the uncertainty in manual labels and results in less generalisable models. To cope with that, we train our EMB model to identify the target boundary from reliable candidate start and end spans (sets of frames) rather than fitting the single pre-fixed manual endpoints, so as to derive consistent video-text correlations from semantically similar activities that are universally interpretable.

Joint content-boundary learning. There are a few recent attempts [22,21,23] on locating video activity jointly by the proposal-based and proposal-free strategies. They mostly explored the interaction of frame’s and segment’s feature representations for better video comprehension. In this work, we study the combination of the two strategies for attention learning of activity temporal boundary conditions beyond feature learning for activity representation. We augment the fixed manual labels by the video segments selected according to their content alignments with query sentences to help improve the robustness of temporal endpoints identification when there is boundary uncertainty.

Temporal boundary uncertainty. Recently, Otani *et al.* [12] quantitatively studied the label uncertainty problems on video activity localisation by collecting multiple boundaries for the same activities from different annotators, the results highlighted the extent of uncertainty in the temporal annotations. However, Otani *et al.* [12] did not explicitly propose a solution to the problem. DeNet [33], on the other hand, addressed it w.r.t. the variety of language descriptions, *i.e.*, the same video activity can be described semantically in different ways. They generated different copies of the same query sentences by perturbing the “modified” phrases (adjective, adverb and *etc.*) so to predict diverse boundaries for the same video activities. Rather than studying the uncertainty from the perspective of semantic description, we analyse the uncertainties in activity temporal boundary annotations, which is intrinsically harder to avoid.

3 Learning Localisation with Uncertainty

Given the feature representations of an untrimmed videos \mathbf{F} composed of T frames, and that of a natural language sentence \mathbf{Q} of L words, the objective of video activity localisation is to identify the temporal boundary of a target moment (S, E) – activity endpoints – so that the video segment $\{\mathbf{f}_t\}_{t=S}^E$ matches




Fig. 2. An overview of *Elastic Moment Bounding* (EMB). (a) The EMB model takes the pretrained 3D-CNN features and the GloVe embeddings as inputs. It consists of an ‘Alignment branch’ (red dashed box) learning the semantic content alignment of video segments and query sentences, and a ‘Bounding branch’ (green dashed box) to predict the temporal endpoints of target activity moments. Both branches are subject to (b) a Guided Attention mechanism with both self- and cross-modalities attention in (c) a Multi-grained Interaction Network (MGIN). (d) The most confidently aligned video segments predicted by the alignment branch are then selected for constructing the elastic boundary to optimise the endpoints predictions from the bounding branch.

with Q in semantics. It is challenging to acquire high-level semantic understandings of either videos or sentences, let alone aligning them to precisely locate the temporal endpoints of a specific activity instance.

In this work, we study the problem of model learning subject to temporal label uncertainty which is inherent to manual video annotation and more importantly not shared in unseen new test videos or language descriptions. To that end, we propose an *Elastic Moment Bounding* (EMB) model (Fig. 2). The EMB model first predicts the per-frame probabilities to be the temporal endpoints of a target moment (Fig. 2’s green dashed box) by a *Multi-grained Interaction Network* (MGIN) (Fig. 2 (c)) incorporating with a *Guided Attention* mechanism (Fig. 2 (b)). EMB then optimises the frame-wise probabilities by mining multiple candidate endpoints beyond the manual annotated ones. The candidate endpoints are discovered by an auxiliary alignment branch (Fig. 2’s red dashed box). The alignment branch explores the visual-textual content alignment at segment-level, which is less sensitive to exact endpoints annotations so more robust to uncertainty. By doing so, we construct an *elastic boundary* interpretable universally for semantically similar activities with endpoints uncertainty.

3.1 Temporal Endpoints Identification

Our elastic moment bounding is a generic formulation deployable in any multi-modal backbone deep networks. Here, we start with the VSLNet [28] and recon-

struct it by introducing a Guided Attention mechanism to form a multi-grained interaction network (MGIN). The overall pipeline of MGIN is shown in Fig. 2 (c) to first encode the video \mathbf{F} and the sentence \mathbf{Q} by attention both within (self) and across modalities, then predict the frame-wise endpoint probabilities by the joint-modal representations fused by context-query attention [18,24,25].

Adopting the convention [28,14,22], we represent video frames by a pre-trained 3D-CNN model [3] as $\mathbf{F} = \{\mathbf{f}_i\}_{i=1}^T \in \mathbb{R}^{D^v \times T}$ and the query sentence by the GloVe embeddings [15] of words $\mathbf{Q} = \{\mathbf{w}_l\}_{l=1}^L \in \mathbb{R}^{D^q \times L}$. To facilitate cross-modal feature interactions, we map both the representations to have the same dimension D by two independent linear projections, *i.e.*, $\mathbf{F} \leftarrow \text{FC}(\mathbf{F}) \in \mathbb{R}^{D \times T}$ and $\mathbf{Q} \leftarrow \text{FC}(\mathbf{Q}) \in \mathbb{R}^{D \times L}$.

