

Open-Vocabulary Video Anomaly Detection

Peng Wu¹, Xuerong Zhou¹, Guansong Pang^{2*}, Yujia Sun³, Jing Liu³, Peng Wang^{1*}, Yanning Zhang¹

¹Northwestern Polytechnical University, ²Singapore Management University, ³Xidian University

{xdwupeng, zxr2333}@gmail.com, gspang@smu.edu.sg, yjsun@stu.xidian.edu.cn

neouma@163.com, {peng.wang, ynzhang}@nwpu.edu.cn

Abstract

Current video anomaly detection (VAD) approaches with weak supervisions are inherently limited to a closed-set setting and may struggle in open-world applications where there can be anomaly categories in the test data unseen during training. A few recent studies attempt to tackle a more realistic setting, open-set VAD, which aims to detect unseen anomalies given seen anomalies and normal videos. However, such a setting focuses on predicting frame anomaly scores, having no ability to recognize the specific categories of anomalies, despite the fact that this ability is essential for building more informed video surveillance systems. This paper takes a step further and explores open-vocabulary video anomaly detection (OVVAD), in which we aim to leverage pre-trained large models to detect and categorize seen and unseen anomalies. To this end, we propose a model that decouples OVVAD into two mutually complementary tasks – class-agnostic detection and class-specific classification – and jointly optimizes both tasks. Particularly, we devise a semantic knowledge injection module to introduce semantic knowledge from large language models for the detection task, and design a novel anomaly synthesis module to generate pseudo unseen anomaly videos with the help of large vision generation models for the classification task. These semantic knowledge and synthesis anomalies substantially extend our model’s capability in detecting and categorizing a variety of seen and unseen anomalies. Extensive experiments on three widely-used benchmarks demonstrate our model achieves state-of-the-art performance on OVVAD task.

1. Introduction

Video anomaly detection (VAD), which aims at detecting unusual events that do not conform to expected patterns, has become a growing concern in both academia and industry communities due to its promising application prospects

in, such as, intelligent video surveillance and video content review. Through several years of vigorous development, VAD has made significant progress with many works continuously emerging.

Traditional VAD can be broadly classified into two types based on the supervised mode, i.e., semi-supervised VAD [17] and weakly supervised VAD [38]. The main difference between them lies in the availability of abnormal training samples. Although they are different in terms of supervised mode and model design, both can be roughly regarded as classification tasks. In the case of semi-supervised VAD, it falls under the category of one-class classification, while weakly supervised VAD pertains to binary classification. Specifically, semi-supervised VAD assumes that only normal samples are available during the training stage, and the test samples which do not conform to these normal training samples are identified as anomalies, as shown in Fig. 1(a). Most existing methods essentially endeavor to learn the one-class pattern, i.e., normal pattern, by means of one-class classifiers [50] or self-supervised learning technique, e.g. frame reconstruction [9], frame prediction [17], jigsaw puzzles [44], etc. Similarly, as illustrated in Fig. 1(b), weakly supervised VAD can be seen as a binary classification task with the assumption that both normal and abnormal samples are available during the training phase but the precise temporal annotation of abnormal events are unknown. Previous approaches widely adopt a binary classifier with the multiple instance learning (MIL) [38] or Top-K mechanism [27] to discriminate between normal and abnormal events. In general, existing approaches of both semi-supervised and weakly supervision VAD restricts their focus to classification and use corresponding discriminator to categorize each video frame. While these practices have achieved significant success on several widely-used benchmarks, they are limited to detecting a closed set of anomaly categories and are unable to handle arbitrary unseen anomalies. This limitation restricts their application in open-world scenarios and poses a risk of increasing missing reports, as many real-world anomalies in actual deployment are not present in the training data.

*Corresponding Authors

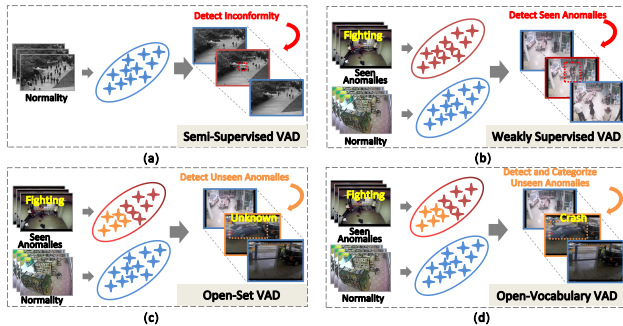


Figure 1. Comparison of different VAD tasks.

To address this issue, a few recent works explore a whole new line of VAD, i.e., open-set VAD [1, 5, 66, 67]. The core purpose of open-set VAD is to train a model with normal and seen abnormal samples to detect unseen anomalies (see Fig. 1(c)). For example, abnormal training samples only includes fighting and shooting events, and it is expected that the trained model can detect abnormal events that occur in the road accident scene. Compared to traditional VAD, open-set VAD breaks out of the close-set dilemma and then possesses ability to deal with open-world problems. Although these works partly reveal their open-world capacity, they fall short in addressing semantic understanding of the abnormal category, which leads to the ambiguous detection process in the open world.

