

HAT: History-Augmented Anchor Transformer for Online Temporal Action Localization

Sakib Reza[Ⓜ], Yuexi Zhang[Ⓜ], Mohsen Moghaddam[Ⓜ], and Octavia Camps[Ⓜ]

Northeastern University, Boston, MA 02115, USA

{reza.s,zhang.yuex,m.moghaddam,o.camps}@northeastern.edu

Abstract. Online video understanding often relies on individual frames, leading to frame-by-frame predictions. Recent advancements such as Online Temporal Action Localization (OnTAL), extend this approach to instance-level predictions. However, existing methods mainly focus on short-term context, neglecting historical information. To address this, we introduce the History-Augmented Anchor Transformer (HAT) Framework for OnTAL. By integrating historical context, our framework enhances the synergy between long-term and short-term information, improving the quality of anchor features crucial for classification and localization. We evaluate our model on both procedural egocentric (PREGO) datasets (EGTEA and EPIC) and standard non-PREGO OnTAL datasets (THUMOS and MUSES). Results show that our model outperforms state-of-the-art approaches significantly on PREGO datasets and achieves comparable or slightly superior performance on non-PREGO datasets, underscoring the importance of leveraging long-term history, especially in procedural and egocentric action scenarios. Code is available at: <https://github.com/sakibreza/ECCV24-HAT/>.

Keywords: Online Video Understanding · Temporal Action Localization · Egocentric Vision · Temporal Transformers

1 Introduction

In today’s digital era, the surge in video content has spurred the demand for innovative solutions to process lengthy, unedited videos. Transformers have become instrumental to understand complex video sequences, particularly in Temporal Action Localization (TAL) [6,29,43,44], which aims to identify and classify actions within videos, marking their start and end times.

Online Action Detection (OAD) identifies ongoing actions in real-time video frames, while Online Detection of Action Start (ODAS) focuses on quickly pinpointing action beginnings, enabling prompt responses. However, both tasks operate at the frame level, lacking contextual understanding at the instance level, limiting their broader applicability. To address this, a new instance-level task, Online Temporal Action Localization (On-TAL) [18], has been introduced. On-TAL, unlike traditional offline TAL, operates online without access to future frames and prohibits alterations to previously generated action proposals, allowing only real-time post-processing.

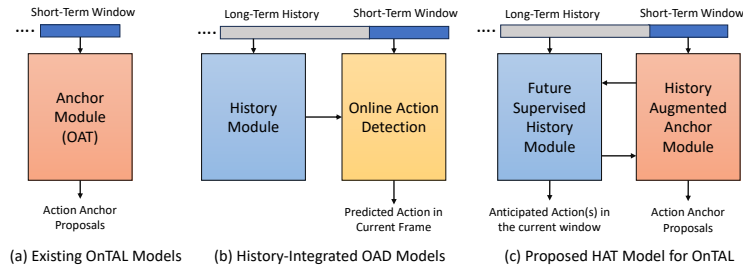


Fig. 1: High-level architecture designs of prior approaches and our method.

Despite its potential, traditional On-TAL methodologies, which focus on aggregating per-frame action predictions into action instances, encounter challenges in accuracy due to issues like tick, fragmentation, and merging of action instances. To address these, the Online Anchor Transformer (OAT) [19] was introduced, using a Transformer architecture with instance-level supervision. However, OAT’s reliance on short-term analysis within a sliding window neglects the potential benefits of incorporating long-term historical context, which is crucial for accurately predicting complex action sequences. While there have been attempts to incorporate historical data in other areas of online action understanding, such as Online Action Detection (OAD) [4, 40], the fusion of historical context with OnTAL, particularly with anchor-based Transformer models, remains largely unexplored. Therefore, this paper aims to bridge this gap by integrating historical context into the OnTAL framework, proposing a novel framework named **History-Augmented Anchor Transformer (HAT)**. Figure 1 illustrates the overarching structure of our proposed framework in comparison to earlier methods. The main contributions of our paper include:

- The first long-term historical context supported anchor-based transformer, tailored for OnTAL task, which efficiently processes extensive historical data and integrates it into the anchor-level feature, enhancing the model’s understanding and contextual relevance.
- A novel history refinement strategy that leverages action anticipation guidance and aligns it with the current temporal context to improve prediction accuracy and relevance.
- A gradient-guided adaptive focal loss function to address the class-imbalance problem effectively, applicable to both foreground versus background and among different foreground classes.
- A comprehensive evaluation across datasets, showcasing state-of-the-art performance on both procedural egocentric (PREGO) datasets (EGTEA and EPIC-Kitchen-100) and standard non-PREGO OnTAL benchmark datasets (THUMOS’14 and MUSES). Our method particularly stands out in procedural tasks and egocentric datasets due to its advanced understanding of action dependencies and its ability to disregard irrelevant past frames.

2 Related Work

Temporal Action Localization TAL detects actions within videos by identifying their start and end points and categorizing them. Initially, researchers used sliding window methods [11, 31]. Recent studies have expanded the scope of video context analysis, including the use of graph networks [38, 42, 45], temporal multi-scale networks [26, 46, 47], multi-scale representation with local self-attention [43], learning based on action sensitivity [29], exploring connections among action proposals [5], incorporating overall video context [48], and utilizing two-stage long-memory transformers [6]. End-to-end training methods [17, 24, 28, 37, 44] are also gaining interest as an alternative approach in TAL.

