

Exploiting Completeness and Uncertainty of Pseudo Labels for Weakly Supervised Video Anomaly Detection

Chen Zhang^{1,2} Guorong Li^{3*} Yuankai Qi⁴ Shuhui Wang^{5,6}
Laiyun Qing³ Qingming Huang^{1,2,3,5} Ming-Hsuan Yang^{7,8}

¹State Key Laboratory of Information Security, Institute of Information Engineering, CAS

²School of Cyber Security, University of Chinese Academy of Sciences

³School of Computer Science and Technology, University of Chinese Academy of Sciences

⁴Australian Institute for Machine Learning, The University of Adelaide

⁵Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, CAS

⁶Peng Cheng Laboratory, ⁷University of California, Merced, ⁸Yonsei University

zhangchen@iie.ac.cn, liguorong@ucas.ac.cn, qykshr@gmail.com,

wangshuhui@ict.ac.cn, {lyqing, qmhuang}@ucas.ac.cn, mhyang@ucmerced.edu

Abstract

Weakly supervised video anomaly detection aims to identify abnormal events in videos using only video-level labels. Recently, two-stage self-training methods have achieved significant improvements by self-generating pseudo labels and self-refining anomaly scores with these labels. As the pseudo labels play a crucial role, we propose an enhancement framework by exploiting completeness and uncertainty properties for effective self-training. Specifically, we first design a multi-head classification module (each head serves as a classifier) with a diversity loss to maximize the distribution differences of predicted pseudo labels across heads. This encourages the generated pseudo labels to cover as many abnormal events as possible. We then devise an iterative uncertainty pseudo label refinement strategy, which improves not only the initial pseudo labels but also the updated ones obtained by the desired classifier in the second stage. Extensive experimental results demonstrate the proposed method performs favorably against state-of-the-art approaches on the UCF-Crime, TAD, and XD-Violence benchmark datasets.

1. Introduction

Automatically detecting abnormal events in videos has attracted increasing attention for its broad applications in intelligent surveillance systems. Since abnormal events are sparse in videos, recent studies are mainly developed within the weakly supervised learning framework [5, 12, 19, 25, 27,

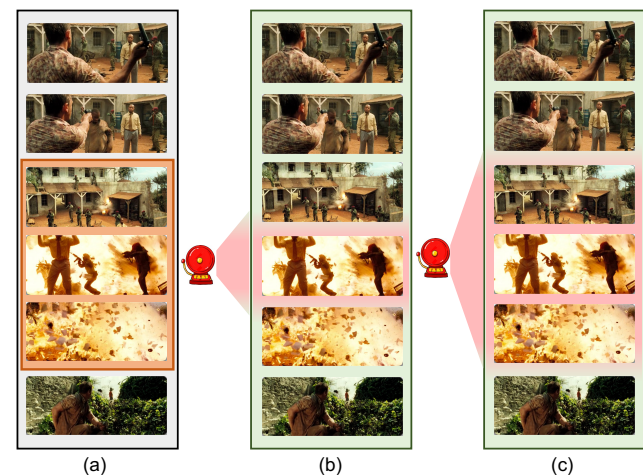


Figure 1. Illustration of the completeness: (a) represents a video that contains multiple abnormal clips (ground truth anomalies are in the orange area). Existing methods tend to focus on the most anomalous clip as shown in (b). We propose to use the multi-head classification module together with a diversity loss to encourage pseudo labels to cover the complete abnormal events as depicted in (c).

29, 32, 34, 37–41], where only video-level annotations are available. However, the goal of anomaly detection is to predict frame-level anomaly scores during test. This results in great challenges for weakly supervised video anomaly detection.

Existing methods broadly fall into two categories: one-stage methods based on Multiple Instance Learning (MIL) and two-stage self-training methods. One-stage MIL-based

*Corresponding author.

methods [19, 27, 29, 39, 41] treat each normal and abnormal video as a negative and positive bag respectively, and clips of a video are the instances of a bag. Formulating anomaly detection as a regression problem, these methods adopt ranking loss to encourage the highest anomaly score in a positive bag to be higher than that in a negative bag. Due to the lack of clip-level annotations, the anomaly scores generated by MIL-based methods are usually less accurate. To alleviate this problem, two-stage self-training methods are proposed [5, 12]. In the first stage, pseudo labels for clips are generated by MIL-based methods. In the second stage, MIST [5] utilizes these pseudo labels to refine discriminative representations. In contrast, MSL [12] refines the pseudo labels via a transformer-based network. Despite progress, existing methods still suffer two limitations. First, the ranking loss used in the pseudo label generator ignores the completeness of abnormal events. The reason is that a positive bag may contain multiple abnormal clips as shown in Figure 1, but MIL is designed to detect only the most likely one. The second limitation is that the uncertainty of generated pseudo labels is not taken into account in the second stage. As the pseudo labels are usually noisy, directly using them to train the final classifier may hamper its performance.

To address these problems, we propose to enhance pseudo labels via exploiting completeness and uncertainty properties. Specifically, to encourage the complete detection of abnormal events, we propose a multi-head module to generate pseudo labels (each head serves as a classifier) and introduce a diversity loss to ensure the distribution difference of pseudo labels generated by the multiple classification heads. In this way, each head tends to discover a different abnormal event, and thus the pseudo label generator covers as many abnormal events as possible. Then, instead of directly training a final classifier with all pseudo labels, we design an iterative uncertainty-based training strategy. We measure the uncertainty using Monte Carlo (MC) Dropout [6] and only clips with lower uncertainty are used to train the final classifier. At the first iteration, we use such uncertainty to refine pseudo labels obtained in the first stage, and in the remaining iterations, we use it to refine the output of the desired final classifier.

The main contributions of this paper are as follows:

- We design a multi-head classifier scheme together with a diversity loss to encourage the pseudo labels to cover as many abnormal clips as possible.
- We design an iterative uncertainty aware self-training strategy to gradually improve the quality of pseudo labels.
- Experiments on UCF-Crime, TAD, and XD-Violence datasets demonstrate the favorable performance compared to several state-of-the-art methods.

