

Mixture of Experts Guided by Gaussian Splatters Matters: A new Approach to Weakly-Supervised Video Anomaly Detection

Giacomo D’Amicantonio^{5*}, Snehashis Majhi^{1,2*}, Quan Kong³, Lorenzo Garattoni⁴, Gianpiero Francesca⁴, François Brémond^{1,2}, Egor Bondarev⁵

¹ INRIA ² Côte d’Azur University ³ Woven by Toyota ⁴ Toyota Motor Europe ⁵ Eindhoven University of Technology

* Joint first authors. Code: <https://github.com/snehashismajhi/GS-MoE>

Abstract

*Video Anomaly Detection (VAD) is a challenging task due to the variability of anomalous events and the limited availability of labeled data. Under the Weakly-Supervised VAD (WSVAD) paradigm, only video-level labels are provided during training, while predictions are made at the frame level. Although state-of-the-art models perform well on simple anomalies (e.g., explosions), they struggle with complex real-world events (e.g., shoplifting). This difficulty stems from two key issues: (1) the inability of current models to address the diversity of anomaly types, as they process all categories with a shared model, overlooking category-specific features; and (2) the weak supervision signal, which lacks precise temporal information, limiting the ability to capture nuanced anomalous patterns blended with normal events. To address these challenges, we propose **Gaussian Splating-guided Mixture of Experts (GS-MoE)**, a novel framework that employs a set of expert models, each specialized in capturing specific anomaly types. These experts are guided by a temporal Gaussian splatting loss, enabling the model to leverage temporal consistency and enhance weak supervision. The Gaussian splatting approach encourages a more precise and comprehensive representation of anomalies by focusing on temporal segments most likely to contain abnormal events. The predictions from these specialized experts are integrated through a mixture-of-experts mechanism to model complex relationships across diverse anomaly patterns. Our approach achieves state-of-the-art performance, with a 91.58% AUC on the UCF-Crime dataset, and demonstrates superior results on XD-Violence and MSAD datasets. By leveraging category-specific expertise and temporal guidance, GS-MoE sets a new benchmark for VAD under weak supervision.*

1. Introduction

Video Anomaly Detection (VAD) in surveillance videos is one of the most challenging tasks in the field of Computer

Vision. With the increasing capabilities of deep-learning models, there have been various approaches to tackle this task. The main focus of recent research in the field of VAD has been to model spatio-temporal dependencies in videos, obtaining meaningful representations of the motion of relevant agents in the scene. In this sense, the transformer architecture has proved to be very effective, forming the backbone of multiple works. While the current state-of-the-art models have achieved reasonable results on publicly available datasets, they still fail to capture subtle anomalies and to detect the temporal window in which they happen.

We identify one of the main reasons for these issues in the formulation of the WSVAD task [31, 36]. Multi-instance learning (MIL) strikes a balance between fully supervised methods, which exhibit good performance but require costly data annotation, and unsupervised methods, which do not require manual annotations but generally result in worse performance. The core idea of MIL is to create bags containing positive and negative data samples (*i.e.*, normal and abnormal videos), labeled only at the video-level. During training, the model assigns a score between 0 and 1 to each snippet, with 0 indicating a normal snippet and 1 indicating an abnormal snippet. The highest-scoring samples in the normal bag are guided towards 0, allowing the model to learn most normal scenarios correctly. On the other hand, the highest-scoring negative samples are pushed towards 1. This leads the model to be supervised, and therefore learn, few and specific instances of anomalous events, ignoring useful information contained in neighboring snippets. Over time, this approach has proved to be powerful but insufficient to train a model to correctly capture the secondary and specific attributes of different anomalous classes. In recent works [8, 38, 40], different auxiliary objectives are identified as priors for the VAD task to optimize the training process.

To address this issue, we propose to model the anomalies in a video as Gaussian distributions, rendering multiple Gaussian kernels in correspondence with peaks detected along the temporal dimension of the scores estimated for




Figure 1. While SoTA methods address the task of WSVAD via the most normal and abnormal snippets in a video, the approach proposed in this paper focuses on learning a more complete representation of anomalous events via Gaussian kernels.

abnormal videos. This technique, called Temporal Gaussian Splatting (TGS), creates a more complete representation of an anomalous event over time, including snippets of the anomaly with lower abnormal scores in the training objective. A side-by-side comparison of the MIL task and the TGS task is shown in Figure 1. The Gaussian kernels are extracted from the abnormal scores produced by the model. An additional challenge is related to the intrinsic differences between abnormal classes. Under the MIL paradigm, the models are trained to learn the difference between normal and abnormal videos, while the specific differences between anomalous classes are overlooked. As a result, these methods mainly focus on coarse-level representations of anomalies that allow us to distinguish between normal and abnormal events, but ignore the fine-grained category-specific cues. Therefore, the more salient anomalies (*i.e.*, such as an explosion) are likely to be easily detected, while subtle anomalies (*i.e.*, shoplifting) are more likely to be confused with normal events. This constitutes a major limitation of most recent methods based on WSVAD. We address this issue via a Mixture-of-Expert (MoE) architecture, in which each expert is trained to model a single anomaly class, enhancing the specific attributes of each anomaly class that are often overlooked. To further leverage the correlations and differences between anomalies, a gate model mediates the predictions of each expert and the more coarse-level anomalous features to learn potential interactions between anomalies.