Online Video Understanding De Geest et al. [8] pioneered the Online Action Detection (OAD) task, focusing on per-frame action classification in streaming video. Various methods, including reinforcement learning [12], future anticipation [20, 34, 39], GRU-based networks [9], minimal RNNs [2], and transformer-based approaches [15, 36, 41], have been utilized for this task. Recent trends favor transformer-based methodologies to leverage historical context for OAD [4, 14, 34, 40].

Shou et al. [30] introduced Online Detection of Action Start (ODAS), focusing on identifying action start frames in streaming video. Gao et al. [13] enhanced ODAS with a two-stage approach and policy gradient learning.

Online Temporal Action Localization (On-TAL) emphasizes online detection of action instances, initially proposed by Kang et al. [18]. Methods like context-aware actionness grouping (CAG) [18], SimOn [33], 2PESNet [21], and OAT model [19] address this task by analyzing streaming videos to generate instance-level action proposals, similar to TAL outputs.

History for Online Action Understanding In online action understanding tasks like OAD, integrating historical context with present information is beneficial. Studies like JOADAA [14] use a transformer encoder to process past information and merge it with current features. LSTR [40] employs transformer decoders for a two-stage compression process to handle extensive historical data efficiently. GateHUB [4] and MAT [34] implement one-stage history compression methods, with GateHUB incorporating a gating mechanism to calibrate attentions and suppress irrelevant frames and backgrounds.

Existing research has explored using long-term historical data in OAD, but similar methods for OnTAL tasks are lacking. To address this, we propose integrating long-term historical context into OnTAL frameworks. However, directly applying techniques from OAD to OnTAL is suboptimal due to task differences: OAD predicts actions at the frame level, while OnTAL identifies discrete action instances.

Thus, we propose a novel history processing and integration method tailored for OnTAL. Our approach compresses long-term historical data efficiently while retaining crucial information for understanding actions at the instance level.

3 Method

3.1 Problem Definition

In this study, we tackle the Online Temporal Action Localization (On-TAL) problem [18], where the task is to identify and classify action instances within an untrimmed video stream, $V = \{v_i\}_{i=1}^T$, consisting of T frames, alongside its corresponding K label sets, $\Psi = \{\psi_k\}_{k=1}^K = \{(s_k, e_k, c_k)\}_{k=1}^K$, where s_k , e_k , and c_k denote the start time, end time, and action category of the k^{th} action instance, respectively. The inherent online constraint of On-TAL stipulates that at any given timestamp t ($1 \leq t \leq T$), only the partial video $V_{1:t} = \{v_i\}_{i=1}^t$ is observable, necessitating the generation of action proposals immediately upon detection of action ends, with the ultimate goal of sequentially recovering Ψ . This paradigm imposes a stringent condition that once an action proposal is generated, it cannot be subsequently revised or removed.

3.2 Proposed Framework Overview

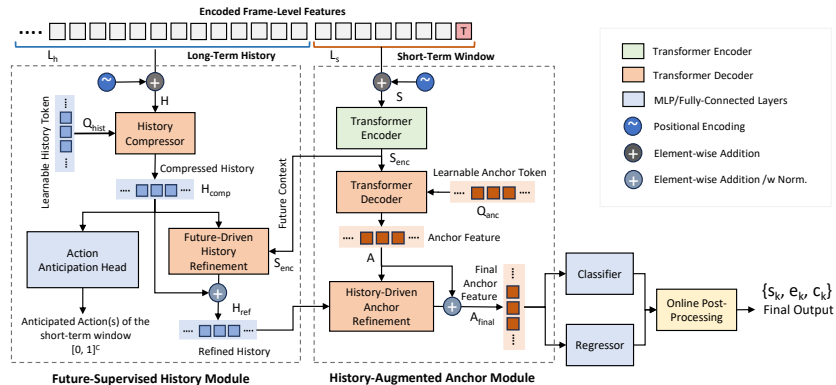


Fig. 2: Proposed History-Augmented Anchor Transformer (HAT) architecture.

Figure 2 illustrates the architecture of the proposed History-Augmented Anchor Transformer (HAT) framework. Its **history module** handles long-term context using a compressor. The compressed history is passed to an action anticipation head, which utilizes window-level supervision for current action anticipation to calibrate the compressor. Simultaneously, the compressed historical data is processed through future-driven history refinement to ensure alignment with the current context. The **history-augmented anchor module** processes short-term window features, utilizing a transformer encoder to capture temporal context and a transformer decoder to generate anchor features. These features, enhanced with the history information, are the augmented anchor features used in the **prediction module**. This module contains a classifier and a regressor

to generate action proposals for anchors, followed by an online post-processing step to eliminate redundant proposals. A detailed description is given next, with additional information in the supplementary material.

3.3 Initial Feature Extraction

Each input video segment undergoes an initial transformation into a frame-level encoded representation using a pre-trained feature encoder (e.g., I3D [3], SlowFast [10]). At a given time point t , these vectors are mapped into a D -dimensional feature space via a linear projection layer, represented as $F = \{f_i\}_{i=t-L_q+1}^t$, where L_q is the total length of the input queue. The feature F is further divided into two distinct subsets to capture different temporal contexts: the long-term history, denoted as $H = \{PE(f_i)\}_{i=t-L_h+1}^{t-L_s}$, and the short-term temporal window, denoted as $S = \{PE(f_i)\}_{i=t-L_s+1}^t$, where PE is the positional encoding function, L_s and $L_h = L_q - L_s$ are the lengths of the short-term and long-term history windows, respectively. The long-term component, $H \in \mathbb{R}^{L_h \times D}$, is processed further in the history module, while the short-term segment, $S \in \mathbb{R}^{L_s \times D}$, is utilized within the anchor module to produce anchor features.

