

Anomize: Better Open Vocabulary Video Anomaly Detection

Fei Li^{1,2#} Wenxuan Liu^{3,4#} Jingjing Chen² Ruixu Zhang¹ Yuran Wang¹
 Xian Zhong⁴ Zheng Wang^{1*}

¹National Engineering Research Center for Multimedia Software, School of Computer Science, Wuhan University

²Shanghai Key Lab of Intelligent Information Processing, School of Computer Science, Fudan University

³State Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

⁴Hubei Key Laboratory of Transportation Internet of Things, Wuhan University of Technology

lifeiwhu@whu.edu.cn, liuw66@pku.edu.cn

Abstract

Open Vocabulary Video Anomaly Detection (OVVAD) seeks to detect and classify both base and novel anomalies. However, existing methods face two specific challenges related to novel anomalies. The first challenge is detection ambiguity, where the model struggles to assign accurate anomaly scores to unfamiliar anomalies. The second challenge is categorization confusion, where novel anomalies are often misclassified as visually similar base instances. To address these challenges, we explore supplementary information from multiple sources to mitigate detection ambiguity by leveraging multiple levels of visual data alongside matching textual information. Furthermore, we propose incorporating label relations to guide the encoding of new labels, thereby improving alignment between novel videos and their corresponding labels, which helps reduce categorization confusion. The resulting Anomize framework effectively tackles these issues, achieving superior performance on UCF-CRIME and XD-VIOLENCE datasets, demonstrating its effectiveness in OVVAD.

1. Introduction

Video Anomaly Detection (VAD) identifies anomaly in videos and is widely used in public safety systems. Traditional VAD methods can be categorized based on the type of training data used. Semi-supervised VAD [3, 22, 49] is trained exclusively on normal samples, detecting anomalies as deviations from learned normal patterns. In contrast, weakly supervised VAD [13, 33, 50] is trained on both normal and anomalous samples but lacks precise temporal labels, treating VAD as a binary classification problem. Both methods focus on detecting specific anomaly within a closed set and exhibit limitations in open-world scenarios.

* Corresponding author. # Contributed equally to this work.




Figure 1. **Challenges Related to Novel Anomalies.** (a) Detection ambiguity: The model struggles to assign accurate anomaly scores to unfamiliar frames containing novel anomalies. (b) Categorization confusion: Novel anomalies are misclassified as visually similar base instances from the training set.

Recent research has explored open-set VAD [1, 9], where anomalies seen in training are considered base cases, while others are treated as novel cases. It trains on normal and base anomalies to detect all anomalies, overcoming the limitations of closed-set detection. However, it struggles with understanding anomaly categories, leading to unclear outputs [42]. Consequently, a leading study [42] has further investigated open vocabulary (OV) VAD, which aims to detect and categorize all anomalies using the same training data as open-set VAD, offering more informative results.

Novel anomalies in OVVAD introduce two challenges that remain unexplored by existing methods: (1) *Detection ambiguity*, where the model often lacks sufficient information to accurately assign anomaly scores to unfamiliar data, as shown in Fig. 1(a). Current methods rely on training or fine-tuning the model, which is inherently limited and cannot adapt to the variability of samples in an open setting. (2) *Categorization confusion*, where novel cases visually similar to base cases are misclassified, as shown in Fig. 1(b). OV tasks generally rely on multimodal alignment for categorization. Since the model tends to extract visual features for novel videos similar to base videos, these features are more likely to align with base label encodings, leading to mis-categorization. Traditional OV methods use pre-trained encoders to encode text, where the input contains labels with

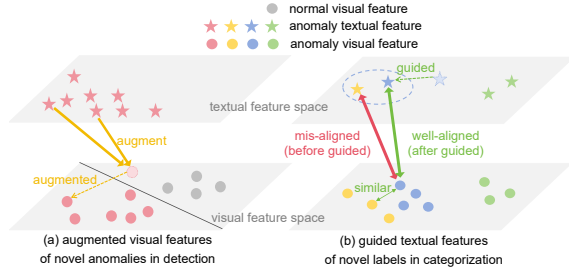


Figure 2. **Feature Visualization of Our Design.** (a) Text augmentation shifts ambiguous frames to the anomalous feature space. In the static stream, text represents anomaly-related nouns (e.g., “abandoned fire starter”), while in the dynamic stream, it denotes label descriptions. (b) Group-guided text encoding improves the alignment of novel anomalies with novel labels, especially for those resembling base samples.

unified templates [16, 29, 44, 45] or embeddings [2, 4, 53]. These methods rely solely on pre-trained encoders without spatial guidance in label encoding, limiting multimodal alignment for novel cases.

To address detection ambiguity, we introduce a **Text-Augmented Dual Stream** mechanism with dynamic and static streams, each focusing on different visual features augmented by corresponding textual information. The dynamic stream captures sequential information through temporal visual encoding, augmented by label descriptions related to dynamic characteristics. The static stream captures scene information through original contrastive language-image pre-training (CLIP)-encoded visual features, augmented by a concept library related to static characteristics. The complementarity between dynamic and static data is crucial: certain anomalies rely on temporal information, such as tailing, while others depend on contextual cues, such as running on a highway. Synergistic training of the streams ensures mutual supplementation and constraints, delivering comprehensive temporal and contextual information, minimizing overfitting to specific anomaly categories, and improving overall performance. Additionally, the augmentation follows common-sense reasoning. To detect anomalies in real-world scenarios, we first define the anomaly and establish correlations between visual data and anomaly texts, providing a reference for detection within the overall context. Similarly, we augment visual features with relevant anomaly text, providing additional information for detection. As shown in Fig. 2(a), novel visual features that cause ambiguous detections are shifted into the anomalous feature space with support from text, helping the model better assess unfamiliar anomalies.

To address categorization confusion, we introduce a **Group-Guided Text Encoding** mechanism, encoding labels using group-based descriptions, with labels sharing similar visual characteristics grouped together. As shown in Fig. 2(b), this mechanism establishes connections be-

tween base and novel data through grouping, positioning the encodings of novel labels close to those of base labels, where videos associated with both base and novel labels are visually similar, thereby enhancing multimodal alignment for novel data. For novel labels not grouped with base labels, the descriptions provide contextual support to pre-trained encoders for text encoding, thus enhancing alignment. Compared to previous methods mainly relying on pre-trained models, our approach strengthens the guidance of the feature space for novel labels, achieving more effective alignment for categorization.