The contributions of this paper are complementary: learning specific representations of anomalous classes allows for more accurate Gaussian kernels, and the Gaussian splatting enables the experts to learn from more subtle anomalous events that would be overlooked otherwise. To summarize, this paper presents:

- A novel formulation of the WSVAD task based on Gaussian kernels extracted from the estimated abnormal scores to generate a more expressive and complete representation of anomalous events. Splatting the kernels along the

temporal dimension allows the model to learn more precise temporal dependencies between snippets and highlight more subtle anomalies;

- A Mixture-of-Expert (MoE) architecture that focuses on individual anomaly types via dedicated class-expert models, allowing a gate model to leverage similarities and diversities between them;
- The impact of the proposed contributions is measured via an extensive set of experiments on the challenging UCF-Crime [30], XD-Violence [35] and MSAD [51] datasets, showing notable improvements in performance w.r.t. previous SoTA methods.

2. Related Work

Weakly-Supervised VAD: In the WSVAD task, anomalous events encompass various classes, each exhibiting distinct characteristics across the spatial and temporal dimensions. The task of WSVAD was introduced in a seminal work by [31]. In the following years, there have been multiple different approaches that addressed the trade-off between the ease of data collection and the performance exhibited by models trained in this task. The limitation of weak labels was addressed by [48] using a graph convolutional network to correct noisy labels and supervise traditional anomaly classifiers. Further, [33] proposed to learn a function of the magnitude of features to improve the classification of normal snippets and, therefore, the detection of abnormal events. The model is based on attention modules and pyramidal convolutions. The idea of improving the quality of weak labels was also explored by [18], which designed a transformer-based method trained to predict abnormal scores both at the snippet and video levels. The video-level predictions are then used to improve the performance of the model at the snippet level. More recently, [43] designed a multi-head classification model that leveraged uncertainty and completeness to produce and refine its own pseudo-labels. [24] proposed a two-stage transformer-based model that generates anomaly-aware position embeddings and then models the short and long-range relationships of anomalous events. Inspired by point-supervision [1], [44] introduced Glance annotations. These annotations enhance the common weak labels by localizing a single frame in which an anomalous event is happening. While reporting very good performance, these annotations require an additional manual-labelling procedure. Recently, [25] proposed a method to include additional data modalities in the anomaly detection process.

Under the MIL paradigm, these variations complicate the model’s ability to effectively differentiate between them. By focusing on the top-k most abnormal snippets of a video, the model is guided towards specific and evident anomalous events, without properly considering the sequence of actions that lead to them and follow them. In fact, some

anomalies occur within short time windows, while others unfold over longer periods; moreover in both cases, the MIL paradigm selects the same amount of abnormal snippets.

Mixture of Experts: This architecture has been introduced by [5] and has since been improved and employed for diverse tasks, from image classification to action recognition [10]. The original MoE design proposed a series of small experts and a separate gate network, all receiving the same input data. Each expert predicts an output, while the gate network assigns a score of importance to them. Since then, this architecture has been improved upon by various works. A common idea across domains is to let a routing network select which portions of the input data, or input tokens, to pass to each expert [6, 14, 26, 29]. A recent work by [27] proposed to weight the input tokens in a different way for each expert.

Gaussian Splatting: It has received a lot of attention in recent years, proving to be very efficient in fields like 3D scene reconstruction [12, 13]. The main idea of Gaussian Splatting is to represent each 3-dimensional point in a scene as a multivariate normal distribution, which allows to render the scene as the sum of the contributions of all the 3-dimensional areas. Gaussian splatting has since been extended to incorporate the temporal dimension in multiple domains, for example, dynamic scene rendering [15, 19] and medical imaging [46].

Our approach utilizes Mixture-of-Experts (MoE) by assigning each expert to a specific anomaly class, enabling fine-grained, category-specific learning often missed in traditional designs. A gate model bridges class-specific experts and coarse features, leveraging anomaly correlations while ensuring balanced utilization. We extend Gaussian Splatting into the temporal domain with Temporal Gaussian Splatting (TGS), capturing nuanced dependencies and integrating subtle, low-scoring snippets into training. Anchored to temporal peaks, TGS mitigates noise, enhances weak supervision, and preserves sharp transitions, avoiding over-smoothing while delivering precise anomaly detection.

3. Methodology

Our novel Gaussian Splatter-guided Mixture-of-Experts (GS-MoE) framework aims to accurately detect complex anomalies using weakly-labeled training videos. GS-MoE leverages two key techniques: **(I) Temporal Gaussian Splatting loss**, to ensure superior separability between normal and anomalous instances under weak-supervision; **(II) Mixture-of-Experts (MoE) architecture**, that learns class-specific representations and detects complex anomalies with high confidence.

3.1. Temporal Gaussian Splatting (TGS)

Our Temporal Gaussian Splatting (TGS) technique provides a novel formulation of the MIL optimization paradigm by




Figure 2. The abnormal scores obtained from the backbone model on a training video at the end of training. The top_k snippets used in the MIL paradigm lead the model to focus on the first and last of the three anomalous events present in the video, overlooking the second anomaly. However, the second anomaly, while not scoring as high as the others, is still detected.




Figure 3. The Gaussian kernels extracted from the abnormal scores shown in Figure 2 are splatted across the width of the detected peaks. This allows the model to learn a more complete representation of the anomalous events in the video.

leveraging Gaussian kernels. The core idea of TGS is to reduce the over-dependency on the most abnormal snippets that is often the result of the classical MIL. An example of such over-dependency is shown in Figure 2. The Top- k abnormal scores are the ones that would normally be used in the loss function in the MIL paradigm:

$$\text{Top-}k(S) = \{score_1, score_2, \dots, score_k\} \quad (1)$$

such that

$$L_{\text{top}_k} = -\underbrace{\frac{1}{N^+} \sum_{i=1}^{N^+} \frac{1}{k} \sum_{j \in \text{top}_k(S^+)} \log \sigma(\text{score}_{ij})}_{L_{\text{top}_k - \text{abn}}} - \underbrace{\frac{1}{N^-} \sum_{i=1}^{N^-} \frac{1}{k} \sum_{j \in \text{top}_k(S^-)} \log (1 - \sigma(\text{score}_{ij}))}_{L_{\text{top}_k - \text{norm}}} \quad (2)$$

where $score_i$ is the score of snippet i of video S . S^+ and S^- denote the set of scores obtained for an abnormal and a normal video, respectively. Similarly, N^+ and N^- are the number of videos in the abnormal and normal classes and σ is the sigmoid function.

