Cross-Domain Learning for Video Anomaly Detection with Limited Supervision

Yashika Jain¹, Ali Dabouei^{2*}, and Min Xu^{2*}

¹ University of Delhi ² Carnegie Mellon University yashikajain201@gmail.com, ali.dabouei@gmail.com, mxul@cs.cmu.edu

Abstract. Video Anomaly Detection (VAD) automates the identification of unusual events, such as security threats in surveillance videos. In real-world applications, VAD models must effectively operate in cross-domain settings, identifying rare anomalies and scenarios not well-represented in the training data. However, existing cross-domain VAD methods focus on unsupervised learning, resulting in performance that falls short of real-world expectations. Since acquiring weak supervision, *i.e.*, video-level labels, for the source domain is cost-effective, we conjecture that combining it with external unlabeled data has notable potential to enhance cross-domain performance. To this end, we introduce a novel weaklysupervised framework for Cross-Domain Learning (CDL) in VAD that incorporates external data during training by estimating its prediction bias and adaptively minimizing that using the predicted uncertainty. We demonstrate the effectiveness of the proposed CDL framework through comprehensive experiments conducted in various configurations on two large-scale VAD datasets: UCF-Crime and XD-Violence. Our method significantly surpasses the state-of-the-art works in crossdomain evaluations, achieving an average absolute improvement of 19.6% on UCF-Crime and 12.87% on XD-Violence.

1 Introduction

Video anomaly detection (VAD) aims to locate anomalous events in the videos [3, 10, 11, 14, 20, 24, 30, 31, 39, 44]. Unlike manual surveillance, which is costly and timeconsuming, video anomaly detection eliminates the need for extensive human effort, saving resources and time. It holds significant potential for playing a vital role in video surveillance by identifying unusual behaviors and activities such as accidents, burglaries, explosions, and other events that signal security threats.

VAD has been extensively studied previously [11, 14, 20, 30, 31, 44]. Owing to the high costs and time associated with obtaining frame-level labels, most approaches formulate the problem as either an unsupervised [10, 14, 20] or weakly-supervised learning setup [11, 30, 31]. In the unsupervised or one-class classification-based) learning setup, only normal videos are used to model the underlying distribution of normal spatiotemporal patterns, and any deviations from the modeled distribution are regarded as anomalies. Despite the convenience of the unsupervised setup, the lack of anomalous

^{*} Corresponding authors.



Fig. 1: Anomaly score comparison on a video of XD-Violence dataset, with and without employing the proposed CDL framework. The model trained without CDL on UCF-Crime as the weakly labeled set consistently yields high anomaly scores. In contrast, the model trained with CDL, using UCF-Crime as the weakly labeled set and HACS as the unlabeled set, is better able to localize the anomalous frames.

Method(s)	Sup. on \mathcal{D}	Target
Acsintoae <i>et al.</i> [1] rGAN [22], MPN [24] zxVAD [3] Ours	unsup. unsup. unsup. weakly-sup.	$\mathcal{D} \\ \mathcal{D}' \\ \mathcal{D} \cup \mathcal{D}' \\ \mathcal{D} \cup \mathcal{D}' $

Table 1: Brief overview of the taxonomy of current works for VAD using a source domain dataset (\mathcal{D}) and a secondary domain dataset (\mathcal{D}') . All these methods do not utilize any labels for training on (\mathcal{D}') and assume distinct distributions for \mathcal{D} and \mathcal{D}' .

videos during training limits the model's ability to learn the specific characteristics of anomalies. This results in limited performance which does not meet real-world expectations. To address this issue, weakly-supervised setup has attracted significant attention. In this setup, merely video-level labels indicating the presence of anomalies within the videos are incorporated as weak supervision to train models capable of making framelevel predictions at inference. Multiple Instance Learning (MIL) [30] is a prominent technique in this domain. By treating each video as a "bag" and each snippet as a "segment", MIL-based algorithms operate under the premise of a worst-case scenario where the segment with the highest predicted probability of being abnormal is considered as the candidate to represent the whole video.

In real-world applications, it is inevitable to encounter environments and scenarios not fully represented in the model's training set. However, it is essential that the model makes correct predictions in such novel situations. For instance, when the training data lacks samples of rare events like "riots" or accidents in novel scenes, the model should be able to characterize such occurrences as anomalous when they occur. Previous works study these novel situations under the cross-domain problem definition [3, 13, 22].

Existing cross-domain VAD methods [3,13,22,24] rely on unsupervised techniques and consequently exhibit limited performance, as demonstrated later in our empirical evaluations in Tables 2 and 3. A solution to this could be the adoption of weaklysupervised techniques for cross-domain VAD. While weakly-supervised approaches have proven promising in single-domain scenarios [11, 30, 31], their effectiveness in cross-domain scenarios has not been extensively explored. Our evaluations in Tables 2 and 3 suggest that directly employing existing weakly-supervised methods to address the cross-domain challenges results in a significant performance drop when tested in scenarios of even similar nature, such as surveillance videos. We argue that this performance gap is due to the following reasons. First, anomalous events, by their very nature, lack a specific pattern or predefined structure. Hence, the definition of anomaly is context-dependent and a naive adaptation of the previous method cannot capture the context-dependencies in multiple domains. Second, anomalous events are relatively infrequent, making VAD a class imbalance problem. This issue becomes more severe when dealing with multiple domains. Third, because of the limited amount of weakly labeled training data, the model's learning capacity to detect novel (open-set) anomalies is also constrained. Due to these challenges, weakly-supervised methods cannot be readily applied to cross-domain or cross-dataset scenarios.

To overcome these challenges and develop a generalized VAD model, substantial amounts of weakly-labeled data are required. However, acquiring even video-level labels for a large number of videos is inefficient and labor-intensive. On the other hand, vast streams of unlabeled videos are generally available. Utilizing the limited weaklylabeled data alongside this abundant unlabeled data provides a notable opportunity to address the aforementioned challenges in cross-domain VAD. Prudent utilization of the unlabeled data can provide valuable insights into the underlying data distribution, leading to improved decision-making and identification of anomalous events.

