

RefineVAD: Semantic-Guided Feature Recalibration for Weakly Supervised Video Anomaly Detection

Junhee Lee^{*1}, ChaeBeen Bang^{*1}, MyoungChul Kim^{*2}, MyeongAh Cho^{†2}

¹Department of Artificial Intelligence, Kyung Hee University

²Department of Software Convergence, Kyung Hee University
{jhlee39, coqls1229, 2018102095, maycho}@khu.ac.kr

Abstract

Weakly-Supervised Video Anomaly Detection aims to identify anomalous events using only video-level labels, balancing annotation efficiency with practical applicability. However, existing methods often oversimplify the anomaly space by treating all abnormal events as a single category, overlooking the diverse semantic and temporal characteristics intrinsic to real-world anomalies. Inspired by how humans perceive anomalies, by jointly interpreting temporal motion patterns and semantic structures underlying different anomaly types, we propose RefineVAD, a novel framework that mimics this dual-process reasoning. Our framework integrates two core modules. The first, Motion-aware Temporal Attention and Recalibration (MoTAR), estimates motion salience and dynamically adjusts temporal focus via shift-based attention and global Transformer-based modeling. The second, Category-Oriented Refinement (CORE), injects soft anomaly category priors into the representation space by aligning segment-level features with learnable category prototypes through cross-attention. By jointly leveraging temporal dynamics and semantic structure, explicitly models both “how” motion evolves and “what” semantic category it resembles. Extensive experiments on WVAD benchmark validate the effectiveness of RefineVAD and highlight the importance of integrating semantic context to guide feature refinement toward anomaly-relevant patterns.

Code —

<https://github.com/VisualScienceLab-KHU/RefineVAD>

Introduction

Video Anomaly Detection (VAD) aims to identify unexpected or abnormal events within video streams and has become increasingly important in applications such as surveillance, public safety, and industrial monitoring. While fully supervised methods achieve precise anomaly localization, they require costly frame-level annotations, limiting their scalability in real-world scenarios. Unsupervised approaches, which learn solely from normal patterns, reduce labeling costs but often suffer from poor general-

ization and high false-positive rates. To balance annotation efficiency with detection performance, Weakly Supervised Video Anomaly Detection (WVAD) has emerged as a compelling alternative. WVAD methods rely exclusively on video-level labels—indicating whether a video contains any anomaly—and typically adopt a Multiple Instance Learning (MIL) paradigm, wherein each video is treated as a bag of segments, and at least one segment in an anomalous video is assumed to contain abnormal content.

However, recognizing abnormal events in videos requires more than coarse video-level classification. Human perception leverages two complementary dimensions when detecting anomalies: **(1) contextual motion dynamics that evolve over time**—such as abrupt or irregular movements—and **(2) prior knowledge of what type of anomaly is occurring**, allowing us to interpret visual patterns semantically. This motivates a hierarchical reasoning framework that first attends to how motion unfolds and then reasons about what kind of abnormal event it resembles. Such reasoning enables more robust and interpretable anomaly detection by aligning perceptual cues with semantic understanding.

Despite recent advancements, existing WVAD methods exhibit two critical limitations that impede this hierarchical reasoning. First, **temporal modeling is often shallow or inflexible**, relying on fixed pooling operations or simple aggregation schemes that fail to adapt to the diverse motion characteristics of real-world anomalies. Since many anomalies are defined by dynamic, non-uniform, or context-specific motion patterns, this temporal rigidity significantly hampers localization accuracy. Second, and more fundamentally, most WVAD frameworks **treat all abnormal events as a single generic class**, ignoring the semantic diversity across anomaly types. While the original VAD task framed anomaly detection as a binary classification problem—determining whether each snippet is normal or abnormal, in practice, anomalies exhibit highly distinctive visual and motion patterns. For instance, fighting involves abrupt, bidirectional motion, while explosions are characterized by sudden flashes and spatial bursts. Ignoring such distinctions limits the model’s ability to capture discriminative features, thereby reducing both interpretability and localization performance. Although some recent works have attempted

^{*}These authors contributed equally.

[†]Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

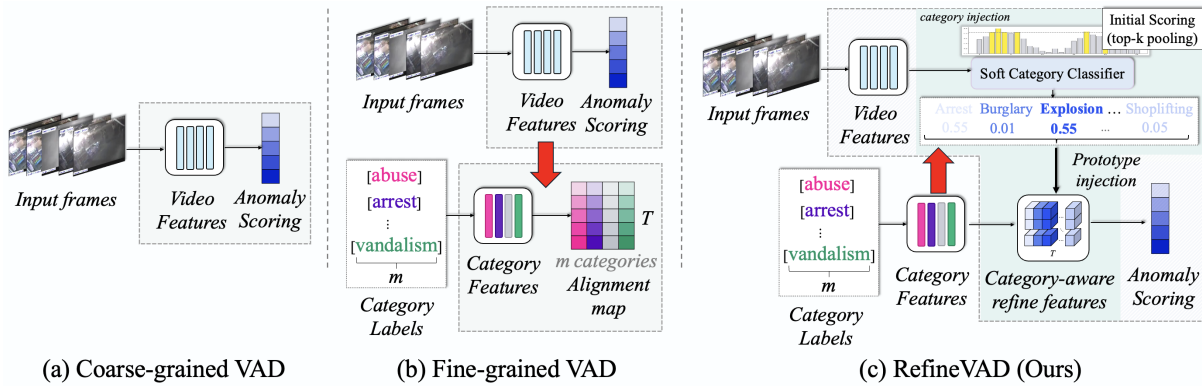


Figure 1: Comparison of WVAD paradigms. (a) Coarse-grained model predicts anomaly scores per snippet. (b) Fine-grained method introduces auxiliary video-level category classification, but these labels are not utilized during anomaly scoring. (c) RefineVAD (Ours): category-aware soft classification guides feature enhancement with learnable prototypes.

to incorporate anomaly categories, these are typically used only for auxiliary video-level classification and not integrated into the anomaly scoring or representation learning processes. Consequently, semantic distinctions remain unexploited during temporal localization, and the model fails to learn category-aware features crucial for precise detection.