With the **Anomize** framework, we achieve notable results across both XD-VIOLENCE [41] and UCF-CRIME [33] datasets. For anomaly detection, we obtain a 2.78% overall improvement on XD-VIOLENCE and 8.21% on novel cases, with UCF-CRIME results comparable to a more complex state-of-the-art model. For categorization, we achieve a 25.61% overall increase in Top-1 accuracy on XD-VIOLENCE and 5.71% on UCF-CRIME, with improvements of 56.53% and 4.49% on novel cases, respectively.

In summary, our contributions are threefold:

- To address detection ambiguity, we discover the importance of providing sufficient informational support. We combine dynamic and static streams to effectively constrain and complement each other. Operating at different levels of visual features, each stream is augmented with corresponding textual information, offering comprehensive support for detection.
- To tackle categorization confusion, we emphasize the importance of establishing connections between labels to guide their encodings. We propose a text encoding mechanism that groups labels based on visual characteristics and generates corresponding descriptions for encodings.
- Our **Anomize** framework targets the challenges of novel anomalies that remain unexplored, offering new insights for OVAD and demonstrating superior performance on two widely-used datasets, particularly for novel cases.

2. Related Work

2.1. Video Anomaly Detection

Semi-Supervised VAD. Existing semi-supervised video anomaly detection (VAD) methods are typically categorized into three groups: one-class classification (OCC), reconstruction-based models, and prediction-based models, all of which are trained exclusively on normal data. OCC models [31, 32, 35, 40, 49] classify anomalies by identifying data points that fall outside a learned hypersphere of normal data. However, defining normality can be ambiguous, which often reduces their effectiveness [23, 46]. Reconstruction-based methods [3, 25, 30, 48, 52] use deep auto-encoders (DAEs) to learn normal patterns, detecting anomalies through high reconstruction errors. However,

DAEs may still reconstruct anomalous frames with low error, weakening detection performance. Prediction models [5, 8, 18, 22, 24, 26], often utilizing GANs, forecast future frames and identify anomalies by comparing predicted frames with actual frames.

Weakly-Supervised VAD. Weakly-supervised video anomaly detection (WSVAD) identifies anomalies using only video-level labels without precise temporal or spatial information. WSVAD methods typically frame the task as a multiple instance learning (MIL) problem [19, 33, 34, 39, 50], where videos are divided into segments, and predictions are aggregated into video-level anomaly scores. Sultani *et al.* [33] first define the WSVAD paradigm using a deep multiple-instance ranking framework. Recent methods focus on optimizing models. Tian *et al.* [34] introduce RTFM, which combines dilated convolutions and self-attention to detect subtle anomalies, while Zaheer *et al.* [50] add a clustering-based normalcy suppression mechanism. Other approaches [13, 27] leverage pre-trained models to gain task-agnostic knowledge. Wu *et al.* [43] propose VadCLIP, which uses CLIP [29] for dual-branch outputs of anomaly scores and labels.

Open-Set VAD. Open-set VAD models are trained on normal behaviors and base anomalies to detect all anomalies, addressing the challenges of open-world environments. Acsintoae *et al.* [1] first introduce open-set VAD, along with a benchmark dataset and evaluation framework. Zhu *et al.* [54] combine evidential deep learning and normalizing flows within a multiple instance learning framework. Hirschorn *et al.* [9] propose a lightweight normalizing flows framework that utilizes human pose graph structures.

Our method provides both detection and categorization results in an open setting, focusing on addressing the challenges related to novel anomalies.

2.2. Open Vocabulary Learning

Recent advancements in pre-trained vision-language models [11, 29] have spurred significant interest in open vocabulary tasks, including object detection [6, 15, 51], semantic segmentation [7, 20, 47], and action recognition [12, 14, 21, 36]. These studies leverage the pre-trained knowledge of multimodal models, demonstrating strong generalization. Wu *et al.* [42] first introduce open vocabulary video anomaly detection (OVVAD) using the pre-trained model CLIP. However, most methods emphasize the visual encoder while neglecting the text encoder, limiting zero-shot capabilities. Our method explores the text encoder and incorporates a guided encoding mechanism to enhance multimodal alignment in OVVAD.

3. Proposed Anomize Method

3.1. Overview

Following Wu *et al.* [42], we define the training sample set as $D = \{(v_i, y_i)\}_{i=1}^{N+A}$, which consists of N normal samples D_n and A abnormal samples D_a . Here, v_i represents video samples, and $y_i \in \mathbb{C}_{\text{base}}$ denotes the corresponding anomaly labels. Each $v_i \in D_a$ contains at least one anomalous frame, while $v_i \in D_n$ consists entirely of normal frames. The complete label set \mathbb{C} includes both base and novel anomaly labels. The objective of OVVAD is to train a model on D to predict frame-level anomaly scores and video-level anomaly labels from \mathbb{C} .

Fig. 3 illustrates the overview of framework. We leverage the encoder of the pre-trained CLIP model for its strong generalization capabilities. Video frames are processed by the CLIP image encoder Φ_{visual} to extract original visual features $x_f \in \mathbb{R}^{n \times d}$, where n is the number of frames and d is the feature dimension. These features are temporally modeled by a lightweight temporal encoder. The original features, augmented by a concept library ConceptLib, pass through the static stream, while temporal features, augmented by label descriptions, pass through the dynamic stream. The prediction from each stream is obtained and aggregated to generate the final frame-level anomaly score $s \in \mathbb{R}^{n \times 1}$. For categorization, a multimodal alignment method is used. A fused visual feature is first generated, and the CLIP text encoder Φ_{text} extracts textual features via the group-guided text encoding mechanism. Frame-level predictions are then obtained through alignment and aggregated for the final video-level result p_{video} .

