

T-DEED: Temporal-Discriminability Enhancer Encoder-Decoder for Precise Event Spotting in Sports Videos

Artur Xarles^{1,2}

Sergio Escalera^{1,2,3}

Thomas B. Moeslund³

Albert Clapés^{1,2}

¹Universitat de Barcelona, Spain

²Computer Vision Center, Spain

³Aalborg University, Denmark

arturxe@gmail.com, sescalera@ub.edu, tbm@create.aau.dk, aclapes@ub.edu

Abstract

*In this paper, we introduce **T-DEED**, a Temporal-Discriminability Enhancer Encoder-Decoder for Precise Event Spotting in sports videos. T-DEED addresses multiple challenges in the task, including the need for discriminability among frame representations, high output temporal resolution to maintain prediction precision, and the necessity to capture information at different temporal scales to handle events with varying dynamics. It tackles these challenges through its specifically designed architecture, featuring an encoder-decoder for leveraging multiple temporal scales and achieving high output temporal resolution, along with temporal modules designed to increase token discriminability. Leveraging these characteristics, T-DEED achieves SOTA performance on the FigureSkating and FineDiving datasets. Code is available at <https://github.com/arturxe2/T-DEED>.*

1. Introduction

In recent years, remarkable progress has been made in the field of video understanding, largely driven by advancements in deep learning and increased computational power. These developments have enabled researchers to move beyond simpler tasks like action recognition [18] in trimmed videos to more complex challenges such as accurate localization of actions within untrimmed videos. Among these tasks, we find Temporal Action Localization (TAL), which represents actions as temporal intervals, and Action Spotting (AS), which uses single keyframes to represent them. While TAL [37] has historically received more attention, AS has recently experienced an increase in research interest, particularly in domains such as sports, as fast-paced actions – common in sports videos – can be better represented by single temporal positions rather than intervals of time. Additionally, Hong et al. [20] extended AS to Precise Event Spotting (PES), distinguishing events from actions and em-

ploying tighter tolerances for evaluation.

This paper specifically focuses on the task of Precise Event Spotting in sports (illustrated in Figure 1), evaluated on FigureSkating [19] and FineDiving [38], two common sports datasets. Following Hong et al. [20], we use tight tolerances to accommodate the fast-paced nature of sporting events, where even a minor temporal deviation of 1-2 frames can lead to missed events. Three main challenges faced by methods addressing this task include: (1) the need for discriminative per-frame representations to differentiate between adjacent frames with high spatial similarity, (2) the necessity of high output temporal resolution to avoid losing prediction precision, and (3) the variability in the amount of temporal context required for different events, influenced by dataset characteristics and event dynamics.

To tackle these challenges we introduce T-DEED, a Temporal-Discriminability Enhancer Encoder-Decoder for PES in sports videos. Introducing skip connections within its encoder-decoder architecture, T-DEED operates across various temporal scales, capturing actions requiring diverse temporal contexts, while restoring the original temporal resolution, thereby addressing challenges (2) and (3). Additionally, it integrates Scalable-Granularity Perception (SGP) [31] based layers to increase discriminability among features within the same temporal sequence, tackling challenge (1). To summarize, our main contributions are:

1. We incorporate residual connections into the SGP layer, enabling the integration of features from multiple temporal scales within the skip connections of our encoder-decoder architecture. This results in the creation of the SGP-Mixer layer, addressing challenge (3).
2. We introduce the SGP-Mixer module within the SGP-Mixer layer, which adapts the SGP module to aggregate information from different temporal scales. This module shares the core principles of SGP to enhance token discriminability while modeling temporal information, therefore tackling challenge (1).




Figure 1. Illustration of the Precise Event Spotting task on the FigureSkating dataset. Red-marked frames contain events that require precise localization and correct classification among possible classes.

3. We conduct extensive ablations of T-DEED components, highlighting the advantages of the encoder-decoder architecture with SGP-based layers to enhance token discriminability. Additionally, T-DEED achieves state-of-the-art performance. Specifically, on FigureSkating, it improves mean Average Precision (mAP) by +1.15 for the FS-Comp split and +3.07 for the FS-Perf split. Similarly, on the FineDiving dataset, T-DEED achieves an improvement of +4.83 mAP.

2. Related work

Over the past decade, deep learning has driven the field of video understanding through a remarkable evolution. Initially focused on simple tasks like classifying short-trimmed videos [1, 6, 16], the field has transitioned to more complex challenges, including Temporal Action Localization (TAL) and Action Spotting (AS). Both TAL and AS focus on temporally locating specific actions within untrimmed videos. TAL specifies temporal intervals for annotations, while AS represents actions with single keyframes, making TAL suitable for prolonged actions and AS more appropriate for fast-occurring actions, with the added benefit of reduced annotation costs. Furthermore, Hong et al. [20] extended the AS task to Precise Event Spotting (PES), introducing a key difference in the required precision of predictions, limited to only a few frames, and distinguishing between actions and events. In the sports domain, AS and PES are particularly popular as they adapt better to its fast nature, leading to the development of many approaches [8, 11, 15, 20, 32, 36] across different sports, including football [10, 14], figure skating [19], diving [38], gymnastics [30], and tennis [45].