To address these challenges, we propose **RefineVAD**, a novel framework that introduces **motion-aware temporal recalibration** and **category-conditioned refinement** into the anomaly detection process. RefineVAD integrates two key components—**MoTAR** and **CORE**—to explicitly models both “*how*” motion evolves and “*what*” semantic category it resembles, enabling structured anomaly reasoning under weak supervision. The Motion-aware Temporal Attention and Recalibration (MoTAR) module captures motion-intensity-aware relationships between segments by dynamically shifting feature channels in proportion to estimated motion saliency. It further employs a lightweight Transformer-based global attention mechanism to model long-range dependencies, allowing the model to focus on critical temporal regions without relying on rigid windowing or fixed aggregation.

The Category-Oriented Refinement (CORE) module incorporates semantic priors by selectively injecting soft category prototypes into the representation space. It first predicts a video-level anomaly category distribution via soft classification, which is then used to compute a weighted combination of learnable category embeddings. Each embedding serves as a prototype encoding the characteristic patterns of a specific anomaly type. These category-aware combined embeddings are injected into the snippet-level features via cross-attention, guiding the model to focus on temporally and semantically meaningful regions during scoring. By combining these two components, RefineVAD effectively disentangles and recomposes multi-modal evidence to enhance anomaly localization. It captures the temporal evolution of motion and aligns it with category-relevant semantics, thereby improving both detection precision and interpretability.

Our contributions are summarized as follows:

- We propose RefineVAD, a novel WVAD framework that jointly models temporal and semantic contexts for anomaly localization under weak supervision.
- We design MoTAR, a motion-aware temporal attention module that adaptively recalibrates temporal feature flows based on motion saliency.
- We introduce CORE, a category-oriented refinement module that injects soft category priors into snippet-level representations via attention-guided semantic alignment.
- We demonstrate that RefineVAD significantly outperforms previous state-of-the-art methods on multiple benchmark datasets through comprehensive ablation studies.

Related Works

Weakly Supervised Video Anomaly Detection. VAD aims to identify events that deviate from typical or expected normal patterns in video streams. Traditional fully supervised learning methods detect anomalies using frame-level annotations, which are expensive and difficult to apply at scale (Zhu, Chen, and Sultani 2021; Majhi et al. 2021; Majhi, Das, and Br mond 2021). On the other hand, Unsupervised VAD (UVAD) methods trained only on normal data struggle to robustly characterize normality and lead to reconstruction and prediction errors (Hasan et al. 2016; Liu et al. 2018; Wu, Liu, and Shen 2019; Cai et al. 2021). To bridge this gap, WVAD has emerged and significantly reduce the annotation cost since it only requires video-level labels to indicate whether anomalous signs are present in the video (Wu et al. 2020; Zaheer et al. 2020; Tian et al. 2021; Chen et al. 2023; Cho et al. 2023; Wan et al. 2020). Most WVAD methods are based on the Multiple-Instance Learning (MIL) framework to utilize anomalous video labels without frame annotations. A foundational approach applies a ranking loss to enforce higher scores for anomalous snippets compared to normal ones (Sultani, Chen, and Shah 2018). Subsequent works build upon this paradigm by introducing techniques such as two-stage pseudo-labeling or surrogate optimization (Feng, Hong, and Zheng 2021). These MIL methods

typically follow a pipeline that extracts pre-trained features (e.g., C3D (Tran et al. 2015)/I3D (Carreira and Zisserman 2017)) and then trains a snippet classifier.

Category and Semantic-Guided VAD. More recent approaches shift towards leveraging semantic priors and category awareness to enhance weakly supervised video anomaly detection. VadCLIP (Wu et al. 2024) applies frozen CLIP (Radford et al. 2021) to WSVAD via a dual-branch design: one branch handles binary classification, while the other aligns frame features with text label embeddings. This enables both coarse- and fine-grained anomaly detection. TPWNG (Yang, Liu, and Wu 2024) leverages CLIP with learnable prompts and event categories as text inputs, generating snippet-level pseudo-labels and training a video classifier in a self-learning loop with temporal context modeling. PEMIL (Chen et al. 2024) formulates WSVAD as prompt learning, injecting semantic priors via anomaly and normal context prompts constructed from class label words. While these methods effectively exploit CLIP’s vision-language capability, they rely on discrete category labels or hand-crafted prompts, which may limit adaptability to ambiguous or overlapping anomalies. Moreover, they often overlook inter-category dependencies or the channel-wise nature of visual representations.