To this end, we propose a weakly-supervised Cross-Domain Learning (CDL) framework for VAD that integrates external, unlabeled data, from the wild with limited weaklylabeled data to provide competitive generalization across the domains. This is achieved by adaptively minimizing the prediction bias over the external data using the estimated prediction variance, which serves as an uncertainty regularization score. In the proposed framework, we first train fine-grained pseudo-label generation models on the weaklylabeled data to obtain sets of segment-level predictions for the external dataset. Second, we compute the variance of the predictions across multiple predictors as a proxy to represent uncertainty associated with the segments in the external data. Third, during the optimization process, involving training on both labeled and external data, we adaptively reweigh the bias on each external data using the uncertainty regularization scores. This dynamic reweighing ensures that segments from the external dataset closer to the source dataset are emphasized during the training, while those with higher uncertainty are down-weighted. Finally, we iteratively regenerate pseudo-labels using the models trained on labeled and pseudo-labeled data, re-estimate the uncertainties, and re-train the model on the union of labeled and external datasets. This iterative process helps refine the pseudo-labels as the training progresses. With this training process, the model learns to generalize to both source and external data, given only supervision on the source data. Figure 1 illustrates the effectiveness of the CDL framework.

To summarize, we make the following contributions:

- We present a practical CDL framework for weakly-supervised VAD, in which unlabeled external videos are employed to enhance the cross-domain generalization of the model.
- We design a novel uncertainty quantification method that enables the adaptive uncertainty-driven integration of external videos into the training set.
- Through extensive experiments and ablation studies on benchmark datasets, we validate the proposed approach, demonstrating state-of-the-art performance in crossdomain settings while retaining a competitive performance on the in-domain data.

2 Related Works

Video Anomaly Detection (VAD). VAD is a well-established problem, with most works formulating it either as unsupervised learning [14, 20, 21, 38, 41] or weakly-supervised learning [28, 30, 31, 40, 45] problem. In unsupervised setups, the training data consists solely of normal videos, with the majority of works encoding normal patterns through techniques like frame reconstruction [14, 36], future frame prediction [20], dictionary learning [21, 41], and one-class classification [16, 23]. Any deviation from the encoded patterns is considered anomalous. Since the model categorizes anything beyond its learned representations as anomalous, it can label novel video actions and scenarios encountered during training but in altered environments as anomalous. Weakly-supervised VAD methods help mitigate these issues by incorporating video-level labels as weak supervision for the model, with the majority of methods utilizing the Multiple Instance Ranking Loss [11, 30, 33, 44].

Given that a VAD model is expected to encounter previously unseen scenarios during deployment, it is of paramount importance for the model to have a high generalization across domains. Previous works refer this as cross-domain [3] or cross-dataset generalization [9]. We provide an overview of the existing works employing external data in VAD in Table 1. Previous works on cross-domain generalization focus on unsupervised methods based on few-shot target-domain scene adaptation. [22,24] employ data from the target domain via meta-learning to adapt to that specific domain. Aich et al. [3] proposed a zero-shot target domain adaptation method that incorporates external data to generate pseudo-abnormal frames. Despite the intriguing setup, these unsupervised cross-domain generalization methods lack explicit knowledge about what constitutes an anomaly, hindering the model's ability to learn the specific characteristics of anomalies. To this end, we propose the use of weakly-supervised learning for crossdomain generalization. We integrate external datasets from diverse domains to enable the cross-domain generalization of a model trained in a weakly-supervised fashion. **Pseudo-Labeling and Self-training.** Pseudo-labeling [4, 27] is a common technique where the model trained on labeled data assigns labels to unlabeled data. Subsequently, the model is trained on both the initially labeled data and the pseudo-labeled data. This self-training strategy [25, 37] operates iteratively, allowing the model to progressively enhance its generalization. In VAD, several works leverage pseudo-labeling and selftraining for generating fine-grained pseudo-labels [11, 19, 39]. However, in contrast to the previous methods, instead of generating pseudo-labels for the weakly labeled data, we leverage pseudo-labels for incorporating the external data.

Uncertainty Estimation. To address pseudo-label noise, prior research in different contexts has explored uncertainty estimation using various approaches, such as data augmentation [5, 29], inference augmentation [12], and model augmentation [43]. While data augmentation is effective for images, it can disrupt temporal relationships in video frames and is not efficient for training on high-cardinality data like videos. On the other hand, inference augmentation methods, such as MC Dropout [12, 39], introduce perturbations during model inference to obtain slightly different predictions, but that is inefficient for training with fixed backbones. In contrast, model augmentation uses different models. Since different models may have varying biases and receptive fields, this would result in diverse predictions. This prediction discrepancy can help quantify





Fig. 2: Overview of the proposed CDL Framework. **CDL Step 0:** The Ranking Loss, $\mathcal{L}_{\text{rank}}$ (Supp Mat. §6), is employed to train two pseudo-label generation models, P_m and P_a , §3.2, on weakly-labeled data, \mathcal{D}_l . **CDL Step** k, k > 0: P_m and P_a are trained iteratively on $\mathcal{D}_l \cup \mathcal{D}_u$, incorporating pseudo-labels for \mathcal{D}_u generated at the end of the previous CDL step. To deal with noise in pseudo-labels, uncertainty regularization scores are estimated using the divergence between the predictions of the two models, §3.4. When optimizing on \mathcal{D}_u , the prediction bias, \mathcal{L}_{bcc} (§3.3), for external data is reweighed using the computed uncertainty regularization scores, §3.5.

uncertainty, making model augmentation well-aligned with our problem. To avoid any manual thresholding for learning from pseudo-labels during training, following [15,43] we use adaptive reweighing of loss with uncertainty values. In [43], Zheng *et al.* quantify uncertainty by estimating discrepancies between predictions made by two classifiers using Kullback–Leibler (KL) divergence. However, given that VAD is a binary classification task, the divergence based on only two outcomes for the posterior probability is not optimally informative. Hence, we propose a method to quantify uncertainty in the high-dimensional feature space instead of the probability space.