Given the inherent similarities between TAL and AS, the methodologies developed for these tasks frequently share common components. However, TAL methods have attracted more attention due to the earlier introduction of the task and a more extensive set of benchmarking datasets [5, 9, 22, 42, 46], placing them a step ahead of AS methods. In contrast, AS methods have been tailored for specific datasets and competitive challenges, exemplified by their development in challenges like SoccerNet Action Spotting and SoccerNet Ball Action Spotting [8, 15]. Notably, E2E-Spot [20] stands out as the only action or event spotting method evaluated across multiple datasets.

Temporal Action Localization. In TAL methods, a common classification divides them into two groups: two-stage methods [3, 13, 17, 28, 39] and one-stage methods [23, 25, 31, 41, 43]. Two-stage methods generate class-agnostic proposals that are later classified into action labels or background, while one-stage models directly localize and classify actions in a single step, offering simplicity and achieving state-of-the-art performance in many TAL and AS scenarios.

Among one-stage methods, early approaches [4, 24] utilized anchor windows to generate action predictions. Later, Yang et al. [40] introduced an anchor-free approach, relying on temporal points instead of anchor windows, highlighting the advantages of both techniques. Building on this approach, current state-of-the-art methods such as ActionFormer [43] and TriDet [31] have exhibited remarkable performance across various datasets. They leverage a feature pyramid network to process features at different temporal scales, a critical aspect for identifying actions that require distinct temporal contexts. The main difference lies in their prediction head and the layers used for feature processing. ActionFormer employs transformer layers, later revealed to suffer from the rank loss problem [12], negatively impacting token discriminability. To alleviate this problem, TriDet proposes a more efficient convolutional-based Scalable-Granularity Perception (SGP) layer, specially designed to increase token discriminability within the same temporal sequence, contributing to an improved overall performance. T-DEED, inspired by TAL trends, focuses on enhancing token discriminability while leveraging multiple temporal scales, with modifications tailored to meet the precision requirements of PES.

Action Spotting. In AS, techniques similar to those in TAL have demonstrated state-of-the-art performance on the SoccerNet challenges [8, 15]. Many methods [11, 32, 36] classify temporal points as either background or actions and refine them through temporal regression using either convolutional [32] or Transformer-based [11, 36] approaches. In contrast, Hong et al. [20] propose a simple end-to-end solution that employs a convolutional backbone with Gate-Shift Modules (GSM) [33] for extracting per-frame features with short-term temporal information, followed by a Gated Recurrent Unit (GRU) [7] layer for long-term temporal in-

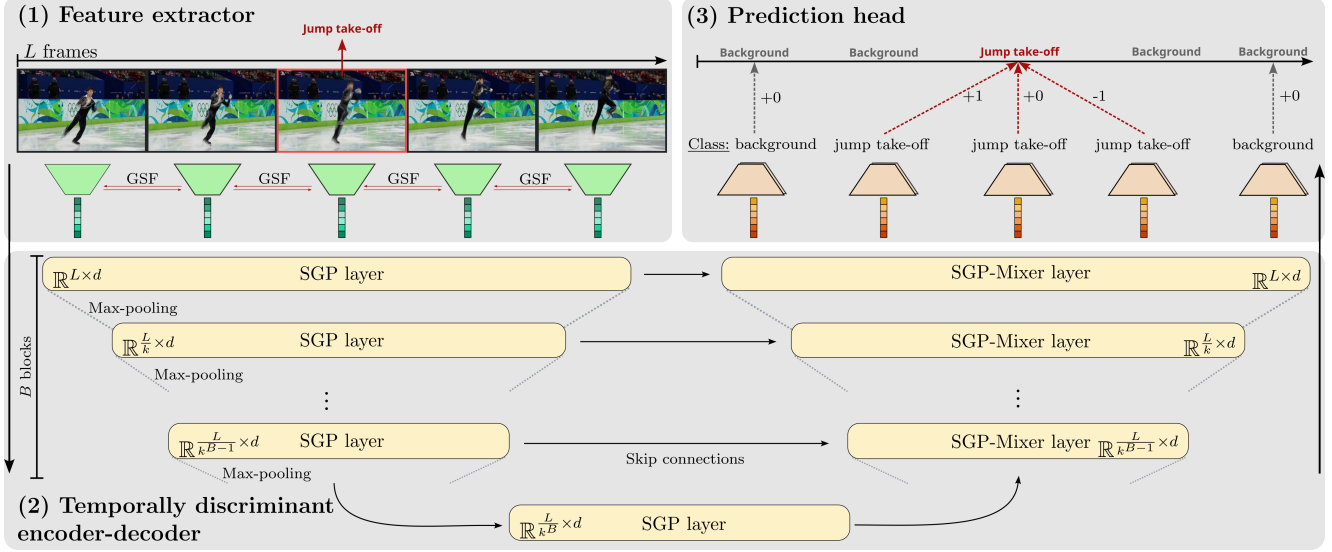


Figure 2. Illustration of **T-DEED** architecture comprising three key components: (1) **Feature extractor** to produce per-frame representations, (2) **Temporally discriminant encoder-decoder** to capture local and global temporal information while enhancing token discriminability, and (3) **Prediction head** to generate per-frame classifications and displacements for refinement.

formation. This approach proves effective in their proposed task of PES across four different datasets [19, 30, 38, 45], and is also adaptable to the coarser AS task in SoccerNet. While many TAL and AS methods [31, 32, 36, 43] rely on pre-extracted features due to their efficiency in training, end-to-end approaches have demonstrated that they can be beneficial in learning more meaningful features in some cases. This is exemplified by Hong et al. [20], particularly in scenarios where precise predictions are essential, such as in the case of PES. Exploiting these advantages, T-DEED also adopts an end-to-end approach.