Proposed Methods

To address the limitations of existing WVAD methods—namely, their static temporal modeling and disregard for semantic category structures—we propose RefineVAD, a modular framework that introduces contextual awareness at two complementary levels: **(1) motion-aware temporal recalibration**, and **(2) category-conditioned semantic refinement**. Rather than relying solely on visual embeddings or treating all anomalies as a single generic class, RefineVAD leverages both structured visual and textual information, dynamically models motion salience over time, and integrates category-aware priors to enhance anomaly representation and scoring.

As illustrated in Figure 2, RefineVAD follows the MIL setup by dividing each input video—annotated with a video-level anomaly label—into T fixed-length segments. Each segment is independently encoded using pretrained visual and textual encoders to extract modality-specific features. These features are then concatenated to form a joint multi-modal representation for each segment, capturing both appearance and semantic context. This joint representation is first processed by the MoTAR (Motion-aware Temporal Attention and Recalibration) module. MoTAR dynamically adjusts feature channel shift ratios based on segment motion intensity to achieve adaptive temporal recalibration based on the degree of motion irregularity. In parallel, it captures long-range temporal dependencies using a lightweight Transformer-based attention mechanism (Vaswani et al. 2017), allowing flexible and adaptive modeling of without rigid pooling or fixed windowing.

Next, the motion-contextual features are refined by the CORE (Category-Oriented Refinement) module, which incorporates semantic priors into the representation space.

CORE first performs a soft classification at the video-level to estimate the probability distribution over anomaly categories. These probabilities are used to compute a weighted combination of learnable category prototypes, each encoding representative patterns of a specific anomaly type. The resulting soft category embedding is then injected into all segment-level features via cross-attention, enabling the model to emphasize segments that are not only temporally salient but also semantically consistent with the predicted anomaly type.

By jointly modeling motion intensity and category relevance, RefineVAD enables the disentanglement and recombination of multi-modal cues to enhance anomaly localization. The final category-aware, temporally contextualized features are passed through a lightweight classifier to compute segment-level anomaly scores, which are subsequently aggregated to produce the video-level anomaly prediction.

Motion-aware Temporal Attention Recalibration

Anomalous events in videos often exhibit irregular motion dynamics—some are abrupt, while others remain deceptively static. To accommodate this variability, we propose MoTAR, a lightweight and effective temporal feature modeling module.

MoTAR is built upon the Temporal Shift Module (TSM) (Lin, Gan, and Han 2019), which introduces temporal interactions by shifting a portion of feature channels forward and backward in time. While TSM is widely used in video understanding due to its efficiency and parameter free design, it applies a fixed shift ratio across all frames, limiting its adaptability to frames with varying motion intensities.

To address this, MoTAR first estimates the motion intensity per-frame using feature differences. As shown in Figure 3, given an input from the vision and text encoder $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{T \times D}$, where T is the number of sequence and D is the feature dimension, we compute $\Delta_t = \mathbf{x}_t - \mathbf{x}_{t-1}$ and obtain $\mathbf{v}_t = \text{Var}(\Delta_t)$ where \mathbf{v}_t reflects local motion intensity. A higher variance indicates more significant motion, suggesting the need for broader temporal context aggregation. This variance-based measure is representation-agnostic, noise-robust, and enables adaptive temporal shifting with minimal computational overhead.

The variance vector is processed by a lightweight MLP-based shift predictor to obtain a shift ratio $r_t \in [0, 1]$:

$$r_t = \sigma(W_3 \cdot \phi(W_2 \cdot \phi(W_1 v_t))) \quad (1)$$

where ϕ denotes the GELU activation and σ is the sigmoid function. The number of channels to shift across features is given by:

$$s_t = \left\lceil r_t \cdot \frac{D}{K} \right\rceil \quad (2)$$

where K is a predefined folding factor. The temporally shifted output of the t -th segment is then constructed as:

$$\mathbf{y}_t = \left[\mathbf{x}_{t-1}^{(1:s_t)} \quad \mathbf{x}_{t+1}^{(s_t:2s_t)} \quad \mathbf{x}_t^{(2s_t:D)} \right] \quad (3)$$

This dynamic construction allows the model to allocate more shifted channels to high-motion frames, while preserving static content in low-variance frames.