3.4 Future-Supervised History Module

The purpose of the future-supervised history module is to compress the long-term historical data (H) to enhance processing efficiency while preserving vital historical information. This module further refines the compressed data by incorporating weak anticipatory supervision of actions and alignment with future contexts. The methodology is detailed next.

History Compressor. The long-term history L_h can be significantly long. Utilizing self-attention mechanisms to encapsulate historical temporal information is computationally intensive as the complexity scales quadratically with L_h . Moreover, integrating this information with anchor features further increases computational demands. Inspired by prior work [40], we introduce a history compression strategy to mitigate these challenges. This method employs a standard transformer decoder integrated with a learnable historical token, Q_{hist} , of size L_{comp} , where $L_{\text{comp}} \ll L_h$. The compression is done through transformer decoder blocks performing the following operations:

$$d(Q, X) = \text{Norm}(X_{\text{attn}} + \text{FFN}(X_{\text{attn}})), \quad (1)$$

$$X_{\text{attn}} = \text{Norm}(Q_{\text{attn}} + \text{CrossAttn}(Q_{\text{attn}}, X)), \quad (2)$$

$$Q_{\text{attn}} = \text{SelfAttn}(Q), \quad (3)$$

where the decoder d integrates multi-head self-attention (SelfAttn), multi-head cross-attention (CrossAttn), layer normalization (Norm), and a feed-forward network (FFN). For this compression operation, we set the query tokens $Q = Q_{\text{hist}}$, with the learnable history token, and the input token $X = H$, with the initial history feature. The architecture incorporates N_c identical decoder blocks,

to generate the compressed history, $H_{\text{comp}} = d_{N_c} \circ \dots \circ d_1(Q_{\text{hist}}, H) \in \mathbb{R}^{L_{\text{comp}} \times D}$.

Action Anticipation Head. To enhance the efficacy of the history compression, it is crucial to retain information from frames that play a significant role in predicting ongoing action instances. In this light, we introduce an innovative concept: the action anticipation head, which processes the compressed history and predicts actions within the current sliding window S (considered as future actions by the history encoder). Our method prioritizes prediction at the window level, rather than at the frame-wise or anchor-level, thereby providing weak supervision through forward anticipation. This novel approach enables the history compressor to calibrate its attention by predicting future actions, allocating increased attention to frames that are pivotal in predicting action instances while diminishing emphasis on less relevant frames.

In our architecture, the condensed historical feature $H_{\text{comp}} \in \mathbb{R}^{L_{\text{comp}} \times D}$ is fed into the action anticipation head to predict forthcoming actions. The anticipation mechanism is built around a simple neural network structure. Initially, a linear layer reduces the feature dimension from D to $\frac{D}{4}$. Subsequently, the features are flattened to prepare them for sequential processing. Finally, two densely connected layers are employed, with the first utilizing the ReLU activation function and the last utilizing the sigmoid function. This architecture results in the generation of predicted action probabilities within the current sliding window S , denoted as $a \in [0, 1]^{1 \times C}$, where C is the number of action classes.

Future-Driven History Refinement. The objective of the future-driven history refinement module is to enhance the intermediate history feature by aligning it with the short-term temporal context, thereby making it more pertinent in understanding current action instances. Its core is built around a transformer decoder (Equation 1) that takes two inputs: the compressed history (H_{comp}) and the short-term temporal features (S_{enc}). This process is executed through a series of transformations with N_r blocks, depicted as $H'_{\text{comp}} = d_{N_r} \circ \dots \circ d_1(H_{\text{comp}}, S_{\text{enc}}) \in \mathbb{R}^{L_{\text{comp}} \times D}$. The cross-attention mechanism in the transformer decoder uses the current short-term context to calibrate the attention weights, ensuring that historical features relevant to the current context are emphasized, while others are de-emphasized. Additionally, to ensure the preservation of the intermediate historical context, we add this refined feature back to the initial compressed history and then apply normalization, resulting in the final refined history, $H_{\text{ref}} = \text{Norm}(H'_{\text{comp}} + H_{\text{comp}})$. This method allows us to enhance historical data with short-term context while retaining the intermediate historical information.

3.5 History-Augmented Anchor Module

Anchor Feature Generation In our framework, we follow the approach used in OAT [19] to generate anchor features. It involves two main stages. First, a transformer encoder, composed of N_e sequential blocks, is utilized to convert

local attributes into a temporal contextual feature, $S_{\text{enc}} \in \mathbb{R}^{L_s \times D}$, which is then supplied to a transformer decoder. This decoder also receives as input a learnable anchor query $Q_{\text{anc}} \in \mathbb{R}^{M \times D}$, where M is the number of pre-defined anchors. As data progresses through its N_d blocks, the decoder produces the anchor feature $A \in \mathbb{R}^{M \times D}$, where each of these M feature vectors represents a pre-defined anchor.

History-Driven Anchor Refinement In this module, we augment the initial anchor features by integrating them with refined historical data, thereby providing a more comprehensive context. To combine two sets of features, rather than using standard approaches such as concatenation and average pooling, we explore more capable techniques employing attention mechanisms. Specifically, cross-attention stands out as a better way to integrate features from different domains, particularly when integrating past context with the current [40]. Here, we utilize a standard Transformer Decoder, consisting of N_a decoder blocks (represented by d in equation 1), to handle the initial anchor feature, A , and the refined historical context, H_{ref} , obtained from our history module. Additionally, we add the enhanced anchor feature with the original anchor feature and apply normalization to derive the final anchor feature representation $A_{\text{final}} \in \mathbb{R}^{M \times D}$. This approach ensures that the model retains original anchor information while augmenting it with historical context. The resulting final anchor feature is subsequently used in both classification and regression tasks to generate action proposals for the pre-defined anchors.