2. Related Work

Semi-Supervised Methods. In the semi-supervised setting, only normal videos are required in the training set. Semi-supervised video anomaly detection methods can be divided into one-class classifier-based methods [24, 33, 36] and reconstruction-based methods [1, 3, 4, 9, 16, 21, 22]. In one-class classifier-based methods, the model constructs a boundary that distinguishes normal events from abnormal events by learning information about normal videos. Xu *et al.* [36] adopt an autoencoder to learn appearance and motion features as well as their associations, and then use multiple one-classifiers to predict anomaly scores based on these three learned feature representations. Sabokrou *et al.* [24] propose to train an end-to-end one-classification model in an adversarial manner. For the problem of anomaly detection in complex scenarios, Wu *et al.* [33] propose to jointly optimize representation learning and one-class classification using convolutional neural networks. Reconstruction-based methods aim to minimize the reconstruction error of the training data, and take the minimum error as the threshold for discriminating abnormality. Some [4, 22] learn the dictionary of normal videos, and video clips that cannot be represented by the dictionary are determined to be abnormal. The others [3, 9] learn the rules of normal video sequences through autoencoders, and events that produce higher reconstruction errors are judged as anomalies. To prevent the model from over-fitting normal videos, later works [1, 16, 21] introduce a memory module for recording normal patterns.

Weakly Supervised Methods. Different from the semi-supervised setting, there are both normal and abnormal videos in the training set for weakly supervised methods, but frame-level annotations are not available. Most of the weakly supervised anomaly detection methods are one-stage MIL-based methods. In [27], the first MIL-based method with ranking loss is proposed, along with a large-scale video anomaly detection dataset. Later, Zhang *et al.* [39] propose an inner bag loss, which is complementary to the outer bag ranking loss. To learn a motion-aware feature that can better detect anomalies, Zhu *et al.* [41] use an attention module to take temporal context into the multi-instance ranking model. Tian *et al.* [29] develop a robust top- k MIL method by training a temporal feature magnitude learning function. To effectively utilize the temporal context, Lv *et al.* [19] propose a high-order context encoding model. Zhang *et al.* [38] aim to address the open event issue existing in anomaly detection tasks by using meta-learning. Sapkota *et al.* [25] construct a submodularity diversified MIL loss in a Bayesian non-parametric way, which can satisfy anomaly detection in more realistic settings with outliers and multimodal scenarios.

Recently, two-stage self-training methods are proposed to generate more accurate and fine-grained anomaly scores,

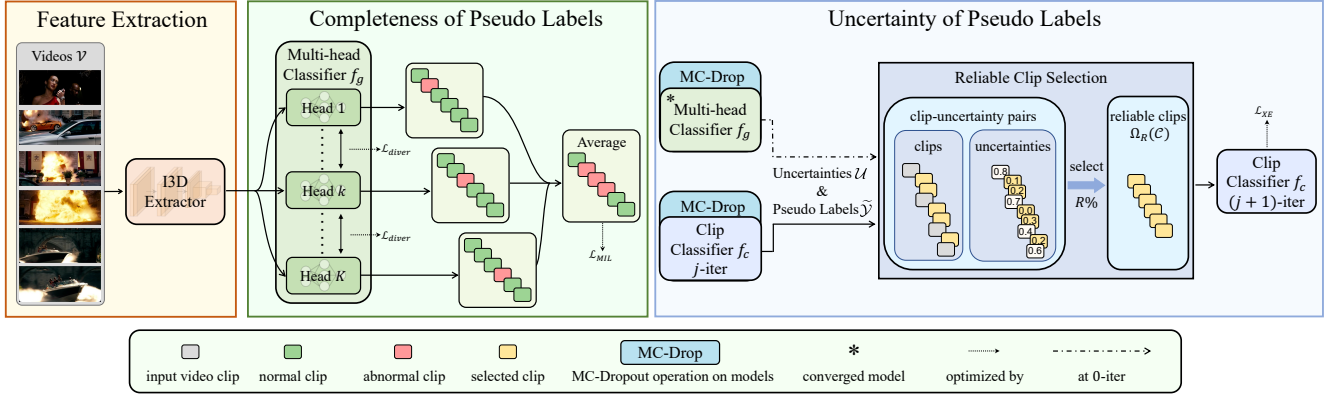


Figure 2. Pipeline of the proposed method: (1) Completeness Enhanced Pseudo Label Generator (Sec. 3.2): First, we use the pre-trained 3D CNN to extract video features. Then the features are fed into the multi-head classifier constrained by a diversity loss to encourage the detection of complete abnormal events. Simultaneously, the MIL ranking loss is used to constrain the anomaly scores of abnormal segments to be larger than those of normal segments. (2) Iterative Uncertainty Aware Pseudo Label Refinement (Sec. 3.3): In the first iteration, we obtain the initial clip-level pseudo labels from the multi-head classifier of the first stage and compute its uncertainty via Monte Carlo Dropout. Then we select reliable clips based on the uncertainty to train a new clip classifier. In the remaining iterations use the new clip classifier to update pseudo labels.

which adopt a two-stage pipeline, *i.e.*, pseudo labels are generated first and then fed into a classification module. Feng *et al.* [5] propose to use the information provided by the multi-instance pseudo label generator to fine-tune the feature encoder to generate more discriminative features specifically for the task of video anomaly detection. Li *et al.* [12] select the sequence consisting of multiple instances with the highest sum of anomaly score as the optimization unit of ranking loss, and gradually reduce the length of the sequence by adopting a self-training strategy to refine the anomaly scores. In addition to the above two categories of methods, there are some fancy ideas for weakly supervised video anomaly detection. Zhong *et al.* [40] reformulate the weakly supervised anomaly detection problem as a supervised learning task under noisy labels. Wu *et al.* [34] propose an audio-visual dataset and design a holistic and localized framework to explicitly model relations of video snippets to learn discriminative feature representations. In [32], Wu *et al.* further explore the importance of temporal relation and discriminative features for weakly supervised anomaly detection.

