

Dual Memory Units with Uncertainty Regulation for Weakly Supervised Video Anomaly Detection

Hang Zhou, Junqing Yu*, Wei Yang*

Huazhong University of Science and Technology, Wuhan, China
{henryzh, yjqing, weiyangcs}@hust.edu.cn

Abstract

Learning discriminative features for effectively separating abnormal events from normality is crucial for weakly supervised video anomaly detection (WS-VAD) tasks. Existing approaches, both video and segment-level label oriented, mainly focus on extracting representations for anomaly data while neglecting the implication of normal data. We observe that such a scheme is sub-optimal, i.e., for better distinguishing anomaly one needs to understand what is a normal state, and may yield a higher false alarm rate. To address this issue, we propose an Uncertainty Regulated Dual Memory Units (UR-DMU) model to learn both the representations of normal data and discriminative features of abnormal data. To be specific, inspired by the traditional global and local structure on graph convolutional networks, we introduce a Global and Local Multi-Head Self Attention (GL-MHSA) module for the Transformer network to obtain more expressive embeddings for capturing associations in videos. Then, we use two memory banks, one additional abnormal memory for tackling hard samples, to store and separate abnormal and normal prototypes and maximize the margins between the two representations. Finally, we propose an uncertainty learning scheme to learn the normal data latent space, that is robust to noise from camera switching, object changing, scene transforming, etc. Extensive experiments on XD-Violence and UCF-Crime datasets demonstrate that our method outperforms the state-of-the-art methods by a sizable margin.

Introduction

Video anomaly detection (VAD) aims to identify the time window of abnormal events in untrimmed videos (Lu, Shi, and Jia 2013; Liu et al. 2018; Sultani, Chen, and Shah 2018; Wu et al. 2020), and has beneficial applications in the context of intelligent video surveillance. Yet collecting fine-grained anomaly videos for training are arduous, weakly supervised approaches, which assign one label, i.e., normal or abnormal, for each panoptic video, have been gaining popularity and achieving state-of-the-art performances (Zhong et al. 2019; Zaheer et al. 2020; Tian et al. 2021). Following weakly supervised setup, multiple instances learning (MIL) is one of the most effective strategies (Wang et al. 2018b;

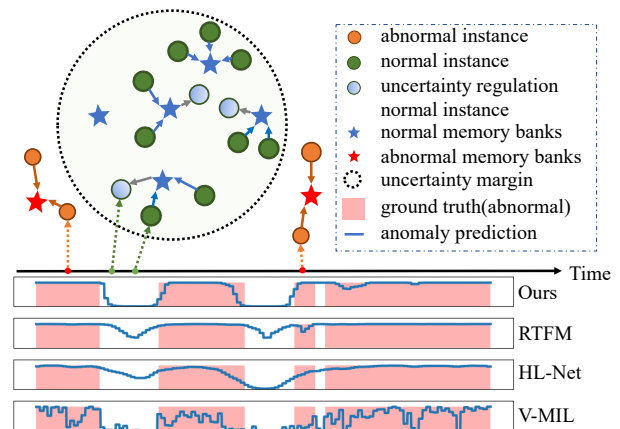


Figure 1: We use two memory banks to store normal and anomaly instances respectively, both memory banks attract close instances. The normal space is regulated to be a Gaussian distribution, while no additional regulation applied for anomaly feature space. Such a strategy can tackle the hard samples better, the bottom curve shows prediction results of a part of the video “Black.Hawk.Down”.

Ilse, Tomczak, and Welling 2018). MIL-based anomaly detection methods consider each video as a bag, where abnormal videos are defined as positive and normal videos as negative. Meanwhile, anomaly snippets are presented as positive instances in positive bags, and normal snippets are negative instances. MIL-based methods then learn the instance-level labels from video bags by selecting top-K instances as abnormal according to discriminative representations (Zhang, Qing, and Miao 2019; Zhu and Newsam 2019; Wu et al. 2021; Wu and Liu 2021).