3.6 Prediction Module and Online Post-Processing

In our framework, we adopt the prediction module and online post-processing techniques as outlined in our baseline work [19]. The prediction module consists of an action classifier and a boundary regressor. Both components use the same processed anchor feature representation A for their input. The function of the action classifier is to identify the specific action (or background) taking place within a given anchor segment. Conversely, the regressor is designed to compute two key measurements: the offset between the end of the target action and the anchor segment’s end, and the ratio of the duration of the target action to the total duration of the anchor segment. In each frame, multiple action proposals are generated and filtered using Non-Maximum Suppression (NMS). The probabilities of the selected proposals are then queued and inputted into the Online Suppression Network (OSN), which determines their inclusion in the final instance set, Ψ . OSN functions as a substitute for the traditional offline NMS method, adhering to online constraints.

3.7 Training Strategy and Loss Function

Adaptive Focal Loss. In the video action understanding datasets, a common issue is the imbalance between action samples and background samples, with

the latter dominating. Chen et. al. [4] proposed a modified focal loss to solve this imbalance by assigning varying levels of emphasis to different instances based on their classification confidence. Specifically, it differentiates between hard-to-classify action instances and background instances by applying distinct focal factors, thereby addressing the imbalance between action and background classes. However, this approach does not address the potential imbalance among foreground action classes. In response, inspired by equalized focal loss used in object detection tasks [22], we introduce a new loss function for OnTAL, named Adaptive Focal Loss (AFL). This new loss function dynamically adjusts the focal factor for each action class during training, informed by the gradients, to better manage class imbalance across the foreground action categories. The Adaptive Focal Loss is formulated as follows:

$$AFL(p_i, y_i) = \sum_{j=0}^C -y_i^j (1 - p_i^j)^{\lambda^j} \log(p_i^j), \quad (4)$$

where p_i is the predicted probability for the i^{th} anchor, y_i is the corresponding ground truth label, and λ^j is the class-specific focal factor for the j^{th} action category. This class-specific focal factor, λ^j , is defined as $\lambda^j = \lambda_b$ for $j = C$ (background), and $\lambda^j = \lambda_b + \lambda_f^j$ for foreground classes, where λ_b is the base focal factor, and $\lambda_f^j = s(1 - r^j)$ is an adaptive component for the foreground classes, with r^j being the ratio of accumulated gradients for positive versus negative samples for the j -th category. The scaling factor s adjusts the influence of λ_f^j . The value of r^j approaches 1 for categories with less imbalance, thereby reducing the additional focal impact, and approaches 0 for categories with greater imbalance, thereby increasing the additional focal impact. Further details on this loss are available in the supplementary material.

Final Loss. In the training phase, we compute the training objective similarly to the established baseline [19], with the exception that we apply our newly introduced adaptive focal loss (AFL) for classification tasks. To determine the overall loss for the framework, we utilize the AFL to compute the classification loss, L_c , alongside the regression losses for offset, L_o , and length adjustment, L_l , which are both calculated using L1 loss. Moreover, for the action anticipation head, we introduce an additional classification loss using AFL, denoted as L_a . Consequently, the overall loss is formulated as $\mathcal{L} = \alpha L_c + \beta(L_o + L_l) + \gamma L_a$, where α , β , and γ represent the weighted coefficients assigned to the classification, regression, and action anticipation losses, respectively. The training and evaluation procedures are carried out in a manner similar to the baseline, with the key difference being that during inference, the output from the action anticipation head is disregarded since it serves solely as a training enhancement.

4 Experiments

4.1 Datasets

Commonly used datasets like THUMOS and MUSES lack action variety within each video, limiting their effectiveness in testing OnTAL models. To overcome

this shortfall, we incorporate two egocentric procedural action datasets, EGTEA and EPIC-Kitchen-100 (EK100) to our testing. These datasets provide longer videos, a broader range of action categories, and a higher frequency of different actions per video. Importantly, the procedural nature of these actions ensures they are contextually linked to previous ones, offering a solid basis for evaluating the OnTAL models’ ability to integrate historical context. Additionally, the unique egocentric viewpoint leads to situations where actions occur outside the camera’s view or are obscured by the user’s hand, requiring the model to deduce actions based on temporal relationships. These characteristics make EGTEA and EK100 perfect for a more comprehensive and realistic evaluation of model performance in dynamic, real-life environments. Thus, we focus on EGTEA and EK100 for validating our proposed framework, while also presenting results for THUMOS and MUSES.

EGTEA [23] includes 28 hours of egocentric video content, comprising 86 sessions from 32 subjects performing kitchen tasks. It features 22 different action classes, with an average of 180 action instances and 15.6 action classes per video.

EPIC-Kitchen [7] is one of the largest egocentric datasets available, containing 100 hours of video from 700 sessions. These videos capture various cooking activities in different kitchens, with an average of 121.5 action instances and 14.5 action classes per video.

THUMOS’14 [16] consists of 200 training videos and 213 test videos, covering 20 action classes related to sports. These are untrimmed videos with an average of 15.5 action instances and only 1.2 action classes per video.

MUSES [27] comprises 3,697 videos, amounting to a total of 716 hours, and featuring 25 action classes. These videos are taken from TV and movie dramas, with an average of 8.5 action instances and 3.3 action classes per video.