At the end of the training, the task encoder is able to detect two out of three anomalies contained in the video as in Figure 2, assigning a very high abnormal score to most

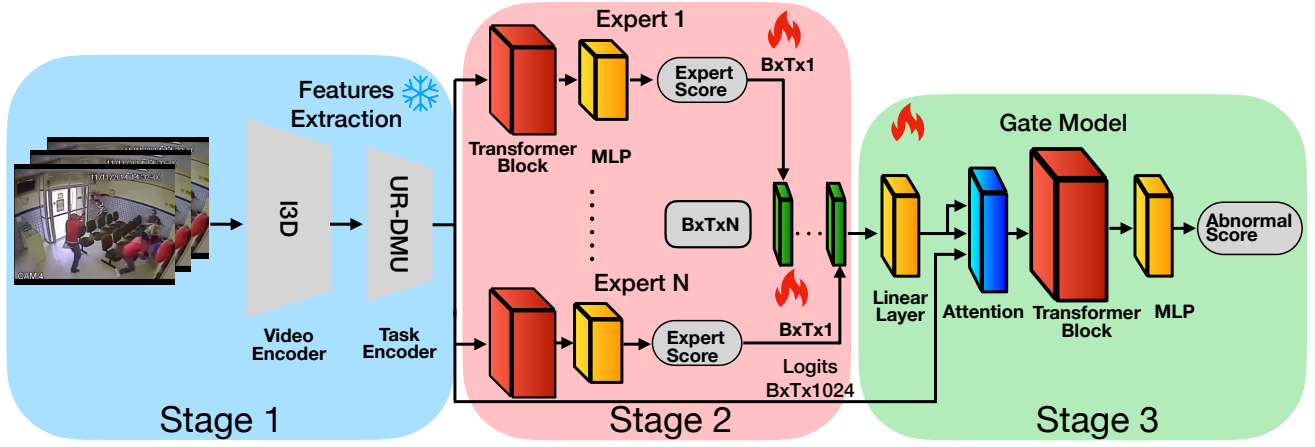


Figure 4. **Overview of the GS-MoE architecture:** First, in the feature extraction stage, the video encoder extracts snippet-level features from the video, and the task encoder refines them in the anomaly-detection latent space. In the second stage, each class-expert is trained only on refined features belonging to its assigned class and to the normal class. In the final stage, the gate model collects the scores assigned by each expert and compares them with the refined features of the task encoder, producing the final abnormal score.

snippets in the first and third anomalies time window. The model is not as confident about the snippets belonging to the second anomaly, due to the fact that during training, it has never been supervised specifically on them, but it assigns them an anomalous score higher than the normal snippets of the video. Additionally, the snippets between anomalies are still considered partially anomalous. We conjecture that it is possible to leverage those situations to generate pseudo-labels that allow a model to be trained on more information, while remaining in the data-annotation boundaries of the WSVAD paradigm. Following [44], we propose a technique called Temporal Gaussian Splatting (TGS) that leverages the peaks in the abnormal scores predicted by the model to represent precisely the temporal windows in which an anomalous event is happening.

Peak detection: Gaussian kernels are extracted in correspondence with local maxima, called "peaks", in the temporal axis of the abnormal scores predicted by a model. The peaks are detected by thresholding the local maxima, selecting only the ones that are above a minimum prominence threshold over the previous two scores and the subsequent two. The width W_i of each peak P_i is determined by $\text{minimum}(v_1, v_2)$, where v_1 is the number of preceding snippets with monotonically increasing scores, and v_2 is the number of following snippets with monotonically decreasing scores.

The set of peaks P , detected for a given video, contains the position of the snippet with the highest abnormal score for each peak in the video. This may lead to the detection of spurious peaks, meaning peaks in the abnormal scores of a video that do not belong to an anomalous event. To mitigate this, the model can be trained for a few iterations with the $L_{\text{topk-norm}}$ component of the standard MIL training objective.

This allows us to identify subtle anomalies that are usu-

ally not included in the top-k snippets described in Equation 2. The kernels obtained from the detected peaks are then rendered over the length of anomalous videos to obtain a more accurate representation of the anomalies along the temporal dimension.