Recently, large language/vision model pre-training [11, 29, 34, 64] has been phenomenally successful across a wide range of downstream tasks [13–15, 24, 25, 28, 47, 48, 58, 65] on account of its learned cross-modal prior knowledge and powerful transfer learning ability, which also allow us to tackle open-vocabulary video anomaly detection (OVVAD). Therefore, in this paper, we propose a novel model built upon large pre-trained vision/language models for OVVAD that aims to detect and categorize seen and unseen anomalies, as shown in Fig. 1(d). Compared to previous VAD, OVVAD has high value to applications as it can provide more informed, fine-grained detection results, but it is more challenging since that 1) it not only needs to detect but also categorize the anomalies; 2) it needs to tackle seen (base) as well as unseen (novel) anomalies. To address these challenges, we explicitly disentangle the OVVAD task into two mutually complementary sub-tasks: one is class-agnostic detection, while another one is class-specific categorization. To improve the class-agnostic detection, we make efforts from two aspects. We first introduce a nearly weight-free temporal adapter (TA) module to model temporal relationships, and then introduce a novel semantic knowledge injection (SKI) module designed to incorporate textual knowledge into visual signals with assistance of large language models. To enhance the class-specific categorization, we take inspirations from the contrastive

language-image pre-training (CLIP) model [29], and use a scalable way to categorize anomalies, i.e., the alignment between textual labels and videos, and furthermore we design a novel anomaly synthesis (NAS) module to generate vision (e.g., images and videos) materials to assist the model better identify novel anomalies. Based on these operations, our model achieves state-of-the-art performance on three popular benchmarks for OVVAD, attaining 86.40% AUC, 66.53% AP and 62.94% AUC on UCF-Crime [38], XD-Violence [51] and UBnormal [1], respectively.

We summarize our contributions as follows:

- We explore video anomaly detection under a challenging yet practically important open-vocabulary setting. To our knowledge, this is the first work for OVVAD.
- We then propose a model built on top of pre-trained large models that disentangles the OVVAD task into two mutually complementary sub-tasks – class-agnostic detection and class-specific categorization – and jointly optimizes them for accurate OVVAD.
- In the class-agnostic detection task, we design a nearly weight-free temporal adapter module and a semantic knowledge injection module for substantially-enhanced normal/abnormal frame detection.
- In the fine-grained anomaly classification task, we introduce a novel anomaly synthesis module to generate pseudo unseen anomaly videos for accurate classification of novel anomaly types.

2. Related Work

Semi-supervised VAD. Mainstream solutions are to build a normal pattern by self-supervised manner (e.g., reconstruction and prediction) or one-class manner. As for the self-supervised manner [8, 54, 56], reconstruction-based approaches [4, 21, 22, 33, 39, 55, 60] typically leverage encoder-decoder frameworks to reconstruct normal events and compute the reconstruction errors, and these events with large reconstruction error are classified as anomalies. Follow-up prediction-based approaches [17, 19] focuses on predicting the future frame with previous video frames and determine whether it is an anomaly frame by calculating the difference between the predicted frame and the actual frame. Recent work [37] combined reconstruction- and prediction-based approaches to improve detection performance. As for one-class models, some works endeavors to learn normal patterns by making use of one-class frameworks [35], e.g., one-class support vector machine and its extension (OCSVM [36], SVDD [50], GODS [45]).

Weakly supervised VAD. In contrast to semi-supervised VAD, weakly supervised VAD [10, 40] consists of normal as well as abnormal samples, which can be regarded as a binary classification task and aims to detect anomalies at frame level under the limitation of temporal annotations. As a pioneer work, Sultani et al. [38] first proposed a large-

scale benchmark and trained a lightweight network with MIL mechanism. Then Zhong et al. [61] proposed a graph convolutional network based approach to capture the similarity relations and temporal relations across frames. Tian et al. [42] introduced self-attention blocks and pyramid dilated convolution layers to capture multi-scale temporal relations. Wu et al. [51, 52] built the largest-scale benchmark that includes audio-visual signals and proposed a multi-task model to deal with coarse- and fine-grained VAD. Zaheer et al. [57] presented a clustering assisted weakly supervised framework with novel normalcy suppression mechanism. Li et al. [16] proposed a transformer-based network with self-training multi-sequence learning. Zhang et al. [59] attempted to exploit the completeness and uncertainty of pseudo labels. The above approaches simply used video or audio inputs encoded by pre-trained models, such as C3D [43] and I3D [3], although a few works [12, 23, 53] introduced CLIP models to weakly-supervised VAD task, they simply used its powerful visual features and ignored the zero-shot ability of CLIP.

Open-set VAD. VAD task naturally exists an open-world requirement. Faced with an open-world requirement, traditional semi-supervised works are more prone to producing large false alarms, while weak-supervised works are effective in detecting known anomalies but could fail in unseen anomalies. Open-set VAD aims to train the model based on normality and seen anomalies, and attempts to detect unseen anomalies. Acsintoae et al. [1] developed the first benchmark called UBnormal for supervised open-set VAD task. Zhu et al. [67] proposed an approach to deal with open-set VAD task by integrating evidential deep learning and normalizing flows into a MIL framework. Besides, Ding et al. [5] proposed a multi-head network based model to learn the disentangled anomaly representations, with each head dedicated to capturing one specific type of anomaly. Compared to our model, these above works mainly devote themselves to open-world detection and overlook anomaly categorization, moreover, these works also fail to take full advantage of pre-trained models.