3. Method

Problem definition. Precise Event Spotting (PES) involves the identification and localization of actions within an untrimmed video X , as illustrated in Figure 1. Given the video input, the objective is to recognize and locate all the events occurring in the video, represented as $E = \{e_1, \dots, e_N\}$. The number of events, denoted as N , may vary across different videos. Each event instance e_i comprises an action class $c_i \in \{1, \dots, C\}$ (where C is the number of distinct event classes) and its corresponding temporal position t_i (i.e. the exact frame where it occurs), forming a pair $e_i = (c_i, t_i)$.

Method overview. Our model, Temporal-Discriminability Enhancer Encoder-Decoder (T-DEED), is designed to increase token discriminability for Precise Event Spotting (PES) while leveraging multiple temporal scales. As illustrated in Figure 2, T-DEED comprises three main blocks: a feature extractor, a temporally discriminant

encoder-decoder, and a prediction head. We process videos through fixed-length clips, each containing L densely sampled frames. The feature extractor, composed of a 2D backbone with Gate-Shift-Fuse (GSF) modules [34], handles the input frames, generating per-frame representations of dimension d , hereby referred to as *tokens*. These tokens undergo further refinement within the temporally discriminant encoder-decoder. This module employs SGP layers which – as shown by Shi et al. [31] – diminish token similarity, thereby boosting discriminability across tokens of the same sequence. The encoder-decoder architecture allows the processing of features across diverse temporal scales, helpful for detecting events requiring different amounts of temporal context. We also integrate skip connections to preserve the fine-grained information from the initial layers. To effectively merge information proceeding from varying temporal scales in the skip connections, we introduce the SGP-Mixer layer, detailed in Section 3.2. This layer employs the same principles as the SGP layer to enhance token discriminability while gathering information from varying range temporal windows. Finally, the output tokens are directed to the prediction head, resembling those commonly used in Action Spotting (AS) literature [11, 32, 36]. It encompasses a classification component to identify whether an event occurs at the given temporal position or in close proximity (within a radius of r_E frames). Additionally, for the positive classifications, a displacement component pinpoints the exact temporal position of ground truth events.

Further details of the proposed method are discussed in the following sections.

3.1. Feature extractor

The feature extractor processes the input frame sequence, $\mathbb{R}^{L \times H \times W \times 3}$ with $H \times W$ denoting the spatial resolution, and produces per-frame feature representations, $\mathbb{R}^{L \times d}$. Following Hong et al. [20], we employ a compact 2D backbone to ensure efficiency and to accommodate longer video sequences. For comparability with previous PES methods, we choose RegNetY [29] as our feature extractor, known for its efficiency. To incorporate local temporal information during the extractor, we utilize GSF [34], which overperforms GSM [33] in action recognition tasks. GSF modules are specifically applied to the latter half of the RegNetY backbone, enabling local spatial modelling before integrating longer-term temporal context.

3.2. Temporally discriminant encoder-decoder

The temporally discriminant encoder-decoder processes the tokens produced by the feature extractor module, which contain mainly spatial information, with the objective of enriching them with essential temporal information – both local and global – to enable precise event predictions. To accomplish this, it incorporates three key components: (1) an encoder-decoder architecture for exploiting information across diverse temporal scales, (2) SGP layers to increase discriminability among tokens within the same sequence, and (3) our novel SGP-Mixer layer, designed to effectively fuse features from varying temporal scales while maintaining the distinctiveness of its output tokens.

Encoder-decoder architecture. Various works [31, 32, 43] have shown the benefits of processing features at different temporal scales in tasks related to action recognition. This is because certain actions inherently require longer temporal context for recognition and localization, while others can be identified with just a few frames. SOTA models in TAL achieve this by employing a feature pyramid, where the temporal dimension is downsampled by a factor of k in the final layers through max-pooling, and predictions are made for each output feature at various temporal scales. However, in PES, where precision is crucial, this is detrimental. As shown in Section 4.5, the further down we go in the feature pyramid, the more diminishes the precision of the output predictions, inevitably impacting model performance. To address this issue while still leveraging different temporal scales, we propose an encoder-decoder architecture. After temporal downscaling during the encoder, akin to the feature pyramid, we restore the original temporal resolution through the decoder, thereby regaining frame-level granularity for the representations. We do this by incorporating skip connections from the encoder to the decoder’s upsampling.