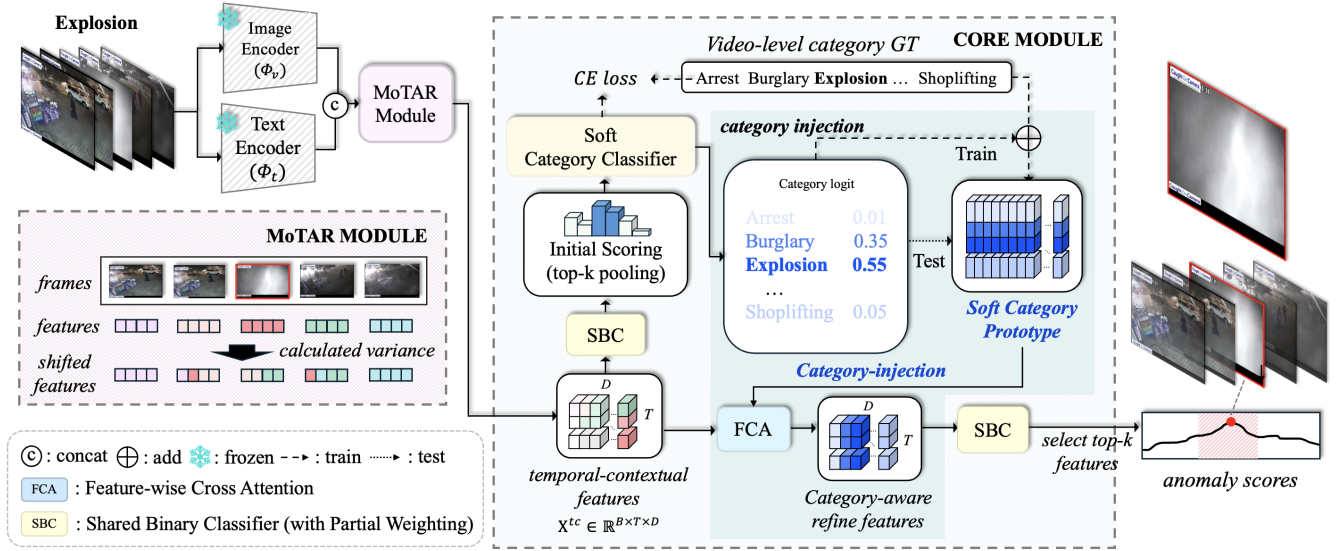


Figure 2: The overall architecture of proposed RefineVAD. Visual and textual features are extracted from each segment and fused before being processed by **MoTAR**, which adaptively recalibrates temporal features based on motion saliency. The resulting features are then refined by **CORE**, which injects soft category priors via cross-attention with learnable prototypes, enabling category-aware anomaly localization.

We further apply a lightweight Transformer encoder over the shifted sequence to incorporate long-range temporal dependencies, producing temporal-contextual features.:

$$\mathbf{X}^{tc} = \text{Transformer}([\mathbf{y}_1, \dots, \mathbf{y}_T]) \quad (4)$$

In summary, MoTAR extracts temporal-contextual feature alignments through adaptive shifting and capturing global temporal context. Its parameter-free shifting ensures computational efficiency, making it suitable for real-time or large-scale video anomaly detection applications.

Detailed numerical results and a comprehensive analysis of the computational efficiency are provided in Appendix F.

Category-Oriented Refinement

To enrich temporal contextual feature with semantic information, we introduce CORE module, a key component of the RefineVAD framework. CORE systematically injects category-specific semantics into segment-level features, thereby guiding the model to better focus on distinctive aspects of each abnormal event through semantic reasoning. The module consists of two main stages. **(1) Soft Category Classification:** We first distinguish the anomaly category associated with a given video sample by applying a soft classification mechanism, which estimates the probability of the sample belonging to each predefined anomaly class. **(2) Category Prototype Injection:** These probabilities are then used as weights to softly combine category-specific prototype embeddings, resulting in a unified semantic vector that captures both dominant and complementary category traits. This vector is further injected into individual temporal features via a cross-attention mechanism, producing category-aware refined features that dynamically align each segment with the most relevant category semantics. By

enriching the representation with this semantically guided context, the model achieves improved discriminability in the final stage of anomaly scoring.

Soft Category Classification. Let $\mathbf{X}^{tc} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ be the temporal-contextual feature sequence generated by the MoTAR module, where each $\mathbf{x}_t^{tc} \in \mathbb{R}^D$ corresponds to the contextualized feature of the t -th segment. Each segment-level feature \mathbf{x}_t^{tc} is processed by relatively simple fully connected (FC) layers consisting of dropout and batch normalization layers for scoring. These roughly estimated anomaly scores are converted into normalized weights which are used to aggregate the segment features along the temporal dimension. Since anomalous regions are relatively sparse within a video, it is more advantageous to focus on high-confidence segments rather than the entire video when distinguishing between categories.

The pooled video feature is passed through a soft category classifier, yielding logits that we reshape into a matrix $\mathbf{z} \in \mathbb{R}^{C \times 2}$, where C is the number of anomaly categories present in the dataset. Each row of this matrix corresponds to the ‘normal’ and ‘abnormal’ scores for a given category. For each category c , we denote the corresponding logit scores for the “normal” and “abnormal” classes as:

$$\mathbf{z}_c = [z_c^n \ z_c^a] \in \mathbb{R}^2. \quad (5)$$

For each category c , we apply a normal vs. anomaly softmax per category to obtain the anomaly probability p_c^a . This type of separation into anomaly and normal cases is motivated by the observation that normal samples encompass a highly diverse range of conditions, making it difficult for the model to learn a coherent representation if they are treated as a single additional anomaly category. Unlike

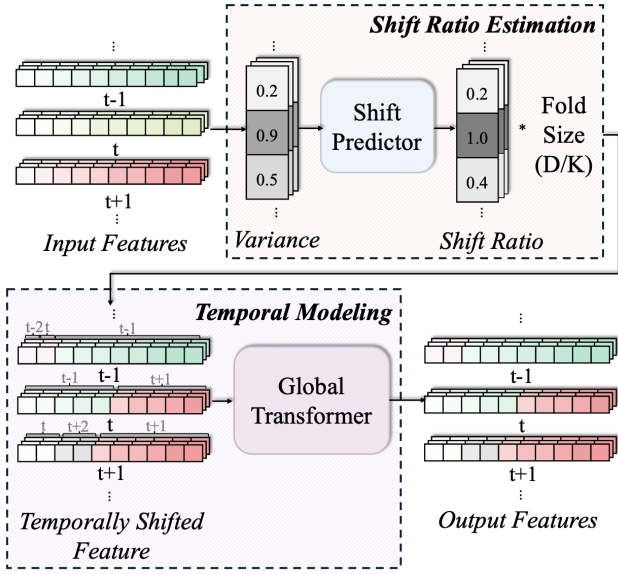


Figure 3: Overview of MoTAR. Motion variance guides adaptive channel shifting, enabling dynamic temporal feature aggregation. A Global Transformer captures long-range dependencies, improving sensitivity to diverse motions.

earlier works that risked treating "normal" as just another anomaly type, we define the normal state by the absence of strong anomaly-specific characteristics, thereby avoiding distortion of learned representations. This is precisely why we designed the model to output the logits in this manner.