Self-Training. Self-training is one of the mainstream techniques in semi-supervised learning [11, 23, 30] and has recently shown important progress for tasks like classification [7, 20, 35] and domain adaptation [14, 42]. For self-training, the training data consists of a small amount of labeled data and a large amount of unlabeled data. The core idea is to use the model trained with labeled data to generate pseudo labels of unlabeled data, and then train the model jointly with labeled data and pseudo labels. The training process is repeated until the model converges. In weakly supervised video anomaly detection, Feng *et al.* [5] propose

a self-training framework in which clip-level pseudo labels generated by a multiple instance pseudo label generator are assigned to all clips of abnormal videos to refine feature encoder. Most similar to our work is the Multi-Sequence Learning method proposed by Li *et al.* [12], which refines anomaly scores by gradually reducing the length of selected sequences by self-training. However, the self-training mechanisms used by these methods do not consider the uncertainty of pseudo labels, leading to a gradually deviating self-training process guided by noisy pseudo labels. In contrast, we develop an uncertainty aware self-training strategy that can reduce the effect of unreliable pseudo labels. We also consider the importance of temporal context for video understanding, in an effort to better refine anomaly scores.

3. Method

The proposed pseudo label enhancement framework based on completeness and uncertainty is shown in Figure 2. We first use a multi-head classifier trained with a diversity loss and MIL ranking loss to generate initial clip-level pseudo labels. Then, we utilize an iterative uncertainty aware pseudo label refinement strategy to gradually improve the quality of pseudo labels to train the final desired classifier. In the following, we first formulate the task of weakly supervised video anomaly detection and then elaborate each component of our method.

3.1. Notations and Preliminaries

Assume that we are given a set of N videos $\mathcal{V} = \{\mathcal{V}_i\}_{i=1}^N$ and the ground-truth labels $\mathcal{Y} = \{\mathcal{Y}_i\}_{i=1}^N \in \{1, 0\}$. $\mathcal{Y}_i = 1$

if an abnormal clip is present in the video and $\mathcal{Y}_i = 0$ otherwise. During training, only video-level labels are available. However, in the testing stage, the goal of this task is to generate frame-level anomaly scores to indicate the temporal location of abnormal events. Following previous MIL-based methods [19, 27, 32, 39, 41], for each input video \mathcal{V}_i , we first divide it into 16-frame non-overlapping T_i clips and use a pre-trained 3D convolutional network to extract features, forming a clip feature sequence $\mathcal{C}_i = \{c_{i,1}, c_{i,2}, \dots, c_{i,T_i}\} \in \mathbb{R}^{T_i \times D_V}$, where T_i is the number of extracted video clip features and D_V is the feature dimension. Since long untrimmed videos may contain different numbers of clips, which is inconvenient for training. Therefore, consistent with [27, 29], the video clip features are combined into S temporal video segments denoted as $\mathcal{X}_i = \{x_{i,1}, x_{i,2}, \dots, x_{i,S}\} \in \mathbb{R}^{S \times D_V}$ by averaging multiple consecutive clip features. We treat an abnormal video \mathcal{V}_i^a as a positive bag and a normal video \mathcal{V}_i^n as a negative bag, and treat each segment $x_{i,s}^a$ or $x_{i,s}^n$ as an instance in the bag.

3.2. Completeness of Pseudo Labels

Inspired by [13], which uses a diversity loss to model action completeness, we design a completeness enhanced pseudo label generator composed of parallel multi-head classifier, together with a diversity loss to detect as many abnormal events as possible in a video. Each head $f_g(\cdot; \phi^k)$ is composed of three fully connected layers parameterized by ϕ^k . Taking the video segment features $\mathcal{X}_i = \{x_{i,s}\}_{s=1}^S$ as input, each head outputs the anomaly scores of each segment, which are further passed through a softmax to generate a score distribution:

$$\hat{\mathcal{Y}}_i^k = \text{softmax}(f_g(\mathcal{X}_i; \phi^k)) \quad (1)$$

where $\hat{\mathcal{Y}}_i^k \in \mathbb{R}^{S \times 1}$ denotes the score distribution of the i -th video from the k -th head. The predicted score distributions of K heads are then enforced to be distinct from each other by a diversity loss, which minimizes the cosine similarity of the distribution between any two heads:

$$\mathcal{L}_{diver} = \frac{1}{Z} \sum_{k=1}^{K-1} \sum_{q=k+1}^K \frac{\hat{\mathcal{Y}}_i^k \cdot \hat{\mathcal{Y}}_i^q}{\|\hat{\mathcal{Y}}_i^k\| \|\hat{\mathcal{Y}}_i^q\|} \quad (2)$$

where $Z = \frac{1}{2}K(K-1)$. A regularization term on the norm of the segment score sequences is used to balance multiple heads and to avoid performance degradation due to dominance by one head:

$$\mathcal{L}_{norm} = \frac{1}{K} \sum_{k=1}^K \left| \|\mathcal{A}^k\| - \|\mathcal{A}^{avg}\| \right| \quad (3)$$

where $\mathcal{A}^k = f_g(\mathcal{X}_i; \phi^k)$ denotes the anomaly scores generated by the k_{th} head and \mathcal{A}^{avg} is the average of

the anomaly scores produced by each head: $\mathcal{A}^{avg} = \frac{1}{K} \sum_{k=1}^K (f_g(\mathcal{X}_i; \phi^k))$.