3 Method

3.1 Problem Definition

In this work, we address a real-world VAD problem, where a weakly-labeled dataset $\mathcal{D}_l = \{(X_l^i, Y_l^i)\}_{i=1}^{n_l}$ and an external unlabeled dataset $\mathcal{D}_u = \{X_u^i\}_{i=1}^{n_u}$ are available for training. Here, n_l and n_u indicate the number of videos in the two datasets, respectively, with $n_u \gg n_l$ due to the convenience of gathering unlabeled video data. The video-level labels of X_l are denoted by $Y_l \in \{0, 1\}$. We do not make any assumption about distributions of \mathcal{D}_l and \mathcal{D}_u , and therefore, they can be drawn from different distributions. We aim to find the model $F(\cdot|\theta)$, parameterized by θ , that provides accurate predictions on weakly-labeled data while adaptively minimizing the prediction bias on the external data using the uncertainty regularization scores. We illustrate the proposed framework in Figure 2.

3.2 Feature Extraction and Temporal Processing

The proposed uncertainty quantification method (Section 3.4) compares two diverse representations of each sample to estimate the uncertainty associated with the segment-

level predictions on external data. To this aim, we employ two different backbones for feature extraction from videos, which are widely used for anomaly detection tasks. The first one is the conventional I3D backbone [6], which extracts segment-level features using 3D convolution, and the other is the CLIP backbone [26], which extracts frame-level features using the frozen CLIP Model's ViT encoder. The contrasting inductive biases of the 3D convolution-based I3D and the transformer-based CLIP help to effectively capture the prediction variance. It is to be noted that only the CLIP backbone is used during inference. We develop two prediction heads, namely the main model, P_m , built on top of the CLIP backbone, and the auxiliary model, P_a , built on top of the I3D backbone.

Video frames are highly correlated in the temporal dimension. To reduce the redundancy in frame-level features extracted by the CLIP backbone, we pool the representations by bilinearly interpolating them to a fixed, empirically determined length, n_s . Each of the n_s interpolated features represents one segment. To ensure consistency, we also fix the length of representations extracted by the I3D backbone. Evaluation in Section 4.6 analyzes the role of n_s on the model's performance. To capture long-range temporal information over the sequence, we employ a lightweight temporal network, *i.e.*, transformer encoder, to implement P_m and P_a .

3.3 Bias Estimation for External Data

Similar to [43], we formulate the prediction bias on external data as:

$$Bias(\mathcal{D}_u) = \mathbb{E}_{X_u}[F(X_u|\theta) - \mathcal{Y}_u],\tag{1}$$

where $F(X_u|\theta)$ represents a set of predicted probability distributions, each one corresponding to a distinct segment of X_u , and \mathcal{Y}_u denotes the set of unknown segment-level labels of X_u . $Bias(\mathcal{D}_u)$ can be re-written as:

$$Bias(\mathcal{D}_u) = \mathbb{E}_{X_u}[F(X_u|\theta) - \hat{\mathcal{Y}}_u] + \mathbb{E}_{X_u}[\hat{\mathcal{Y}}_u - \mathcal{Y}_u], \tag{2}$$

where $\hat{\mathcal{Y}}_u$ denotes the set of segment-level pseudo-labels for X_u . $\hat{\mathcal{Y}}_u$ can be generated by performing inference on the model trained on \mathcal{D}_l . The first term in Equation 2 denotes the difference between the predicted posterior probability and the pseudo-labels, while the second term denotes the error between the pseudo-labels and the ground-truth labels. While minimizing the prediction bias, due to the lack of ground truth supervision, we employ a self-training mechanism, considering $\hat{\mathcal{Y}}_u$ as the soft labels, thereby treating the second term as a constant and minimizing the first term. Specifically, we use the binary cross-entropy (BCE) loss, \mathcal{L}_{bce} , given by:

$$\mathcal{L}_{\text{bce}} = -\hat{\mathcal{Y}}_u \log(F(X_u|\theta)) - (1 - \hat{\mathcal{Y}}_u) \log(1 - F(X_u|\theta)),$$
(3)

to estimate the prediction bias associated with each video segment, for both P_m and P_a .

3.4 Uncertainty Estimation

Since \mathcal{D}_u and \mathcal{D}_l do not necessarily share the same distribution, the generated pseudolabels are noisy. This noise can adversely affect the subsequent training process as it causes bias to further magnify and propagate within the model. This issue, known as Confirmation Bias [4], is often mitigated by quantifying the uncertainty associated with pseudo-labels and then incorporating this uncertainty into the training process to compensate for the noise. As discussed in Section 2, we opt to address the confirmation bias by computing uncertainty using model augmentation. To quantify uncertainty through model augmentation, following [43], we estimate prediction variance, which is formulated as:

$$Var(\mathcal{D}_u) = \mathbb{E}_{X_u}[(F(X_u|\theta) - \mathcal{Y}_u)^2].$$
(4)

Due to the lack of ground-truth labels, Equation 4 can be approximated as:

$$Var(\mathcal{D}_u) \approx \mathbb{E}_{X_u}[(F(X_u|\theta) - \hat{\mathcal{Y}}_u)^2].$$
(5)

When optimizing the prediction bias in Equation 2, the variance in Equation 5 will also be minimized, potentially resulting in inaccurate quantification of the true prediction variance. To address this, we adopt an alternative approximation, expressed as:

$$Var(\mathcal{D}_u) \approx \mathbb{E}_{X_u}[\left(P_m(X_u|\theta_{P_m}) - P_a(X_u|\theta_{P_a})\right)^2].$$
(6)

Since VAD is a binary classification task, the probability distributions corresponding to each segment have limited support. Consequently, estimating prediction variance using only the predicted anomaly scores, as in Equation 6, may not be robust. Hence, instead of measuring the divergence between the predicted posterior probabilities for the two classes, we propose quantifying pseudo-label uncertainty in the high-dimensional space. To this end, we compute the cosine similarity between the segments in each set of the representations, Z_m and Z_a , obtained from the penultimate layer of P_m and P_a , respectively. Here, $Z_m = {\mathbf{z}_m^1, \mathbf{z}_m^2, \dots, \mathbf{z}_m^{n_s}}$ and $Z_a = {\mathbf{z}_a^1, \mathbf{z}_a^2, \dots, \mathbf{z}_a^{n_s}}$.