The category-wise anomaly probabilities serve as modulation signals that guide how the input representation is refined, enabling the model to softly encode the relevance of different anomaly categories based on their activation levels. Instead of selecting only the most likely category, as is typical in standard classification, the model derives a set of importance weights from the full anomaly probability distribution, reflecting its confidence across all categories. Formally, we define the weights of each category as:

$$w_c = \begin{cases} \frac{1}{C}, & \text{if } S = \emptyset, \\ \frac{\exp(z_c^a)}{\sum_{k=1}^C \exp(z_k^a)}, & \text{if } S \neq \emptyset. \end{cases} \quad (6)$$

where $S = \{c \mid p_c^a \geq \tau\}$ is the set of categories whose anomaly probability exceeds a threshold τ . This allows the model to naturally separate normal inputs from anomaly categories without introducing a dedicated "normal" prototype, enabling more faithful and effective learning that better reflects how anomalies are identified in real-world settings.

Category Prototype Injection. These weights are then used to softly combine the corresponding category embeddings, which act as prototypes representing each category, resulting in a feature representation that captures both shared and distinctive anomaly characteristics. Prototypes are learnable

and arranged row-wise in the matrix:

$$\mathbf{E} = \begin{bmatrix} \mathbf{e}_1^T \\ \vdots \\ \mathbf{e}_C^T \end{bmatrix} \in \mathbb{R}^{C \times d_{\text{emb}}}. \quad (7)$$

To encode contextual category semantics, we compute a soft category embedding vector \mathbf{v} as a weighted sum over learnable category embeddings:

$$\mathbf{v} = \sum_{c=1}^C w_c \mathbf{e}_c, \quad (8)$$

where w_c represents the model's estimated anomaly probability for the c -th category, and \mathbf{e}_c is its corresponding embedding. This formulation allows the model to softly blend semantic cues across categories, capturing both confident and ambiguous class-related anomaly traits. During training, we enhance the soft category embedding \mathbf{v} by adding the ground-truth class embedding, resulting in $\mathbf{v}_{\text{train}} = \mathbf{v} + \mathbf{e}_y$. This encourages class-aware specialization in the learned representations.

The embedding \mathbf{v} is incorporated into the temporal-contextual features via a multi-head cross-attention mechanism. Specifically, \mathbf{v} serves as the global query, while the temporal-contextual features act as both keys and values. This operation refines each segment feature \mathbf{x}_t^{lc} into a category-aware refine feature \mathbf{x}_t^{ca} :

$$\mathbf{x}_t^{\text{ca}} = \text{CrossAttn}(\mathbf{v}, \mathbf{x}_t^{\text{lc}}, \mathbf{x}_t^{\text{lc}}), \quad (9)$$

effectively aligning low-level feature patterns with high-level category cues.

Finally, anomaly scoring is performed on the aligned features described above. This scoring part shares parameters with the earlier rough scoring function that was used for pooling features to be passed into the category classifier, enabling some degree of mutual learning during training.

We combine a Top- k MIL ranking loss (Sultani, Chen, and Shah 2018) \mathcal{L}_{MIL} to promote truly anomalous segments over normal ones, and an improved GMM-based smoothing loss (Wang and Chen 2025) \mathcal{L}_{GMM} . The original method fitted a Gaussian mixture by computing per-video mean and variance across all segments; we enhance it by injecting our weighted category embedding into the mixture to more clearly reflect category-specific traits. We also use a category classification loss \mathcal{L}_{cat} , where Eq. 5 outputs a two-logit vector z_c for each category c . For a video with ground-truth category c^* , we assign the target $t_c = [0, 1]$ if $c = c^*$ and $t_c = [1, 0]$ otherwise, and define $\mathcal{L}_{\text{cat}} = \sum_{c=1}^C \text{BCE}_2(\sigma(z_c), t_c; w_c)$, with optional class weights w_c to mitigate imbalance. This encourages the semantic embeddings to remain category-discriminative while retaining anomaly relevance.

The overall training objective is $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MIL}} + \lambda_1 \mathcal{L}_{\text{GMM}} + \lambda_2 \mathcal{L}_{\text{cat}}$.