Vision-language attention representation. We deploy attentive encoding [20,9] for both the visual and textual representations to explore the dependencies among elements in both. In general, to encode a target sequence $\mathbf{X}^t \in \mathbb{R}^{D \times L^t}$ of L^t elements with the help of a reference sequence $\mathbf{X}^r \in \mathbb{R}^{D \times L^r}$ in size L^r , we first compute an attention matrix \mathbf{A} indicating the pairwise target-reference correlations, then represent each target element by its correlated references:

$$\mathbf{A} = \text{FC}(\mathbf{X}^t)^\top \text{FC}(\mathbf{X}^r) / \sqrt{D} \in \mathbb{R}^{L^t \times L^r} \quad (1)$$

$$g(\mathbf{X}^t, \mathbf{X}^r) = \mathbf{X}^t + \text{FC}(\mathbf{X}^r) \text{Softmax}(\mathbf{A})^\top \in \mathbb{R}^{D \times L^t}. \quad (2)$$

An attention layer formulated in Eq. (2) is parameterised by three independent fully-connected layers. Our MGIN shown in Fig. 2 (c) is constructed by both self-attention within modalities: $\mathbf{F} \leftarrow g(\mathbf{F}, \mathbf{F})$, $\mathbf{Q} \leftarrow g(\mathbf{Q}, \mathbf{Q})$ for context exploration and cross-attention between modalities: $\mathbf{F} \leftarrow g(\mathbf{F}, \mathbf{Q})$, $\mathbf{Q} \leftarrow g(\mathbf{Q}, \mathbf{F})$ to learn the semantic correlations between video frames and query words.

Guided attention. To effectively locate the temporal endpoints of activities, it is essential for the model to be aware of not only what is shown in each individual frame but also what’s different before and after it. As a simple example, the starting point of an activity ‘person puts on shoes’ should not be arbitrary frames involving shoes-like objects in-between the period but be consistent like when the shoes first appear to interact with the person. Therefore, we propose a content-guided attention module (Fig. 2 (b)’s bottom) to explicitly encode the preceding and subsequent content information of each frame into its representation:

$$\begin{aligned} \mathbf{F}_{\text{pre}} &= \{\text{MaxPool}(\{\mathbf{f}_i\}_{i=1}^t)\}_{t=1}^T \in \mathbb{R}^{D \times T}, \\ \mathbf{F}_{\text{sub}} &= \{\text{MaxPool}(\{\mathbf{f}_i\}_{i=t}^T)\}_{t=1}^T \in \mathbb{R}^{D \times T}, \\ \tilde{\mathbf{F}} &= \text{Conv2d}(\{\mathbf{F}, \mathbf{F}_{\text{pre}}, \mathbf{F}_{\text{sub}}\}) \in \mathbb{R}^{D \times T}. \end{aligned} \quad (3)$$

The feature $\text{MaxPool}(\{\mathbf{f}_i\}_{i=1}^t) \in \mathbb{R}^D$ in Eq. 3 aggregate all the frames before \mathbf{f}_t by max-pooling as its preceding content representation. Similarly, the subsequent content of the t -th frame is obtained by $\text{MaxPool}(\{\mathbf{f}_i\}_{i=t}^T)$. Both the preceding \mathbf{F}_{pre} and subsequent \mathbf{F}_{sub} content features are then stacked and assembled with the frame-wise representations \mathbf{F} by a 2D convolution layer. After that, the

content-guided representations of video frames $\hat{\mathbf{F}}$ are used for attentive encoding (Eq. (2)) both within $\mathbf{F} \leftarrow g(\tilde{\mathbf{F}}, \hat{\mathbf{F}})$ and across modalities $\mathbf{F} \leftarrow g(\tilde{\mathbf{F}}, \mathbf{Q})$.

Boundary prediction. Given a video $\mathbf{F} \in \mathbb{R}^{D \times T}$ and sentence $\mathbf{Q} \in \mathbb{R}^{D \times L}$ representations, we estimate the frame-wise endpoint probabilities by computing context-query attention [18,24,25], same as the baseline [28]. It is defined as

$$\begin{aligned} (\mathbf{p}^s, \mathbf{p}^e) &= \text{Softmax}(\text{LSTM}(\hat{\mathbf{F}} \odot \mathbf{h})), \text{ where } \mathbf{h} = \sigma(\text{Conv1d}(\hat{\mathbf{F}} \parallel \mathbf{q})) \in \mathbb{R}^{1 \times T}, \\ \hat{\mathbf{F}} &= H(\mathbf{F}, \mathbf{Q}) = \text{FC}(\mathbf{F} \parallel \mathbf{X}^{v2q} \parallel \mathbf{F} \odot \mathbf{X}^{v2q} \parallel \mathbf{F} \odot \mathbf{X}^{q2v}) \in \mathbb{R}^{D \times T}; \text{ and} \\ \mathbf{A} &= \frac{\text{FC}(\mathbf{F})^\top \text{FC}(\mathbf{Q})}{\sqrt{D}}, \quad \mathbf{X}^{v2q} = \mathbf{Q} \mathbf{A}^{r \top}, \quad \mathbf{X}^{q2v} = \mathbf{F} \mathbf{A}^r \mathbf{A}^{c \top}. \end{aligned} \quad (4)$$