To obtain a set of stabilized, segment-level uncertainty regularization scores within a bounded range from the computed cosine similarity, we introduce the following function. Let $S = {s^1, s^2, ..., s^{n_s}}$ be the set of surrogate variances that we use as proxies for the uncertainty of segments. The surrogate variance is computed as:

$$\mathbf{s}^{j} = e^{\tau(\langle \mathbf{z}_{\mathbf{m}}^{j}, \mathbf{z}_{a}^{j} \rangle - 1)},\tag{7}$$

where s^j indicates the uncertainty regularization score for the j^{th} segment, $\langle \mathbf{z}_m^j, \mathbf{z}_a^j \rangle$ indicates the cosine similarity, and τ denotes the temperature parameter.

Higher uncertainty regularization scores indicate the similar encoding of data between the models, implying less uncertainty in the predicted labels, while, lower scores imply high uncertainty in the predicted labels. Empirical evidence in Section 4.4 demonstrates a significant negative correlation between uncertainty regularization scores and Binary Cross-Entropy (BCE) loss between the predicted labels and ground truths. This affirms that the proposed uncertainty regularization score effectively serves as a proxy for the quality of pseudo-labels.

3.5 Training Process

CDL Step 0. We initially train P_m and P_a separately on the labeled set, optimizing both of them using the Ranking Loss, \mathcal{L}_{rank} , discussed in Supp. Mat. Sec. 6. We then perform

inference on the trained models to generate the sets of soft segment-level pseudo-labels for training on \mathcal{D}_u .

CDL Step > 0. Following the generation of the sets of pseudo-labels for \mathcal{D}_u , we enter an iterative pseudo-label refinement phase, where we train P_m and P_a on $\mathcal{D}_l \cup \mathcal{D}_u$ for multiple CDL steps. Each CDL step comprises a fixed number of epochs. In each epoch, we regenerate the sets of segment-level uncertainty regularization scores. To enable the uncertainty-driven learning from external data, similar to [43], we use the estimated uncertainty regularization scores, S, as automatic thresholds as this dynamically adjusts learning from noisy labels by scaling the prediction bias associated with external data based on S. This helps filter out unreliable predictions while prioritizing highly confident predictions. To encourage lower prediction variance, which would in turn lead to increased pseudo-label quality, we explicitly add the prediction variance to the optimization objective corresponding to the external data, \mathcal{L}_{ext} , as:

$$\mathcal{L}_{\text{ext}} = \mathbb{E}_{X_u} [\frac{1}{Var(\mathcal{D}_u)} \cdot Bias(\mathcal{D}_u) + Var(\mathcal{D}_u)].$$
(8)

Equation 8 is rewritten with the approximated terms as:

$$\mathcal{L}_{\text{ext}} = \mathbb{E}_{X_u} [S \cdot \mathcal{L}_{\text{bce}} - \lambda_3 \cdot \langle Z_m, Z_a \rangle].$$
(9)

Alternatively, Equation 9 can be rewritten as:

$$\mathcal{L}_{\text{ext}} = \frac{1}{n_s \cdot n_u} \sum_{i=1}^{n_u} \sum_{j=1}^{n_s} \left(S^{i,j} \cdot \mathcal{L}_{\text{bce}}^{i,j} - \lambda_3 \cdot \langle \mathbf{Z}_{\text{m}}^{i,j}, \mathbf{Z}_a^{i,j} \rangle \right), \tag{10}$$

where λ_3 is a hyper-parameter to balance the losses. Similar to CDL step 0, to optimize the training on \mathcal{D}_l , we use \mathcal{L}_{rank} . The total optimization objective for training on $\mathcal{D}_l \cup \mathcal{D}_u$ can be expressed as:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{rank}} + \lambda_4 \cdot \mathcal{L}_{\text{ext}},\tag{11}$$

where λ_4 is a trade-off parameter for \mathcal{L}_{ext} . We employ the optimization objective defined in Equation 11 during training on $\mathcal{D}_l \cup \mathcal{D}_u$ for each epoch within every CDL step. After each CDL step is completed, we respensate the set of soft segment-level pseudo-labels using the models trained on $\mathcal{D}_l \cup \mathcal{D}_u$. This iterative refinement process repeats k times, where k is a hyper-parameter determining the number of CDL steps. With each CDL step, the models' performance gets further refined as the pseudo-labels get iteratively improved.

3.6 Inference - Extending Segment-level Scores to Frame-level Scores

During inference, we compute segment-level anomaly scores for the videos using P_m . Since we encounter long-untrimmed videos with varying numbers of frames, for extending the segment-level anomaly score to the frame level, for each video, we divide the total number of frames n_f by the number of segments n_s to obtain the number of frames per segment, n_{fs} . We assign the anomaly score of each segment to its consecutive frames. The first segment corresponds to the first n_{fs} frames, and so forth until the $(n_s - 1)^{th}$ segment. For the last segment, its anomaly score is assigned to any remaining frames, potentially exceeding n_{fs} , if there is a remainder.

4 Experiments

We evaluate the proposed method on the major video anomaly datasets, UCF-Crime (UCF) [30] and XD-Violence (XDV) [35]. Additionally, we use 11,000 videos from the HACS [42] dataset as a source of external data. We provide detailed information about the datasets in Supp. Mat. §7. In §4.1, we discuss the implementation details. In §4.2, we discuss the inherent noise in the test annotations of benchmark datasets. We proceed to compare the proposed framework with prior works in cross-domain scenarios (§4.3.1) and open-set scenarios (§4.3.2). Subsequently, in §4.4, we demonstrate a strong correlation between the quality of pseudo labels and the computed uncertainty scores. We then explore the evolution of these uncertainty scores through the training process in §4.5. Finally, in §4.6, we conduct ablation studies and hyper-parameter analysis to analyze the impact of individual components of the proposed framework.