3.2. Lightweight Temporal Encoder

We utilize the frozen Φ_{visual} for visual features to leverage its zero-shot capabilities. However, since CLIP is pre-trained on image-text pairs, it lacks temporal modeling for video. Recent methods [14, 37, 38] commonly introduce a temporal encoder. However, this often leads to performance degradation on novel cases, as the additional parameters in the encoder may become specialized for the training set, leading to overfitting. Therefore, we adopt a lightweight long short-term memory (LSTM) [10] for temporal modeling, resulting in the temporal visual feature $x_{\text{tem}} \in \mathbb{R}^{n \times d}$:

$$x_{\text{tem}} = \text{LSTM}(x_f). \quad (1)$$

Other parameter-efficient models may also be suitable, as discussed in the supplementary material.

3.3. Group-Guided Text Encoding

Previous methods mainly rely on the generalization capabilities of pre-trained models without task-specific guidance, often leading to categorization confusion. We introduce a group-guided text encoding mechanism to address this.

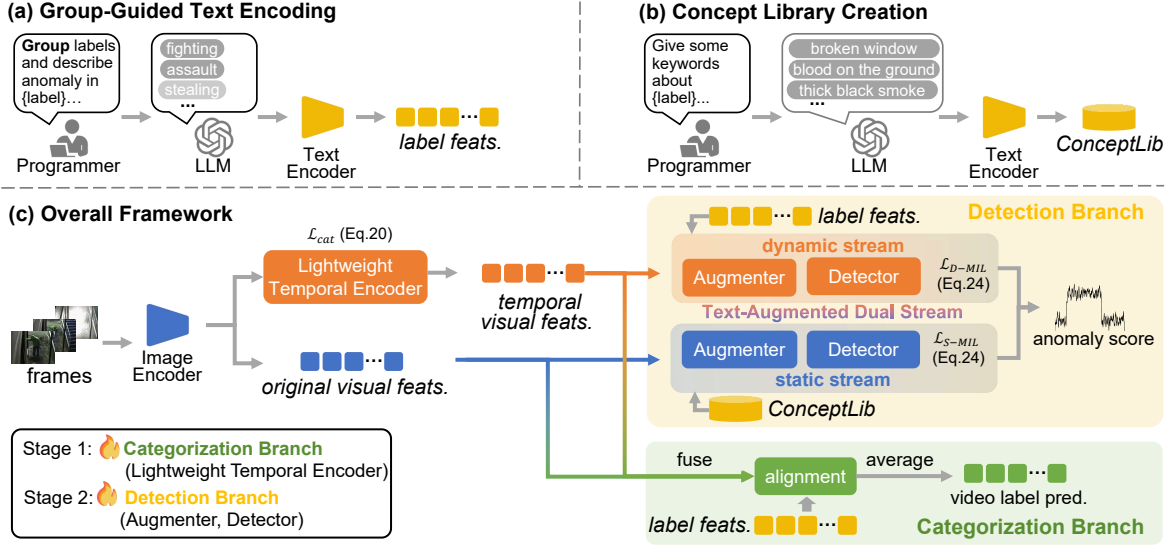


Figure 3. **Overview of Our Anomize Framework.** (a) Process for obtaining label features via the Group-Guided Text Encoding mechanism. (b) Creation of the concept library ConceptLib for anomaly detection. (c) The framework processes anomaly labels and video frames to generate frame-level anomaly scores and detected labels. Scoring is performed using a Text-Augmented Dual Stream mechanism, where each stream receives corresponding text and visual features, and the fused scores are produced as output. For labeling, the model aligns label features from the Group-Guided Text Encoding mechanism with the fused original and temporal visual encodings. Both the text and image encoders, pre-trained on CLIP, remain frozen without further optimization.

We leverage large language models (LLM), specifically GPT-4 [28], for textual encoding. We first use the prompt $\text{prompt}_{\text{group}}$ to group labels, ensuring that corresponding videos in each group exhibit high visual similarity. Then, we apply the prompt $\text{prompt}_{\text{desc}}$ to generate text descriptions for each label based on the grouping. These descriptions capture shared elements while emphasizing unique characteristics within each group, ensuring the encodings remain similar yet distinguishable:

$$\text{result}_{\text{group}} = \text{GPT}(\text{prompt}_{\text{group}}, \mathbb{C}), \quad (2)$$

$$\text{result}_{\text{desc}} = \text{GPT}(\text{prompt}_{\text{desc}}, \text{result}_{\text{group}}), \quad (3)$$

where $\text{result}_{\text{group}}$ and $\text{result}_{\text{desc}}$ represent the label groups and their descriptions. The descriptions are then passed into the frozen CLIP text encoder Φ_{text} to obtain encodings $t_{\text{desc}} \in \mathbb{R}^{c \times d}$, where c is the number of anomaly labels:

$$t_{\text{desc}} = \Phi_{\text{text}}(\text{result}_{\text{desc}}). \quad (4)$$

These encodings are used for multimodal alignment. For the visual features, we combine the temporal and original encodings, preserving the knowledge captured by the pre-trained model:

$$x_{\text{fused}} = x_{\text{tem}} + \alpha x_f, \quad (5)$$

where α is a scalar weight. The prediction probabilities for video frames are expressed as:

$$p_{\text{frame}} = \frac{x_{\text{fused}} t_{\text{desc}}^{\top}}{\|x_{\text{fused}}\| \|t_{\text{desc}}\|}, \quad (6)$$

where $p_{\text{frame}} \in \mathbb{R}^{n \times c}$ represents the probability distribution over c anomaly labels for each frame. To obtain the video-level prediction, we select the top M probabilities for each label and average these top values, where M is the total number of frames divided by 16:

$$p_{\text{avg}} = \sigma \left[\frac{1}{M} \sum_{k=1}^M \text{Top}M(p_{\text{frame}}[j]) \right]_{j=1}^c, \quad (7)$$

where $p_{\text{avg}} \in \mathbb{R}^c$ is the average probabilities for each label after applying the softmax function $\sigma(\cdot)$. Finally, the video-level prediction p_{video} is determined as the label with the highest average probability:

$$p_{\text{video}} = \arg \max_j p_{\text{avg}}^{(j)}. \quad (8)$$

3.4. Augmenter

In our method, both streams utilize text to enhance visual encodings via a unified augmenter. The augmenter takes visual encodings e_{visual} and textual encodings e_{text} as input. These are processed by a multi-head attention layer $\text{MHA}(\cdot)$, where e_{visual} acts as the query and e_{text} acts as the key and value. This operation extracts the most relevant textual features for supplementation, denoted as e_{refine} :

$$e_{\text{refine}} = \text{MHA}(e_{\text{visual}}, e_{\text{text}}, e_{\text{text}}). \quad (9)$$