Under the action of the diversity loss and norm regularization, the anomaly scores generated by multiple heads can achieve maximum differentiation and detect different anomalous segments. Finally, \mathcal{A}^{avg} is followed by a sigmoid function to obtain the predicted segment-level labels ranging from 0 to 1:

$$\hat{\mathcal{Y}}_i = \text{sigmoid}(\mathcal{A}^{avg}). \quad (4)$$

where $\hat{\mathcal{Y}}_i = \{\hat{y}_{i,1}, \hat{y}_{i,2}, \dots, \hat{y}_{i,S}\}$ represents the predicted segment-level labels of the i -th video. For the abnormal video \mathcal{V}_i^a , the predicted labels are denoted as $\hat{\mathcal{Y}}_i^a = \{\hat{y}_{i,1}^a, \hat{y}_{i,2}^a, \dots, \hat{y}_{i,S}^a\}$. For the normal video \mathcal{V}_i^n , the predicted labels are denoted as $\hat{\mathcal{Y}}_i^n = \{\hat{y}_{i,1}^n, \hat{y}_{i,2}^n, \dots, \hat{y}_{i,S}^n\}$.

Like [27], the ranking loss is used to constrain the highest anomaly score of abnormal videos to be higher than that of normal videos:

$$\max_{\hat{y}_{i,s}^a \in \hat{\mathcal{Y}}_i^a} \hat{y}_{i,s}^a > \max_{\hat{y}_{i,s}^n \in \hat{\mathcal{Y}}_i^n} \hat{y}_{i,s}^n \quad (5)$$

To maximize the separability between positive and negative instances, a hinge-based ranking loss is used:

$$\mathcal{L}_{MIL} = \max \left(0, 1 - \max_{\hat{y}_{i,s}^a \in \hat{\mathcal{Y}}_i^a} \hat{y}_{i,s}^a + \max_{\hat{y}_{i,s}^n \in \hat{\mathcal{Y}}_i^n} \hat{y}_{i,s}^n \right) \quad (6)$$

Finally, the completeness enhanced pseudo label generator is trained with loss as follows:

$$\mathcal{L}_{f_g} = \mathcal{L}_{MIL} + \alpha \mathcal{L}_{diver} + \alpha \mathcal{L}_{norm} \quad (7)$$

where α is the hyper-parameter to balance the losses.

3.3. Uncertainty of Pseudo Labels

Instead of directly using the clip-level pseudo labels obtained in the first stage (Sec. 3.2) to train the final desired clip classifier f_c , we propose an uncertainty aware self-training strategy to mine clips with reliable pseudo labels. Specifically, we introduce the uncertainty estimation leveraging Monte Carlo (MC) Dropout [6] so that clips with low uncertainty (*i.e.*, reliable) pseudo labels are selected for training f_c . This process is conducted several iterations. Throughout these iterations, the pseudo labels are continuously refined, eventually generating high-quality fine-grained pseudo labels to train the final desired clip classifier f_c . Note that the pseudo labels are initially obtained in the first stage and are then updated by f_c .

Uncertainty Estimation. We use MC-Dropout [6] to estimate the uncertainty of clip-level pseudo labels. For training video clips $\mathcal{C}_i = \{c_{i,t}\}_{t=1}^{T_i}$, we perform M stochastic forward passes through the model f trained with dropout. In the first iteration, we use the multi-head classifier in the

first stage as the trained model (*i.e.*, $f = f_g$). In the remaining iterations, $f = f_c$. Each pass generates clip-level pseudo labels as follows:

$$\hat{\mathcal{Y}}_i^m = \text{sigmoid}\left(f(\mathcal{C}_i; \tilde{W}^m)\right) \quad (8)$$

where \tilde{W}^m denotes the m^{th} sampled masked model parameters and $\hat{\mathcal{Y}}_i^m = \{\hat{y}_{i,t}^m\}_{t=1}^{T_i}$. The clip-level pseudo labels $\tilde{\mathcal{Y}}_i = \{\tilde{y}_{i,t}\}_{t=1}^{T_i}$ used as the supervision for training the clip classifier are given by the predictive mean:

$$E(\hat{\mathcal{Y}}_i) = \frac{1}{M} \sum_{m=1}^M (\hat{\mathcal{Y}}_i^m) \quad (9)$$

The prediction uncertainties $\mathcal{U}_i = \{u_{i,t}\}_{t=1}^{T_i}$ of $\hat{\mathcal{Y}}_i$ are given by the diagonal entries (*i.e.*, the variance of elements in $\hat{\mathcal{Y}}_i$) in the covariance matrix $\text{Cov}(\hat{\mathcal{Y}}_i)$:

$$\text{Cov}(\hat{\mathcal{Y}}_i) = \frac{1}{M} \sum_{m=1}^M \hat{\mathcal{Y}}_i^{m\top} \hat{\mathcal{Y}}_i^m - E(\hat{\mathcal{Y}}_i)^\top E(\hat{\mathcal{Y}}_i) \quad (10)$$

Iterative Reliable Clip Mining. As the goal is to train a reliable model with low-uncertainty pseudo labels, we mine reliable clips after uncertainty estimation. For each video, we rank the uncertainties \mathcal{U}_i of its pseudo labels from small to large and remain the clips with the least $R\%$ uncertainty scores, where $R\%$ represents the sample ratio. In this way, reliable video clips and corresponding clip-level pseudo labels can be mined.

Since participating in training with only clip-level features ignores the contextual relationship of the video, we use a long-term feature memory [31] to model the temporal relationship between video clips. The clip features of each video are stored in a memory pool. After obtaining the selected reliable clip, we retrieve the window size w clip features $\mathcal{H}_{i,t} = [c_{i,t-w}, \dots, c_{i,t-1}]$ before the current clip $c_{i,t}$ from the memory pool. We obtain $\tilde{\mathcal{H}}_{i,t}$ by performing mean pooling on $\mathcal{H}_{i,t}$, and then concatenate the current clip feature $c_{i,t}$ with $\tilde{\mathcal{H}}_{i,t}$ into a new temporal clip feature $\bar{c}_{i,t}$. Thus we can use all reliable temporal features set $\Omega_R(\mathcal{C})$ and reliable clip-level pseudo labels $\Omega_R(\tilde{\mathcal{Y}})$ to train the clip classifier f_c based on the binary cross entropy loss:

$$\mathcal{L}_{XE} = \sum_{\bar{c}_{i,t} \in \Omega_R(\mathcal{C})} -(\tilde{y}_{i,t} \log(f_c(\bar{c}_{i,t})) + (1 - \tilde{y}_{i,t}) \log(1 - f_c(\bar{c}_{i,t}))) \quad (11)$$

where $\tilde{y}_{i,t} \in \Omega_R(\tilde{\mathcal{Y}})$ represents the pseudo label of the t_{th} clip in the i_{th} video. Then we can obtain the clip classifier trained in the current iteration, and perform uncertainty estimation and reliable sample selection in the next iteration to further train the desired clip classifier.