Specifically, the encoder begins by complementing the tokens with a learnable encoding that specifies its temporal

position. Furthermore, it comprises B encoder blocks, each consisting of an SGP layer to enhance discriminability while capturing temporal context, and a max-pooling operation to reduce the temporal dimension by a factor of k . An additional SGP layer is applied in the neck of the encoder-decoder, before passing the features to the decoder. In the decoder, the original temporal resolution is restored through B decoder blocks. Each of these blocks incorporates an SGP-Mixer layer, which extends the SGP layer to increase the temporal dimension by a factor of k while integrating information coming from the skip connections.


SGP layer. The SGP layer, as introduced by Shi et al. [31], addresses the rank-loss problem [12] commonly encountered in Transformers [35], thus improving token discriminability within sequences. As depicted in the top of Figure 3, it replaces the self-attention module in the Transformer layer with an SGP module illustrated at the bottom of the figure. This module comprises two primary branches: an instant-level branch and a window-level branch. The instant-level branch aims to boost token discriminability by increasing its distance from the clip-level average token, while the window-level branch captures temporal information from multiple receptive fields. Furthermore, the SGP layer substitutes one of the layer normalization modules with group normalization.

Given its advantageous characteristics, which contributed to Tridet’s SOTA results in TAL, we believe it can be even more beneficial in PES, where precise frame-level predictions are crucial and can benefit from improved discriminability between concurrent tokens, as discussed in Section 4.5.

SGP-Mixer layer. As shown in the top of Figure 4, the SGP-Mixer layer extends the original SGP layer to accommodate two inputs with distinct temporal scales. This adaptation becomes necessary in our architecture due to the reception of an input feature $z \in \mathbb{R}^{L/k^j \times d}$ from the preceding layer, and features $x \in \mathbb{R}^{L/k^{(j-1)} \times d}$ from the skip connection, where $j \in \{1, \dots, B\}$ denotes the depth within the decoder blocks. Both features undergo layer normalization before entering the SGP-Mixer module, where a combination of both features is generated. Finally, the output features are further processed through identical components within an SGP layer.

The fusion of both features is done in the SGP-Mixer module, shown at the bottom of the Figure 4. First, it up-samples the features from the previous layer to match the temporal dimensions ($\mathbb{R}^{L/k^j} \rightarrow \mathbb{R}^{L/k^{(j-1)}}$) using linear interpolation. Following the principles of the SGP layer, it has two instant-level branches – one for each input feature – to improve token discriminability. This is particularly im-

SGP Layer



SGP Module

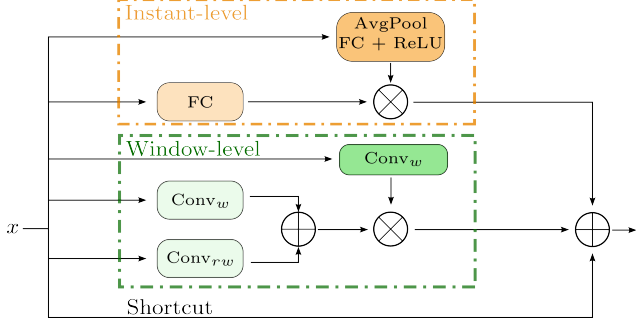


Figure 3. Illustration of the SGP layer structure, with its module comprising an instant-level branch to boost token discriminability and a window-level branch for temporal modeling.

portant for the features in the previous layer, since upsampling may produce similar tokens in adjacent temporal positions. In addition, it includes two window-level branches to aggregate information from both features while capturing different temporal contexts. In each branch, one feature evolves to mix information across different temporal receptive fields, which is later gated by the other feature. Shortcuts are also introduced, but, unlike SGP, feature aggregation from different branches involves concatenation and linear projection, which we find more effective than simple addition, as will be shown later in Section 4.5.


3.3. Prediction head

Following common AS approaches [11, 32, 36], we include a prediction head consisting of a classification head and a displacement head. The classification head uses a linear layer and a softmax activation to project the output of the temporally discriminant encoder-decoder, $\mathbb{R}^{L \times d}$, onto $\hat{y}^c \in \mathbb{R}^{L \times (C+1)}$, representing the probability of each temporal position containing each of the events or a background class. Similarly, the displacement head uses a linear layer to project the same output to $\hat{y}^d \in \mathbb{R}^{L \times 1}$, representing the displacement toward the ground truth event if an actual event is present at the corresponding temporal position.

3.4. Training details

The model is trained using a combination of a classification loss (\mathcal{L}_c) and a displacement loss (\mathcal{L}_d). For classification, per-frame cross-entropy loss is employed weighting the positive classes by a factor of w for all events. Mean squared error is used for displacement. The final loss (\mathcal{L}) for a given clip is the sum of both losses:

SGP-Mixer Layer



SGP-Mixer Module

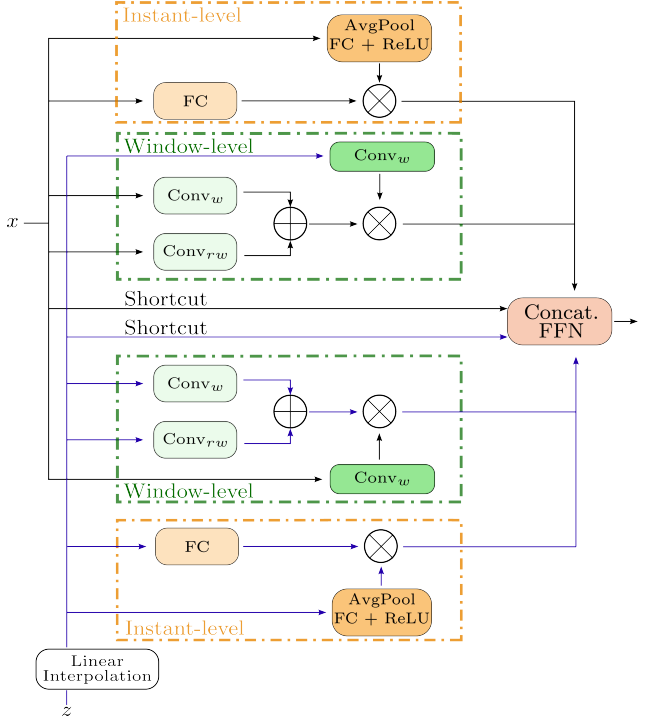


Figure 4. Illustration of the SGP-Mixer layer structure integrating an SGP-Mixer module to aggregate features of different temporal scales. This module follows the SGP principles, incorporating instant-level and window-level branches to boost token discriminability while merging the features.

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_d = \frac{1}{L} \sum_{l=1}^L CE_w(y_l^c, \hat{y}_l^c) + MSE(y_l^d, \hat{y}_l^d), \quad (1)$$

where y_l^c represents the one-hot encoding of the event in frame l , \hat{y}_l^c denotes the classification probabilities at that frame, y_l^d indicates the actual displacement to a ground-truth event (if an event is within the detection radius r_E), and \hat{y}_l^d represents the predicted displacement.

During training, we perform typical data augmentation techniques such as random cropping, random horizontal flip, gaussian blur, color jitter, and mixup [44].

3.5. Inference

At inference time, the data augmentation techniques are disabled and we use clips with 50% of overlapping. Moreover, to reduce the number of candidate events, Soft Non-Maximum Suppression [2] is applied.

Table 1. Comparison with SOTA methods on FigureSkating (FS-Comp and FS-Perf splits) and FineDiving datasets, highlighting the best results in bold and the second-top performance with underlining. For each model, we report feature extractor sizes in MegaFlops (MF) and input modalities including RGB images and Optical Flow (OF).

Model	Size	Input	FS-Comp		FS-Perf		FineDiving	
			$\delta = 1$	2	1	2	1	2
E2E-Spot [20]	200MF	RGB	81.00	<u>93.50</u>	85.10	95.70	68.40	85.30
E2E-Spot [20]	800MF	RGB	84.00	-	83.60	-	64.60	-
E2E-Spot [20]	800MF	RGB + OF [21]	83.40	94.90	83.30	<u>96.00</u>	66.40	84.80
T-DEED	200MF	RGB	85.15	91.70	<u>86.79</u>	96.05	<u>71.48</u>	<u>87.62</u>
T-DEED	800MF	RGB	<u>84.77</u>	92.86	88.17	95.87	73.23	88.88

4. Results

In this section, we detail the evaluation setup for our proposed method and present experimental results demonstrating its superiority over current SOTA methods in PES on the application domain of sports data. Additionally, we conduct ablations on key components of our proposal.

4.1. Datasets

We conduct experiments on two challenging datasets: FigureSkating [19] and FineDiving [38]. FigureSkating comprises 11 videos featuring 371 short-program performances from Winter Olympics and World Championships, annotated with 4 event classes. Annotations cover 0.23% of the frames. Evaluation is performed on two different splits: FS-Comp and FS-Perf, as detailed in Hong et al. [20]. FineDiving [38] consists of 3000 diving clips annotated with 4 different event classes, covering 2.2% of the frames. Both datasets' videos have a frame rate between 25 and 30 frames per second.

4.2. Evaluation

Following standard practices, we train on the train split, use the validation split for early stopping, and evaluate on the test split for all datasets. Model performance is assessed using the mAP metric, which calculates the mean of Average Precisions across different events. We employ a tight version of the metric with a tolerance of $\delta = 1$ frame and a loose metric with $\delta = 2$ frames.

4.3. Implementation details

We train T-DEED using clips of $L = 100$ frames with a batch size of 8 clips. Each epoch comprises 5000 clips, randomly sampled from the training videos. Models are trained for 50 epochs using AdamW optimizer [26], with a base learning rate of $8e-04$, with 3 linear warmup epochs, and cosine decay. Positive classes in the cross-entropy loss are weighted with $w = 5$. We evaluate two versions of our feature extractor, denoted as RegNetY-200MF and RegNetY-800MF with hidden dimensions set to $d = 368$ and $d = 768$, respectively, and a downscaling factor of

$k = 2$. Additional dataset-sensitive model hyperparameters are outlined in the supplementary material.