Notably, MIST (Feng, Hong, and Zheng 2021) develops a two-stage self-training procedure by generating pseudo labels using MIL approaches and emphasizing anomalous regions with a self-guided attention boosted feature encoder. RTFM (Tian et al. 2021) uses a feature magnitude learning function to recognise the positive instances, effectively improving the robustness of the MIL approach. Nevertheless, existing weakly supervised approaches focus on learning discriminative features for anomalous instances, while neglecting the implication of normal data, and hence yield

*indicates co-corresponding author.

high false alarm rate. As shown in the bottom of Figure 1, RTFM and HL-Net both lead to alarm falsely.

To address the problems and better distinguish anomaly, we propose a dual memory units module with uncertainty regulation to store normal and abnormal prototypes respectively and learn their distinctive distributions, as shown in the top of Figure 1. Previous approaches use only one memory for normality, while we observe such strategy can not handle hard samples well, e.g., the action of running exists in both normal exercising videos and criminal arresting videos, and hence yield a high false alarm rate. Dual memory units consider more peripheral factors compared to a single memory, and may reduce false alarms. Specifically, we encourage the anomaly memory bank to intake more anomaly information from videos annotated as anomaly, while the normal memory bank to obtain normal data pattern from both normal and abnormal videos. Following RTFM’s paradigm which indicates that normal features are in low magnitude while anomaly features are on the contrary, we observe that normal features exhibit fluctuations due to noises coming from camera switching, subject changing, etc., and further format low magnitude normal features with Gaussian distribution. We adopt a normal data uncertainty learning (NUL) scheme to learn both the mean feature and uncertainty variance, and consequentially generate a latent normal space. We can segregate the normality and anomaly according to the latent space and reduce the false alarm rate with noise perturbations. Moreover, inspired by the traditional global and local structure on graph convolutional networks, we introduce a Global and Local Multi-Head Self Attention (GL-MHSA) module for the Transformer network to obtain more expressive embeddings for capturing associations in videos.

In summary, the main contributions of this paper are as follows:

- We design a Global and Local Multi-Head Self Attention model to capture different associations of anomaly videos for enhancing the expressive embeddings.
- We propose the dual memory units, which use both normal and abnormal memory banks for better distinguishing patterns of normality and anomaly. We also introduce a dual memory separation loss to learn more robust and discriminative features.
- We propose a normal data uncertainty learning scheme to enhance the normal pattern representation ability and better separate the anomaly from latent normality.

Related Work

Weakly Supervised Video Anomaly Detection. VAD is an important task in computer vision (Pang et al. 2021a), and mainstream methods either adopt an unsupervised learning strategy with only normal data or weakly supervised learning setup with video-level annotated data. Unsupervised methods with only normal data can not well handle complex environments, and stimulated by progress in weakly supervised learning in video action localization (Nguyen et al. 2018; Lee, Uh, and Byun 2020; Lee et al. 2021; Yang, Qin, and Huang 2022), WS-VAD gradually gains popularity. (Sultani, Chen, and Shah 2018) introduce a large real-world

surveillance dataset and proposed a vanilla MIL (V-MIL) framework with ranking loss. (Zhong et al. 2019) consider the normal label to contain noise and apply pseudo label generation method by graph convolutional network to tackle noise specifically. Recent work mostly adopt the MIL training strategy. Most improvements expand the receptive field and enhance the instance embedding (Wu and Liu 2021), e.g., RTFM proposes Multi-scale Temporal Network (MTN) to capture the long and short-range temporal dependencies. Other approaches focus on multi-modal setup, for instance (Wei et al. 2022; Pang et al. 2021b) use both audio and video information for VAD.