3. Method

Problem Statement. The studied problem, OVVD, can be formally stated as follows. Suppose we are given a set of training samples $\mathcal{X} = \{x_i\}_{i=1}^{N+A}$, where $\mathcal{X}_n = \{x_i\}_i^N$ is the set of normal samples and $\mathcal{X}_a = \{x_i\}_{i=N+1}^{N+A}$ is the set of abnormal samples. For each sample x_i in \mathcal{X}_a , it has a corresponding video-level category label y_i , $y_i \in C_{base}$. Here, C_{base} represents the set of base (seen) anomaly categories, and \mathcal{C} is the union of C_{base} and C_{novel} , where C_{novel} stands for the set of novel (unseen) anomaly categories. Based on the training samples \mathcal{X} , the objective is to train a model capable of detecting and categorizing both base and novel anomalies. Specifically, the goal of model is

to predict anomaly confidence for each frame, and identify the anomaly category if anomalies are present in the video.

3.1. Overall Framework

Traditional methods based on close-set classifications are less likely to deal with VAD under the open-vocabulary scenario. To this end, we leverage language-image pre-training models, e.g., CLIP, as the foundation thanks to its powerful zero-shot generalization ability. As illustrated in Fig. 2, given a training video, we first feed it into image encoder of CLIP Φ_{CLIP-v} to obtain frame-level features x_f with shape of $n \times c$, where n is the number of video frames, and c is the feature dimension. Then these features pass through TA module, SKI module and a detector to produce frame-level anomaly confidence p , this pipeline is mainly applied to class-agnostic detection task. On the other hand, for class-specific categorization, we take inspirations from other open-vocabulary works across different vision tasks [31, 46, 63] and use cross-modal alignment mechanism. Specifically, we first generate a video-level aggregated feature across frame-level features, then also generate textual features/embeddings of anomaly categories, finally, we estimate the anomaly category based on alignments between video-level features and textual features. Moreover, we introduce NAS module to generate potential novel anomalies with the assistance of large language models (LLM) and AI-generated content models (AIGC) for novel category identification achievement.

3.2. Temporal Adapter Module

Temporal dependencies plays a vital role in VAD [49, 62]. In this work, we employ the frozen image encoder of CLIP to attain vision features, but it lacks consideration of temporal dependencies since CLIP is pre-trained on image-text pairs. To bridge the gap between images and videos, the use of a temporal transformer [13, 25] has emerged as a routine practice in recent studies. However, such a paradigm suffers from a clear performance degradation on novel categories [13, 32], the possible reason is that additional parameters in temporal transformer could specialise on the training set, thus harming the generalisation towards novel categories. Therefore, we design a nearly weight-free temporal adapter for temporal dependencies, which is built on top of classical graph convolutional networks. Mathematically, it can be presented as follows,

$$x_t = LN(\text{softmax}(H) x_f) \quad (1)$$

where LN is the layer normalization operation, H is the adjacency matrix, the softmax normalization is used to ensure the sum of each row of H equals to one. Such a design is used to capture contextual dependencies based on positional distance between each two frames. The adjacency matrix is

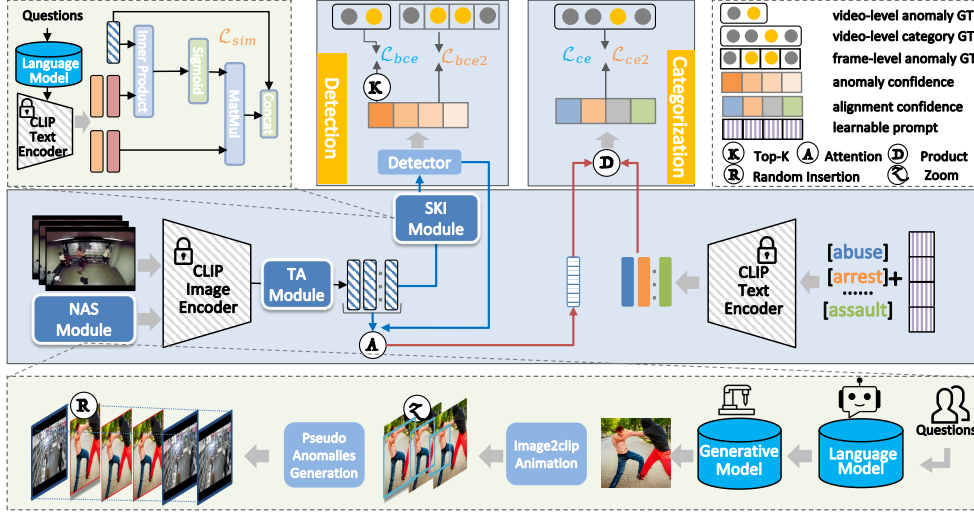


Figure 2. Overview of our proposed framework.

calculated as follows:

$$H_{(i,j)} = \frac{-|i-j|}{\sigma} \quad (2)$$

the proximity relation between i^{th} and j^{th} frames only determined by their relative temporal position. σ is a hyper-parameter to control the range of influence of distance relation. According to this formula, the closer the temporal distance between two frames, the higher proximity relation the score, otherwise the lower. Notably, across TA module, only layer normalization involves few parameters.