4.1 Implementation Details

We implement the proposed method using PyTorch. We extract CLIP and I3D features at a fixed frame rate of 30 FPS. CLIP features are extracted from the frozen CLIP model's image encoder (ViT-B/32). For the hyper-parameters, in the open-set scenarios, we empirically set the value of n_s to 64, τ to 1.25, λ_1 and λ_2 to 5e - 4, λ_3 to 1e - 3, and λ_4 to 700. Ablation studies for selecting n_s and λ_3 are included in Section 4.6. We use the Adam optimizer with a weight decay of 1e - 3, and we set a learning rate of 3e - 5 for the transformer encoder and 5e - 4 for the fully connected layers. We use a batch size of 64. In both P_m and P_a , we explicitly encode positional information in the segments using sinusoidal positional encodings [32]. We train on the weakly-labeled source dataset for 200 epochs, followed by training on the union of weakly-labeled and external datasets for 40 CDL steps, each CDL step comprising 4 epochs. Additional information regarding hyper-parameters is provided in Supp. Material Section 8.

Model Architecture. Both P_m and P_a consist of a transformer encoder layer with four heads, followed by four fully connected layers, each consisting of 4096, 512, 32, and 1 neurons, respectively. In both the models, for all the layers except the last, we use ReLU [2] activation while for the last layer, we use Sigmoid activation.

Evaluation Setup. To reduce bias, we perform each experiment three times with different seeds and average the results. In open-set experiments, we repeat each experiment three times, using different sets of anomaly classes each time.

Evaluation Metric. Following previous works on UCF-Crime [30], we adopt the framelevel area under the ROC curve (AUC) to evaluate on UCF-Crime. In line with previous works on XD-Violence [35], we use the frame-level area under the Precision-Recall curve (PRAUC), also known as Average Precision (AP), to evaluate on XDV.

4.2 Noise in the Test Annotations of Benchmark Datasets

Our manual inspection reveals that the frame-level testing annotations of the UCF-Crime (UCF) [30] and XD-Violence (XDV) [35] datasets, which are commonly used for benchmarking VAD models, exhibit significant noise. This noise largely stems from the fact that the original annotations do not consistently label the frames leading up to

Table 2: Comparison with prior works on XDV, considering UCF-Crime as the source data. Asterisk (*) indicates that evaluations were conducted by us using the official code. Dagger (†) indicates that evaluations were conducted by our implementation due to the lack of an official implementation.

	Methods	Features	UCF AUC(%)	UCF-R AUC(%)	XDV AP(%)
Crear Demain	rGAN [22]	-	64.35^{*}	65.19^{*}	37.74^{*}
(Unsupervised)	MPN [24]	-	65.67^{*}	67.98^{*}	38.89^{*}
(Onsupervised)	zxVAD [3]	-	68.74^{\dagger}	69.39^{\dagger}	40.68^{\dagger}
	Sultani et al. [30]	I3D	80.70	84.63^{*}	53.88^{*}
	MIST [11]	I3D	82.30	86.17^{*}	50.33^{*}
	RTFM []	I3D	84.03	86.47^{*}	37.30^{*}
	S3R [34]	I3D	85.99	87.11^{*}	49.84^{*}
Non Cross-Domain	CU-Net [39]	I3D	86.22	88.15^{*}	37.98^{*}
	MGFN [8]	I3D	86.98	87.33^{*}	32.16^{*}
	SSRL [18]	I3D	87.43	87.02^{*}	51.60^{*}
	CLIP-TSA [17]	CLIP	87.58	73.20^{*}	44.33^{*}
	Ours (No ext. data)	CLIP	84.49	89.96	58.13
Cross-Domain	Ours (UCF + HACS)	CLIP	84.63	90.53	65.14
(Weakly-Supervised)	Ours (UCF + XDV)	CLIP	84.73	90.26	68.37

the primary anomalous events and their subsequent consequences as anomalous. For instance, in a video assigned a label like "shooting", we assert that frames showing the person holding the gun and frames illustrating the injured victim should also be marked as anomalous. This perspective aligns with the fundamental goal of VAD, which is to identify all anomalous frames within a video, irrespective of the video's primary label. However, it should also be noted that in the original annotations, for some videos, certain frames related to the video's primary anomaly label are also not marked anomalous.

To address this, we re-annotate the test set of UCF-Crime by assigning each video to three independent annotators. We then combine their annotations to generate more accurate frame-level labels. Compared to the original annotations where 7.58% of the total frames are labeled as anomalous, the proposed annotations label 16.55% of the total frames as anomalous. The proposed annotations are available here³. We provide a comparison of the proposed and original annotations here⁴. For the remainder of this paper, we refer the re-annotated test set of the UCF-Crime dataset as UCF-R.

4.3 Comparison with Prior Works

4.3.1 Cross-Domain Scenarios

While the UCF-Crime [30] and XD-Violence [35] datasets share similar definitions of what constitutes anomalies, that definition differs from those of smaller datasets like

³ https : / / drive . google . com / drive / folders / 1IVjQQFHXVcsaT63HUjpfk8C5KH6HsQ7t?usp=drive_link

⁴ https://rb.gy/4vkr1r

Table 3: Comparison with prior works on UCF-Crime, considering XDV as the source data. Asterisk (*) indicates that evaluations were conducted by us using the official code. Dagger (†) indicates that evaluations were conducted by our implementation due to the lack of an official implementation.