Gaussian kernels G_i are then initialized with a unitary value for the snippets corresponding to each peak P_i detected in the abnormal scores of the video. To further represent the duration of the anomaly, the kernel values corresponding to snippets that are within the width W_i of the respective peak are also set to 1 if their abnormal score is higher than the difference between the peak score and the standard deviation of the normal distribution centered in the peak:

$$G_{i,t} = \begin{cases} 1, & \text{if } t = P_i, \\ 1, & \text{if } s_t \geq s_{P_i} - \sigma_i \wedge t \in W_i, \forall t \in [1, T] \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where s_t is the abnormal score assigned to snippet t and σ_i is the standard deviation of the normal distribution centered in peak i . This allows to treat each anomaly separately, which is beneficial for the WSVAD task due to the fact that different anomalies have different characteristics along the temporal dimension. Computing the Gaussian kernels in this way represents an improvement upon the top-k formulation, allowing the model to learn from the entirety of an anomalous event instead of its most abnormal snippets. Each kernel is splatted via:

$$f_i(t) = G_{i,t} \cdot \exp\left(-\frac{\|t - P_i\|^2}{2\sigma_i^2}\right), \forall t \in [1, T] \quad (4)$$

where T is the length of the video and σ_i is the standard deviation of the scores around the peak centered in P_i within the width W_i . Finally, the pseudo-labels \hat{y} are generated by rendering each of the K extracted kernels over the

length of the video:

$$\hat{y} = \|(\sum_{i=1}^k f_i(t))\| \quad (5)$$

An example of such pseudo-label (Temporal Gaussian Splatting) is shown in Figure 3. The generated pseudo-labels contain a target abnormal score between 0 and 1 for each snippet in the video, allowing the model to learn the severity of each abnormal snippet. This represents a relevant improvement over the standard MIL training objective, where only the top-k snippets are pushed towards 1 in the training objective, as in Equation 2. Instead, the TGS loss function used to train the experts and the MoE is formulated as:

$$L_{TGS} = L_{topk-norm} + BCE(y, \hat{y}) \quad (6)$$

3.2. Mixture of Experts (MoE)

Our Mixture-of-Experts (MoE) architecture, illustrated in Figure 4, combines three stages. The first is task-agnostic and task-aware feature extraction, the second has class-specific expertise, and the third provides a novel gate mechanism. This multi-stage framework directly addresses the challenges of weak supervision and anomaly diversity, enabling precise detection of complex anomalous patterns through enriched representations and specialized models.

Stage 1: Enhanced Temporal-Spatial Feature Extraction. The synergy of I3D and UR-DMU forms the backbone of feature extraction. The widely-used I3D model provides task-agnostic general features, capturing basic video dynamics. However, these features lack the specificity needed for detecting intricate spatial and temporal anomalies. To address this, UR-DMU [50] acts as a task-aware feature extractor tailored to anomaly detection. Trained initially with the standard MIL loss [49] and fine-tuned using our novel Temporal Gaussian Splatting (TGS) loss (Equation 6), UR-DMU extracts highly informative features by leveraging temporal consistency. These richer features filter motion dynamics and fine-grained temporal patterns crucial for distinguishing between normal and abnormal events. While UR-DMU can coarsely differentiate normal and abnormal events, it cannot fully address the complexities of different anomaly types.

Stage 2: Class-Specific Anomaly Detection with Expert Models. To overcome the limitations of coarse anomaly detection, our framework incorporates multiple expert models optimized with our TGS loss (Equation 6), each dedicated to identifying specific types of anomalies. This design introduces a crucial level of specialization, allowing the framework to capture the unique attributes of individual anomaly classes. Each expert model is composed of a transformer block with four self-attention heads and a MLP with

GELU activation [9], which maps the extracted features to an anomaly score for its respective class. These expert models leverage the enriched UR-DMU features, expanding the boundaries of the latent anomaly space. These fine-grained, specialized experts enable the model to detect subtle or complex anomalies that may blend seamlessly with normal events, a feat that generic models fail to achieve.

Stage 3: Collaborative Integration with the Gate Model.

In the final stage of the framework, the scores generated by the expert models are passed to the gate model, which acts as collaborative integration mechanism. This step ensures that the individual strengths of the experts are harnessed to create a unified representation capable of robust anomaly detection. The gate model comprises of three components: **(a) Score Refinement:** The expert scores are concatenated and projected into a higher-dimensional space, enriching the representation of class-specific anomaly logits. This projection enables the gate model to handle the intricate variations across anomaly classes effectively. **(b) Bi-Directional Cross-Attention Module:** To bridge the gap between fine-grained class-specific logits and coarse abnormal logits from the task encoder, the gate model incorporates a bi-directional cross-attention mechanism. This module learns the correlations and contrasts between the expert predictions and the coarse anomaly-aware features, allowing the gate model to leverage detailed, class-specific insights as well as the more general anomaly-aware features. **(c) Final Prediction:** The refined and integrated features are processed through a transformer block followed by a four-layer MLP, similar to the architecture of the experts, to produce the final anomaly scores. This step ensures that the latent space representation is expressive and well-suited for capturing diverse anomaly patterns.

4. Experiments

Datasets. We conduct our experiments on two widely-used Weakly-Supervised Video Anomaly Detection (WSVD) datasets, namely UCF-Crime [30] and XD-Violence [35]. We further experiment on the recent MSAD dataset [51]. Importantly, for all datasets, the training videos are annotated with only video-level labels, without access to frame-level annotations.

Evaluation Metrics. We adhere to the evaluation protocols established in prior works [22, 30, 35, 37]. To ensure comprehensive evaluation, we utilize multiple indicators, such as frame-level Average Precision (AP), Abnormal AP (AP_A) for XD-Violence and Area Under the Curve (AUC), Abnormal AUC (AUC_A) for UCF-Crime dataset. The AP and AUC metrics show the method robustness towards both normal and anomaly videos. However, AP_A and AUC_A allows to exclude normal videos where all snippets are labeled as normal and retain only the abnormal videos