Algorithm 1 Completeness-and-Uncertainty Aware Pseudo Label Enhancement

Input: A set of N videos $\mathcal{V} = \{\mathcal{V}_i\}_{i=1}^N$ and video-level labels $\mathcal{Y} = \{\mathcal{Y}_i\}_{i=1}^N$.

Output: Clip classifier f_c .

- 1: // Completeness Enhanced Pseudo Label Generator.
 - 2: Extract T_i clip features for each video \mathcal{V}_i as $\{c_{i,t}\}_{t=1}^{T_i}$.
 - 3: Combine $\{c_{i,t}\}_{t=1}^{T_i}$ into S segment features $\{x_{i,s}\}_{s=1}^S$.
 - 4: Training multi-head classifier f_g with $\{x_{i,s}\}_{s=1}^S$ and \mathcal{Y}_i via Eq. 7.
 - 5: Obtain the trained model $f = f_g$.
 - 6: // Uncertainty Aware Pseudo Label Refinement.
 - 7: **while** not converged **do**
 - 8: Leverage MC-Dropout for the trained model f .
 - 9: Get pseudo labels via Eq. 9, $\tilde{\mathcal{Y}}_i = E(\hat{\mathcal{Y}}_i)$.
 - 10: Get uncertainties via Eq. 10, $\mathcal{U}_i = \text{diag}(\text{Cov}(\hat{\mathcal{Y}}_i))$.
 - 11: $(\mathcal{U}_i)_{\text{sorted}} \leftarrow \text{SORT}(\mathcal{U}_i)$. // ascending sort
 - 12: Remain clips with the least $R\%$ uncertainty scores.
 - 13: Model temporal relationship for the remained clips.
 - 14: All reliable clip set $\Omega_R(\mathcal{C}) = \Omega_R(\{\mathcal{C}_i\}_{i=1}^N)$.
 - 15: All reliable label set $\Omega_R(\tilde{\mathcal{Y}}) = \Omega_R(\{\tilde{\mathcal{Y}}_i\}_{i=1}^N)$.
 - 16: Training f_c with $\Omega_R(\mathcal{C})$ and $\Omega_R(\tilde{\mathcal{Y}})$ via Eq. 11.
 - 17: Obtain the trained model $f = f_c$.
 - 18: **end while**
 - 19: Return clip classifier f_c for inference.
-

3.4. Model Training and Inference

Training. We increase completeness of pseudo labels with Eq. 7 to cover as many abnormal clips as possible in the first stage. In the second stage, we mine reliable video clips with uncertainty estimation to train the clip classifier using Eq. 11 and gradually refine clip-level pseudo labels through multiple iterations. Algorithm 1 shows the main steps of the training process.

Inference. Given a test video, we directly utilize the clip classifier f_c to predict anomaly scores.

4. Experimental Results

We perform experiments on three publicly available datasets including UCF-Crime [27], TAD [19] and XD-Violence [34].

4.1. Datasets and Evaluation Metrics

Datasets. UCF-Crime is a large-scale benchmark dataset for video anomaly detection with 13 anomaly categories. The videos are captured from diverse scenes, such as streets, family rooms, and shopping malls. The dataset contains 1610 training videos annotated with video-level labels and 290 test videos with frame-level annotation. TAD is a re-

cently released dataset in traffic scenarios. It covers 7 real-world anomalies in 400 training videos and 100 test videos. In line with the weak supervision setting in [27], the training set is annotated with video-level labels, and the test set provides frame-level labels. XD-Violence is the largest dataset currently used for weakly supervised anomaly detection. Its videos are collected through multiple channels, such as movies, games, and car cameras. It contains 3954 training videos with video-level labels and 800 test videos with frame-level labels, covering 6 anomaly categories. Moreover, it provides audio-visual signals, enabling anomaly detection by leveraging multimodal cues.

Evaluation Metrics. Similar to the previous works [5, 27, 39, 40], we choose the area under the curve (AUC) of the frame-level receiver operating characteristic curve to evaluate the performance of our proposed method on UCF-Crime and TAD datasets. For XD-Violence, following [12, 29, 32, 34], we use average precision (AP) as the metric.

4.2. Implementation Details

Feature Extractor. Consistent with existing methods [29], we use pre-trained I3D [2] model to extract clip features from 16 consecutive frames. Then we divide each video into 32 segments as the input to the multi-head classifier, *i.e.*, $S = 32$. For the XD-violence dataset, following the setting in [32, 34], we extracted audio features by leveraging the VGGish network [8].

Training Details. The head number K of the multi-head classifier is 2 for all datasets. Each head consists of three fully connected (FC) layers with 512, 128, and 1 units, respectively. The first and third fully connected layers are followed by ReLU activation and sigmoid activation, respectively. We use a dropout function between FC layers with a dropout rate of 0.6. The classifier is trained using the Adadelta optimizer with a learning rate of 0.1. The parameter α is set to 10 for UCF-Crime, and 0.1 for TAD and XD-Violence. The number of stochastic forward passes M is set to 50 for all datasets when generating the initial pseudo labels and uncertainty scores. During pseudo label refinement, the clip classifier consists of three FC layers and 80% of the dropout regulation is used between FC layers. It is trained using the Adam optimizer with a learning rate of $1e - 4$ and a weight decay of $5e - 4$. At the end of each iteration, we set $M = 5$ to obtain pseudo labels and uncertainty scores.