	Methods	Features	XDV AP(%)	UCF-R AUC(%)
	rGAN [22]	-	40.10^{*}	59.82*
Cross-Domain (Unsup.)	MPN [24]	-	44.79*	60.35*
	zxVAD [3]	-	47.53 [†]	63.61 [†]
	Sultani et al. [30]	I3D	73.20	71.23*
	RTFM [31]	I3D	77.81	70.46^{*}
Non Cross Domain	MGFN [8]	I3D	80.11	69.12*
Non Cross-Domain	S3R [34]	I3D	80.26	69.04*
	CLIP-TSA [17]	CLIP	80.67	67.58*
	Ours (No ext. data)	CLIP	75.13	76.39
Cross Domain (Waskly Sun)	Ours (XDV + UCF)	CLIP	77.04	88.06
Cross-Domain (weakly-Sup.)	Ours (XDV + HACS)	CLIP	78.61	88.50

ShanghaiTech [20], CUHK-Avenue [21], UCSD Pedestrian [7], UBnormal [1], where anomalies are more subtle. For instance, running is considered anomalous in UBnormal but not in XD-Violence. Due to these divergent notions of anomalies across datasets, we conduct cross-domain experiments by simultaneously evaluating on the UCF-Crime and XD-Violence datasets, given their more aligned anomaly definitions.

UCF-Crime as the Weakly-Labeled Source Set, XDV as the Cross-Domain Set. Table 2 summarizes the results for this scenario. First, we observe that the proposed method achieves state-of-the-art results on XDV and UCF-R even without utilizing any external data (without CDL). We believe this is due to the inductive bias of previous methods towards the noisy annotations of UCF-Crime. Next, we observe that the addition of external data, HACS and XDV, leads to a significant enhancement in the performance of the cross-domain dataset, XDV, by 11.26% and 14.49%, respectively, compared to the previous state-of-the-art baseline. Additionally, there is also a marginal improvement in the performance of the source set upon integration of external datasets.

XDV as the Weakly-Labeled Source Set, UCF-Crime as the Cross-Domain Set. Table 3 summarizes the results for this scenario. Notably, the proposed method achieves state-of-the-art performance on the cross-domain dataset, UCF-R, even without the utilization of any external data during training. This is attributed to the simplicity of the proposed architecture compared to other baselines. The proposed architecture prevents overfitting to the source dataset, thereby increasing its generalizability to the crossdomain dataset. Additionally, integrating external data further enhances performance on both the cross-domain and source sets. Specifically, leveraging the CDL framework with UCF-Crime and HACS as external datasets boosts UCF-R's AUC by 18.94% and 19.39% respectively, compared to previous state-of-the-art baselines. We also observe that the proposed method's performance is inferior on XDV. We attribute this to the noise in the annotations of XDV's test set.

Table 4: Comparison with prior works in open-set setting on UCF-Crime dataset; *c* denotes the no. of anomalous classes included for weakly-supervised training.

_	UCF (AUC%)					UCF-R (A	UC%)
c	Wu et al. [35]	RTFM [31]	Zhu et al. [46]	Ours (w/o CDL)	Ours (CDL)	Ours (w/o CDL)	Ours (CDL)
1	73.22	75.91	76.73	75.17	77.45	84.32	85.39
3	75.15	76.98	77.78	81.51	82.57	86.84	87.69
6	78.46	77.68	78.82	82.97	83.44	87.85	88.21
9	79.96	79.55	80.14	83.02	83.37	89.22	89.82

These results highlight that the proposed CDL framework is capable of effectively exploiting external data with vast domain gaps to achieve a significant cross-domain generalization. It's noteworthy that the performance gain observed with the proposed CDL framework remains consistent across all tested datasets, suggesting that the performance improvement is not dependent on any specific source or external dataset.

4.3.2 Open-Set Scenarios

In Table 4, we evaluate the proposed framework's performance on the UCF-Crime dataset in a realistic open-set scenario, where the model is evaluated on both, previously seen and unseen anomaly classes. To simulate this scenario, we randomly include *c* anomalous classes in the weakly-labeled set, while the remaining anomalous classes are placed in the unlabeled set. In both the weakly-supervised source set and the unlabeled set, the number of normal videos equals the number of anomalous videos. We evaluate two model configurations; one trained solely on the weakly-labeled set (without CDL) and the other on the union of weakly-labeled and unlabeled sets using the CDL Framework.

On UCF-Crime, the proposed model, without CDL, surpasses the state-of-the-art baselines for c > 1. This highlights its efficacy in open-set settings. While, with CDL, the model surpasses the baselines across all values of c by a considerable margin.

For both UCF-Crime and UCF-R, when unlabeled data is incorporated, we observe a consistent performance gain across all values of *c*, suggesting the effectiveness of the CDL framework across varying amounts of weakly-labeled and unlabeled data.

4.4 Correlation between Uncertainty Scores and BCE Loss (Proxy to Label Quality)

To assess the efficacy of the proposed uncertainty quantification method as a proxy for pseudo-label quality, we compute the non-parametric Spearman correlation between estimated uncertainty regularization scores and BCE loss between the predicted pseudo-labels and the corresponding ground truths. For this experiment, we consider UCF-Crime as the weakly-labeled source set and XDV as the external set. In Figure 3(a), with $\lambda_3 = 1e - 3$, CDL step 1 onwards, a consistently high negative correlation (-0.46 in CDL step 6, with a p-value < 1e-5) emerges, indicating the robustness of the proposed uncertainty quantification framework. Conversely, setting λ_3 to 0 results in a sustained positive correlation, signifying sub-optimal pseudo-labels in the absence of cosine similarity loss term.

13



Fig. 3: (a) Correlation between uncertainty scores and BCE loss computed between the estimated scores and ground truth. When $\lambda_3 = 1e - 3$, as expected, a consistently high negative correlation emerges, demonstrating the effectiveness of the proposed uncertainty quantification method as a reliable proxy for pseudo-label quality. (b) Cumulative Distribution Function (CDF) plots illustrating the progression of average uncertainty regularization scores for each video during training. CDL step 20 has a higher concentration of scores around 1 compared to CDL step 2, while CDL step 2 has a higher tendency for scores to have elevated values, indicating more confident pseudo-label predictions. (c) Ablation study on the coefficient of the cosine similarity loss term, λ_3 . (d) Ablation study on the number of segments, n_s .