In Eq. (4), we predict the frame-wise endpoint probabilities by two stacked LSTM. This is based on fusing the two modalities \mathbf{F} and \mathbf{Q} by function H then rescale the per-frame fused feature $\hat{\mathbf{F}} \in \mathbb{R}^{D \times T}$ using their estimated likelihood $\mathbf{h} \in \mathbb{R}^{1 \times T}$ of being foreground to suppress any distractions from redundant frames. Matrix $\mathbf{A} \in \mathbb{R}^{T \times L}$ consists of frame-to-word correlation scores; \mathbf{A}^r and \mathbf{A}^c are its row and column-wise softmax normalised copies. The \mathbf{q} are the sentence-level representations from weighted sum of words [2]; $(\cdot \parallel \cdot)$ stands for concatenation (broadcast if necessary) while \odot is the Hadamard Product.

3.2 Elastic Moment Bounding

Given the uncertainty and ambiguity in manually annotated activity temporal boundaries, it is ineffective to decide heuristically and universally which frames and how many of them should be taken as the candidate endpoints ($\tilde{\mathbf{S}}, \tilde{\mathbf{E}}$) for different video activities. To address this problem, we formulate an auxiliary alignment branch in the model to learn the visual-textual content mapping per each video segment. It serves as an additional self-learning ‘‘annotator’’ to expand the given single pair of manually annotated boundaries into candidate endpoints proposal sets tailored for individual activities.

Elastic boundary construction. As shown in Fig. 2’s red dashed box, we first generate a 2D feature map [29] by enumerating pairwise start-end frames to represent $K = T \times T$ video segments $\mathbf{V} = \{\mathbf{v}_k\}_{k=1}^K \in \mathbb{R}^{D \times K}$ as the proposals for a target moment. We flatten the 2D map here for clarity. The k -th proposal with the temporal boundary of (t_k^s, t_k^e) is represented by max-pooling the frames it is composed of $\mathbf{v}_k = \text{MaxPool}(\{\mathbf{f}_t \mid t \in [t_k^s, t_k^e]\})$. The segment-wise representations will then be fed into an independent MGIN equipped with *boundary-guided* attention modules (Fig. 2 (b)’s top) for visual encoding. Similar as in the *content-guided* attention for video frames, we explicitly assemble the frame-wise boundary features with the content representations of video segments to encourage boundary-sensitive content alignment:

$$\begin{aligned} \mathbf{V}_{\text{sta}} &= \{\mathbf{f}_{t_k^s}\}_{k=1}^K \in \mathbb{R}^{D \times K}, \quad \mathbf{V}_{\text{end}} = \{\mathbf{f}_{t_k^e}\}_{k=1}^K \in \mathbb{R}^{D \times K}, \\ \tilde{\mathbf{V}} &= \text{Conv2d}(\{\mathbf{V}, \mathbf{V}_{\text{sta}}, \mathbf{V}_{\text{end}}\}) \in \mathbb{R}^{D \times K}. \end{aligned} \quad (5)$$

The features \mathbf{V}_{sta} and \mathbf{V}_{end} in Eq. (5) are the representations of the start and end frames for each of the K proposals. They are stacked and assembled with the segment-wise content features \mathbf{V} to derive the boundary guided segment representations $\tilde{\mathbf{V}}$ by a 2D convolution layer. Such boundary-guided attention share a similar spirit with temporal pyramid pooling [32], that is to explicitly encode the temporal structure into segment’s representation so to be sensitive to its boundary. $\tilde{\mathbf{V}}$ is then used for attentive encoding (Eq. (2)) within $\mathbf{V} \leftarrow g(\tilde{\mathbf{V}}, \tilde{\mathbf{V}})$ and across $\mathbf{V} \leftarrow g(\tilde{\mathbf{V}}, \mathbf{Q})$ modalities.