3.3. Semantic Knowledge Injection Module

Human often make use of prior knowledge when perceiving the environment, for example, we can infer the presence of a fire based on the smell and smoke without directly seeing the flames. Building on this idea, we propose SKI module to explicitly introduce additional semantic knowledge for assisting visual detection. As depicted in Fig. 2, for normal events in videos, we prompt the large-scale language models, e.g., ChatGPT [2] and SparkDesk¹, with a fixed template, to obtain about common scenarios and actions, such as, street, park, shopping hall, walking, running, working, etc. Likewise, we generate additional words related to anomaly scenes, including terms like explosion, burst, fire-light, etc. Finally, we obtain several phrase lists denoted by M_{prior} , which consists of noun words (scenes) and verb words (actions). With M_{prior} in hands, we exploit the text encoder of CLIP to extract textual embeddings as the semantic knowledge, which is show as follows,

$$F_{text} = \Phi_{CLIP-t}(\Phi_{token}(M_{prior})) \quad (3)$$

¹<https://xinghuo.xfyun.cn>

where $F_{text} \in \mathcal{R}^{l \times c}$, Φ_{CLIP-t} denotes the text encoder of CLIP, and Φ_{token} refers to the language tokenizer that converts words into vectors.

Then, towards the goal of effectively incorporating these semantic knowledge into visual information to boost anomaly detection, we design a cross-modal injection strategy. This strategy encourage visual signals to seek related semantic knowledge and integrate it into the process. Such an operation is demonstrated as,

$$F_{know} = \text{sigmoid}(x_t F_{text}^T) F_{text} / l \quad (4)$$

where $F_{know} \in \mathcal{R}^{n \times c}$, and we employ sigmoid instead of softmax to ensure that visual signals can encompass more relevant semantic concepts.

Finally, we concatenate F_{know} and x_t creating an input that contains both visual information and integrated semantic knowledge. We feed this input into a binary detector to generate anomaly confidence for class-agnostic detection.

3.4. Novel Anomaly Synthesis Module

While current pre-trained vision-language models, such as, CLIP, possess impressive zero-shot capacities, their zero-shot performance on various downstream tasks, especially video-related ones, remains far from satisfactory. For the same reason, our model, which is built on these pre-trained vision-language models, is trained on base anomalies and normal samples, making it susceptible to a generalization deficiency when faced with novel anomalies. With the advent of large generative models, generating samples as pseudo training data has become a feasible solution [20, 26]. Consequently, we propose NAS module to generate a series of pseudo novel anomalies based solely on potential anomaly categories. We then leverage these samples to fine-

tune the proposed model for improved categorization and detection of novel anomalies. On the whole, NAS module consists of three key processes:

1) Initially, we prompt LLMs (e.g., ChatGPT, ERNIE Bot [41]) with pre-defined templates $prompt_{gen}$ like *generate ten shorter scene descriptions about "Fighting" in real world* to produce textual descriptions of potential novel categories. We then employ AIGC models, e.g., DALL·E mini [30], Gen-2 [7], to generate corresponding images or short videos. This can be represented as follows,

$$V_{gen} = \Phi_{AIGC}(\Phi_{LLM}(prompt_{gen})) \quad (5)$$

where V_{gen} is the combination of generated images (I_{gen}) and short videos (S_{gen}).

2) Subsequently, for I_{gen} , we draw inspiration from [18] and introduce a simple yet effective animation strategy to convert single images into video clips that simulates scene changes. Specifically, given an image, we employ the center crop mechanism with different crop ratios to select corresponding image regions, then resize these regions back to original size and cascade them to create new video clips S_{cat} .

3) Finally, to mimic real-world situation where anomaly videos are generally long and untrimmed, we introduce the third step, pseudo anomaly synthesis, by inserting S_{cat} or S_{gen} into randomly selected normal videos. Moreover, the insertion position is also randomly chosen. This process yields the final pseudo anomaly samples \mathcal{V}_{nas} . Refer to supplementary materials for detailed descriptions and results.

With \mathcal{V}_{nas} in hands, we fine-tune our model, which was initially trained on \mathcal{X} , to enhance its generalization capacities for novel anomalies.

3.5. Objective Functions

3.5.1 Training stage without pseudo anomaly samples

For class-agnostic detection, following previous VAD works [27, 49], we use the Top-K mechanism to select the top K high anomaly confidences in both abnormal and normal videos. We compute the average values of these selections and feed the average values into the sigmoid function as the video-level predictions. Here, we set $K = n/16$ for abnormal videos and $K = n$ for normal videos. Finally, we compute binary cross entropy L_{bce} between video-level prediction and binary labels.

In regard to class-specific categorization, we compute the similarity between aggregated video-level features and textual category embeddings to derive video-level classification predictions. We also use a cross entropy loss function to compute the video-level categorization loss L_{ce} . Given that OVVD is a weakly supervised task, we can not obtain video-level aggregated features directly from frame-level annotations. Following [49], we employ a soft attention

based aggregation, as shown below,

$$x_{agg} = softmax(p)^\top x_t \quad (6)$$

For textual category embeddings, we are inspired by CoOp[63] and append the learnable prompt to original category embeddings.

For the parameters of SKI module, namely F_{text} , we aim for explicit optimization during the training stage. We intend to distinguish between normal knowledge embeddings and abnormal knowledge embeddings. For normal videos, we expect their visual features have higher similarities with normal knowledge embeddings and lower similarities with abnormal knowledge embeddings. To this end, we first extract the similarity matrix between each video and textual knowledge embeddings, and then select the top 10% highest scores for each frame and compute the average value, finally, we apply the cross-entropy-base loss L_{sim-n} . For abnormal videos, we anticipate the high similarities between abnormal knowledge embeddings and abnormal video-frame features. Since precise frame-level annotations are absent under weak supervision, we employ a hard attention based selection mechanism know as Top-K to locate abnormal regions. The same operations are then performed to compute the loss L_{sim-a} .