4.4. Comparison to SOTA

In Table 1, we compare our proposed T-DEED in two variations, with smaller and larger feature extractors, against previous SOTA models across four different configurations and datasets. We present mAP scores with $\delta = 1$ and $\delta = 2$. For the Figure Skating dataset, in both FS-Comp and FS-Perf configurations, our models exhibit improvements on the tight metric, with gains of up to +1.15 and +3.07, respectively. However, performance on the loose metric, with a 2-frame tolerance, is comparable for FS-Comp and slightly worse in FS-Perf. We attribute this to the use of dilation in the E2E-Spot model for the FigureSkating dataset, which may be beneficial when using larger tolerances. Additionally, in the FineDiving dataset, T-DEED clearly outperforms the SOTA on both metrics, with improvements of up to +4.87 on the tight metric and +3.58 on the loose metric.

4.5. Ablations

In this section, we incrementally ablate the different components of our approach. We start with a base model consisting of the feature extractor without GSF, producing features with only spatial information, and plain layers of the different temporal approaches. Initially, we consider only the classification head (i.e., $\mathcal{L} = \mathcal{L}_c$), resulting in per-frame classifications. For each approach, we evaluate multiple configurations with different hyperparameters (e.g., number of layers) and present results with the best configuration, reporting the mAP with a tolerance of $\delta = 1$. Each experiment is trained with two different seeds and we report the average for robustness. Ablations are conducted using two datasets: FineDiving and FigureSkating on the FS-Comp split.

Enhancing token discriminability. As previously discussed, tasks like PES require tokens with sufficient discriminability to prevent precision loss in predictions due to similar representations for spatially similar adjacent frames.

Table 2. Ablation of T-DEED’s main components using mAP with $\delta = 1$, highlighting the best results in bold.


Experiments	FS-Comp mAP ($\delta = 1$)	FineDiving mAP ($\delta = 1$)
<i>(a) Temporal module</i>		
Transformer	69.68	59.84
GRU	72.04	59.09
SGP	73.14	60.07
<i>(b) Skip connection</i>		
w/o	74.29	63.01
sum	73.27	60.14
concat	76.78	61.90
SGP-Mixer (sum)	76.88	61.34
SGP-Mixer	77.96	63.67

Here, we evaluate the discriminability of three commonly used layers for modeling temporal information: Transformer [35], GRU [7], and SGP [31]. We measure the discriminability of their output tokens by averaging the cosine similarity between each token and the mean token of the sequence, as described in Shi et al. [31].


Figure 5 illustrates the rank-loss problem in Transformers, where token similarity increases across layers, therefore losing discriminability. In contrast, GRU layers, typically requiring fewer layers, show slight improvement in similarity. However, SGP layers exhibit the lowest similarity (i.e., highest discriminability) among output tokens, demonstrating its effectiveness in enhancing token discriminability. Additionally, this enhanced discriminability leads to the best performance among the different temporal modules, as shown in Table 2a. This observation highlights the advantages of employing the SGP layer, particularly in precision-demanding tasks like PES.

Defining multiple temporal scales. Employing multiple temporal scales while processing videos has proven effective across various TAL and AS methodologies [31, 32, 43]. Here, we evaluate our encoder-decoder architecture, operating on multiple temporal scales, using diverse mixture approaches in skip connections. Specifically, we assess five variations: (1) no skip connections, (2) addition, (3) concatenation and linear projection, (4) a modified SGP-Mixer aggregating branches information with summation, and (5) our proposed SGP-Mixer layer. Further details of each approach can be found in the supplementary material.

Results of these approaches are presented in Table 2b, where we observe that all encoder-decoder-based approaches, with or without skip connections, outperform the baseline using a single scale. This underscores the importance of capturing information from multiple temporal scales, effectively achieved in our encoder-decoder. However, further analysis reveals issues with approaches uti-



(a) FS-Comp discriminability analysis.



(b) FineDiving discriminability analysis.

Figure 5. **Temporal module discriminability analysis.** Cosine similarity after backbone (BB), post-positional encoding (PE), and at each temporal layer is displayed. Additionally, mAP performance with $\delta = 1$ is reported.

lizing addition for feature aggregation within skip connections, as they tend to yield inferior results compared to alternative methods. We hypothesize that this problem may be due to the attribution of equal weight to features from skip connections and previous layers when adding them. This could be particularly critical in the top layers of the architecture, where information passed from skip connections may be too primitive. Finally, our proposed SGP-Mixer layer stands as the best approach for both datasets in aggregating information within the skip connections. This emphasizes the advantages of our carefully designed SGP-Mixer layer in aggregating information across multiple temporal scales while enhancing token discriminability.

Introducing the displacement head. In Table 3a we evaluate various methods for addressing the imbalance between frames containing actual events and background frames in PES. These methods include label dilation, which aims to detect events within a radius around the ground truth but may impact prediction precision, and the utilization of a displacement head, as detailed in Section 3.3. Results indicate that the prediction head offers more benefits compared to label dilation, enabling wider range of event detection without sacrificing prediction precision.

Table 3. Further ablations and analysis of T-DEED using mAP with $\delta = 1$.