Model	Encoder	UCF-Crime		XD-Violence	
		AUC	AUC _A	AP	AP _A
SoTA Methods With Multi-modal Features					
M.A. [52]	C3D	79.10	62.18	-	-
HL-Net [35]	I3D	82.44	-	-	-
HSN [23]	I3D	85.45	-	-	-
MACIL-SD [41]	I3D+audio	-	-	83.40	-
UR-DMU [49]	I3D+audio	-	-	81.77	-
TPWNG [39]	CLIP	87.79	-	83.68	-
VadCLIP [37]	CLIP	88.02	70.23	84.15	-
SoTA Methods With RGB only Features					
MIL [30]	C3D	75.41	54.25	75.68	78.61
	I3D	77.42	-	-	-
TCN [45]	C3D	78.66	-	-	-
GCN [47]	TSN	82.12	59.02	78.64	-
MIST [7]	I3D	82.30	-	-	-
Dance-SA [28]	TRN	85.00	-	-	-
RTFM [32]	I3D	84.30	62.96	77.81	78.57
CLAV [4]	I3D	86.10	-	-	-
UR-DMU [49]	I3D	86.97	70.81	81.66	83.94
SSRL [16]	I3D	87.43	-	-	-
MSL [17]	V-Swin	85.30	-	78.28	-
WSAL [21]	I3D	85.38	67.38	-	-
ECU [42]	V-Swin	86.22	-	-	-
MGFN [3]	V-Swin	86.67	-	-	-
UMIL [22]	CLIP	86.75	68.68	-	-
TSA [11]	CLIP	87.58	-	82.17	-
GS-MoE (Ours) I3D		91.58	83.86	82.89	85.74
		(+3.56%) (+13.63%)			

Table 1. State-of-the-art comparisons on UCF-Crime and XD-Violence datasets. The best results are written in **bold**.

Model	MSAD			
	AUC	AUC _A	AP	AP _A
RTFM[32]	86.65	-	-	-
MGFN[3]	84.96	-	-	-
TEVAD[2]	86.82	-	-	-
UR-DMU[49]	85.02	-	-	-
UR-DMU *	85.78	67.95	67.35	75.30
GS-MoE	87.72	69.54	68.91	76.68

Table 2. State-of-the-art comparisons on MSAD. * indicates our own implementation. The best results are written in **bold**.

containing both normal and anomalous snippets. This poses a more meaningful challenge to the model’s ability to accurately localize anomalies.

4.1. State-of-the-art Comparison

In our experiments, the proposed GS-MoE model outperforms prior state-of-the-art (SoTA) approaches across multiple metrics, as summarized in Table 1. On the challenging **UCF-Crime dataset**, GS-MoE achieves an AUC of

Baseline	TGS	MoE		AUC(%)		AP _A (%)	
		Experts	Gate	UCF	XD-V	UCF	XD-V
✓	-	-	-	86.97	94.07	45.65	82.91
✓	✓	-	-	87.84	94.13	46.01	83.39
✓	✓	✓	-	89.53	94.29	47.17	84.16
✓	✓	✓	✓	91.58	94.52	51.63	85.74

Table 3. Impact of each component in GS-MoE framework on UCF-Crime and XD-Violence datasets.




Figure 5. Category-wise performance comparison with UR-DMU.

91.58%, surpassing the previous best model, VadCLIP [37], by 3.56%. This significant improvement illustrates the effectiveness of our model in detecting complex video anomalies in real-world datasets. Additionally, when considering the performance on the abnormal videos (AUC_A) only, GS-MoE achieves a score of 83.86%, which constitutes a remarkable 13.63% improvement over the second-best approach, UR-DMU [49], at 70.81%. This result supports one key hypothesis of our work: different types of anomalies require class-specific fine-representations for more effective detection. UR-DMU performance remains limited due to feature-magnitude based optimization which overlooks the subtle cues and enhances the sharp cues. However, the proposed TGS loss promotes both subtle and sharp cues to take part in the separability optimization. Further, the mixture-of-experts architecture is capable of capturing these class-specific representations, leading to substantial performance gains, especially on complex anomalies.

On the **XD-Violence** dataset, GS-MoE achieves an AP score of 82.89%, which is competitive with the best-performing TSA [11] model (82.89%). Moreover, when focusing on anomalous videos only, GS-MoE achieves an AP_A score of 85.74%, outperforming the second-best approach, UR-DMU [49], which achieved an AP_A score of 83.94%. Since the AP metric considers both normal and anomaly videos for evaluation, the performance gets elevated by accurately predicting many normal videos.

As a result, methods performing well on the AP metric may still struggle in anomaly detection. The proposed method outperforms previous SoTA in the AP_A metric, reinforcing its utility in real-world scenarios. On the recently released **MSAD**, GS-MoE surpasses the available baseline models by up to 2.74% on the AUC metric and establishes a new SoTA on the dataset. We report the other metrics as well, in order to provide a baseline for subsequent works.

Expert	Abuse	Arrest	Arson	Assault	Burglary	Explosion	Fighting	RoadAcc.	Robbery	Shooting	Shoplifting	Stealing	Vandalism
Mask	50.02	50.51	49.27	50.72	49.49	49.92	49.95	49.91	50.04	49.20	49.39	50.52	49.87
W/o Mask	86.37	55.48	61.73	63.12	53.65	57.04	65.14	65.22	72.37	60.89	54.73	77.62	57.43

Table 4. Category-wise performance comparison on UCF-Crime dataset between the UR-DMU baseline model and GS-MoE without the expert model for a given class. Masking the relevant experts results in an almost random output from the gate model.




Figure 6. Category-wise t-SNE feature distribution comparison between the baseline, the experts and the gate model.