Overall, during the training phase, we employ three loss functions, with the total loss function given as:

$$L_{train} = L_{bce} + L_{ce} + L_{sim} \quad (7)$$

where L_{sim} is the sum of L_{sim-n} and L_{sim-a} .

3.5.2 Fine-tuning stage with pseudo anomaly samples

After obtaining \mathcal{V}_{nas} from NAS module, we proceed with fine-tuning our model. \mathcal{V}_{nas} is synthetic, providing us with frame-level annotations and allowing us to optimize our model with full supervisions for detection. For categorization, L_{ce2} remains the same as L_{ce} , with the key difference being that labels are available not only for base categories but also for potential novel categories. For detection, L_{bce2} is the binary cross entropy loss at the frame level.

Finally, the total loss function during the fine-tuning phase is shown as:

$$L_{tune} = L_{bce2} + L_{ce2} + \lambda(L_{bce} + L_{ce}) \quad (8)$$

4. Experiments

4.1. Experiment Setup

Datasets. UCF-Crime [38] is a large-scale VAD dataset for surveillance scenes, containing 13 types of abnormal events. 800 normal videos and 810 abnormal videos are provided for training, and the remaining 140 normal videos

and 150 abnormal videos for test. **XD-Violence** [51] is the largest VAD benchmark to date, it contains 6 anomalous categories with 3954 videos for training and 800 videos for test. To align with our model, which supports single-category identification, we exclude videos with multiple categories in XD-Violence. **UBnormal** [1] is a synthesized benchmark which defines seven types of normal events and 22 types of abnormal events. During training, only 7 abnormal categories are visible, while 12 abnormal categories are used for test.

Evaluation metrics. OVVD entails detecting and categorizing anomalies. To assess detection performance, we employ standard metrics from previous works [38, 51]. For UCF-Crime and UBnormal, we use the area under the curve of the frame-level receiver operating characteristic (AUC) to evaluate performance. For XD-Violence, we utilize AUC of the frame-level precision-recall curve (AP). For classification, we report the video-level TOP1 accuracy for abnormal test videos on both UCF-Crime and XD-Violence. UBnormal lacks category labels, so we exclusively report AUC results. During the test phase, we provide these metrics for the entire set of categories, as well as separately for base and novel categories, on both UCF-Crime and XD-Violence.

Implementation Details. The proposed model is implemented using PyTorch and trained on single RTX3090 GPU. The frozen image encoder and text encoder stem from pre-trained CLIP(ViT-B/16) [6] model. The detector is a modified feed-forward network (FFN) layer in Transformer with ReLU replaced by GeLU. In line with existing works, we process 1 out of 16 frames for each video, and during the training phase, the maximum video length is set to 256. For model optimization, we use AdamW optimizer to train the model with learning rate of $1e^{-4}$ and train epoch of 20. The batch size is set to 64, consisting of an equal number of normal and abnormal samples. During the fine-tune phase with pseudo novel anomalies, the learning rate is set to $1e^{-5}$ on UBnormal and $5e^{-6}$ on UCF-Crime and XD-Violence. The fine-tuning process spans 10 epochs, with each batch containing 10 pseudo novel anomaly videos and 10 base anomaly videos. σ is set to 0.07 across all situations, and λ is set as $1e^{-1}$ on UCF-Crime, $1e^0$ on XD-Violence and UBnormal, respectively.

4.2. Comparison with State-of-the-Arts

In Tab. 1 to Tab. 3, we report comparison results with existing approaches on three public benchmarks. Since prior approaches are designed for close-set VAD, our focus is primarily on the comparison results for open-set detection. For the sake of fairness, most of comparison approaches are re-implemented with the same visual feature as our approach. The symbol † indicates that these approaches follow traditional VAD works and use the entire training set, which includes novel anomaly samples. Consequently, the per-

Mode	Method	AUC(%)	AUC _b (%)	AUC _n (%)
Semi	SVM baseline	50.10	N/A	N/A
	OCSVM[36]	63.20	N/A	N/A
	Hasan et al.[9]	51.20	N/A	N/A
Weak	Sultani et al.†[38]	84.14	N/A	N/A
	Wu et al.†[51]	84.57	N/A	N/A
	AVVD†[52]	82.45	N/A	N/A
	RTFM†[42]	85.66	N/A	N/A
	DMU†[62]	86.75	N/A	N/A
	UMIL†[23]	86.75	N/A	N/A
	CLIP-TSA†[12]	87.58	N/A	N/A
Weak	Zhu et al.*[67]	78.82	N/A	N/A
	Sultani et al.[38]	78.25	86.31	80.12
	Wu et al.[51]	82.24	90.62	84.13
	RTFM[42]	84.47	92.54	85.87
	DMU[62]	85.14	93.52	86.24
	Ours	86.40	93.80	88.20

Table 1. AUC Comparisons on UCF-Crime.

Mode	Method	AP(%)	AP _b (%)	AP _n (%)
Semi	SVM baseline	50.80	N/A	N/A
	OCSVM[36]	28.63	N/A	N/A
	Hasan et al.[9]	31.25	N/A	N/A
Weak	Sultani et al.†[38]	75.18	N/A	N/A
	Wu et al.†[51]	80.00	N/A	N/A
	RTFM†[42]	78.27	N/A	N/A
	AVVD†[52]	78.10	N/A	N/A
	DMU†[62]	82.41	N/A	N/A
	CLIP-TSA†[12]	82.17	N/A	N/A
Weak	Zhu et al.*[67]	64.40	N/A	N/A
	Sultani et al.[38]	52.26	51.25	54.64
	Wu et al.[51]	55.43	52.94	64.10
	RTFM[42]	58.99	55.72	65.97
	DMU[62]	63.90	60.12	71.63
	Ours	66.53	57.10	76.03

Table 2. AP Comparisons on XD-Violence.