Self-training strategy generate labels from coarse to fine level (Zou et al. 2019; Pang et al. 2020) and has been explored in semi-supervised learning. In VADs, (Pang et al. 2020) present an end-to-end and self-training method for anomaly regression without manually labeled training data. (Feng, Hong, and Zheng 2021) propose a multiple instance pseudo label generator and self-guided attention boosted encoder to refine discriminative representations. (Li, Liu, and Jiao 2022) apply multi-sequence learning with Transformer to refining abnormal scores through self-training method. Both types of methods consider normal snippets are noisy and lack normal data learning, while learning anomalous patterns excessively without normal contrast suffers from overfitting. In contrast, we adopt the MIL setup and use memory banks to store the normal and abnormal patterns. Moreover, we apply normal data uncertainty regulation to learn the latent normal space and enlarge the distance between normality and anomaly.

Memory Networks. Memory network is an efficient neural network to store different prototypes and has been applied in various tasks (Yang et al. 2021b; Luo et al. 2021). (Gong et al. 2019) propose a memory augmented deep autoencoder to suppress the reconstruction of anomaly. (Stephen and Menon 2020) follows the MemAE framework and presents a new readable and updatable memory scheme, where memory banks also update at test time. (Liu et al. 2021b) design a multi-level memory module in an autoencoder with skip connections to memorize normal patterns for optical flow reconstruction. All the above methods use memory banks to learn from normality. In this work, We introduce dual memory units to store both normal and abnormal prototypes, and we design a dual loss to guide memory banks in prototype learning.

Data Uncertainty Learning. Data uncertainty learning captures the data variance due to noise and randomness (Chang et al. 2020; Li et al. 2022; Lee et al. 2021). (Chang et al. 2020) propose the data uncertainty learning scheme to learn the feature embedding and its uncertainty in face recognition. In video temporal action localization, (Lee et al. 2021) formulate the background frames as out-of-distribution (OOD) samples and proposed an uncertainty model to learn the uncertainty. Different from (Lee et al. 2021), we presume that the anomaly data is OOD, and the normal data follows a Gaussian distribution. The latent space of the normal uncertainty pattern is effective for unknown normal samples. Meanwhile, it also helps the model to separate normality and anomaly.