Category-Wise Performance Analysis: To bring additional analytical insights on the complex anomaly performance, Figure 5 provides an anomaly category-wise performance comparison between GS-MoE and the baseline UR-DMU method on the UCF-Crime dataset. Notably, significant performance boosts are recorded for complex categories like ‘‘Arson’’, ‘‘Assault’’, ‘‘Fighting’’, ‘‘Stealing’’ and ‘‘Burglary’’, up to +24.3%. These performance gains corroborate the benefits of GS-MoE in detecting complex video anomalies. Figure 6 shows the t-SNE plot [34] of the logits obtained at the first and third stages of GS-MoE for the anomalous videos in the test set. The plot in Figure 6a, obtained from the baseline UR-DMU, shows a low degree of separability. The class diversification performed by the experts and shown in Figure 6b demonstrates the capability of GS-MoE to learn enhanced class representations.

4.2. Qualitative Results

As shown in Figure 7, the Gaussian kernels extracted from the abnormal score contain a precise representation of the anomalous events present in videos of the UCF-Crime dataset. The kernel temporal activation (heatmaps) demonstrate the capabilities of this approach. By correctly distinguishing the peaks of the anomalous events and from the spurious peaks, the model is trained to predict high anomaly scores for the associated anomalous snippets. In the ‘‘Assault-010’’ video sample, two peaks are detected in the abnormal score and the TGS finds a small variance for both, leading to a steep normal distribution for each of them. On the other hand, in the ‘‘Arson-011’’ and ‘‘Explosion-033’’ samples, the TGS creates much longer distributions by lead-

ing the model to estimate a large variance and producing a long time-window for the anomaly.

4.3. Ablation Studies

Component Impact: Extensive ablation studies are conducted to evaluate the impact of each contribution to the final performance of GS-MoE, as shown in Table 3. Fine-tuning the baseline UR-DMU model with the TGS loss in Equation 6 leads to a performance increase of +1.77% on the AUC metric of UCF-Crime, while the AP_A of XD-Violence increases by +0.48%. These results show that the new formulation of the WSVAD task is beneficial to existing methods as well. The class-experts outperform the fine-tuned baseline by +0.79% on UCF-Crime. Notably, the AP_A increases on both datasets, leading to +1.16% for UCF-Crime and +0.76% on XD-Violence, further supporting the idea that different classes of anomaly should be treated separately. Adding the gate model to the framework brings the largest performance increment. For UCF-Crime, the AUC increases by +2.05% and the AP_A by +4.46%. On XD-Violence, we observe relatively smaller improvements, increasing AUC by +0.23% and AP_A by +1.68%.

Datasets	W. TA features	W/o. TA features
UCF-Crime (AUC)	91.58	90.98
XD-Violence (AP_A)	85.74	81.45

Table 5. Evaluation of the importance of the task-aware (TA) features for the gate model on the key metrics of the UCF-Crime and XD-Violence datasets.

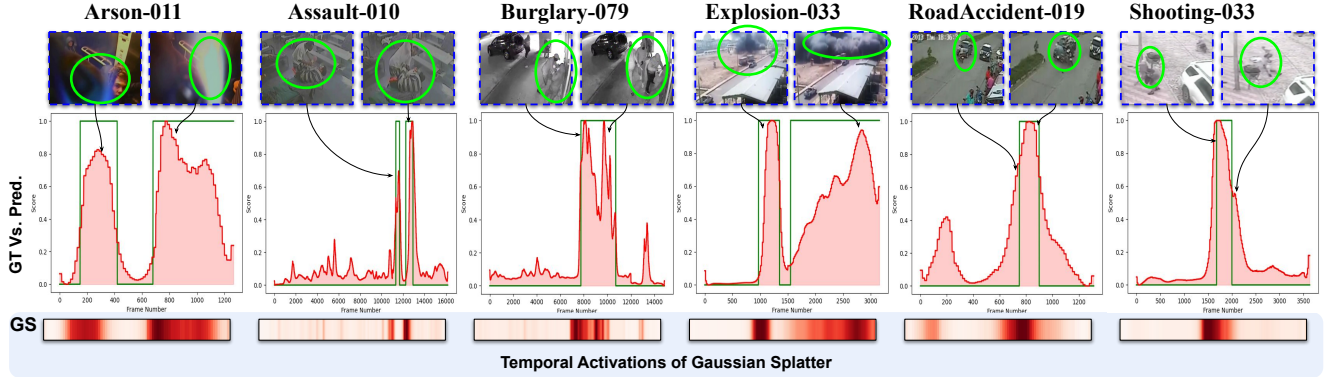


Figure 7. Visualization of sample frames and ground truth (green shed) vs. prediction scores (red shed) for various cases in Row-1 and Row-2. For each plot in Row-2, the X and Y axis denotes the number of frames and corresponding anomaly scores. Row-3 shows the temporal activation (heatmaps) learned by Gaussian splatter (GS).

Task-Aware Features: In order to further analyze this performance increment, the gate model was trained with and without the task-aware features. The results of this experiment are shown in Table 5. The task-aware features seem to have a key role in the performance on the AP_A metric of XD-Violence. In fact, the Gate model trained with the task-aware features outperforms the other configuration by 4.29% on this setting, and by 0.6% on UCF-Crime.

Class-Experts Impact: The relevance of the expert models on the performance of the gate model is measured with the class-wise AUC score obtained by masking the respective class expert on the UCF-Crime dataset. The results of this experiment are shown in Table 4. By masking the experts, the measured AUC hovers around 50% for each class. On the other hand, the gate model predictions are much improved when the relevant expert score is included, leading to a significant performance boost.