A fully connected layer $\text{FC}(\cdot)$ linearly projects the visual encoding, which is concatenated with the refined textual features and passed through a multi-layer perceptron

MLP(\cdot) for dimensionality reduction. This results in the augmented output e_{aug} :

$$e_{\text{aug}} = \text{MLP}([e_{\text{refine}}; \text{FC}(e_{\text{visual}})]), \quad (10)$$

3.5. Text-Augmented Dual Stream

In open settings, models may struggle to assess unfamiliar anomalies due to limited information. We propose a Text-Augmented Dual Stream mechanism with complementary dynamic and static streams, each augmented by relevant text to provide sufficient support for detection.

Since video anomaly detection (VAD) relies on temporal cues, we employ a dynamic stream to predict anomaly scores $s_{\text{dyn}} \in \mathbb{R}^{n \times 1}$ based on refined visual features $f_{\text{aug}} \in \mathbb{R}^{n \times d}$, derived from temporal visual features and augmented by label descriptions via the augmenter in Sec. 3.4:

$$f_{\text{aug}} = \text{Augmenter}(x_{\text{tem}}, t_{\text{desc}}), \quad (11)$$

$$s_{\text{dyn}} = \text{Sigmoid}(\text{FC}(f_{\text{aug}} + \text{MLP}(f_{\text{aug}}))), \quad (12)$$

where $\text{Sigmoid}(\cdot)$ converts predictions to $[0, 1]$.

Since the dynamic stream is limited in scene context, we employ a static stream with original visual features augmented by anomaly-relevant concept data.

Specifically, we create a concept library ConceptLib containing key features related to anomalies. These features are generated by Φ_{text} from various nouns describing significant characteristics of the anomalies:

$$\text{ConceptLib} = \Phi_{\text{text}}(\text{GPT}(\text{prompt}_{\text{conc}}, \mathbb{C})), \quad (13)$$

where $\text{prompt}_{\text{conc}}$ is a prompt for relevant nouns. We then compute the cosine similarity between the visual feature of frame i , $x_f^{(i)}$, and concept features $h \in \text{ConceptLib}$:

$$\text{sim}(x_f^{(i)}, h) = \frac{x_f^{(i)} h^\top}{\|x_f^{(i)}\| \|h\|}. \quad (14)$$

The top K relevant concept features $h_f^{(i)} \in \mathbb{R}^{K \times d}$ and their scores $s_f^{(i)} \in \mathbb{R}^K$ are selected:

$$h_f^{(i)}, s_f^{(i)} = \text{TopK}(\text{sim}(x_f^{(i)}, h)). \quad (15)$$

The selected features are then weighted by their scores and concatenated to form refined textual features for the video:

$$h_f^{\text{new}} = \left[\sigma(s_f^{(i)}) \odot h_f^{(i)} \right]_{i=1}^n. \quad (16)$$

Next, the refined features $h_f^{\text{new}} \in \mathbb{R}^{n \times k \times d}$ and the original visual features x_f are passed through the augmenter to generate the augmented encoding $x_{\text{aug}} \in \mathbb{R}^{n \times d}$:

$$x_{\text{aug}} = \text{Augmenter}(x_f, h_f^{\text{new}}). \quad (17)$$

Similar to the dynamic stream, x_{aug} is fed into a detector for anomaly score prediction s_{sta} in the static stream:

$$s_{\text{sta}} = \text{Sigmoid}(\text{FC}(x_{\text{aug}} + \text{MLP}(x_{\text{aug}}))). \quad (18)$$

Finally, the outputs of the two streams are aggregated for the overall anomaly score prediction s :

$$s = \beta s_{\text{dyn}} + (1 - \beta) s_{\text{sta}}, \quad (19)$$

where β is a tunable parameter balancing the contributions of the dynamic and static streams.

3.6. Objective Functions

First Training Stage. In the first stage, we focus on video anomaly categorization to train the LSTM while freezing other modules to prevent optimization conflicts. We use cross-entropy loss \mathcal{L}_{ce} for categorization. To prevent overfitting to normal data due to class imbalance, we add a separation loss \mathcal{L}_{sep} to enhance the distinction between normal and anomalous predictions. The loss for the first stage is:

$$\mathcal{L}_{\text{cat}} = \mathcal{L}_{\text{ce}} + \mathcal{L}_{\text{sep}}, \quad (20)$$

$$\mathcal{L}_{\text{ce}} = -\frac{1}{N} \sum_{i=1}^N g_i \log(p_{\text{avg}, i}), \quad (21)$$

$$\mathcal{L}_{\text{sep}} = -\frac{1}{N} \sum_{i=1}^N |\max(p_{\text{avg}, i}[1:] - p_{\text{avg}, i}[0])| + 1, \quad (22)$$

where $p_{\text{avg}, i}$ denotes the predicted probabilities for the i -th video, and the normal label is at the first index. g_i is the one-hot ground truth, and N denotes the batch size.