Experiments	FS-Comp mAP ($\delta = 1$)	FineDiving mAP ($\delta = 1$)
<i>(a) Displacement head</i>		
w/o & dilation = 0	74.74	63.67
w/o & dilation = 1	77.96	61.97
$r_E = 1$	78.20	67.49
$r_E = 2$	77.12	68.31
<i>(b) Feature pyramids</i>		
Tridet	68.31	64.28
<i>(c) Feature extractor</i>		
w/ gsm	81.00	67.95
w/ gsm (half)	81.05	67.47
w/ gsf	80.43	67.87
w/ gsf (half)	81.25	68.40
<i>(d) Clip length</i>		
$L = 25$ frames	71.02	66.61
$L = 50$ frames	76.87	64.84
$L = 100$ frames	81.25	68.40
$L = 200$ frames	79.19	65.29
<i>(e) Postprocessing</i>		
NMS [27]	81.27	68.40
SNMS [2]	85.15	71.48

Feature pyramids. A common approach in TAL to process multiple temporal scales is the use of feature pyramids, which resembles using only the encoder of our approach. In Table 3b, we explore adapting the SGP Feature Pyramid proposed in Shi et al. [31] to our task. This method generates predictions at each temporal scale and integrates a displacement head to locate them. However, we observe a performance decrease compared to our encoder-decoder architecture. Further analysis of the predictions produced at each layer of the pyramid, depicted in Figure 6, reveals a decrease in performance as we descend the pyramid towards more high-level but lower-resolution features. This suggests that generating predictions at lower resolutions might compromise the accuracy. Our encoder-decoder architecture is able to model high-level information while recovering the original temporal resolution, thus avoiding losing the temporal precision that is critical in PES.

Feature extractor ablations. In Table 3c, we explore the impact of integrating different temporal shift modules into our 2D backbone to capture local temporal context. Specifically, we consider two modules, GSM [33] and GSF [34], known for their effectiveness in action recognition tasks. We assess two variants: one applying these modules across all backbone layers, and another limited to the latter half to promote stronger spatial modeling before temporal inte-




Figure 6. Per-layer mAP analysis using the Feature Pyramid Network. mAP is reported at each layer of the SGP feature pyramid, accumulating predictions from previous layers, for the FineDiving and FigureSkating datasets.

gration. Following [20], we apply these modules to $\frac{1}{4}$ of the channels of the residual blocks. Results indicate that incorporating local temporal modeling into the backbone is especially useful in the FigureSkating dataset, while FineDiving shows comparable results to those without temporal modules. Notably, GSF applied to the latter half of the backbone yields the best performance for both datasets.

Clip length analysis. Regarding optimal clip length, Table 3d evaluates various number of frames. In FigureSkating, performance improves with increasing clip length, plateauing at 100 frames. Similarly, results for FineDiving also indicate 100 frames as the optimal choice.

Postprocessing analysis. In Table 3e, we illustrate the impact of employing two postprocessing techniques: Non-Maximum Suppression (NMS) [27] and Soft Non-Maximum Suppression (SNMS) [2]. The optimal results are achieved with a 1-frame window for NMS and a 3-frame window for SNMS. Notably, SNMS consistently improves mAP for both datasets.

5. Conclusion

This work presented T-DEED, a model designed to address Precise Event Spotting across various sports datasets. Ablation studies underscore the importance of processing videos in multiple temporal scales and enhancing token discriminability for precise predictions. To address these challenges, we integrate an encoder-decoder architecture and propose the SGP-Mixer layer, aimed at aggregating information at various temporal scales within skip connections while improving token discriminability. Additionally, T-DEED achieves SOTA performance on the FigureSkating and FineDiving datasets.

Acknowledgements. This work has been partially supported by the Spanish project PID2022-136436NB-I00 and by ICREA under the ICREA Academia programme.