Given the segment-level video representations \mathbf{V} , we fuse them with the sentence features by H defined in Eq. (4), and re-arrange it to be a 2D feature map then predict the per-proposal alignment scores by a 2D convolution layer:

$$\mathbf{p}^a = \sigma(\text{Conv2d}(H(\mathbf{V}, \mathbf{Q}))) \text{ s.t. } p_k^a \in (0, 1) \forall k \in [1, K]. \quad (6)$$

The segment-wise alignment scores \mathbf{p}^a activated by the Sigmoid function σ is then supervised by the temporal overlaps between every proposals and the manual boundary:

$$\begin{aligned} \alpha_k &= \text{IoU}((t_k^s, t_k^e), (S, E)) \\ \mathbf{y}_k^a &= \begin{cases} 1, & \text{if } \alpha_k \geq \tau_u \\ 0, & \text{if } \alpha_k < \tau_l \\ \alpha_k, & \text{otherwise} \end{cases} \\ \mathcal{L}_{\text{align}}(\mathbf{V}, \mathbf{Q}, S, E) &= \text{BCE}(\mathbf{y}^a, \mathbf{p}^a). \end{aligned} \quad (7)$$

The notations τ_u and τ_l are the upper and lower overlap thresholds to control the flexibility of video-text alignment, which are set to 0.7 and 0.3 respectively as in [29]. With the learned segment-wise alignment scores \mathbf{p}^a , we take the boundary $(t_{k^*}^s, t_{k^*}^e)$ of the most confident proposal with the greatest predicted score $p_{k^*}^a \geq p_k^a \forall k \in [1, K]$ as the pseudo boundary and construct the corresponding candidate endpoint sets by:

$$\tilde{\mathbf{S}} = [\min(t_{k^*}^s, S), \max(t_{k^*}^s, S)], \quad \tilde{\mathbf{E}} = [\min(t_{k^*}^e, E), \max(t_{k^*}^e, E)]. \quad (8)$$

We customise the candidate endpoint sets for every individual activity by exploring the content alignments between video segments and query sentences, *i.e.*, elastic boundary. This is intuitively more reliable than applying label smoothing globally [22,23] without considering video context and language semantics.

Reliability vs. flexibility. Introducing too many candidate endpoints that are semantically irrelevant to the query sentences is prone to distracting the model from learning effective visual-textual correlations, especially at the early stage of training a randomly initialised model which is likely to yield inaccurate pseudo boundaries $(t_{k^*}^s, t_{k^*}^e)$. Therefore, we balance the reliability and flexibility of our elastic boundary by a controllable threshold τ :

$$k^* = \arg \max_k \mathbf{p}^a \text{ s.t. } \alpha_k \geq \tau. \quad (9)$$

The α_k in Eq. (9) implies the overlap between the k -th proposal and the manual boundary, whilst the threshold τ serving as a tradeoff between flexibility and

reliability so that only the sufficiently overlapped proposals will be selected for constructing the elastic boundary in Eq. (8).

Learning from elastic boundary. With the elastic boundary $(\tilde{\mathbf{S}}, \tilde{\mathbf{E}})$, we formulate the boundary supervision signals to maximise the sum of the candidate endpoint’s probabilities obtained in Eq. (4):

$$\mathcal{L}_{\text{bound}}(\mathbf{F}, \mathbf{Q}, S, E) = -\log\left(\sum_{t \in \tilde{\mathbf{S}}} p_t^s\right) - \log\left(\sum_{t \in \tilde{\mathbf{E}}} p_t^e\right). \quad (10)$$

Comparing with the commonly adopted frame-wise supervision which trains \mathbf{p}^s and \mathbf{p}^e to be one-hot [28,14], we provide in Eq. (10) a more flexible boundary to the target moments so that the model can learn in a data-driven manner to select the endpoints beyond the manual boundary and ignore the unconcerned actions involved.

3.3 Model Training and Inference

Inference. We consider two scenarios when predicting the boundary of video activity: **(a) DET:** following the standard protocol of the task [5,8], we predict a determined boundary enclosed by a single start and end frames according to the outputs of bounding branch in a maximum likelihood manner

$$\hat{S} = \arg \max_t \mathbf{p}^s, \quad \hat{E} = \arg \max_t \mathbf{p}^e, \quad (11)$$

where \hat{S} and \hat{E} are the predicted start and end frame indices of a video that are corresponding to a given query. **(b) ELA:** considering the uncertain nature of temporal boundary, it is more intuitive to estimate the endpoints of video activity by temporal spans rather than specific frames. Our model is able to predict also an elastic boundary in a similar way as in training:

$$\hat{\mathbf{S}} = [\min(t_{k^*}^s, \hat{S}), \max(t_{k^*}^s, \hat{S})], \quad \hat{\mathbf{E}} = [\min(t_{k^*}^e, \hat{E}), \max(t_{k^*}^e, \hat{E})]. \quad (12)$$

In Eq. (12), we denote $\hat{\mathbf{S}}$ and $\hat{\mathbf{E}}$ in bold to indicate a set of candidate endpoints, and differentiate them from the determined boundary in Eq. (11). The $(t_{k^*}^s, t_{k^*}^e)$ is the boundary of the most confident proposals selected from the alignment branch without constraint on their overlaps to the ground-truth (Eq. (9)).