Class experts vs cluster experts: In practical applications, anomalies often span multiple classes, making it challenging to train a predefined set of specialized experts. To address this issue, we trained GS-MoE using cluster-based experts rather than class-specific experts. To form the data clusters, we calculated the average task-aware features for each anomalous video in the UCF-Crime training set and applied the K-Means algorithm [20] to group them. Each expert was then trained using videos from a single cluster combined with normal videos, resulting in k specialized expert models. This approach enabled us to evaluate the model’s performance in real-world scenarios where the number of classes is undefined. The results are reported in Table 6. In this setting, GS-MoE is able to outperform current SoTA models by 0.56% clustering the anomalous training videos in 7 clusters and using 7 experts while performing on par with other SoTA models using fewer experts. These results highlight the capabilities of GS-MoE in a real-world use case where the number of anomalous events is not predefined.

Model	AUC
UR-DMU [49]	86.97
TSA [11]	87.58
TPWNG [39]	87.79
VadCLIP [37]	88.02
GS-MoE (5 clusters / 5 experts)	87.35
GS-MoE (6 clusters / 6 experts)	88.03
GS-MoE (7 clusters / 7 experts)	88.58
GS-MoE (class experts)	91.58

Table 6. Comparison between the performance of GS-MoE with varying number of experts.

5. Conclusion

We proposed GS-MoE to provide a novel formulation for weakly-supervised video anomaly detection by leveraging Temporal Gaussian Splatting to overcome the limitations of previous methods. More specifically, we address the over-dependency on the most abnormal snippets for separability optimization. Our framework utilizes a mixture-of-experts architecture that learns category-specific fine-grained representations, building a correlation between coarse abnormal cues and fine-grained cues to learn a more compact representation for each category. Extensive experiments on challenging datasets across various metrics show GS-MoE consistently outperforms SoTA methods with significant performance gains. In the future, we aim to leverage LLMs to provide more explainability to anomaly classes.

Acknowledgements: This work was supported by Toyota Motor Europe (TME) and the French government, through the 3IA Cote d’Azur Investments managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

- [1] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016. 2
- [2] Weiling Chen, Keng Teck Ma, Zi Jian Yew, Minhoe Hur, and David Aik-Aun Khoo. Tevad: Improved video anomaly detection with captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5549–5559, 2023. 6
- [3] Yingxian Chen, Zhengzhe Liu, Baoheng Zhang, Wilton Fok, Xiaojuan Qi, and Yik-Chung Wu. Mgfn: Magnitude-contrastive glance-and-focus network for weakly-supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 387–395, 2023. 6
- [4] MyeongAh Cho, Minjung Kim, Sangwon Hwang, Chaewon Park, Kyungjae Lee, and Sangyoun Lee. Look around for anomalies: Weakly-supervised anomaly detection via context-motion relational learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12137–12146, 2023. 6
- [5] David Eigen, Marc’Aurelio Ranzato, and Ilya Sutskever. Learning factored representations in a deep mixture of experts. *arXiv preprint arXiv:1312.4314*, 2013. 3
- [6] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022. 3
- [7] Jia-Chang Feng, Fa-Ting Hong, and Wei-Shi Zheng. Mist: Multiple instance self-training framework for video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14009–14018, 2021. 6
- [8] Mariana-Iuliana Georgescu, Antonio Barbalau, Radu Tudor Ionescu, Fahad Shahbaz Khan, Marius Popescu, and Mubarak Shah. Anomaly detection in video via self-supervised and multi-task learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12742–12752, 2021. 1
- [9] Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *CoRR*, abs/1606.08415, 2016. 5
- [10] Gagan Jain, Nidhi Hegde, Aditya Kusupati, Arsha Nagrani, Shyamal Buch, Prateek Jain, Anurag Arnab, and Sujoy Paul. Mixture of nested experts: Adaptive processing of visual tokens. *arXiv preprint arXiv:2407.19985*, 2024. 3
- [11] Hyekang Kevin Joo, Khoa Vo, Kashi Yamazaki, and Ngan Le. Clip-tsa: Clip-assisted temporal self-attention for weakly-supervised video anomaly detection. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 3230–3234. IEEE, 2023. 6, 8
- [12] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 3
- [13] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, pages 29–43. Wiley Online Library, 2021. 3
- [14] Dmitry Lepikhin, HyounJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020. 3
- [15] Deqi Li, Shi-Sheng Huang, Zhiyuan Lu, Xinran Duan, and Hua Huang. St-4dgs: Spatial-temporally consistent 4d gaussian splatting for efficient dynamic scene rendering. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery. 3
- [16] Guoqiu Li, Guanxiong Cai, Xingyu Zeng, and Rui Zhao. Scale-aware spatio-temporal relation learning for video anomaly detection. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 6
- [17] Shuo Li, Fang Liu, and Licheng Jiao. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1395–1403, 2022. 6
- [18] Shuo Li, Fang Liu, and Licheng Jiao. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1395–1403, 2022. 2
- [19] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8508–8520, 2024. 3
- [20] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 8
- [21] Hui Lv, Chuanwei Zhou, Zhen Cui, Chunyan Xu, Yong Li, and Jian Yang. Localizing anomalies from weakly-labeled videos. *IEEE transactions on image processing*, 30:4505–4515, 2021. 6
- [22] Hui Lv, Zhongqi Yue, Qianru Sun, Bin Luo, Zhen Cui, and Hanwang Zhang. Unbiased multiple instance learning for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8022–8031, 2023. 5, 6
- [23] Snehashis Majhi, Rui Dai, Quan Kong, Lorenzo Garattoni, Gianpiero Francesca, and Francois Bremond. Human-scene network: A novel baseline with self-rectifying loss for weakly supervised video anomaly detection. *Computer Vision and Image Understanding*, 241:103955, 2024. 6
- [24] Snehashis Majhi, Rui Dai, Quan Kong, Lorenzo Garattoni, Gianpiero Francesca, and François Brémont. Oectst: Outlier-embedded cross temporal scale transformer for weakly-supervised video anomaly detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 8574–8583, 2024. 2
- [25] Snehashis Majhi, Giacomo D’Amicantonio, Antitza Dantcheva, Quan Kong, Lorenzo Garattoni, Gianpiero Francesca, Egor Bondarev, and Francois Bremond. Just dance with pi! a poly-modal inductor for weakly-supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24265–24274, 2025. 2