Mode	Method	AUC(%)
Semi	Georgescu et al.[8]	59.30
Weak	Georgescu et al.[8]+anomalies	61.30
	Sultani et al.[38]	50.30
	Wu et al.[51]	53.70
	RTFM[42]	60.94
	DMU[62]	59.91
	Ours	62.94

Table 3. AUC Comparisons on UBnormal.

formance of these approaches outperforms models trained without novel anomaly samples. This underscores the considerable challenge presented by OVVD from a detection perspective. Regarding the comparison between our approach and other approaches on OVVD task, we observe

TA	SKI	NAS	AUC(%)	AUC _b (%)	AUC _n (%)
×	×	×	84.79	92.75	86.73
✓	×	×	85.14	93.22	86.79
×	✓	×	85.04	92.96	86.89
✓	✓	×	85.81	93.85	87.62
✓	✓	✓	86.40	93.80	88.20

Table 4. Ablations studies with different designed module on UCF-Crime for detection.

that our approach demonstrates distinct advantages over state-of-the-art counterparts. In fact, our model performs on par with the best competitors that make use of the complete training dataset. For example, our approach surpasses the top-performing model, DMU[62], by 1.26% AUC on UCF-Crime, 2.63% AP on XD-Violence, and 3.03% AUC on UBnormal. Particularly, when it comes to novel categories, our approach exhibits a clear performance advantage compared to other approaches. Notably, Zhu et al. [67] is the first work to tackle open-set VAD, where the symbol * indicates that its category division setup differs from ours. We report its detection results under settings as identical as possible, with the number of novel categories matching ours. Our model outperforms it by a substantial margin on both UCF-Crime and XD-Violence datasets.

4.3. Ablation Studies

4.3.1 Contribution of TA module

As aforementioned, TA module is devised to capture temporal dependencies, thus enhancing class-agnostic detection abilities. To verify the effectiveness of TA module, we conduct experiments and present ablation results in Tab. 4 to Tab. 6. It can be found that the inclusion of TA module, our model achieves a significant performance improvement across various datasets and metrics. More importantly, unlike previous transformer-like temporal modeling modules [13, 25] used on other open-vocabulary tasks, this nearly weight-free designed module also shows a clear gain for novel anomaly categories, e.g., adding TA module results in an improvement of 14.47% AP on XD-Violence.

4.3.2 Contribution of SKI module

In this section, we investigate the contribution of SKI module to class-agnostic detection. As reported in Tab. 4 to Tab. 6, SKI module boosts detection performance on all datasets regardless of whether TA module is introduced or not. Similar to TA module, SKI can also clearly improve performance for novel anomaly categories. The difference with TA module is that SKI module leverage LLMs to explicitly introduce semantic knowledge into visual signals and knowledge helps better distinguish between normal and abnormal events.

TA	SKI	NAS	AP(%)	AP _b (%)	AP _n (%)
×	×	×	53.11	54.84	53.69
✓	×	×	60.13	59.38	68.16
×	✓	×	56.62	53.03	63.92
✓	✓	×	65.60	61.40	73.67
✓	✓	✓	66.53	57.10	76.03

Table 5. Ablations studies with different designed module on XD-Violence for detection.

TA	SKI	NAS	AUC(%)	AP(%)
×	×	×	60.51	65.18
✓	×	×	61.18	67.36
×	✓	×	61.93	66.36
✓	✓	×	61.67	67.43
✓	✓	✓	62.94	68.07

Table 6. Ablations studies with different designed module on UB-normal for detection.

	ACC(%)	ACC _b (%)	ACC _n (%)
w/o NAS	37.86	43.14	34.83
Finetune_N	39.29	37.25	40.45
Finetune_N+B(Ours)	41.43	49.02	37.08
w/o NAS	59.60	91.98	15.18
Finetune_N	62.03	82.06	34.55
Finetune_N+B(Ours)	64.68	89.31	30.90

Table 7. Ablations studies on UCF-Crime and XD-Violence for categorization.

Test⇒ Train↓	UCF-Crime		XD-Violence	
	AUC(%)	ACC(%)	AP(%)	ACC(%)
UCF-Crime	86.05	45.00	63.74	47.90
XD-Violence	82.42	40.71	82.86	88.96

Table 8. Cross-dataset results on UCF-Crime and XD-Violence.

4.3.3 Contribution of NAS module

From Tab. 4 to Tab. 6, we can see that, for all test categories and novel anomaly categories, NAS module can obtain a significant performance gain for the class-agnostic detection. For base categories, it causes a relatively small performance degradation, which is also observed in other open-vocabulary tasks. We argue that the introduction of pseudo novel samples makes the model pay more attention to these generated novel samples, thus partially diminishing the importance of base categories. Moreover, Tab. 7 reveals that NAS module also obtains a significant performance gain for the class-specific categorization, especially for novel anomaly categories. Besides, we also found that only using generated novel samples results in a clear performance drop for base anomaly categories during the fine-tuning phase. This illustrates while generated anomaly samples benefit the generalization abilities of our model, it is essential to adopt reasonable and effective fine-tuning schemes.




Figure 3. Qualitative results of our model on testing videos. Colored window denotes ground-truth anomalous region.