Training. In addition to $\mathcal{L}_{\text{bound}}$ and $\mathcal{L}_{\text{align}}$, we follow the baseline to learn \mathbf{h} in Eq. (4) by a binary cross-entropy loss to highlight foreground video content:

$$\mathcal{L}_{\text{high}}(\mathbf{F}, \mathbf{Q}, S, E) = \text{BCE}(\mathbf{y}^h, \mathbf{h}), \quad \mathbf{y}_t^h = \mathbb{1}[\min(\tilde{\mathbf{S}}) \leq t \leq \max(\tilde{\mathbf{E}})]. \quad (13)$$

Note that, the boundary $(\min(\tilde{\mathbf{S}}), \max(\tilde{\mathbf{E}}))$ is also extended as in [28] to encourage the model to be sensitive to subtle visual changes around the temporal endpoints. The overall loss function of EMB is then formulated as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{bound}} + \lambda_2 \mathcal{L}_{\text{align}} + \lambda_3 \mathcal{L}_{\text{high}} \quad (14)$$

The EMB model is optimised end-to-end by stochastic gradient descent. Its overall training process is summarised in Alg. 1.

Algorithm 1 Elastic Moment Bounding

Input: An untrimmed video F , a query sentence Q , a temporal boundary (S, E) .**Output:** An updated video activity localisation model.

Encode frames by content-guided attention via Eq. (2)(3);

Fuse frames with query and predict per-frame endpoint probabilities via Eq. (4);

Construct 2D feature map of proposals;

Encode proposals by boundary-guided attention via Eq. (2)(5);

Fuse proposals with query and predict proposal-query alignment scores via Eq. (6);

Construct the elastic boundary via Eq. (8)(9);

Optimise model weights by minimising \mathcal{L} via Eq. (14).

4 Experiments

Datasets. We evaluated the proposed EMB model on three widely adopted video activity localisation benchmark datasets: **(1)** TACoS [16,17], **(3)** CharadesSTA [5,19] and **(2)** ActivityNet-Captions [10,8]. Their different data characteristics are summarised in Table 1. Among the three datasets, the raw videos in TACoS have the longest durations while that of its MoIs are shortest in contrast, which means that the video activities are temporally covering less than 2% of the complete videos on average. Therefore, the videos in TACoS contain a lot of redundancy in terms of every MoIs. On the other hand, the ActivityNet-Captions is very different from TACoS whose video activities temporally cover much larger proportions of the videos ($\sim 30\%$) than the other two.

Performance Metrics. We followed the common practices [28,22,14] to measure the quality of our video activity localisation results by their average recall rate at different temporal IoU thresholds (IoU@ m). The predicted boundary (\hat{S}, \hat{E}) of a MoI is considered correct if its IoU with the manual temporal label (S, E) is greater than the thresholds m which are predefined as $m = \{0.3, 0.5, 0.7\}$. Besides, we also reported the mean IoU (mIoU) of all predictions with their corresponding ground-truth to show the average overlaps between the predicted and manual boundaries. For our elastic boundary, we enumerate all the start-end pairs from \hat{S} and \hat{E} (Eq. (12)) respectively. If a manual boundary’s overlap to any of the combinations is greater than the IoU threshold, we consider it is correctly predicted.

Table 1. Statistics of datasets. L^v and L^m are the average lengths of videos and MoIs, respectively. L^q is the average number of words in query sentences.

Datset	#Train	#Val	#Test	L^v	L^q	L^m
TACoS [16]	10,146	4,589	4,083	287.14s	10.1	5.45s
ANet [10]	37,421	17,031	17,505	117.61s	14.8	36.18s
Charades [5]	12,408	-	3,720	30.59s	7.2	8.22s

Implementation Details. We adopted the video features provided by our baseline model [28] and the 300D GloVe [15] embeddings to encode the video and text inputs, respectively. We downsampled videos to have 128 frames at most by max-pooling and zero-padded the shorter ones. The outputs of all the hidden layers were 128D as in [28] and the multi-head variant [20] of the attention layer in Eq. (1) was used with 8 heads followed by layer normalisation and random dropout at 0.2. Cosine positional embeddings were applied to the inputs. The EMB model was trained for 100 epochs with a batch size of 16. It was optimised by an Adam optimiser using a linearly decaying learning rate of $5e-4$ and gradient clipping of 1.0. In the alignment branch, we downsampled the videos to have 16 clips by the max-pooling of every 8 continuous frames for constructing the 2D feature maps of video segments to avoid over-dense proposals. The threshold τ in Eq. (9) was initiated to be 1 and progressively decreased to 0.5. The weights of losses were empirically set to $\lambda_1 = \lambda_2 = 1$ and $\lambda_3 = 5$ in all the datasets.