4.3. Comparisons with Prior Work

Results on UCF-Crime. Table 1 summarizes the performance comparison between the proposed method and other methods on the UCF-Crime test set. From the results, we can see that our proposed method outperforms all the previous semi-supervised methods [9, 10, 17, 28] and weakly su-

Supervised	Methods	Feature	AUC(%)
Semi	Hasan <i>et al.</i> [9]	AE	50.60
	Ionescu <i>et al.</i> [10]	-	61.60
	Lu <i>et al.</i> [17]	Dictionary	65.51
	Sun <i>et al.</i> [28]	-	72.70
Weakly	Binary classifier	C3D	50.00
	Sultani <i>et al.</i> [27]	C3D	75.41
	Zhang <i>et al.</i> [39]	C3D	78.66
	Zhu <i>et al.</i> [41]	AE	79.00
	GCN [40]	TSN	82.12
	CLAWS [37]	C3D	83.03
	Wu <i>et al.</i> [34]	I3D	82.44
	MIST [5]	I3D	82.30
	Wu <i>et al.</i> [32]	I3D	84.89
	WSAL [19]	TSN	85.38
	RTFM [29]	I3D	84.30
	MSL [12]	I3D	85.30
	BN-SVP [25]	I3D	83.39
	Ours	I3D	86.22

Table 1. Comparison with other methods on UCF-Crime.

Supervised	Methods	Feature	AUC(%)
Semi	Luo <i>et al.</i> [18]	TSN	57.89
	Liu <i>et al.</i> [15]	-	69.13
Weakly	Sultani <i>et al.</i> [27]	C3D	83.27
	Sultani <i>et al.</i> [27]	TSN	85.95
	Sultani <i>et al.</i> [27]	I3D	88.34
	Zhu <i>et al.</i> [41]	TSN	83.08
	MIST [5]	I3D	89.26
	WSAL [19]	TSN	89.64
	RTFM [29]	I3D	89.64
Ours	I3D	91.66	

Table 2. Comparison with other methods on TAD.

pervised methods [5, 12, 19, 25, 27, 29, 32, 34, 37, 39–41]. Using the same I3D-RGB features, an absolute gain of 0.92% is achieved in terms of the AUC when compared to the best previous two-stage method [12].

Results on TAD. The performance comparisons on TAD are shown in Table 2. Compared with the previous semi-supervised approaches [15, 18] and weakly supervised methods [5, 19, 27, 29, 41], our method achieves the superior performance on AUC. Remarkably, with the same I3D features, our method is 2.40% and 2.02% better than two-stage MIST [5] and one-stage RTFM [29].

Results on XD-Violence. As shown in Table 3, our method also achieves favorable performance when compared with previous semi-supervised approaches [9, 26] and weakly su-

Supervised	Methods	Feature	AP(%)
Semi	OCSVM [26]	I3D+VGGish	27.25
	Hasan <i>et al.</i> [9]	I3D+VGGish	30.77
Weakly	SVM baseline	I3D+VGGish	50.78
	Sultani <i>et al.</i> [27]	I3D+VGGish	73.20
	Wu <i>et al.</i> [34]	I3D	75.41
	Wu <i>et al.</i> [34]	I3D+VGGish	78.64
	Wu <i>et al.</i> [32]	I3D+VGGish	75.90
	RTFM [29]	I3D	77.81
	MSL [12]	I3D	78.28
	Ours	I3D	78.74
	Ours	I3D+VGGish	81.43

Table 3. Comparison with other methods on XD-Violence.

Baseline	Completeness	Uncertainty	AUC(%)
✓			82.86
✓	✓		84.89
✓		✓	84.69
✓	✓	✓	86.22

Table 4. Ablation results on UCF-Crime.

pervised methods [12, 27, 29, 32, 34] on XD-Violence. Using the same I3D features, our method achieves new state-of-the-art performance of 78.74% AP. Following Wu *et al.* [34], we show the results when using the fused features of I3D and VGGish. Our method gains clear improvements against Wu *et al.* [34] by 2.79% in terms of AP. Moreover, the comparison results between using only I3D features and using multi-modal features (*i.e.*, I3D and VGGish) demonstrate that audio can provide more useful information for anomaly detection.

4.4. Ablation Study

We conduct multiple ablation studies on the UCF-Crime dataset to analyze how each component in our method influences the overall performance and show the results in Table 4. We start with the baseline that directly uses the MIL-based method as a pseudo label generator and then selects all the pseudo labels to train a clip-level classifier.

Effectiveness of Completeness of Pseudo Labels. To demonstrate the necessity of exploiting the completeness property, we compare our method that uses a multi-head classifier constrained by diversity loss as pseudo label generator (the 2nd row) with the baseline (the 1st row) on the UCF-Crime dataset. The results show a significant improvement of 2.03% in terms of AUC, which proves that taking the completeness of pseudo labels into consideration can effectively enhance the quality of pseudo labels and thus improve the anomaly detection performance.



Figure 3. Performance comparison of training strategy between Uncertainty Aware Self-Training and Standard Self-Training.

K	1	2	3	4	5
AUC(%)	84.69	86.22	85.05	85.90	85.10

Table 5. Performance comparison of different head numbers.

Effectiveness of Uncertainty of Pseudo Labels. To investigate the effect of exploiting the uncertainty property, we conduct experiments (the 3rd row) with only iterative uncertainty aware pseudo label refinement strategy added to the baseline. From the results, we can see that uncertainty property can bring a performance gain of 1.83% in AUC, which indicates that using an uncertainty aware self-training strategy to gradually improve the quality of pseudo labels is important.

Analysis on the Uncertainty Aware Self-Training. To look deeper into the proposed uncertainty aware self-training strategy, we also make a comparison with standard self-training mechanism that randomly selects samples from the new labeled set. As shown in Figure 3, with the number of iterations increasing, the performance of the uncertainty aware self-training strategy increases rapidly and then stabilizes gradually, whereas the performance improvement of the standard self-training strategy is much smaller. This shows that adopting the uncertainty aware self-training strategy to select reliable samples in each iteration is more beneficial for training the desired classifier.