Experiments

Experimental Setup

We evaluate our method on two widely used benchmark datasets for WSVAD: UCF-Crime(Sultani, Chen, and Shah

Approach	Method	Source	UCF-AUC(%)	XD-AP(%)
semi-sup.	SVM baseline	-	50.10	50.80
	OCSVM (1999)	NeurIPS	63.20	28.63
	Hansan et al. (2016)	CVPR	51.20	31.25
weakly-sup.	Ju et al. (2022)	ECCV	84.72	76.57
	Sultani et al. (2018)	CVPR	84.14	75.18
	Wu et al. (2020)	ECCV	84.57	80.00
	AVVD (2022)	TMM	82.45	78.10
	RTFM (2021)	ICCV	85.66	78.27
	DMU (2023)	AAAI	86.75	81.66
	UMIL (2023)	CVPR	86.75	-
	CLIP-TSA (2023)	ICIP	87.58	82.17
	HSN (2024)	CVIU	85.45	-
	IFS-VAD (2024)	TCSVT	85.47	83.14
	VadCLIP (2024)	AAAI	88.02	84.57
	PEMIL (2024)	CVPR	86.83	<u>88.21</u>
	ReFLIP (2024)	TCSVT	88.57	85.81
	CMHKF (2025a)	ACL	-	86.57
	Ex-VAD (2025)	ICML	88.29	86.52
	π -VAD (2025)	CVPR	90.33	85.37
	RefineVAD (Ours)			<u>88.92</u>

Table 1: Comparison of state-of-the-art methods on UCF-Crime and XD-Violence datasets. Best result is **bolded** and second best result is underlined.

2018) and XD-Violence(Wu et al. 2020). **UCF-Crime**. For evaluation, we follow standard protocols. On UCF-Crime, we use the Area Under the Curve (AUC) and AUC for anomaly videos (AnoAUC) at the frame level. On XD-Violence, we use frame-level Average Precision (AP) and report mean Average Precision (mAP) across IoU thresholds from 0.1 to 0.5, along with the average score (AVG) for segment-level performance.

Implementation Details

We divide each video into 32 snippets and use pre-trained CLIP(ViT-L/14) (Joo et al. 2023) as both the image and text encoder. For textual representation, InternVideo2.5 (Wang et al. 2025b) is utilized. The encoders are kept frozen during training. For the XD-Violence dataset, we train the model for a maximum of 30 epochs using the AdamW optimizer (LR: 3e-3, Batch Size: 64), where each mini-batch consisted of 32 normal and 32 abnormal samples. The loss weights were set to $\lambda_1 = 0.1$ and $\lambda_2 = 0.2$. The implementation used PyTorch on a single NVIDIA A5000 GPU.

Comparison on WVAD Benchmarks

We compare our performance with current state-of-the-art (SOTA) methods on UCF-Crime and XD-Violence datasets. According to Table 1, our model achieves 88.92% AUC on UCF-Crime and surpasses all previous weakly supervised learning models on XD-Violence with 88.66% AP. Regarding the segment-level evaluation in Table 2, our method achieves competitive mAP across IoU thresholds on UCF-Crime, ranking among the top existing approaches. Unlike

Method	mAP@IOU (%)					
	0.1	0.2	0.3	0.4	0.5	AVG
Random Baseline	0.21	0.14	0.04	0.02	0.01	0.08
Sultani et al. (2018)	5.73	4.41	2.69	1.93	1.44	3.24
AVVD (2022)	10.27	7.01	6.25	3.42	3.29	6.05
VadCLIP (2024)	11.72	7.83	6.40	4.53	2.93	6.68
ITC (2024)	13.54	9.24	7.45	5.46	3.79	7.90
ReFLIP (2024)	14.23	10.34	9.32	7.54	6.81	9.62
Ex-VAD (2025)	16.51	12.35	9.41	7.82	4.65	10.15
RefineVAD (Ours)	20.90	13.17	8.14	4.41	3.03	<u>9.93</u>

Table 2: Fine-grained comparisons on UCF-Crime.

MoTAR	CORE		AUC(%)
	Category-Injection	Soft-Classification	
			84.60
✓			85.43
	✓		87.28
✓	✓		87.85
✓	✓	✓	88.89

Table 3: Ablation studies of our module on the UCF-Crime.

prior works that explicitly optimize snippet-level classification, our framework predicts only a single video-level category, which is then used to guide temporal localization. While this design improves category-level agreement, it places our model at a structural disadvantage under higher IoU thresholds, where the metric emphasizes frame-accurate boundary localization rather than alignment with video-level predictions. The slight performance gap at high IoU thus arises from this evaluation-objective mismatch rather than from the capability of our approach. On XD-Violence, our method remains competitive and surpasses previous weakly supervised methods by 0.19%, validating the benefit of adaptive category-aware feature refinement via MoTAR and CORE.

Ablation Studies

To assess the contribution of each component, we conduct an ablation study on UCF-Crime (Table 3). Starting from the Base model (MLP with MIL), we sequentially add MoTAR, Category-Injection, and Soft Classification.

Introducing MoTAR provides a modest AUC gain (from 84.60% to 85.43%), showing the utility of motion-aware temporal adjustment. Notably, Category-Injection brings a much larger improvement, raising the AUC to 87.28% and delivering the most significant single-step boost. Finally, adding Soft Classification—where embeddings are formed not from a single top category but from a weighted sum over category probabilities—further refines the representation. Combining all components yields the best performance at 88.89%, confirming that each module is effective and that their integration provides complementary benefits for weakly supervised video anomaly detection.

Method	Category GT	AP (%)
Full Training	✓	88.66
Freezing CORE Module (Category Classifier & Embeddings; Trained only on UCF)	✗	87.52
Direct Cross-Domain Transfer (Train: UCF-Crime, Test: XD-Violence, No Fine-Tuning)	✗	77.56

Table 4: Semantic transfer of UCF-Crime category classifier and embedding on XD-Violence.