4.3.4 Analysis of cross-dataset ability

To further investigate the zero-shot abilities of our model, we conducted experiments where we train our model under the cross-dataset setup. In this case, we take UCF-Crime and XD-Violence as examples. These datasets have some overlapping categories but completely different sources, with UCF-Crime developed from surveillance videos and XD-Violence collected from movies and online videos. From the evaluation results in Tab. 8, we can draw the following conclusions: First, our model achieves better performance with the whole training samples. Second, the cross-dataset test results show that our model can compete with or outperform current approaches on both UCF-Crime and XD-Violence, further validating the favorable generalization abilities of the proposed model.

4.4. Qualitative Results

We first present qualitative detection results on three datasets in Fig. 3, where the top column denotes UCF-Crime, the first three in the bottom column denote XD-Violence, and the rest denotes UBnormal. As we can see, whether base or novel categories, our method produces high anomaly confidence in anomaly regions, even there are multiple discontinuous abnormal regions in a video. Besides, we present confusion matrices of anomaly categorization in Fig. 4, it is not hard to see that there are some anomaly categories that our model cannot effectively identify, either base or novel, especially on UCF-Crime, such results indicate OVVD is a unique and challenging task, especially for the anomaly categorization. Refer to supplementary materials for more ablation studies and qualitative results.

5. Conclusion

In this paper, we present a new model built on top of pre-trained large models for open-vocabulary video anomaly detection task under weak supervision. Owing to the chal-




Figure 4. Confusion matrices of anomaly categorization.

lenging nature of open-vocabulary video anomaly detection, current video anomaly detection approaches face difficulties in working efficiently. To address these unique challenges, we explicitly disentangle open-vocabulary video anomaly detection into the class-agnostic detection and class-specific classification sub-tasks. We then introduce several ad-hoc modules: temporal adapter and semantic knowledge injection modules mainly aim at promoting detection for both base and novel anomalies, novel anomaly synthesis module generates several potential pseudo novel sample to assist the proposed model in perceiving novel anomalies more accurately. Extensive experiments on three public datasets demonstrate the proposed model performs advantageously on open-vocabulary video anomaly detection task. In the future, generating more vivid pseudo anomaly samples in the form of videos with the assistance of AIGC models is yet to be explored.

6. Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62306240, U23B2013), China Postdoctoral Science Foundation (No. 2023TQ0272), National Key R&D Program of China (No.2020AAA0106900), and the Fundamental Research Funds for the Central Universities (No. D5000220431).

References

- [1] Andra Acsintoae, Andrei Florescu, Mariana-Iuliana Georgescu, Tudor Mare, Paul Sumedrea, Radu Tudor Ionescu, Fahad Shahbaz Khan, and Mubarak Shah. Ubnormal: New benchmark for supervised open-set video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20143–20153, 2022. 2, 3, 6
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 4
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 3
- [4] Yang Cong, Junsong Yuan, and Ji Liu. Sparse reconstruction cost for abnormal event detection. In *CVPR 2011*, pages 3449–3456. IEEE, 2011. 2
- [5] Choubo Ding, Guansong Pang, and Chunhua Shen. Catching both gray and black swans: Open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7388–7398, 2022. 2, 3
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [7] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023. 5
- [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. 2, 6
- [9] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 733–742, 2016. 1, 6
- [10] Chao Huang, Chengliang Liu, Jie Wen, Lian Wu, Yong Xu, Qiuping Jiang, and Yaowei Wang. Weakly supervised video anomaly detection via self-guided temporal discriminative transformer. *IEEE Transactions on Cybernetics*, 2022. 2
- [11] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2
- [12] Hyekang Kevin Joo, Khoa Vo, Kashu 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. 3, 6
- [13] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 105–124. Springer, 2022. 2, 3, 7
- [14] Chen Ju, Zeqian Li, Peisen Zhao, Ya Zhang, Xiaopeng Zhang, Qi Tian, Yanfeng Wang, and Weidi Xie. Multi-modal prompting for low-shot temporal action localization. *arXiv preprint arXiv:2303.11732*, 2023.
- [15] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11144–11154, 2023. 2
- [16] 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. 3
- [17] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection—a new baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6536–6545, 2018. 1, 2
- [18] Yu Liu, Huai Chen, Lianghua Huang, Di Chen, Bin Wang, Pan Pan, and Lisheng Wang. Animating images to transfer clip for video-text retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1906–1911, 2022. 5
- [19] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13588–13597, 2021. 2
- [20] Zuhao Liu, Xiao-Ming Wu, Dian Zheng, Kun-Yu Lin, and Wei-Shi Zheng. Generating anomalies for video anomaly detection with prompt-based feature mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24500–24510, 2023. 4
- [21] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal event detection at 150 fps in matlab. In *Proceedings of the IEEE international conference on computer vision*, pages 2720–2727, 2013. 2
- [22] Weixin Luo, Wen Liu, and Shenghua Gao. Remembering history with convolutional lstm for anomaly detection. In *2017 IEEE International Conference on Multimedia and Expo (ICME)*, pages 439–444. IEEE, 2017. 2
- [23] Hui Lv, Zhongqi Yue, Qianru Sun, Bin Luo, Zhen Cui, and Hanwang Zhang. Unbiased multiple instance learning for weakly supervised video anomaly detection. *arXiv preprint arXiv:2303.12369*, 2023. 3, 6