4.5 Progression of Uncertainty Scores

To assess the evolution of uncertainty regularization scores through the training process, in Figure 3(b), we plot the Cumulative Distribution Function (CDF) of average uncertainty regularization scores for external videos across the first epoch of three different CDL steps. We conduct this experiment considering UCF-Crime as the weakly-labeled source set and XDV as the external set. We observe that in CDL step 1, 16.65% of the uncertainty scores fall within the range [0, 0.1]. As training progresses to CDL steps 2 and 20, this proportion decreases to 13.06% and 11.39%, respectively. Meanwhile, the proportion of uncertainty scores in the range [0.9, 1] increases from 35.11% in CDL step 1 to 56.70% in CDL step 2 and further to 57.68% in CDL step 20. This trend indicates a discernible shift towards higher uncertainty scores as training progresses, suggesting an improvement in model confidence due to increased pseudo-label quality.

4.6 Ablation Studies and Hyper-parameter Analysis

For the sake of consistency, we conduct all ablation studies on UCF-Crime in an openset setting, with c = 1. However, it should be noted that for different training setups, hyper-parameters are tuned separately as well.

Impact of Various Components of the CDL Framework. We assess the effectiveness of each component of the CDL framework by adding them sequentially. The results are summarized in Table 5. We consider training on c = 1 anomaly class in a weaklysupervised fashion as our baseline. The remaining c-1 anomalous classes are placed in the external set. We first observe that integrating external data into the source set without accounting for pseudo-label uncertainty ($S^{i,j} = 1, \forall i, j$) and without minimizing cosine similarity between representations ($\lambda_3 = 0$) yields a 0.35% gain in AUC, highlighting the effectiveness of external data in improving the model's performance. Next, we study the impact of uncertainty-aware integration of external data, *i.e.*, adaptively





Ext. data	Unc. Coef.	CosSim. Los	s AUC(%)
×	X	X	84.32
1	×	×	84.67
1	1	×	84.80
1	1	1	85.39

Fig. 4: Ablation study on the impact of the size of external data.

Table 5: Ablation study of various components on the UCF-R dataset in an open-set setting (c = 1).

reweighing the prediction bias of external data with the computed uncertainty values and with λ_3 set to 0. This results in a gain of 0.13% in AUC, demonstrating the superiority of uncertainty-driven integration compared to the standard integration. Finally, we assess the impact of adding the cosine similarity loss term during uncertainty-aware training. This further leads to a significant boost of 0.59%, validating its effectiveness.

Impact of Cosine Similarity Loss. In Figure 3(c), we explore the impact of varying the coefficient of the cosine similarity loss on the model's performance. We observe a gradual increase in AUC as λ_3 increases from 1e-9 to 1e-3. This could be due to the effect of cosine similarity loss getting more pronounced with higher values of λ_3 . However, beyond 1e-3, there is a rapid decline in AUC, likely due to the dominance of the cosine similarity loss over other losses when its coefficient is high. Therefore, we select 1e-3 as the optimal choice for λ_3 .

Impact of Number of Segments. In Figure 3(d), we observe that the performance consistently improves as no. of segments, n_s , increases from 16 to 64, but it begins to decline rapidly afterward. Therefore, we set n_s as 64.

Impact of the Size of External Data. To determine the optimal number of unlabeled external videos from the HACS dataset to integrate into the weakly-labeled training set of UCF-Crime, we conduct an ablation study, depicted in Figure 4. We observe that increasing the size of the external set increases the performance on XDV. However, this increase tends to plateau after the inclusion of 11,000 videos. Consequently, we do not include additional videos beyond the 11,000 threshold.

5 Conclusion

In this work, we demonstrated the effectiveness of integrating external, unlabeled data with weakly-labeled source data to enhance the cross-domain generalization of VAD models. To enable this integration, we proposed a weakly-supervised CDL (Cross-Domain Learning) framework that adaptively minimizes the prediction bias on external data by scaling it with the prediction variance, which serves as an uncertainty regularization score. The proposed method outperforms baseline models significantly in cross-domain and open-set settings while retaining competitive performance in in-domain settings.

Acknowledgement

This work was supported in part by U.S. NIH grants R01GM134020 and P41GM103712, NSF grants DBI-1949629, DBI-2238093, IIS-2007595, IIS-2211597, and MCB-2205148. This work was supported in part by Oracle Cloud credits and related resources provided by Oracle for Research, and the computational resources support from AMD HPC Fund. We thank Eshaan Mandal and Bhavay Malhotra for their assistance, which has been instrumental in completing this work.

References

- Acsintoae, A., Florescu, A., Georgescu, M., Mare, T., Sumedrea, P., Ionescu, R.T., Khan, F.S., Shah, M.: Ubnormal: New benchmark for supervised open-set video anomaly detection. In: CVPR (2022)
- 2. Agarap, A.F.: Deep learning using rectified linear units (relu) (2019)
- Aich, A., Peng, K.C., Roy-Chowdhury, A.K.: Cross-domain video anomaly detection without target domain adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2579–2591 (January 2023)
- Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: IJCNN (2020)
- Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In: ICLR (2020)
- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: CVPR. pp. 6299–6308 (2017)
- Chan, A.B., Vasconcelos, N.: Modeling, clustering, and segmenting video with mixtures of dynamic textures. IEEE Transactions on Pattern Analysis and Machine Intelligence pp. 909– 926 (2008)
- Chen, Y., Liu, Z., Zhang, B., Fok, W., Qi, X., Wu, Y.C.: Mgfn: magnitude-contrastive glanceand-focus network for weakly-supervised video anomaly detection. In: Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence (2023)
- Cho, M., Kim, M., Hwang, S., Park, C., Lee, K., Lee, S.: Look around for anomalies: Weaklysupervised anomaly detection via context-motion relational learning. In: CVPR. pp. 12137– 12146 (2023)
- Cong, Y., Yuan, J., Liu, J.: Sparse reconstruction cost for abnormal event detection. In: CVPR. pp. 3449–3456 (2011)
- Feng, J.C., Hong, F.T., Zheng, W.S.: MIST: Multiple instance self-training framework for video anomaly detection. In: CVPR. pp. 14009–14018 (2021)
- Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: ICML. pp. 1050–1059 (2016)
- Georgescu, M.I., Ionescu, R.T., Khan, F.S., Popescu, M., Shah, M.: A background-agnostic framework with adversarial training for abnormal event detection in video. IEEE Transactions on Pattern Analysis and Machine Intelligence pp. 4505–4523 (2022)
- 14. Hasan, M., Choi, J., Neumann, J., Roy-Chowdhury, A.K., Davis, L.S.: Learning temporal regularity in video sequences. In: CVPR (June 2016)
- Huang, K., Sresht, V., Rai, B., Bordyuh, M.: Uncertainty-aware pseudo-labeling for quantum calculations. In: UAI. pp. 853–862 (2022)