Second Training Stage. In the second stage, we focus on anomaly detection to train the static and dynamic streams while freezing other modules. Following Wu *et al.* [43], we apply the MIL loss. Specifically, we first average the top M frame anomaly scores to obtain the video-level prediction \hat{q}_i , then compute the loss $\mathcal{L}_{X-\text{MIL}}$ for each stream using binary cross-entropy to quantify the difference between predictions and binary labels q_i , where $q_i = 1$ denotes an anomaly and $X \in \{D, S\}$ denotes the type of stream. Additionally, we apply a weight w_i in this phase to tackle data imbalance by increasing the penalty for incorrect scores related to anomalous videos. The loss is defined as follows:

$$\mathcal{L}_{\text{det}} = \mathcal{L}_{D-\text{MIL}} + \mathcal{L}_{S-\text{MIL}}, \quad (23)$$

$$\mathcal{L}_{X-\text{MIL}} = -\frac{1}{N} \sum_{i=1}^N w_i (q_i \log(\hat{q}_i) + (1 - q_i) \log(1 - \hat{q}_i)). \quad (24)$$

Method	AP	AP _b	AP _n	AUC	AUC _b	AUC _n
Zhu et al.* [54]	64.40	-	-	78.82	-	-
Sultani et al. [33]	52.26	51.25	54.64	78.25	86.31	80.12
Wu et al. [41]	55.43	52.94	64.10	82.24	90.62	84.13
RTFM [34]	58.99	55.72	65.97	84.47	92.54	85.87
Wu et al. [42]	66.53	57.10	76.03	86.40	93.80	88.20
Ours	69.31	57.37	84.24	<u>84.49</u>	<u>93.00</u>	<u>87.05</u>

Table 1. **Detection Metrics (%) Comparisons for XD-VIOLENCE (left) and UCF-CRIME (right).** The best results are highlighted in bold, our method is shaded in gray, the symbol * indicates different category divisions, and underlined values represent the second-best results.

4. Experimental Results

4.1. Datasets and Implementation Details

Datasets. We evaluate the performance of Anomize on two widely-used benchmark datasets. XD-VIOLENCE [41] is the largest dataset focused on violent events in videos, containing 3,954 training videos and 800 testing videos. The videos are collected from movies and YouTube, capturing six types of anomalous events across diverse scenarios. UCF-CRIME [33] is a large-scale dataset with 1,610 untrimmed surveillance videos for training and 290 for testing, totaling 128 hours. This dataset includes 13 types of anomalous events, spanning both indoor and outdoor settings and providing broad coverage of real-world scenarios.

Evaluation Metrics. For anomaly detection, we employ standard metrics from previous works [33, 41]. Specifically, for UCF-CRIME, we compute AUC, which captures the trade-off between true positive and false positive rates. For XD-VIOLENCE, we report AP, reflecting the balance between precision and recall. For anomaly categorization, we report Top-1 accuracy on anomalous test videos from both datasets. These metrics are provided for all categories combined, as well as separately for base and novel categories, denoted by the subscripts b and n, respectively.

Implementation Details. We implement our model in PyTorch and train it on an RTX 4090 with a 256-frame limit. Using the AdamW optimizer [17] with a learning rate of 2×10^{-5} and a batch size of 32, we train for 16 and 64 epochs in two phases. We use the pre-trained CLIP (ViT-B/16) model. The MLP module contains 2 fully connected layers with GeLU activation. The fusion weight α is 1 during training and 2 during testing. K is 25 for XD-VIOLENCE and 5 for UCF-CRIME. Score weight β is 1 on XD-VIOLENCE (0 for base categories) and 0.5 on UCF-CRIME (0 for novel categories). Loss weight w_i follows the normal-to-anomaly ratio per iteration.

4.2. Comparison with State-of-the-Art Methods

In Tab. 1, we compare the performance of our method with prior VAD methods, ensuring that all methods use the same

Method	ACC	ACC _b	ACC _n	ACC	ACC _b	ACC _n
Wu et al. [42]	64.68	89.31	30.90	41.43	49.02	37.08
Ours	90.29	92.37	87.43	47.14	56.86	41.57

Table 2. **Top-1 Accuracy (%) Comparisons on XD-VIOLENCE (left) and UCF-CRIME (right).**

visual features from CLIP and adopt an open-set setting. On XD-VIOLENCE, we achieve the best performance, with an increase of 2.78% overall and 8.21% on novel cases, demonstrating the effectiveness of our method in reducing detection ambiguity. On UCF-CRIME, our method achieves competitive results, likely due to the lightweight temporal encoder and segmented training, which limits optimization for the detection branch. Since other studies focus solely on traditional VAD without categorization, we compare our method’s categorization performance with a leading study [42] in Tab. 2. Our method shows significant improvements, with a 25.61% gain on XD-VIOLENCE and 5.71% on UCF-CRIME, as well as further improvements of 56.53% and 4.49% on novel cases, highlighting its effectiveness in reducing categorization confusion.

4.3. Ablation Studies

Effectiveness of Lightweight Temporal Encoder. The experiments confirm the importance of temporal information in video-level tasks. Tab. 3 shows that the dynamic stream, with temporal encoding, provides useful temporal cues and complements the static stream effectively. While relying solely on temporal information performs poorly due to noise, the dynamic stream becomes more effective with loss weighting and text augmentation. Tab. 4 shows that adding the temporal encoder improves performance on base cases, but without further guidance, the lightweight encoder still introduces confusion for novel anomalies.

Effectiveness of Group-Guided Text Encoding Mechanism. Comparisons in Tab. 4 between the second and third rows, or the fourth and fifth rows, on both datasets show that textual encodings based on group descriptions outperform the baseline, demonstrating the importance of our text encoding mechanism.

Effectiveness of Text Augmentation. As shown in Tab. 3, text augmentation in both dynamic and static streams generally reduces detection ambiguity by compensating for the limitations of visual features. The dynamic stream with only text augmentation shows a slight drop on UCF-CRIME, suggesting noise in the temporal encodings. However, when combined with the loss weight, it demonstrates the importance of text, as shown in Rows 6 and 7.