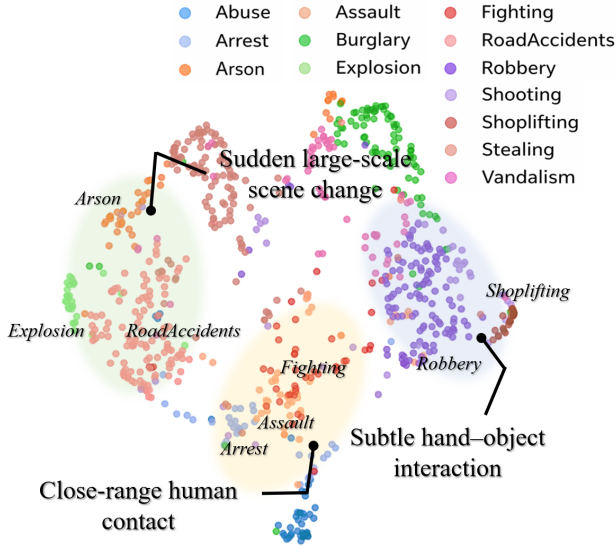


Figure 4: t-SNE visualization of logit features, where semantically similar categories form meaningful clusters.

Cross-Dataset Semantic Transfer Evaluation

To evaluate the transferability of the learned category semantics, we perform cross-dataset experiments on XD-Violence, as summarized in Table 4. Full training on XD-Violence serves as the upper-bound reference, achieving 88.66% AP. When freezing the CORE module (category classifier and semantic embeddings learned solely from UCF-Crime) and training the remaining modules, the model still attains 87.52% AP, indicating the UCF-Crime semantic space effectively encodes transferable category-level structure despite dataset differences. Finally, a strict zero-shot test—evaluating a UCF-Crime trained model directly on XD-Violence—achieves a robust 77.56% AP, demonstrating that the prototype-based semantics retain discriminative power across domains with mismatched category taxonomies and distinct visual distributions.

Collectively, these results show that our semantic representations extend beyond dataset-specific labels. Unlike prior category-driven approaches that required collecting new annotations for each target dataset and often degraded under mismatch, our compositional semantic space is label-efficient, scalable, and reusable without additional supervision, enabling practical cross-dataset deployment in weakly supervised video anomaly detection.




Figure 5: Example frame (top) and corresponding predicted scores with ground truth (bottom). The blue curve represents the predicted scores, and the red shaded regions indicate the ground-truth anomalous intervals.

Category Logits Visualization

To examine whether the predicted logit vectors are effectively separated according to semantic categories, we visualize the outputs of our soft category classifier using t-SNE. The logit vector corresponds to the category-wise activation output from the classifier, and similar vectors tend to be positioned close together in the embedding space.

As shown in the Figure 4, samples with similar semantic characteristics form distinct clusters. For instance, Arson, Explosion, and RoadAccidents are grouped together, where the classifier tends to focus on sudden background changes or scene-level phenomena. In contrast, Arrest, Assault, and Fighting are clustered with a focus on the interactions among multiple people. On the other hand, Shoplifting and Robbery form a third group, where the model mainly attends to the actions of a single individual.

This observation suggests that different anomaly categories require attention to different aspects of the scene (e.g., background changes, group activity, or individual behavior), with similar categories exhibiting similar logit activation distributions. These results validate the effectiveness of category-aware anomaly detection, where adapting the decision-making process based on the anomaly type leads to more accurate detection.

Conclusion

In this work, we proposed RefineVAD, a weakly-supervised framework that integrates motion-aware temporal recalibration and category-conditioned feature refinement to enhance anomaly localization. By jointly modeling motion patterns and semantic categories, RefineVAD captures both generic and category-specific anomaly cues, enabling more precise detection. Our soft classification mechanism guides the model toward meaningful temporal and semantic evidence, leading to improved performance across benchmarks. In addition, the resulting category-aware feature space provides improved interpretability and demonstrates strong transferability across datasets. This approach also lays a foundation for future extensions toward fine-grained anomaly taxonomy, adaptive prototype refinement, and label-efficient cross-domain generalization.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT)(RS-2024-00456589) and Korea Planning & Evaluation Institute of Industrial Technology(KEIT) grant funded by the Korea government(MOTIE) (No. RS-2024-00442120, Development of AI technology capable of robustly recognizing abnormal and dangerous situations and behaviors during night and bad weather conditions)