- 16 Y. Jain et al.
- Ionescu, R.T., Khan, F.S., Georgescu, M.I., Shao, L.: Object-centric auto-encoders and dummy anomalies for abnormal event detection in video. In: CVPR. pp. 7842–7851 (2019)
- Joo, H.K., Vo, K., Yamazaki, K., Le, N.: Clip-tsa: Clip-assisted temporal self-attention for weakly-supervised video anomaly detection. In: ICIP. pp. 3230–3234 (2023)
- Li, G., Cai, G., Zeng, X., Zhao, R.: Scale-aware spatio-temporal relation learning for video anomaly detection. In: ECCV. pp. 333–350 (2022)
- Li, S., Liu, F., Jiao, L.: Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. AAAI pp. 1395–1403 (2022)
- Liu, W., Luo, W., Lian, D., Gao, S.: Future frame prediction for anomaly detection-a new baseline. In: CVPR. pp. 6536–6545 (2018)
- Lu, C., Shi, J., Jia, J.: Abnormal event detection at 150 fps in matlab. In: ICCV. pp. 2720– 2727 (2013)
- 22. Lu, Y., Yu, F., Reddy, M.K.K., Wang, Y.: Few-shot scene-adaptive anomaly detection. In: ECCV. pp. 125–141 (2020)
- Luo, W., Liu, W., Gao, S.: Remembering history with convolutional lstm for anomaly detection. In: ICME. pp. 439–444 (2017)
- Lv, H., Chen, C., Cui, Z., Xu, C., Li, Y., Yang, J.: Learning normal dynamics in videos with meta prototype network. In: CVPR. pp. 15425–15434 (2021)
- McClosky, D., Charniak, E., Johnson, M.: Effective self-training for parsing. In: NAACL. pp. 152–159 (2006)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763 (2021)
- Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In: ICLR (2021)
- Sapkota, H., Yu, Q.: Bayesian nonparametric submodular video partition for robust anomaly detection. In: CVPR. pp. 3212–3221 (2022)
- Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: simplifying semi-supervised learning with consistency and confidence. In: NeurIPS (2020)
- Sultani, W., Chen, C., Shah, M.: Real-world anomaly detection in surveillance videos. In: CVPR. pp. 6479–6488 (2018)
- Tian, Y., Pang, G., Chen, Y., Singh, R., Verjans, J.W., Carneiro, G.: Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In: ICCV. pp. 4975– 4986 (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS. p. 6000–6010 (2017)
- Wan, B., Fang, Y., Xia, X., Mei, J.: Weakly supervised video anomaly detection via centerguided discriminative learning. ICME pp. 1–6 (2020)
- Wu, J.C., Hsieh, H.Y., Chen, D.J., Fuh, C.S., Liu, T.L.: Self-supervised sparse representation for video anomaly detection. In: ECCV. pp. 729–745 (2022)
- 35. Wu, P., Liu, j., Shi, Y., Sun, Y., Shao, F., Wu, Z., Yang, Z.: Not only look, but also listen: Learning multimodal violence detection under weak supervision. In: ECCV (2020)
- Xu, D., Ricci, E., Yan, Y., Song, J., Sebe, N.: Learning deep representations of appearance and motion for anomalous event detection. In: BMVC. pp. 8.1–8.12 (2015)
- Yarowsky, D.: Unsupervised word sense disambiguation rivaling supervised methods. In: ACL. p. 189–196 (1995)
- Yu, G., Wang, S., Cai, Z., Zhu, E., Xu, C., Yin, J., Kloft, M.: Cloze test helps: Effective video anomaly detection via learning to complete video events. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 583–591 (2020)

Cross-Domain Learning for Video Anomaly Detection with Limited Supervision

- Zhang, C., Li, G., Qi, Y., Wang, S., Qing, L., Huang, Q., Yang, M.H.: Exploiting completeness and uncertainty of pseudo labels for weakly supervised video anomaly detection. In: CVPR. pp. 16271–16280 (2023)
- 40. Zhang, J., Qing, L., Miao, J.: Temporal convolutional network with complementary inner bag loss for weakly supervised anomaly detection. In: ICIP. pp. 4030–4034 (2019)
- Zhao, B., Li, F.F., Xing, E.: Online detection of unusual events in videos via dynamic sparse coding. In: CVPR. pp. 3313–3320 (2011)
- 42. Zhao, H., Torralba, A., Torresani, L., Yan, Z.: Hacs: Human action clips and segments dataset. In: ICCV. pp. 8668–8678 (2019)
- 43. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. International Journal of Computer Vision (IJCV) (2021)
- Zhong, J.X., Li, N., Kong, W., Liu, S., Li, T.H., Li, G.: Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection. In: CVPR. pp. 1237– 1246 (2019)
- Zhu, Y., Newsam, S.D.: Motion-aware feature for improved video anomaly detection. In: BMVC. p. 270 (2019)
- Zhu, Y., Bao, W., Yu, Q.: Towards open set video anomaly detection. In: ECCV. pp. 395–412 (2022)