Effectiveness of Integrating Dynamic and Static Streams. Tab. 3 shows that integrating the two streams is generally more effective than using them independently, as

M_{sta}			M_{dyn}			ST	XD-VIOLENCE			UCF-CRIME		
visual	text	w_i	visual	text	w_i		AP (%)	AP_b (%)	AP_n (%)	AUC (%)	AUC_b (%)	AUC_n (%)
✓	×	×	×	×	×	✓	54.75	46.30	76.61	48.47	47.72	50.05
✓	✓	×	×	×	×	✓	59.65	56.41	75.29	52.21	59.18	48.45
✓	✓	✓	×	×	×	✓	59.26	57.20	74.28	84.26	92.06	86.86
×	×	×	✓	×	×	✓	47.19	36.87	69.79	26.93	17.39	25.73
×	×	×	✓	✓	×	✓	54.98	57.79	69.40	23.06	14.68	19.74
×	×	×	✓	✓	✓	✓	54.16	56.03	68.19	83.21	91.94	85.43
×	×	×	✓	×	✓	✓	48.55	51.15	59.93	80.97	90.71	83.63
✓	✓	×	✓	✓	×	✓	64.92	52.93	79.91	52.37	59.44	48.61
✓	✓	✓	✓	✓	✓	×	58.22	58.58	72.31	84.52	92.62	86.52
✓	✓	✓	✓	✓	✓	✓	69.31	57.37	84.24	84.49	93.00	87.05

Table 3. **Effectiveness of Dynamic M_{dyn} and Static M_{sta} Streams with Text-Augmented Visual Data, Additional Loss Weight w_i , and Segmented Training (ST) on XD-VIOLENCE (left) and UCF-CRIME (right).** In M_{sta} , the visual data is the original visual feature output by CLIP, while in M_{dyn} , it is derived from a lightweight temporal encoder.

E_{tem}	M_{group}	F_{fus}	\mathcal{L}_{sep}	ST	XD-VIOLENCE			UCF-CRIME		
					ACC (%)	ACC_b (%)	ACC_n (%)	ACC (%)	ACC_b (%)	ACC_n (%)
×	×	×	×	✓	45.92	74.81	6.28	35.71	56.86	23.60
✓	×	×	×	✓	53.42	92.37	0	25.71	70.59	0
✓	✓	×	×	✓	56.07	92.75	5.76	27.86	68.63	4.49
✓	×	✓	×	✓	56.51	94.66	4.19	27.14	74.51	0
✓	✓	✓	×	✓	90.07	91.98	87.43	46.43	56.86	40.45
✓	✓	✓	✓	×	89.95	91.98	86.91	46.43	52.94	42.70
✓	✓	✓	✓	✓	90.29	92.37	87.43	47.14	56.86	41.57

Table 4. **Effectiveness of the Lightweight Temporal Encoder E_{tem} , Group-Guided Text Encoding Mechanism M_{group} , Fusion Function F_{fus} , Additional Separation Loss \mathcal{L}_{sep} , and Segmented Training (ST) on XD-VIOLENCE (left) and UCF-CRIME (right).**

they complement and constrain each other, which is evident from the comparisons in the third, sixth, and last rows.

Effectiveness of Additional Loss Design. Tab. 3 emphasizes the importance of loss weight w_i , especially when training dynamic and static streams together. On XD-VIOLENCE, adding w_i slightly degrades single-stream training but significantly benefits integrated streams. Besides, Tab. 4 shows that adding separation loss \mathcal{L}_{sep} improves performance, effectively addressing class imbalance.

Effectiveness of Segmented Training. Results in Tab. 3 and Tab. 4 show that segmented training generally outperforms single-phase training, especially for novel cases, confirming that single-phase training may cause optimization conflicts and increase overfitting risks. On UCF-CRIME, single-phase training shows slightly better categorization performance on novel categories, likely due to randomness caused by optimization conflicts, as reflected in the poor performance on base categories.

4.4. Qualitative results

Fig. 4 presents qualitative results for anomaly detection, featuring two base and two novel categories from each dataset to cover all label groups. The comparison between the blue predicted score lines and the pink rectangles for ground truth demonstrates the effectiveness of our model in detecting anomalies. Notably, our method demonstrates

Test Train	XD-VIOLENCE		UCF-CRIME	
	AP (%)	ACC (%)	AUC (%)	ACC (%)
XD-VIOLENCE	69.31	90.29	81.69	45.71
UCF-CRIME	66.16	85.87	84.68	47.14

Table 5. **Cross-Dataset Detection and Categorization Results.**

its capability to handle novel cases with minimal detection ambiguity, highlighting the strong support of the text-augmented dual stream.

Fig. 5 shows similarity matrices of textual encodings, where labels indicated by the same dot are grouped together. Fig. 5(a) and (c) show results from encoding text using only label text, where visually similar anomalies lack corresponding textual similarity, revealing the limitations of relying solely on pre-trained models. Fig. 5(b) and (d) present results after being guided, where textual similarities within groups are enhanced to better align with visual similarities. This demonstrates the benefit of establishing label connections to guide encoding.

4.5. Adaptability and Generalization

Analysis of Cross-Dataset Ability. Tab. 5 shows that training on one dataset in an open-set manner and testing on another achieves results comparable to direct training on the target dataset, highlighting the adaptability of our method.



Figure 4. **Qualitative Results for Anomaly Detection.** The first and second rows present results on XD-VIOLENCE and UCF-CRIME respectively. Red boxes and rectangles highlight the ground-truth anomalous frames, while blue lines represent predicted anomaly scores.

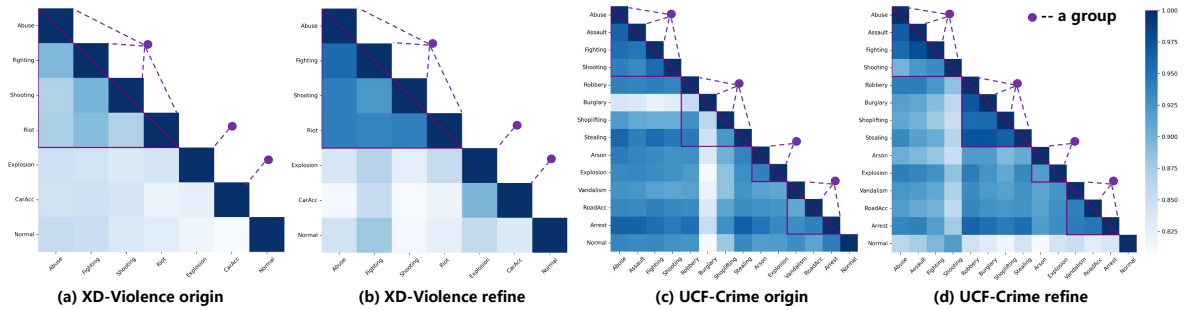


Figure 5. **Similarity Matrices of Textual Encoding.** (a) and (c) depict results using encodings from the original label data, while (b) and (d) show improvements achieved with the group-guided text encoding mechanism.