- [24] Sauradip Nag, Xiatian Zhu, Yi-Zhe Song, and Tao Xiang. Zero-shot temporal action detection via vision-language prompting. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, pages 681–697. Springer, 2022. 2
- [25] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 1–18. Springer, 2022. 2, 3, 7
- [26] Minheng Ni, Zitong Huang, Kailai Feng, and Wangmeng Zuo. Imaginarynet: Learning object detectors without real images and annotations. *arXiv preprint arXiv:2210.06886*, 2022. 4
- [27] Yujiang Pu, Xiaoyu Wu, and Shengjin Wang. Learning prompt-enhanced context features for weakly-supervised video anomaly detection. *arXiv preprint arXiv:2306.14451*, 2023. 1, 5
- [28] Jie Qin, Jie Wu, Pengxiang Yan, Ming Li, Ren Yuxi, Xuefeng Xiao, Yitong Wang, Rui Wang, Shilei Wen, Xin Pan, et al. Freeseg: Unified, universal and open-vocabulary image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19446–19455, 2023. 2
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [30] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 5
- [31] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Densclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. 3
- [32] Hanoona Rasheed, Muhammad Uzair Khattak, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Fine-tuned clip models are efficient video learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6545–6554, 2023. 3
- [33] 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. *arXiv preprint arXiv:2306.12041*, 2023. 2
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [35] Mohammad Sabokrou, Mohammad Khaloeei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3379–3388, 2018. 2
- [36] Bernhard Schölkopf, Robert C Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. *Advances in neural information processing systems*, 12, 1999. 2, 6
- [37] Chenrui Shi, Che Sun, Yuwei Wu, and Yunde Jia. Video anomaly detection via sequentially learning multiple pretext tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10330–10340, 2023. 2
- [38] 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, 5, 6
- [39] Shengyang Sun and Xiaojin Gong. Hierarchical semantic contrast for scene-aware video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22846–22856, 2023. 2
- [40] Shengyang Sun and Xiaojin Gong. Long-short temporal co-teaching for weakly supervised video anomaly detection. *arXiv preprint arXiv:2303.18044*, 2023. 2
- [41] Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. Ernie 2.0: A continual pre-training framework for language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8968–8975, 2020. 5
- [42] 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. 3, 6
- [43] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 3
- [44] Guodong Wang, Yunhong Wang, Jie Qin, Dongming Zhang, Xiuguo Bao, and Di Huang. Video anomaly detection by solving decoupled spatio-temporal jigsaw puzzles. In *European Conference on Computer Vision*, pages 494–511. Springer, 2022. 1
- [45] Jue Wang and Anoop Cherian. Gods: Generalized one-class discriminative subspaces for anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8201–8211, 2019. 2
- [46] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. 3
- [47] Zejia Weng, Xitong Yang, Ang Li, Zuxuan Wu, and Yungang Jiang. Transforming clip to an open-vocabulary video model via interpolated weight optimization. *arXiv preprint arXiv:2302.00624*, 2023. 2
- [48] Jianzong Wu, Xiangtai Li, Shilin Xu Haobo Yuan, Henghui Ding, Yibo Yang, Xia Li, Jiangning Zhang, Yunhai Tong,

- Xudong Jiang, Bernard Ghanem, et al. Towards open vocabulary learning: A survey. *arXiv preprint arXiv:2306.15880*, 2023. 2
- [49] Peng Wu and Jing Liu. Learning causal temporal relation and feature discrimination for anomaly detection. *IEEE Transactions on Image Processing*, 30:3513–3527, 2021. 3, 5
- [50] Peng Wu, Jing Liu, and Fang Shen. 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, 2019. 1, 2
- [51] 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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 322–339. Springer, 2020. 2, 3, 6
- [52] Peng Wu, Xiaotao Liu, and Jing Liu. Weakly supervised audio-visual violence detection. *IEEE Transactions on Multimedia*, pages 1674–1685, 2022. 3, 6
- [53] 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 (AAAI)*, 2024. 3
- [54] Qingsen Yan, Tao Hu, Yuan Sun, Hao Tang, Yu Zhu, Wei Dong, Luc Van Gool, and Yanning Zhang. Towards high-quality hdr deghosting with conditional diffusion models. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2023. 2
- [55] Zhiwei Yang, Jing Liu, Zhaoyang Wu, Peng Wu, and Xiaotao Liu. Video event restoration based on keyframes for video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14592–14601, 2023. 2
- [56] Guang Yu, Siqu 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. 2
- [57] 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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 358–376. Springer, 2020. 3
- [58] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. 2
- [59] 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. 3
- [60] Yiru Zhao, Bing Deng, Chen Shen, Yao Liu, Hongtao Lu, and Xian-Sheng Hua. Spatio-temporal autoencoder for video anomaly detection. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1933–1941, 2017. 2
- [61] 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. 3
- [62] 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. 3, 6, 7
- [63] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 3, 5
- [64] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. Anomalyclip: Object-agnostic prompt learning for zero-shot anomaly detection. *arXiv preprint arXiv:2310.18961*, 2023. 2
- [65] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11175–11185, 2023. 2
- [66] Jiawen Zhu, Choubo Ding, Yu Tian, and Guansong Pang. Anomaly heterogeneity learning for open-set supervised anomaly detection. *arXiv preprint arXiv:2310.12790*, 2023. 2
- [67] Yuansheng Zhu, Wentao Bao, and Qi Yu. Towards open set video anomaly detection. In *European Conference on Computer Vision*, pages 395–412. Springer, 2022. 2, 3, 6, 7