4.1 Comparisons to the State-of-the-art

As shown in Table 2, the determined boundary yielded by EMB (DET) outperforms the baseline VSLNet [28] by non-negligible margins on all tests. The more recent IVG [14] shares the same baseline as EMB. The notable performance advantages of EMB over both of them demonstrate its non-trivial improvements. Furthermore, EMB surpasses the state-of-the-art methods on TACoS against all the performance metrics while remaining its competitiveness on the other two datasets. Among the three datasets, TACoS poses the hardest test with the longest average untrimmed videos and the shortest activity moments (see Table 1). That is, TACoS exhibits more realistic scenarios for activity localisation test. In this context, EMB shows its advantage most clearly when the untrimmed videos are longer whilst the video MoIs are sparse and far between.

Moreover, table 2 shows also the clear performance advantages of the elastic boundary predicted by our EMB (ELA) model over a wide range of the state-of-the-art methods. When constructing the elastic boundaries in inference, over 80% of the predictions pairs yielded by the alignment and bounding branches are consistent with each other ($\text{IoU} > 0.5$). Therefore, the performance improvements we obtained is not due to over-dense sampling of the potential boundaries. For fairer comparisons, we took the adjacent frames before and after the endpoints predicted by VSLNet to generate multiple candidate boundaries for its evaluation. The number of frames is set to be 10% of the moment length so that the density of candidate boundaries is consistent with ours. Although clear performance gains are observed, the improvements from such a *global* shifting strategy are less competitive to our per-sample *adaptive* designs due to missing considerations of sample-dependent bias.

4.2 Ablation Study

We conducted comprehensive ablation studies based on the EMB’s determined predictions to provide in-depth analyses and better understandings.

Table 2. Performance comparisons to the state-of-the-art models on three video activity localisation benchmark datasets. The first and second best results are highlighted in red and blue, respectively. The ‘DET’ modifier of EMB stands for the determined boundary predicted in Eq. (11) while ‘ELA’ is the elastic boundary (Eq. (12)). The symbol † denotes the reproduced results of our baseline model under the strictly identical setups using the code from authors and ★ indicates multi-candidate predictions.

Method	TACoS [16]				Charades-STA [5]				ActivityNet-Captions [10]			
	mIoU	IoU@m			mIoU	IoU@m			mIoU	IoU@m		
		0.3	0.5	0.7		0.3	0.5	0.7		0.3	0.5	0.7
VSLNet [28]	24.11	29.61	24.27	20.03	45.15	64.30	47.31	30.19	43.19	63.16	43.22	26.16
IVG [14]	28.26	38.84	29.07	19.05	48.02	67.63	50.24	32.88	44.21	63.22	43.83	27.10
2D-TAN [29]	-	37.29	25.32	-	-	-	39.70	23.31	-	59.45	44.51	26.54
LGI [13]	-	-	-	-	51.38	72.96	59.46	35.48	41.13	58.52	41.51	23.07
DPIN [21]	-	46.74	32.92	-	-	-	47.98	26.96	-	62.40	47.27	28.31
DRN [27]	-	-	23.17	-	-	-	53.09	31.75	-	-	45.45	24.36
SCDM [26]	-	26.11	21.17	-	-	-	54.44	33.43	-	54.80	36.75	19.86
BPNet [23]	19.53	25.93	20.96	14.08	46.34	65.48	50.75	31.64	42.11	58.98	42.07	24.69
CPNet [11]	28.69	42.61	28.29	-	52.00	-	60.27	38.74	40.65	-	40.56	21.63
CPN [31]	34.63	48.29	36.58	21.25	51.85	72.94	56.70	36.62	45.70	62.81	45.10	28.10
DeNet [33]	-	-	-	-	-	-	59.75	38.52	-	61.93	43.79	-
CBLN [31]	-	38.98	27.65	-	-	-	61.13	38.22	-	66.34	48.12	27.60
SMIN [22]	-	48.01	35.24	-	-	-	64.06	40.75	-	-	48.46	30.34
VSLNet [†] [28]	28.15	39.07	27.59	16.65	47.33	67.26	50.46	31.53	42.26	57.75	41.10	25.58
EMB (DET)	35.49	50.46	37.82	22.54	53.09	72.50	58.33	39.25	45.59	64.13	44.81	26.07
VSLNet ^{†★}	30.61	41.14	30.09	18.97	53.88	71.59	57.98	41.64	49.49	65.83	49.68	32.00
EMB [★] (ELA)	48.36	63.31	52.49	37.02	62.16	79.73	69.22	51.40	56.25	73.72	58.65	40.74

Components analysis. We investigated the individual contributions of different components in our EMB model to its improvements over the baseline model [28]. As shown in Fig. 3, both our elastic boundary learning objective (Eq. (10)) and the multi-grained interaction network brought clear benefits to the baseline. Such results demonstrate the effectiveness to learn the temporal endpoints of video activities with higher flexibility so to tolerant the uncertainty of manual labels. Besides, they also imply the superiority of our visual encoders which conduct both within and cross-modal attention learning and complement the boundary and content information of video segments mutually.