4.2 Implementation Details

In our initial extraction of video features for EGTEA, we use the I3D model pre-trained on the Kinetics dataset [3]. The video frames were processed at their original frame rate of 24 FPS and a stride of 12, yielding two feature vectors per second. For the EPIC-Kitchens dataset, we adopted SlowFast features in line with the settings utilized by state-of-the-art TAL literature. For the THUMOS’14 dataset, we used the two-stream TSN approach [35], which was trained on the Kinetics dataset, following the methodology of prior research. Lastly, for the MUSES dataset, we used the I3D feature extractor again, aligning with the standards established by previous studies on MUSES.

In our framework configuration, we utilize a transformer with $D = 1024$ feature dimensions. In our architecture, transformer decoders have $N_c = N_d = N_a = 5$ blocks with 4 heads each, except the future-driven refinement decoder with $N_r = 2$ blocks with 4 heads each. The transformer encoder has $N_e = 3$ blocks with 8 heads. Regarding dataset-specific parameters, for the EGTEA and EPIC-Kitchens datasets, we set the history length $L_h = 48$ and the short-term window length $L_s = 16$, which respectively correspond to 24 seconds of historical data and 8 seconds of short-term data. Additionally, we employ $M = 6$ anchor

sizes: 2, 4, 6, 8, 12, 16. For the THUMOS’14 dataset, the history and short-term lengths are set to $L_h = 256$ and $L_s = 64$, with $M = 6$ anchor sizes of 4, 8, 16, 32, 48, 64. For the MUSES dataset, we set $L_h = 300$, $L_s = 150$, and use $M = 7$ anchor sizes: 4, 9, 18, 37, 75, 112, 150. In the adaptive focal loss function, we define the parameters as follows: $\lambda_b = 0.025$ and $s = 0.05$. Additionally, the weighted coefficients for the final loss are specified as $\alpha = \beta = 1$ and $\gamma = 0.2$. For other hyperparameters and the training of both the main and suppression networks, we adhere to the settings specified in OAT [19]. Additionally, our model operates under the online execution setting as mentioned in [19], achieving an inference speed of 147 FPS on an Nvidia GeForce RTX 4090 GPU.

4.3 Comparison with the State of the Art

Since OnTAL is a relatively new and less investigated area, there are only a few existing models available for comparison. Our model is primarily evaluated against the current leading OnTAL model, OAT. Additionally, we have included the performance metrics of other available models for comparison, as reported in the existing literature.

Table 1: Comparison of mAP (%) across various tIOU thresholds with OnTAL baseline on the Procedural Egocentric (PREGO) Datasets.

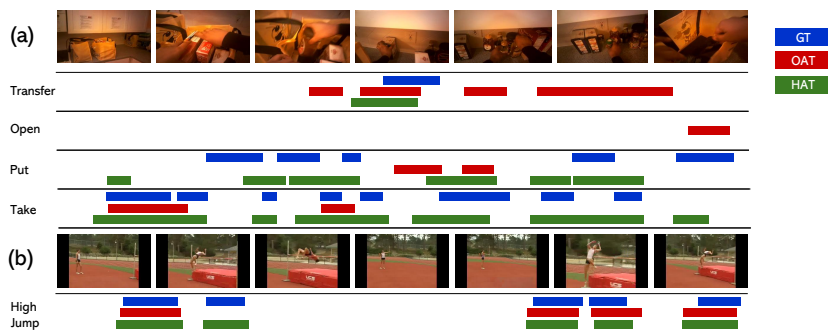
Dataset	Model	0.1	0.2	0.3	0.4	0.5	Avg.
EGTEA	OAT [19]	24.9	23.1	20.6	17.1	12.2	19.6
	HAT (Ours)	27.5	25.4	22.6	18.3	13.5	21.5
EK-100	OAT [19]	17.8	16.3	14.3	12.8	10.1	14.2
	HAT (Ours)	18.3	17.0	15.8	13.9	11.5	15.3

In Table 1, we present the performance of our model on the Procedural Egocentric (PREGO) datasets, specifically the EGTEA and EPIC-KITCHENS datasets. Our findings indicate a notable improvement over the current state-of-the-art OnTAL model, OAT. This enhancement underlines the effectiveness of our proposed history integration technique in understanding procedural and egocentric instance-level actions. The results convincingly demonstrate that our model holds a significant advantage in handling PREGO scenarios.

On the other hand, Table 2 shows our model’s performance on non-PREGO datasets, namely THUMOS and MUSES. Here, our model displays performance levels that are comparable to, or in some instances slightly superior to, existing models (with some exceptions). These results suggest that while our model is still effective in non-PREGO contexts, the impact of historical information is more pronounced and beneficial in PREGO situations. This distinction underscores the specialized advantage of our model in scenarios requiring nuanced, instance-level action understanding within procedural and egocentric contexts.

Table 2: Comparison of mAP (%) across various tIOU thresholds on the Non-PREGO Standard OnTAL Datasets. (*Reconstructed result)

Dataset	Model	0.3	0.4	0.5	0.6	0.7	Avg.
THUMOS'14	CAG-QIL [18]	44.7	37.6	29.8	21.9	14.5	29.7
	2PESNet [21]	47.4	39.8	31.4	21.8	14.0	30.9
	SimOn [33]	57.0	47.5	37.3	26.6	16.0	36.9
	OAT [19]	63.0	56.7	47.1	36.3	20.0	44.6
	HAT (Ours)	62.0	57.0	48.0	36.5	20.7	44.8
MUSES	CAG-QIL [18]	8.5	6.5	4.2	2.8	1.9	4.8
	OAT* [19]	16.7	13.0	10.0	6.3	3.2	9.8
	HAT (Ours)	19.1	14.7	10.1	6.3	3.7	10.8

**Fig. 3:** Comparison of qualitative results of HAT with OAT [19]. Blue is the ground truth, red is the result of OAT, and green is the result of HAT.