References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. [2](#)
- [2] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *Proceedings of the IEEE international conference on computer vision*, pages 5561–5569, 2017. [5](#), [8](#)
- [3] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2911–2920, 2017. [2](#)
- [4] Shyamal Buch, Victor Escorcia, Bernard Ghanem, Li Fei-Fei, and Juan Carlos Niebles. End-to-end, single-stream temporal action detection in untrimmed videos. In *Proceedings of the British Machine Vision Conference 2017*. British Machine Vision Association, 2019. [2](#)
- [5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. [2](#)
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. [2](#)
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. [2](#), [7](#)
- [8] Anthony Cioppa, Silvio Giancola, Vladimir Somers, Floriane Magera, Xin Zhou, Hassan Mkhallati, Adrien Deliege, Jan Held, Carlos Hinojosa, Amir M Mansourian, et al. Soccernet 2023 challenges results. *arXiv preprint arXiv:2309.06006*, 2023. [2](#)
- [9] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pages 720–736, 2018. [2](#)
- [10] Adrien Deliege, Anthony Cioppa, Silvio Giancola, Meisam J Seikavandi, Jacob V Dueholm, Kamal Nasrollahi, Bernard Ghanem, Thomas B Moeslund, and Marc Van Droogenbroeck. Soccernet-v2: A dataset and benchmarks for holistic understanding of broadcast soccer videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4508–4519, 2021. [2](#)
- [11] Julien Denize, Mykola Liashuha, Jaonary Rabarisoa, Astrid Orcesi, and Romain Hérault. Comedian: Self-supervised learning and knowledge distillation for action spotting using transformers. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 530–540, 2024. [2](#), [3](#), [5](#)
- [12] Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International Conference on Machine Learning*, pages 2793–2803. PMLR, 2021. [2](#), [4](#)
- [13] Victor Escorcia, Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Daps: Deep action proposals for action understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 768–784. Springer, 2016. [2](#)
- [14] Silvio Giancola, Mohieddine Amine, Tarek Dghaily, and Bernard Ghanem. Soccernet: A scalable dataset for action spotting in soccer videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1711–1721, 2018. [2](#)
- [15] Silvio Giancola, Anthony Cioppa, Adrien Deliege, Floriane Magera, Vladimir Somers, Le Kang, Xin Zhou, Olivier Barnich, Christophe De Vleeschouwer, Alexandre Alahi, et al. Soccernet 2022 challenges results. In *Proceedings of the 5th International ACM Workshop on Multimedia Content Analysis in Sports*, pages 75–86, 2022. [2](#)
- [16] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017. [2](#)
- [17] Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1914–1923, 2016. [2](#)
- [18] Samitha Herath, Mehrtash Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image and vision computing*, 60:4–21, 2017. [1](#)
- [19] James Hong, Matthew Fisher, Michaël Gharbi, and Kayvon Fatahalian. Video pose distillation for few-shot, fine-grained sports action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9254–9263, 2021. [1](#), [2](#), [3](#), [6](#)
- [20] James Hong, Haotian Zhang, Michaël Gharbi, Matthew Fisher, and Kayvon Fatahalian. Spotting temporally precise, fine-grained events in video. In *European Conference on Computer Vision*, pages 33–51. Springer, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [8](#)
- [21] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. [6](#)
- [22] Haroon Idrees, Amir R Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017. [2](#)
- [23] Chuming Lin, Chengming Xu, Donghao Luo, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yan-

- wei Fu. Learning salient boundary feature for anchor-free temporal action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3320–3329, 2021. 2
- [24] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 988–996, 2017. 2
- [25] Xiaolong Liu, Qimeng Wang, Yao Hu, Xu Tang, Shiwei Zhang, Song Bai, and Xiang Bai. End-to-end temporal action detection with transformer. *IEEE Transactions on Image Processing*, 31:5427–5441, 2022. 2
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [27] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *18th international conference on pattern recognition (ICPR’06)*, pages 850–855. IEEE, 2006. 8
- [28] Zhiwu Qing, Haisheng Su, Weihao Gan, Dongliang Wang, Wei Wu, Xiang Wang, Yu Qiao, Junjie Yan, Changxin Gao, and Nong Sang. Temporal context aggregation network for temporal action proposal refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 485–494, 2021. 2
- [29] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10428–10436, 2020. 4
- [30] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Finegym: A hierarchical video dataset for fine-grained action understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2616–2625, 2020. 2, 3
- [31] Dingfeng Shi, Yujie Zhong, Qiong Cao, Lin Ma, Jia Li, and Dacheng Tao. Tridet: Temporal action detection with relative boundary modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18857–18866, 2023. 1, 2, 3, 4, 7, 8
- [32] Joao VB Soares, Avijit Shah, and Topojoy Biswas. Temporally precise action spotting in soccer videos using dense detection anchors. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 2796–2800. IEEE, 2022. 2, 3, 4, 5, 7
- [33] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift networks for video action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1102–1111, 2020. 2, 4, 8
- [34] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift-fuse for video action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 3, 4, 8
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4, 7
- [36] Artur Xarles, Sergio Escalera, Thomas B Moeslund, and Albert Clapés. Astra: An action spotting transformer for soccer videos. In *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports*, pages 93–102, 2023. 2, 3, 5
- [37] Huifen Xia and Yongzhao Zhan. A survey on temporal action localization. *IEEE Access*, 8:70477–70487, 2020. 1
- [38] Jinglin Xu, Yongming Rao, Xumin Yu, Guangyi Chen, Jie Zhou, and Jiwen Lu. Finediving: A fine-grained dataset for procedure-aware action quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2949–2958, 2022. 1, 2, 3, 6
- [39] Mengmeng Xu, Chen Zhao, David S Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10156–10165, 2020. 2
- [40] Le Yang, Houwen Peng, Dingwen Zhang, Jianlong Fu, and Junwei Han. Revisiting anchor mechanisms for temporal action localization. *IEEE Transactions on Image Processing*, 29:8535–8548, 2020. 2
- [41] Le Yang, Junwei Han, Tao Zhao, Nian Liu, and Dingwen Zhang. Structured attention composition for temporal action localization. *IEEE Transactions on Image Processing*, 2022. 2
- [42] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 126:375–389, 2018. 2
- [43] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *European Conference on Computer Vision*, pages 492–510. Springer, 2022. 2, 3, 4, 7
- [44] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [45] Haotian Zhang, Cristobal Sciutto, Maneesh Agrawala, and Kayvon Fatahalian. Vid2player: Controllable video sprites that behave and appear like professional tennis players. *ACM Transactions on Graphics (TOG)*, 40(3):1–16, 2021. 2, 3
- [46] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8668–8678, 2019. 2