4.5. Hyperparameter Analysis

Effect of the head number K . In Table 5, we show the comparative result of varying the head number K of the multi-head classifier on the UCF-Crime dataset. When $K = 1$, the pseudo label generator is a single-head classifier, which limits the ability of pseudo labels to cover complete anomalies. As the head number increase from 2 to 5, the performance of multi-head classifier models all exceed that of single-head classifier model and our method yields the best performance when $K = 2$. Since the performance difference between models with head numbers from 2 to 5 is not so significant, we set $K = 2$ for all datasets.

α	0.01	0.1	1	10	100
AUC(%)	84.56	84.43	84.82	86.22	75.06

Table 6. Performance comparison of different diversity loss weights.

R	0.1	0.3	0.5	0.7	0.9
AUC(%)	85.13	85.25	86.22	85.58	85.71

Table 7. Performance comparison of different sample ratios.

Effect of the diversity loss weight α . To explore the influence of the diversity loss weight, we conduct experiments on the UCF-Crime dataset and report the AUC with different diversity loss weights. As shown in Table 6, we report the results as α increases from 0.01 to 100. The AUC can be consistently improved as α grows from 0.01 to 10 and then decreases when α increases to 100, which means 10 is sufficient to model the completeness of pseudo labels on the UCF-Crime dataset. For the TAD and XD-Violence datasets, the weight of the diversity loss α is set to 0.1.

Effect of self-training iterations j . In Figure 3, we report the AUC of our proposed method with different iterations on the UCF-Crime dataset. In general, we observe the performance to improve rapidly initially as j increases, and gradually converge in 3 – 5 iterations. Considering efficiency and performance comprehensively, we set $j = 3$ for UCF-Crime and XD-Violence datasets. Due to the relatively small scale of the TAD dataset, we set $j = 5$ for TAD.

Effect of the sample ration R . Table 7 reports the experimental results evaluated with different sample ratios on the UCF-Crime dataset. We observe that setting R to 0.5 is an optimum choice. As R increases, the AUC improves rapidly and then declines gradually. According to our analysis, the sampling ratio of reliable samples should depend on the scale of the training set. With a large-scale training set, it is more likely to generate high-quality pseudo labels, so a larger sampling ratio can be set. On the contrary, a smaller sampling ratio should be set. Therefore, for the large-scale XD-Violence dataset, we set $R = 0.7$. For the small-scale TAD dataset, R is set to 0.3.

4.6. Qualitative Results

In order to further prove the effectiveness of our method, we visualize the anomaly score results on UCF-Crime, TAD, and XD-Violence datasets in Figure 4. As shown in Figure 4(a) and Figure 4(b), our method can predict relatively accurate anomaly scores for the multi-segment abnormal event (Explosion) in the UCF-Crime dataset and the long-term abnormal event (PedestrianOnRoad) in the TAD

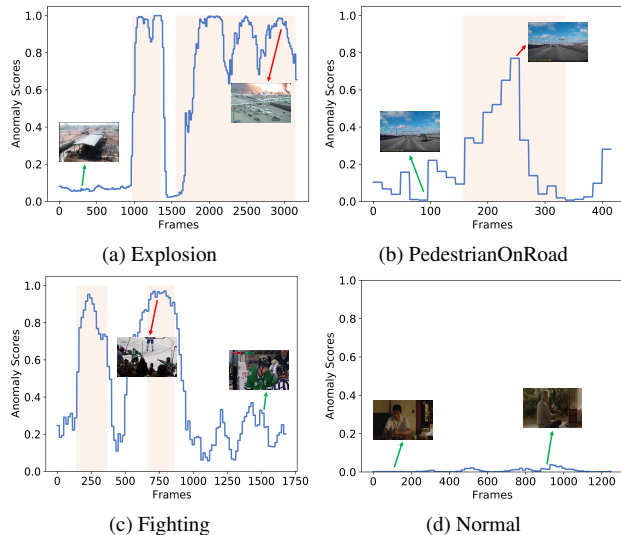


Figure 4. Qualitative results on UCF-Crime (Explosion), TAD (PedestrianOnRoad) and XD-Violence (Fighting, Normal) test videos. The pink square area is the interval where abnormal events occur in the video, and the blue curve indicates the predicted anomaly scores of each video frame.

dataset. Figure 4(c) and Figure 4(d) depict the anomaly scores of the abnormal event (Fighting) and normal event in the XD-Violence dataset, our method can completely detect two anomalous intervals in the abnormal video and predict anomaly scores close to 0 for the normal video.

5. Conclusions

In this paper, we focus on enhancing the quality of pseudo labels and propose a two-stage self-training method that exploits completeness and uncertainty properties. First, to enhance the completeness of pseudo labels, a multi-head classification module constrained by a diversity loss is designed to generate pseudo labels that can cover as many anomalous events as possible. Then, an iterative uncertainty aware self-training strategy is employed to select reliable samples to train the clip classifier. The output of the clip classifier is gradually refined through multiple iterations, resulting in high-quality pseudo labels to train the final desired clip classifier. Extensive experiments show the effectiveness of our method.

Acknowledgements. This work was supported in part by the Italy–China Collaboration Project TALENT under Grant 2018YFE0118400; in part by the National Natural Science Foundation of China under Grant 62236008, 61836002, 62272438, 61976069, U21B2038, 61931008, 62022083, and 61872333; and in part by the Fundamental Research Funds for Central Universities.