In addition to standard quantitative evaluation methods, we conducted a manual qualitative analysis to determine the effectiveness of our proposed HAT model. Figure 3(a) presents a case from the procedural ego-centric (PREGO) dataset EGTEA. Here, it is evident that the baseline model predicted a significantly higher number of false positives and missed action instances, failing to capture the broader context that our model provides. In Figure 3(b), we examine an example from the non-PREGO dataset THUMOS. In this scenario, the performance of the baseline and our model are nearly identical, with our model showing slight improvements in certain aspects. This qualitative assessment corroborates the findings from our previous quantitative analysis.

4.4 Model Analysis

In this section, we delve into the intricacies of the model by conducting ablation studies on the EGTEA dataset (split-1). These studies are designed to assess the effectiveness of various components within the proposed framework.

Impact of the History Module. In this experiment, we evaluated the effectiveness of the Future-Supervised History Module (FSHM) by sequentially eliminating its components and observing the effects on performance across various tIOU thresholds. The experimental results are summarized in Table 3.

Removing the action anticipation head ('w/o AH'), which leaves only the history compressor and future-driven history refinement active, results in a notable performance decline across most tIOU thresholds. This indicates the critical role of the action anticipation head in calibrating the attention mechanism of the history compressor. In the scenario 'w/o FDHR,' where the Future-Driven History Refinement component is omitted, there is a less severe performance decline, showing that while important, FDHR's impact is slightly less than that of the action anticipation head. When both the action anticipation head and the future-driven history refinement are removed ('w/o AH & FDHR'), leaving only the history compressor, there is an even more noticeable performance decrease, though still not as drastic as removing the entire module. Finally, completely eliminating the FSHM ('w/o FSHM') leads to a significant drop in performance across all tIOU thresholds, typically by about 2-3%, which underlines the effectiveness of the proposed future-supervised history component in our framework.

Comparison with other SOTA History Modules. We conducted an analysis to assess the effectiveness of our history module compared to other state-of-the-art (SOTA) history processing methods. Since this is the first-ever exploration of historical context processing specifically for OnTAL, and no prior methods have been designed for OnTAL, we opted to compare it with two SOTA history modules developed for OAD tasks: GateHUB and LSTR. Table 4 presents a comparison of our proposed future-supervised history module against alternative methods for processing historical data. Unlike the history module in LSTR, which employs a two-stage history compression method, our approach utilizes a single-stage compression combined with an action anticipation head and future-driven history refinement. The results indicate that our method performs better in the specific context of the On-

Table 3: Impact of history module on EGTEA dataset(Split-1). FSHM = Future-Supervised History Module, AH = Action Anticipation Head, FDHR = Future-Driven History Refinement.

Model	0.1	0.2	0.3	0.4	0.5	Avg.
w/ FSHM (Ours)	27.3	25.3	21.6	16.9	12.8	20.8
w/o AH	25.8	24.0	20.8	16.9	11.5	19.8
w/o FDHR	26.1	24.1	20.8	16.8	12.0	20.0
w/o AH & FDHR	24.7	23.0	20.1	16.2	11.7	19.2
w/o FSHM	24.2	22.7	20.5	16.3	10.9	18.9

Table 4: Comparison of our history processing method against other SOTA on EGTEA Dataset (Split-1).

Model	0.1	0.2	0.3	0.4	0.5	Avg.
No History	24.2	22.7	20.5	16.3	10.9	18.9
LSTR [40]	25.3	23.3	20.4	16.4	12.0	19.5
GHU [4]	26.2	24.0	21.0	16.6	11.7	19.9
FSHM (Ours)	27.3	25.3	21.6	16.9	12.8	20.8

TAL task. In GateHUB’s history module, a one-stage history compression method is utilized, similar to ours, but it incorporates a gating mechanism to adjust the focus during the compression process, aiming to filter out irrelevant and background frames. However, our results demonstrate the superior performance of our approach, highlighting the effectiveness of our action anticipation guided attention calibration compared to their gating mechanism, especially within the context of the PREGO OnTAL task. Consequently, we can assert that our future-supervised history module is more effective for OnTAL applications than other SOTA history processing methods designed for OAD.

Impact of the Adaptive Focal Loss In Table 5, we compare our introduced adaptive focal loss function with traditional cross-entropy, standard focal loss (with a constant focal factor for all classes), and the background suppression focal (BSF) loss as utilized in GateHub (with distinct fixed focal factors for foreground and background). Our findings show that cross-entropy, regular focal loss, and BSF loss all result in lower mAP across various tIOUs when compared to our adaptive focal loss function. This underscores the effectiveness of applying adaptive focusing factors to the foreground classes based on the gradient ratio of positive-to-negative samples during training, in addition to applying a distinct focal factor for the background class. This approach effectively addresses the issue of class imbalance between foreground and background, as well as among the foreground classes themselves. Whereas the BSF loss previously addressed only the foreground-background class imbalance through two separate focusing factors, it failed to address the imbalance among the foreground classes. The superior performance of our adaptive focal loss demonstrates its capacity to more effectively focus on the less confident, imbalanced classes during training, leading to improved outcomes.