Candidate endpoints mining. We evaluated the advantages of mining candidate endpoints adaptively over several heuristic strategies without the MGIN design: (1) boundary extension [28], (2) smoothing by a gaussian kernel [22,23] and (3) single-frame endpoints (baseline). As shown in Fig. 4, simply improving the boundary’s flexibility without considering their reliability (“Extend”) tends to degrade the model’s performances on both datasets. Boundary smoothing by a gaussian kernel (“Kernel”) is sometimes beneficial but less stable than our adaptive designs because their candidates were determined according to only the duration of MoIs without considering the video context and query’s unambiguity.

Evolving threshold. We studied the effects of threshold’s evolving schemes to our elastic boundary constructions (Eq. (9)). Fig. 5 shows the curves of schemes and their corresponding performances. The model trained with a constant threshold yielded the worst results in most cases while the ‘Sigmoid’ scheme




Fig. 3. Effectiveness of different proposed components. The elastic moment bounding formulation is denoted as component “A” while the multi-grained interaction network is component “B”.




Fig. 4. Effects of multi-candidate mining strategies in training. “Fix”: single-frame boundaries. Multiple endpoints are generated by extension (Extend), a gaussian kernel (Kernel), or our elastic bounding.

is always the best. This is because the ‘Sigmoid’ scheme maintains a persistently high threshold at the early training stages to avoid introducing distractions when the alignment branch is under-trained, then drops rapidly to involve more diverse candidate endpoints when the alignment branch is reliable.




Fig. 5. Effects of constructing elastic boundary subject to an evolving threshold on Charades-STA.




Fig. 6. Effectiveness of guided attention mechanism by comparing with the conventional attention modules [20].

Guided attention. We validated the effectiveness of our guided attention mechanism by replacing it in our MGIN encoder by the conventional attention modules proposed in [20]. From the comparison results shown in Fig. 6, the models trained with guided attention outperformed their counterparts which learned the video representations without interacting information in multiple granularities. Such results imply the complementary of segment’s content and boundary information, which encourage the video feature representations to be sensitive to redundancy and activity transitions.

Qualitative case study. We provide several video examples in Fig. 7 which are showing video activities corresponding to semantically similar sentence descriptions. However, their manual boundary are inconsistent, demonstrating the uncertainty in temporal boundaries. Specifically, the manual boundary for Q_1 starts from grabbing the food right before putting it into the mouth while Q_2 skipping the action of “grabbing” and starts when the person takes a bite. The



Fig. 7. Cases of video activities with similar semantics but inconsistent manual boundaries. The manual and predicted boundary are shown in red and green, respectively.

Q_3 involves even more redundancy which covers the actions of taking a plate from a desk and blending foods by a folk. In contrast, the predictions made by our model are more consistent on interpreting the action of “eat” in different videos, *i.e.*, always starts from delivering food to the mouth. This is accomplished by learning with highly flexible boundaries instead of fitting rigid and ambiguous manual endpoints which are prone to visual-textual miscorrelations.

5 Conclusion

In this work, we introduced a new Elastic Moment Bounding (EMB) approach to learn a more robust model for identifying video activity temporal endpoints with the inherent uncertainty in training labels. EMB is based on modelling elastic boundary tailored to optimise learning more flexibly the endpoints of every target activity moment with the knowledge that the given training labels are uncertain with inconsistency. EMB learns a more accurate and robust visual-textual correlation generalisable to activity moment localisation in more naturally prolonged unseen videos where activity of interests are fractionally small and harder to detect. Comparative evaluations and ablation studies on three activity localisation benchmark datasets demonstrate the competitiveness and unique advantages of EMB over the state-of-the-art models especially when the untrimmed videos are long and activity moments are short.

Acknowledgements

This work was supported by the China Scholarship Council, Vision Semantics Limited, the Alan Turing Institute Turing Fellowship, Adobe Research and Zhejiang Lab (NO. 2022NB0AB05).