References

- [1] Ruichu Cai, Hao Zhang, Wen Liu, Shenghua Gao, and Zhifeng Hao. Appearance-motion memory consistency network for video anomaly detection. In *AAAI*, pages 938–946, 2021. [2](#)
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. [6](#)
- [3] Yong Shean Chong and Yong Haur Tay. Abnormal event detection in videos using spatiotemporal autoencoder. In *ISNN*, pages 189–196, 2017. [2](#)
- [4] Yang Cong, Junsong Yuan, and Ji Liu. Sparse reconstruction cost for abnormal event detection. In *CVPR*, pages 3449–3456, 2011. [2](#)
- [5] Jia-Chang Feng, Fa-Ting Hong, and Wei-Shi Zheng. MIST: Multiple instance self-training framework for video anomaly detection. In *CVPR*, pages 14009–14018, 2021. [1](#), [2](#), [3](#), [6](#)
- [6] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, pages 1050–1059, 2016. [2](#), [4](#)
- [7] Kirill Gavrilyuk, Mihir Jain, Ilia Karmanov, and Cees GM Snoek. Motion-augmented self-training for video recognition at smaller scale. In *CVPR*, pages 10429–10438, 2021. [3](#)
- [8] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*, pages 776–780, 2017. [6](#)
- [9] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *CVPR*, pages 733–742, 2016. [2](#), [6](#), [7](#)
- [10] Radu Tudor Ionescu, Fahad Shahbaz Khan, Mariana-Iuliana Georgescu, and Ling Shao. Object-centric auto-encoders and dummy anomalies for abnormal event detection in video. In *CVPR*, pages 7842–7851, 2019. [6](#)
- [11] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICMLW*, page 896, 2013. [3](#)
- [12] Shuo Li, Fang Liu, and Licheng Jiao. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. In *AAAI*, pages 1395–1403, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [13] Daochang Liu, Tingting Jiang, and Yizhou Wang. Completeness modeling and context separation for weakly supervised temporal action localization. In *CVPR*, pages 1298–1307, 2019. [4](#)
- [14] Hong Liu, Jianmin Wang, and Mingsheng Long. Cycle self-training for domain adaptation. In *NeurIPS*, pages 22968–22981, 2021. [3](#)
- [15] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection—a new baseline. In *CVPR*, pages 6536–6545, 2018. [6](#)
- [16] 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 *CVPR*, pages 13588–13597, 2021. [2](#)
- [17] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal event detection at 150 fps in matlab. In *ICCV*, pages 2720–2727, 2013. [6](#)
- [18] Weixin Luo, Wen Liu, and Shenghua Gao. A revisit of sparse coding based anomaly detection in stacked rnn framework. In *ICCV*, pages 341–349, 2017. [6](#)
- [19] Hui Lv, Chuanwei Zhou, Zhen Cui, Chunyan Xu, Yong Li, and Jian Yang. Localizing anomalies from weakly-labeled videos. *IEEE TIP*, 30:4505–4515, 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [20] Subhabrata Mukherjee and Ahmed Awadallah. Uncertainty-aware self-training for few-shot text classification. In *NeurIPS*, pages 21199–21212, 2020. [3](#)
- [21] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning memory-guided normality for anomaly detection. In *CVPR*, pages 14372–14381, 2020. [2](#)
- [22] Huamin Ren, Weifeng Liu, Søren Ingvor Olsen, Sergio Escalera, and Thomas B Moeslund. Unsupervised behavior-specific dictionary learning for abnormal event detection. In *BMVC*, pages 28–1, 2015. [2](#)
- [23] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S. Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *ICLR*, 2021. [3](#)
- [24] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *CVPR*, pages 3379–3388, 2018. [2](#)
- [25] Hitesh Sapkota and Qi Yu. Bayesian nonparametric submodular video partition for robust anomaly detection. In *CVPR*, pages 3212–3221, 2022. [1](#), [2](#), [6](#)
- [26] Bernhard Schölkopf, Robert C. Williamson, Alexander J. Smola, John Shawe-Taylor, and John C. Platt. Support vector method for novelty detection. In *NeurIPS*, pages 582–588, 1999. [6](#), [7](#)
- [27] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *CVPR*, pages 6479–6488, 2018. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [28] Che Sun, Yunde Jia, Yao Hu, and Yuwei Wu. Scene-aware context reasoning for unsupervised abnormal event detection in videos. In *ACM MM*, pages 184–192, 2020. [6](#)
- [29] 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 *ICCV*, pages 4975–4986, 2021. [1](#), [2](#), [4](#), [6](#), [7](#)
- [30] Chen Wei, Kihyuk Sohn, Clayton Mellina, Alan Yuille, and Fan Yang. CReST: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In *CVPR*, pages 10857–10866, 2021. [3](#)
- [31] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, pages 284–293, 2019. [5](#)
- [32] Peng Wu and Jing Liu. Learning causal temporal relation and feature discrimination for anomaly detection. *IEEE TIP*, 30:3513–3527, 2021. [1](#), [3](#), [4](#), [6](#), [7](#)

- [33] Peng Wu, Jing Liu, and Fang Shen. A deep one-class neural network for anomalous event detection in complex scenes. *IEEE TNNLS*, 31(7):2609–2622, 2020. [2](#)
- [34] 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 *ECCV*, pages 322–339, 2020. [1](#), [3](#), [5](#), [6](#), [7](#)
- [35] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, pages 10687–10698, 2020. [3](#)
- [36] Dan Xu, Elisa Ricci, Yan Yan, Jingkuan Song, and Nicu Sebe. Learning deep representations of appearance and motion for anomalous event detection. In *BMVC*, pages 8.1–8.12, 2015. [2](#)
- [37] 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 *ECCV*, pages 358–376, 2020. [1](#), [6](#)
- [38] Chen Zhang, Guorong Li, Qianqian Xu, Xinfeng Zhang, Li Su, and Qingming Huang. Weakly supervised anomaly detection in videos considering the openness of events. *IEEE TITS*, 23(11):21687–21699, 2022. [1](#), [2](#)
- [39] Jiangong Zhang, Laiyun Qing, and Jun Miao. Temporal convolutional network with complementary inner bag loss for weakly supervised anomaly detection. In *ICIP*, pages 4030–4034, 2019. [1](#), [2](#), [4](#), [6](#)
- [40] 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 *CVPR*, pages 1237–1246, 2019. [1](#), [3](#), [6](#)
- [41] Yi Zhu and Shawn D. Newsam. Motion-aware feature for improved video anomaly detection. In *BMVC*, page 270, 2019. [1](#), [2](#), [4](#), [6](#)
- [42] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, pages 5982–5991, 2019. [3](#)