Setting	ACC (%)	ACC _b (%)	ACC _n (%)
orig.	90.29	92.37	87.43
+ shoplifting	90.51	92.37	88.21
+ arrest	89.96	92.37	86.73
+ arson	88.77	92.37	84.08
+ assault	87.28	87.40	87.11
+ labels	85.21	86.64	83.25
orig.	47.14	56.86	41.57
+ riot	65.69	56.86	68.09
+ labels	45.00	54.90	39.33

Table 6. **Impact of Data Addition on Top-1 Accuracy for XD-VIOLENCE (top) and UCF-CRIME (bottom).** “+ shoplifting” denotes the addition of both videos and labels, while “+ labels” indicates adding new labels only.

Analysis of Open Vocabulary Ability. The open vocabulary ability is demonstrated by stable categorization performance after adding novel anomaly data, which we further validate by evaluating top-1 accuracy with the added data, as shown in Tab. 6. The additional data for one dataset is sourced from another. On XD-VIOLENCE, incorporating both same-group (e.g., arson) and different-group data (e.g., shoplifting) leads to stable performance, though assault shows the clearest decline due to confusion

with similar labels like fighting. On UCF-CRIME, incorporating riot leads to notable improvement, with most instances accurately categorized despite being grouped with four other labels. Although the addition of new labels introduces some confusion, the performance remains robust across both datasets.

5. Conclusion

We propose the Anomize framework to address detection ambiguity and categorization confusion in open vocabulary video anomaly detection (OVVAD). By augmenting visual encodings with anomaly-related text through a dual-stream mechanism, Anomize improves the detection of unfamiliar samples and resolves ambiguity. Additionally, a group-guided text encoding mechanism enhances multi-modal alignment and reduces categorization confusion. Experiments on XD-VIOLENCE and UCF-CRIME demonstrate the effectiveness of our method.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (Grants 62171325, 62271361), the Hubei Provincial Key Research and Development Program (Grant 2024BAB039), and the supercomputing system at Wuhan University.

References

- [1] Andra Acsintoae, Andrei Florescu, Mariana-Iuliana Georgescu, Tudor Mare, Paul Sumedrea, Radu Tudor Ionescu, Fahad Shahbaz Khan, and Mubarak Shah. Ub-normal: New benchmark for supervised open-set video anomaly detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 20111–20121, 2022. 1, 3
- [2] Dibyadip Chatterjee, Fadime Sener, Shugao Ma, and Angela Yao. Opening the vocabulary of egocentric actions. In *Adv. Neural Inf. Process. Syst.*, 2023. 2
- [3] Yang Cong, Junsong Yuan, and Ji Liu. Sparse reconstruction cost for abnormal event detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 3449–3456, 2011. 1, 2
- [4] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 14064–14073, 2022. 2
- [5] Xinyang Feng, Dongjin Song, Yuncong Chen, Zhengzhang Chen, Jingchao Ni, and Haifeng Chen. Convolutional transformer based dual discriminator generative adversarial networks for video anomaly detection. In *Proc. ACM Int. Conf. Multimedia*, pages 5546–5554, 2021. 3
- [6] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *Proc. Int. Conf. Learn. Represent.*, 2022. 3
- [7] Kunyang Han, Yong Liu, Jun Hao Liew, Henghui Ding, Jiajun Liu, Yitong Wang, Yansong Tang, Yujiu Yang, Jiashi Feng, Yao Zhao, and Yunchao Wei. Global knowledge calibration for fast open-vocabulary segmentation. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 797–807, 2023. 3
- [8] Yi Hao, Jie Li, Nannan Wang, Xiaoyu Wang, and Xinbo Gao. Spatiotemporal consistency-enhanced network for video anomaly detection. *Pattern Recognit.*, 121:108232, 2022. 3
- [9] Or Hirschorn and Shai Avidan. Normalizing flows for human pose anomaly detection. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 13499–13508, 2023. 1, 3
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997. 3
- [11] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proc. Int. Conf. Mach. Learn.*, pages 4904–4916, 2021. 3
- [12] Chengyou Jia, Minnan Luo, Xiaojun Chang, Zhuohang Dang, Mingfei Han, Mengmeng Wang, Guang Dai, Sizhe Dang, and Jingdong Wang. Generating action-conditioned prompts for open-vocabulary video action recognition. In *Proc. ACM Int. Conf. Multimedia*, pages 4640–4649, 2024. 3
- [13] Hyekang Kevin Joo, Khoa Vo, Kashu Yamazaki, and Ngan Le. CLIP-TSA: clip-assisted temporal self-attention for weakly-supervised video anomaly detection. In *Proc. IEEE Int. Conf. Image Process.*, pages 3230–3234, 2023. 1, 3
- [14] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *Proc. Eur. Conf. Comput. Vis.*, pages 105–124, 2022. 3
- [15] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 11144–11154, 2023. 3
- [16] Jooyeon Kim, Eulrang Cho, Sehyung Kim, and Hyunwoo J. Kim. Retrieval-augmented open-vocabulary object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 17427–17436, 2024. 2
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. Int. Conf. Learn. Represent.*, 2015. 6
- [18] Sangmin Lee, Hak Gu Kim, and Yong Man Ro. BMAN: bidirectional multi-scale aggregation networks for abnormal event detection. *IEEE Trans. Image Process.*, 29:2395–2408, 2020. 3
- [19] Shuo Li, Fang Liu, and Licheng Jiao. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. In *Proc. AAAI Conf. Artif. Intell.*, pages 1395–1403, 2022. 3
- [20] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted CLIP. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 7061–7070, 2023. 3
- [21] Kun-Yu Lin, Henghui Ding, Jiaming Zhou, Yi-Xing Peng, Zhilin Zhao, Chen Change Loy, and Wei-Shi Zheng. Rethinking clip-based video learners in cross-domain open-vocabulary action recognition. *arXiv:2403.01560*, 2024. 3
- [22] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection - A new baseline. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 6536–6545, 2018. 1, 3
- [23] Wenxuan Liu, Shilei Zhao, Xiyu Han, Aoyu Yi, Kui Jiang, Zheng Wang, and Xian Zhong. Pixel-refocused navigated trimargin for semi-supervised action detection. In *Proc. ACM Int. Conf. Multimedia Workshop*, pages 23–31, 2024. 2
- [24] Zhian Liu, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. A hybrid video anomaly detection framework via memory-augmented flow reconstruction and flow-guided frame prediction. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 13568–13577, 2021. 3
- [25] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal event detection at 150 FPS in MATLAB. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 2720–2727, 2013. 2
- [26] Yiwei Lu, K. Mahesh Kumar, Seyed Shahabeddin Nabavi, and Yang Wang. Future frame prediction using convolutional VRNN for anomaly detection. In *Proc. IEEE Int. Conf. Adv. Video Signal Based Surveill.*, pages 1–8, 2019. 3
- [27] Hui Lv, Zhongqi Yue, Qianru Sun, Bin Luo, Zhen Cui, and Hanwang Zhang. Unbiased multiple instance learning for weakly supervised video anomaly detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 8022–8031, 2023. 3