Table 5: Comparison of different loss functions on EGTEA dataset (Split-1).

Loss	0.1	0.2	0.3	0.4	0.5	Avg.
Adaptive Focal Loss (Ours)	27.3	25.3	21.6	16.9	12.8	20.8
BG-Sup. Focal Loss [4]	27.2	24.4	21.1	16.7	11.5	20.2
Regular Focal Loss [25]	26.6	24.1	21.3	15.9	11.8	19.9
Cross-Entropy Loss [1]	25.1	22.9	20.2	16.6	12.1	19.4

Qualitative Analysis of the History Module To assess the effectiveness of our future-supervised history module beyond quantitative metrics, we conducted a qualitative analysis by examining the attention weights. This evaluation aims to determine if the history compressor within the module can successfully downplay irrelevant frames while highlighting those crucial to the current action. A representative example in Figure 4 demonstrates this by showcasing the subject performing a “cutting” action on a tomato. In this context, we divided the historical frames into two groups: those receiving higher attention than the median, and those receiving lower. The frames receiving more focus depict pertinent ac-




Fig. 4: Frame attention analysis for an “cutting” a tomato example from EGTEA dataset. The frames with related actions such as “taking” the tomato from the freezer, “taking” a knife, “putting” the tomato on a plate and the initial “cutting” action receive more attention (upper row).

tivities such as “taking” the tomato from the freezer, “taking” a knife, “putting” the tomato on a plate and the initial “cutting” action, illustrating the compressor’s ability to highlight relevant action-related history. Conversely, the less emphasized frames mostly include irrelevant egocentric movements and unrelated background details, underscoring the module’s capability to filter out irrelevant information to the ongoing action.

5 Conclusions and Future Work

In conclusion, we present a novel approach to Online Temporal Action Localization (OnTAL) by introducing a History-Augmented Anchor Transformer, which innovatively combines long-term historical information with anchor-level features. We designed a unique history processing method, incorporating action anticipation-guided history compression and future context-driven alignment, to enhance instance-level action prediction performance. Additionally, We introduced an adaptive focal loss function that addresses the class imbalance issue between foreground and background, as well as among different foreground classes. Our comprehensive analysis across four datasets demonstrates the effectiveness of our approach on the OnTAL task, particularly highlighting the influence of our history processing and integration in procedural egocentric (PREGO) contexts.

Our framework, while capable of utilizing long-term history, is restricted by a fixed history length. This creates a challenge as determining the optimal length for this history involves extensive manual adjustments. Furthermore, different instances within the same use case may require varying lengths of history—one may need a shorter span, while another may require a longer one. Although our system can filter out irrelevant information and focus on pertinent details within this fixed span, it lacks the ability to dynamically modify the history span. It would be ideal if the history component could automatically adjust its length to capture the long-term history and allocate attention to specific frames as needed, similar to the adaptive attention span techniques used in language models [32]. This adjustment would not only enable more effective long-term context capture but also reduce overall computational demands.

Acknowledgments

This work is supported by the U.S. National Science Foundation (NSF) Grants No. FW-HTF-2128743, IIS-2302838, IIS-1814631, and CNS-2038493, ONR grant N00014-21-1-2431 from NCI, and U.S. Department of Homeland Security grant 22STESE00001-03-02. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security, the NSF and the ONR.

References

1. Aldrich, J.: Ra fisher and the making of maximum likelihood 1912-1922. *Statistical science* **12**(3), 162–176 (1997)
2. An, J., Kang, H., Han, S.H., Yang, M.H., Kim, S.J.: Miniroad: Minimal rnn framework for online action detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10341–10350 (2023)
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*. pp. 6299–6308 (2017)
4. Chen, J., Mittal, G., Yu, Y., Kong, Y., Chen, M.: Github: Gated history unit with background suppression for online action detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19925–19934 (2022)
5. Chen, P., Gan, C., Shen, G., Huang, W., Zeng, R., Tan, M.: Relation attention for temporal action localization. *IEEE Transactions on Multimedia* **22**(10), 2723–2733 (2019)
6. Cheng, F., Bertasius, G.: Tallformer: Temporal action localization with a long-memory transformer. In: *European Conference on Computer Vision*. pp. 503–521. Springer (2022)
7. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision* pp. 1–23 (2022)
8. De Geest, R., Gavves, E., Ghodrati, A., Li, Z., Snoek, C., Tuytelaars, T.: Online action detection. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14*. pp. 269–284. Springer (2016)
9. Eun, H., Moon, J., Park, J., Jung, C., Kim, C.: Learning to discriminate information for online action detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 809–818 (2020)
10. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6202–6211 (2019)
11. Gao, J., Yang, Z., Chen, K., Sun, C., Nevatia, R.: Turn tap: Temporal unit regression network for temporal action proposals. In: *Proceedings of the IEEE international conference on computer vision*. pp. 3628–3636 (2017)