- [26] Basil Mustafa, Carlos Riquelme, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. Multimodal contrastive learning with limoe: the language-image mixture of experts. *Advances in Neural Information Processing Systems*, 35:9564–9576, 2022. 3
- [27] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. From sparse to soft mixtures of experts, 2024. 3
- [28] Didik Purwanto, Yie-Tarn Chen, and Wen-Hsien Fang. Dance with self-attention: A new look of conditional random fields on anomaly detection in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 173–183, 2021. 6
- [29] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021. 3
- [30] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6479–6488, 2018. 2, 5, 6
- [31] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6479–6488, 2018. 1, 2
- [32] Yu Tian, Guansong Pang, Yuanhong Chen, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4975–4986, 2021. 6
- [33] Yu Tian, Guansong Pang, Yuanhong Chen, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4975–4986, 2021. 2
- [34] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. 7
- [35] Peng Wu, Jing Liu, Yujia Shi, Yujia Sun, Fangtao Shao, Zhaoyang Wu, and Zhiwei Yang. Not only look, but also listen: Learning multimodal violence detection under weak supervision. In *European Conference on Computer Vision*, pages 322–339. Springer, 2020. 2, 5, 6
- [36] Peng Wu, Xiaotao Liu, and Jing Liu. Weakly supervised audio-visual violence detection. *IEEE Transactions on Multimedia*, 25:1674–1685, 2022. 1
- [37] Peng Wu, Xuerong Zhou, Guansong Pang, Lingru Zhou, Qingsen Yan, Peng Wang, and Yanning Zhang. Vadclip: Adapting vision-language models for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6074–6082, 2024. 5, 6, 8
- [38] Qingsen Yan, Tao Hu, Yuan Sun, Hao Tang, Yu Zhu, Wei Dong, Luc Van Gool, and Yanning Zhang. Towards high-quality hdr dehazing with conditional diffusion models. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023. 1
- [39] Zhiwei Yang, Jing Liu, and Peng Wu. Text prompt with normality guidance for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18899–18908, 2024. 6, 8
- [40] Guang Yu, Siqi Wang, Zhiping Cai, En Zhu, Chuanfu Xu, Jianping Yin, and Marius Kloft. Cloze test helps: Effective video anomaly detection via learning to complete video events. In *Proceedings of the 28th ACM international conference on multimedia*, pages 583–591, 2020. 1
- [41] Jiashuo Yu, Jinyu Liu, Ying Cheng, Rui Feng, and Yuejie Zhang. Modality-aware contrastive instance learning with self-distillation for weakly-supervised audio-visual violence detection. In *Proceedings of the 30th ACM international conference on multimedia*, pages 6278–6287, 2022. 6
- [42] Chen Zhang, Guorong Li, Yuankai Qi, Shuhui Wang, Laiyun Qing, Qingming Huang, and Ming-Hsuan Yang. Exploiting completeness and uncertainty of pseudo labels for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16271–16280, 2023. 6
- [43] Chen Zhang, Guorong Li, Yuankai Qi, Shuhui Wang, Laiyun Qing, Qingming Huang, and Ming-Hsuan Yang. Exploiting completeness and uncertainty of pseudo labels for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16271–16280, 2023. 2
- [44] Huaxin Zhang, Xiang Wang, Xiaohao Xu, Xiaonan Huang, Chuchu Han, Yuehuan Wang, Changxin Gao, Shanjun Zhang, and Nong Sang. Glancevad: Exploring glance supervision for label-efficient video anomaly detection. *arXiv preprint arXiv:2403.06154*, 2024. 2, 4
- [45] Jiangong Zhang, Laiyun Qing, and Jun Miao. Temporal convolutional network with complementary inner bag loss for weakly supervised anomaly detection. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 4030–4034. IEEE, 2019. 6
- [46] Shuai Zhang, Huangxuan Zhao, Zhenghong Zhou, Guanjuan Wu, Chuansheng Zheng, Xinggang Wang, and Wenyu Liu. Togs: Gaussian splatting with temporal opacity offset for real-time 4d dsa rendering, 2024. 3
- [47] Jia-Xing Zhong, Nannan Li, Weijie Kong, Shan Liu, Thomas H. Li, and Ge Li. Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6
- [48] Jia-Xing Zhong, Nannan Li, Weijie Kong, Shan Liu, Thomas H Li, and Ge Li. Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1237–1246, 2019. 2
- [49] Hang Zhou, Junqing Yu, and Wei Yang. Dual memory units with uncertainty regulation for weakly supervised video anomaly detection. *arXiv preprint arXiv:2302.05160*, 2023. 5, 6, 8
- [50] Hang Zhou, Junqing Yu, and Wei Yang. Dual memory units with uncertainty regulation for weakly supervised video

anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3769–3777, 2023. [5](#)

- [51] Liyun Zhu, Lei Wang, Arjun Raj, Tom Gedeon, and Chen Chen. Advancing video anomaly detection: A concise review and a new dataset. In *The Thirty-eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. [2](#), [5](#)
- [52] Yi Zhu and Shawn Newsam. Motion-aware feature for improved video anomaly detection. *arXiv preprint arXiv:1907.10211*, 2019. [6](#)