- [28] OpenAI. GPT-4 technical report. *arXiv:2303.08774*, 2023. 4
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proc. Int. Conf. Mach. Learn.*, pages 8748–8763, 2021. 2, 3
- [30] Nicolae-Catalin Ristea, Florinel-Alin Croitoru, Radu Tudor Ionescu, Marius Popescu, Fahad Shahbaz Khan, and Mubarak Shah. Self-distilled masked auto-encoders are efficient video anomaly detectors. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 15984–15995, 2024. 2
- [31] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 3379–3388, 2018. 2
- [32] Bernhard Schölkopf, Robert C. Williamson, Alexander J. Smola, John Shawe-Taylor, and John C. Platt. Support vector method for novelty detection. In *Adv. Neural Inf. Process. Syst.*, pages 582–588, 1999. 2
- [33] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 6479–6488, 2018. 1, 2, 3, 6
- [34] 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 *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 4955–4966, 2021. 3, 6
- [35] Jue Wang and Anoop Cherian. GODS: generalized one-class discriminative subspaces for anomaly detection. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 8200–8210, 2019. 2
- [36] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv:2109.08472*, 2021. 3
- [37] Syed Talal Wasim, Muzammal Naseer, Salman H. Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Videogrounding-dino: Towards open-vocabulary spatio-temporal video grounding. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 18909–18918, 2024. 3
- [38] Zejia Weng, Xitong Yang, Ang Li, Zuxuan Wu, and Yu-Gang Jiang. Open-vclip: Transforming CLIP to an open-vocabulary video model via interpolated weight optimization. In *Proc. Int. Conf. Mach. Learn.*, pages 36978–36989, 2023. 3
- [39] Peng Wu and Jing Liu. Learning causal temporal relation and feature discrimination for anomaly detection. *IEEE Trans. Image Process.*, 30:3513–3527, 2021. 3
- [40] Peng Wu, Jing Liu, and Fang Shen. A deep one-class neural network for anomalous event detection in complex scenes. *IEEE Trans. Neural Networks Learn. Syst.*, 31(7):2609–2622, 2020. 2
- [41] 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 *Proc. Eur. Conf. Comput. Vis.*, pages 322–339, 2020. 2, 6
- [42] Peng Wu, Xuerong Zhou, Guansong Pang, Yujia Sun, Jing Liu, Peng Wang, and Yanning Zhang. Open-vocabulary video anomaly detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 18297–18307, 2024. 1, 3, 6
- [43] 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 *Proc. AAAI Conf. Artif. Intell.*, pages 6074–6082, 2024. 3, 5
- [44] Tao Wu, Shuqiu Ge, Jie Qin, Gangshan Wu, and Limin Wang. Open-vocabulary spatio-temporal action detection. *arXiv:2405.10832*, 2024. 2
- [45] Zuxuan Wu, Zejia Weng, Wujian Peng, Xitong Yang, Ang Li, Larry S. Davis, and Yu-Gang Jiang. Building an open-vocabulary video CLIP model with better architectures, optimization and data. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(7):4747–4762, 2024. 2
- [46] Haiyang Xie, Zhengwei Yang, Huilin Zhu, and Zheng Wang. Striking a balance: Unsupervised cross-domain crowd counting via knowledge diffusion. In *Proceedings of the 31st ACM international conference on multimedia*, pages 6520–6529, 2023. 2
- [47] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. SAN: side adapter network for open-vocabulary semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(12):15546–15561, 2023. 3
- [48] Zhiwei Yang, Jing Liu, Zhaoyang Wu, Peng Wu, and Xiaotao Liu. Video event restoration based on keyframes for video anomaly detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 14592–14601, 2023. 2
- [49] Muhammad Zaigham Zaheer, Jin-Ha Lee, Marcella Astrid, and Seung-Ik Lee. Old is gold: Redefining the adversarially learned one-class classifier training paradigm. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 14171–14181, 2020. 1, 2
- [50] Muhammad Zaigham Zaheer, Arif Mahmood, Marcella Astrid, and Seung-Ik Lee. CLAWS: clustering assisted weakly supervised learning with normalcy suppression for anomalous event detection. In *Proc. Eur. Conf. Comput. Vis.*, pages 358–376, 2020. 1, 3
- [51] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 14393–14402, 2021. 3
- [52] Yuanhong Zhong, Xia Chen, Jinyang Jiang, and Fan Ren. A cascade reconstruction model with generalization ability evaluation for anomaly detection in videos. *Pattern Recognit.*, 122:108336, 2022. 2
- [53] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *Int. J. Comput. Vis.*, 130(9):2337–2348, 2022. 2
- [54] Yuansheng Zhu, Wentao Bao, and Qi Yu. Towards open set video anomaly detection. In *Proc. Eur. Conf. Comput. Vis.*, pages 395–412, 2022. 3, 6