12. Gao, J., Yang, Z., Nevatia, R.: RED: reinforced encoder-decoder networks for action anticipation. In: British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017. BMVA Press (2017)
13. Gao, M., Xu, M., Davis, L.S., Socher, R., Xiong, C.: Startnet: Online detection of action start in untrimmed videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5542–5551 (2019)
14. Guermal, M., Ali, A., Dai, R., Brémond, F.: Joadaa: joint online action detection and action anticipation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 6889–6898 (2024)
15. Guo, H., Ren, Z., Wu, Y., Hua, G., Ji, Q.: Uncertainty-based spatial-temporal attention for online action detection. In: European Conference on Computer Vision. pp. 69–86. Springer (2022)
16. Idrees, H., Zamir, A.R., Jiang, Y.G., Gorban, A., Laptev, I., Sukthakar, R., Shah, M.: The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding* **155**, 1–23 (2017)
17. Kang, H., Kim, H., An, J., Cho, M., Kim, S.J.: Soft-landing strategy for alleviating the task discrepancy problem in temporal action localization tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6514–6523 (2023)
18. Kang, H., Kim, K., Ko, Y., Kim, S.J.: Cag-qil: Context-aware actionness grouping via q imitation learning for online temporal action localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13729–13738 (2021)
19. Kim, Y.H., Kang, H., Kim, S.J.: A sliding window scheme for online temporal action localization. In: European Conference on Computer Vision. pp. 653–669. Springer (2022)
20. Kim, Y.H., Nam, S., Kim, S.J.: Temporally smooth online action detection using cycle-consistent future anticipation. *Pattern Recognition* **116**, 107954 (2021)
21. Kim, Y.H., Nam, S., Kim, S.J.: 2pesnet: Towards online processing of temporal action localization. *Pattern Recognition* **131**, 108871 (2022)
22. Li, B., Yao, Y., Tan, J., Zhang, G., Yu, F., Lu, J., Luo, Y.: Equalized focal loss for dense long-tailed object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6990–6999 (2022)
23. Li, Y., Liu, M., Rehg, J.M.: In the eye of beholder: Joint learning of gaze and actions in first person video. In: Proceedings of the European conference on computer vision (ECCV). pp. 619–635 (2018)
24. Lin, C., Xu, C., Luo, D., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Fu, Y.: Learning salient boundary feature for anchor-free temporal action localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3320–3329 (2021)
25. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
26. Liu, Q., Wang, Z.: Progressive boundary refinement network for temporal action detection. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 11612–11619 (2020)
27. Liu, X., Hu, Y., Bai, S., Ding, F., Bai, X., Torr, P.H.: Multi-shot temporal event localization: a benchmark. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12596–12606 (2021)

28. Phan, T., Vo, K., Le, D., Doretto, G., Adjeroh, D., Le, N.: Zeetad: Adapting pre-trained vision-language model for zero-shot end-to-end temporal action detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 7046–7055 (2024)
29. Shao, J., Wang, X., Quan, R., Zheng, J., Yang, J., Yang, Y.: Action sensitivity learning for temporal action localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13457–13469 (October 2023)
30. Shou, Z., Pan, J., Chan, J., Miyazawa, K., Mansour, H., Vetro, A., Giro-i Nieto, X., Chang, S.F.: Online detection of action start in untrimmed, streaming videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 534–551 (2018)
31. Shou, Z., Wang, D., Chang, S.F.: Temporal action localization in untrimmed videos via multi-stage cnns. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1049–1058 (2016)
32. Sukhbaatar, S., Grave, É., Bojanowski, P., Joulin, A.: Adaptive attention span in transformers. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 331–335 (2019)
33. Tang, T.N., Park, J., Kim, K., Sohn, K.: Simon: A simple framework for online temporal action localization. arXiv preprint arXiv:2211.04905 (2022)
34. Wang, J., Chen, G., Huang, Y., Wang, L., Lu, T.: Memory-and-anticipation transformer for online action understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13824–13835 (2023)
35. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: European conference on computer vision. pp. 20–36. Springer (2016)
36. Wang, X., Zhang, S., Qing, Z., Shao, Y., Zuo, Z., Gao, C., Sang, N.: Oadtr: Online action detection with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7565–7575 (2021)
37. Xu, M., Perez Rúa, J.M., Zhu, X., Ghanem, B., Martinez, B.: Low-fidelity video encoder optimization for temporal action localization. *Advances in Neural Information Processing Systems* **34**, 9923–9935 (2021)
38. Xu, M., Zhao, C., Rojas, D.S., Thabet, A., Ghanem, B.: G-tad: Sub-graph localization for temporal action detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10156–10165 (2020)
39. Xu, M., Gao, M., Chen, Y.T., Davis, L.S., Crandall, D.J.: Temporal recurrent networks for online action detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5532–5541 (2019)
40. Xu, M., Xiong, Y., Chen, H., Li, X., Xia, W., Tu, Z., Soatto, S.: Long short-term transformer for online action detection. *Advances in Neural Information Processing Systems* **34**, 1086–1099 (2021)
41. Yang, L., Han, J., Zhang, D.: Colar: Effective and efficient online action detection by consulting exemplars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3160–3169 (2022)
42. Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., Gan, C.: Graph convolutional networks for temporal action localization. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7094–7103 (2019)
43. Zhang, C.L., Wu, J., Li, Y.: Actionformer: Localizing moments of actions with transformers. In: European Conference on Computer Vision. pp. 492–510. Springer (2022)

44. Zhao, C., Liu, S., Mangalam, K., Ghanem, B.: Re2tal: Rewiring pretrained video backbones for reversible temporal action localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10637–10647 (2023)
45. Zhao, C., Thabet, A.K., Ghanem, B.: Video self-stitching graph network for temporal action localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13658–13667 (2021)
46. 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)
47. Zhao, Z., Wang, D., Zhao, X.: Movement enhancement toward multi-scale video feature representation for temporal action detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13555–13564 (2023)
48. Zhu, Z., Tang, W., Wang, L., Zheng, N., Hua, G.: Enriching local and global contexts for temporal action localization. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13516–13525 (2021)