References

1. Anne Hendricks, L., Wang, O., Shechtman, E., Sivic, J., Darrell, T., Russell, B.: Localizing moments in video with natural language. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 5803–5812 (2017)
2. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: Proceedings of the International Conference on Learning Representations (ICLR) (2015)
3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6299–6308 (2017)
4. Chen, J., Chen, X., Ma, L., Jie, Z., Chua, T.S.: Temporally grounding natural sentence in video. In: Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 162–171 (2018)
5. Gao, J., Sun, C., Yang, Z., Nevatia, R.: Tall: Temporal activity localization via language query. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 5267–5275 (2017)
6. Ge, R., Gao, J., Chen, K., Nevatia, R.: Mac: Mining activity concepts for language-based temporal localization. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 245–253. IEEE (2019)
7. Ghosh, S., Agarwal, A., Parekh, Z., Hauptmann, A.: ExCL: Extractive Clip Localization Using Natural Language Descriptions. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 1984–1990. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019), <https://www.aclweb.org/anthology/N19-1198>
8. Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C.: Activitynet: A large-scale video benchmark for human activity understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 961–970 (2015). <https://doi.org/10.1109/CVPR.2015.7298698>
9. Huang, J., Liu, Y., Gong, S., Jin, H.: Cross-sentence temporal and semantic relations in video activity localisation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7199–7208 (2021)
10. Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Niebles, J.C.: Dense-captioning events in videos. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017)
11. Li, K., Guo, D., Wang, M.: Proposal-free video grounding with contextual pyramid network. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). vol. 35, pp. 1902–1910 (2021)
12. Mayu Otani, Yuta Nakahima, E.R., Heikkilä, J.: Uncovering hidden challenges in query-based video moment retrieval. In: Proceedings of the British Machine Vision Conference (BMVC) (2020)
13. Mun, J., Cho, M., Han, B.: Local-global video-text interactions for temporal grounding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10810–10819 (2020)
14. Nan, G., Qiao, R., Xiao, Y., Liu, J., Leng, S., Zhang, H., Lu, W.: Interventional video grounding with dual contrastive learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2765–2775 (2021)

15. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 1532–1543 (2014), <http://www.aclweb.org/anthology/D14-1162>
16. Regneri, M., Rohrbach, M., Wetzels, D., Thater, S., Schiele, B., Pinkal, M.: Grounding action descriptions in videos. *Transactions of the Association for Computational Linguistics* **1**, 25–36 (2013)
17. Rohrbach, M., Regneri, M., Andriluka, M., Amin, S., Pinkal, M., Schiele, B.: Script data for attribute-based recognition of composite activities. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 144–157. Springer (2012)
18. Seo, M., Kembhavi, A., Farhadi, A., Hajishirzi, H.: Bidirectional attention flow for machine comprehension. arXiv preprint arXiv:1611.01603 (2016)
19. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in homes: Crowdsourcing data collection for activity understanding. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 510–526. Springer (2016)
20. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Proceedings of the Conference on Neural Information Processing Systems (NeurIPS). pp. 5998–6008 (2017)
21. Wang, H., Zha, Z.J., Chen, X., Xiong, Z., Luo, J.: Dual path interaction network for video moment localization. In: Proceedings of the ACM International Conference on Multimedia (MM). pp. 4116–4124 (2020)
22. Wang, H., Zha, Z.J., Li, L., Liu, D., Luo, J.: Structured multi-level interaction network for video moment localization via language query. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7026–7035 (2021)
23. Xiao, S., Chen, L., Zhang, S., Ji, W., Shao, J., Ye, L., Xiao, J.: Boundary proposal network for two-stage natural language video localization. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). vol. 35, pp. 2986–2994 (2021)
24. Xiong, C., Zhong, V., Socher, R.: Dynamic coattention networks for question answering. arXiv preprint arXiv:1611.01604 (2016)
25. Yu, A.W., Dohan, D., Le, Q., Luong, T., Zhao, R., Chen, K.: Fast and accurate reading comprehension by combining self-attention and convolution. In: Proceedings of the International Conference on Learning Representations (ICLR). vol. 2 (2018)
26. Yuan, Y., Ma, L., Wang, J., Liu, W., Zhu, W.: Semantic conditioned dynamic modulation for temporal sentence grounding in videos. In: Proceedings of the Conference on Neural Information Processing Systems (NeurIPS). pp. 534–544 (2019)
27. Zeng, R., Xu, H., Huang, W., Chen, P., Tan, M., Gan, C.: Dense regression network for video grounding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10287–10296 (2020)
28. Zhang, H., Sun, A., Jing, W., Zhou, J.T.: Span-based localizing network for natural language video localization. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 6543–6554. Association for Computational Linguistics, Online (Jul 2020), <https://www.aclweb.org/anthology/2020.acl-main.585>
29. Zhang, S., Peng, H., Fu, J., Luo, J.: Learning 2d temporal adjacent networks for moment localization with natural language. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). vol. 34, pp. 12870–12877 (2020)
30. Zhang, S., Su, J., Luo, J.: Exploiting temporal relationships in video moment localization with natural language. In: Proceedings of the ACM International Conference on Multimedia (MM). pp. 1230–1238 (2019)

31. Zhao, Y., Zhao, Z., Zhang, Z., Lin, Z.: Cascaded prediction network via segment tree for temporal video grounding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4197–4206 (2021)
32. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2914–2923 (2017)
33. Zhou, H., Zhang, C., Luo, Y., Chen, Y., Hu, C.: Embracing uncertainty: Decoupling and de-bias for robust temporal grounding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8445–8454 (2021)