References

- Cai, R.; Zhang, H.; Liu, W.; Gao, S.; and Hao, Z. 2021. Appearance-motion memory consistency network for video anomaly detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 938–946.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308.
- Chen, J.; Li, L.; Su, L.; Zha, Z.-J.; and Huang, Q. 2024. Prompt-enhanced multiple instance learning for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18319–18329.
- Chen, W.; Ma, K. T.; Yew, Z. J.; Hur, M.; and Khoo, D. A.-A. 2023. TEVAD: Improved video anomaly detection with captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5549–5559.
- Cho, M.; Kim, M.; Hwang, S.; Park, C.; Lee, K.; and Lee, S. 2023. 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*, 12137–12146.
- Dev, P. P.; Hazari, R.; and Das, P. 2024. Reflip-vad: Towards weakly supervised video anomaly detection via vision-language model. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Feng, J.-C.; Hong, F.-T.; and Zheng, W.-S. 2021. Mist: Multiple instance self-training framework for video anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14009–14018.
- Hasan, M.; Choi, J.; Neumann, J.; Roy-Chowdhury, A. K.; and Davis, L. S. 2016. Learning temporal regularity in video sequences. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 733–742.
- Huang, C.; Shi, Y.; Wen, J.; Wang, W.; Xu, Y.; and Cao, X. 2025. Ex-VAD: Explainable Fine-grained Video Anomaly Detection Based on Visual-Language Models. In *Forty-second International Conference on Machine Learning*.
- Joo, H. K.; Vo, K.; Yamazaki, K.; and Le, N. 2023. Clip-tsa: Clip-assisted temporal self-attention for weakly-supervised video anomaly detection. In *2023 IEEE International Conference on Image Processing (ICIP)*, 3230–3234. IEEE.
- Ju, C.; Han, T.; Zheng, K.; Zhang, Y.; and Xie, W. 2022. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, 105–124. Springer.
- Lin, J.; Gan, C.; and Han, S. 2019. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7083–7093.
- Liu, T.; Lam, K.-M.; and Bao, B.-K. 2024. Injecting text clues for improving anomalous event detection from weakly labeled videos. *IEEE Transactions on Image Processing*.
- Liu, W.; Luo, W.; Lian, D.; and Gao, S. 2018. Future frame prediction for anomaly detection—a new baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6536–6545.
- Lv, H.; Yue, Z.; Sun, Q.; Luo, B.; Cui, Z.; and Zhang, H. 2023. Unbiased multiple instance learning for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8022–8031.
- Majhi, S.; Dai, R.; Kong, Q.; Garattoni, L.; Francesca, G.; and Bremond, F. 2024. Human-scene network: A novel baseline with self-rectifying loss for weakly supervised video anomaly detection. *Computer Vision and Image Understanding*, 241: 103955.
- Majhi, S.; D’Amicantonio, G.; Dantcheva, A.; Kong, Q.; Garattoni, L.; Francesca, G.; Bondarev, E.; and Brémond, F. 2025. Just Dance with pi! A Poly-modal Inductor for Weakly-supervised Video Anomaly Detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 24265–24274.
- Majhi, S.; Das, S.; and Brémond, F. 2021. DAM: Dissimilarity attention module for weakly-supervised video anomaly detection. In *2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 1–8. IEEE.
- Majhi, S.; Das, S.; Brémond, F.; Dash, R.; and Sa, P. K. 2021. Weakly-supervised joint anomaly detection and classification. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, 1–7. IEEE.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Schölkopf, B.; Williamson, R. C.; Smola, A.; Shawe-Taylor, J.; and Platt, J. 1999. Support vector method for novelty detection. *Advances in neural information processing systems*, 12.
- Sultani, W.; Chen, C.; and Shah, M. 2018. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6479–6488.
- Tian, Y.; Pang, G.; Chen, Y.; Singh, R.; Verjans, J. W.; and Carneiro, G. 2021. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4975–4986.

- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 4489–4497.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wan, B.; Fang, Y.; Xia, X.; and Mei, J. 2020. Weakly supervised video anomaly detection via center-guided discriminative learning. In *2020 IEEE international conference on multimedia and expo (ICME)*, 1–6. IEEE.
- Wang, G.; Song, S.; He, W.; and Zheng, Y. 2025a. CMHKF: Cross-Modality Heterogeneous Knowledge Fusion for Weakly Supervised Video Anomaly Detection. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 31594–31607.
- Wang, Y.; and Chen, S. 2025. Learning Event Completeness for Weakly Supervised Video Anomaly Detection. *arXiv preprint arXiv:2506.13095*.
- Wang, Y.; Li, X.; Yan, Z.; He, Y.; Yu, J.; Zeng, X.; Wang, C.; Ma, C.; Huang, H.; Gao, J.; Dou, M.; Chen, K.; Wang, W.; Qiao, Y.; Wang, Y.; and Wang, L. 2025b. InternVideo2.5: Empowering Video MLLMs with Long and Rich Context Modeling. *arXiv:2501.12386*.
- Wu, P.; Liu, J.; and Shen, F. 2019. A deep one-class neural network for anomalous event detection in complex scenes. *IEEE transactions on neural networks and learning systems*, 31(7): 2609–2622.
- Wu, P.; Liu, J.; Shi, Y.; Sun, Y.; Shao, F.; Wu, Z.; and Yang, Z. 2020. Not only look, but also listen: Learning multimodal violence detection under weak supervision. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, 322–339. Springer.
- Wu, P.; Liu, X.; and Liu, J. 2022. Weakly supervised audiovisual violence detection. *IEEE Transactions on Multimedia*, 25: 1674–1685.
- Wu, P.; Zhou, X.; Pang, G.; Zhou, L.; Yan, Q.; Wang, P.; and Zhang, Y. 2024. Vadclip: Adapting vision-language models for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6074–6082.
- Yang, Z.; Liu, J.; and Wu, P. 2024. Text prompt with normality guidance for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18899–18908.
- Zaheer, M. Z.; Mahmood, A.; Astrid, M.; and Lee, S.-I. 2020. Claws: Clustering assisted weakly supervised learning with normalcy suppression for anomalous event detection. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, 358–376. Springer.
- Zhong, Y.; Zhu, R.; Yan, G.; Gan, P.; Shen, X.; and Zhu, D. 2024. Inter-clip feature similarity based weakly supervised video anomaly detection via multi-scale temporal MLP. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Zhou, H.; Yu, J.; and Yang, W. 2023. Dual memory units with uncertainty regulation for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3769–3777.
- Zhu, S.; Chen, C.; and Sultani, W. 2021. Video anomaly detection for smart surveillance. In *Computer Vision: A Reference Guide*, 1315–1322. Springer.