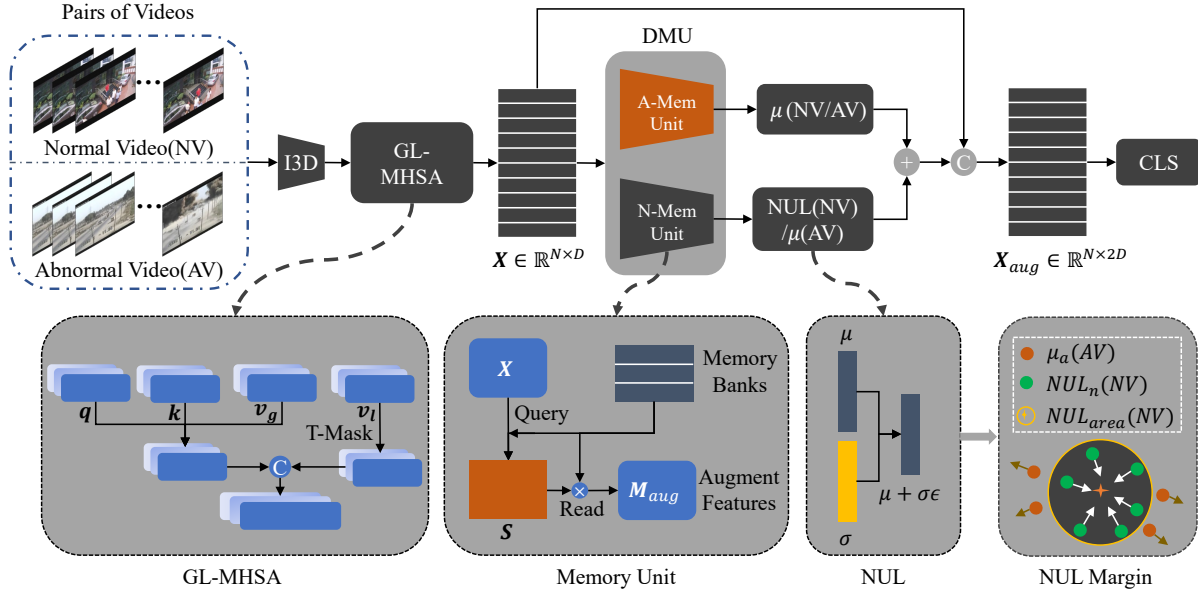


Figure 2: The framework of our UR-DMU model consists of three parts: global and local feature learning (GL-MHSA), dual memory units (DMU), and normal data uncertainty learning (NUL). GL-MHSA extracts better expressive embeddings, DMU stores both normal and abnormal patterns for discrimination, and NUL constrains the normal data as a Gaussian distribution to handle uncertainty perturbation.

Method

Here we introduce the Uncertainty Regulated Dual Memory Units (UR-DMU) method to store both normal and abnormal prototypes and enforce additional prior distributions in details. As illustrated in Figure 2, UR-DMU consists of three modules: the Global and Local Multi-Head Self Attention (GL-MHSA) for feature learning, the Dual Memory Units (DMU) for storing prototypes, and the Normal Data Uncertainty Learning (NUL) for normality latent embedding learning. In the following sections, we briefly introduce the formations of our problem first, and then describe the above three modules respectively, and finally show the training and testing process.

Problem Formation

Following the MIL setup for WS-VAD tasks, we define normal videos as $\{v_i^n\}_{i=1}^N$ and abnormal videos as $\{v_i^a\}_{i=1}^N$. Each anomaly video is a bag $Y^a = 1$ with at least one abnormal clip, similarly a normal video is marked as $Y^n = 0$ with all normal clips. Our goal is to score each v_i according to $Y_i = f(v_i)$ with the model $f(\cdot)$. It is a common practice to extract snippet features using pretrained backbones (Tran et al. 2015; Carreira and Zisserman 2017; Liu et al. 2021c) and conduct future processes on extracted features, in this paper we choose I3D model pre-trained on Kinetics dataset as our backbone, and extract RGB features of consecutive 16-frames (Wu et al. 2020). The extracted video features are denoted as $\mathbf{X}_{i3d} \in \mathbb{R}^{N \times F}$, where N represent the number of snippets, and F is the dimension of extracted features.

Global and Local Feature Learning

Learning long and short temporal dependencies of anomalous features is crucial for VAD (Wang et al. 2018a; Wang and Gupta 2018). Traditional approaches adopt the graph convolution network for extracting global and local associations (Wu et al. 2020). Recently, the transformer structure shows its superiority in handling long-range associations between temporal signals (Vaswani et al. 2017; Dosovitskiy et al. 2021; Liu et al. 2021a; Yang et al. 2021a). For better capturing long and short temporal dependencies simultaneously, we borrow the idea of global and local structure in GCN and develop the Global and Local Multi-Head Self Attention (GL-MHSA) for transformer. Different from the conventional MHSA, we add an extra encoder layer of $\mathbf{V}_l = \mathbf{X}_{i3d} \mathbf{W}_l$ to learn local features with a temporal mask:

$$\mathbf{T}_m(i, j) = -\frac{|i - j|}{e^\tau}. \quad (1)$$

Then we normalize the mask with $\text{softmax}(\mathbf{T}_m)$, where τ is the sensitivity factor for balancing the local distance. While computing Q, K, V , we obtain the global and local feature $\mathbf{X} \in \mathbb{R}^{N \times D}$ from Multilayer Perceptron (MLP) and Layer-Norm (LN), which can be formulated as follows:

$$\begin{aligned} \mathbf{X} &= \text{MH}([\text{softmax}(\frac{\mathbf{QK}^T}{\sqrt{D}})\mathbf{V}_g; \text{softmax}(\mathbf{T}_m)\mathbf{V}_l]), \\ \mathbf{X} &= \text{LN}(\text{MLP}(\text{LN}(\mathbf{X})) + \mathbf{X}). \end{aligned} \quad (2)$$

Dual Memory Units

Previous approaches use memory units to store prototypes of normal data for better scoring. We observe that a single unit

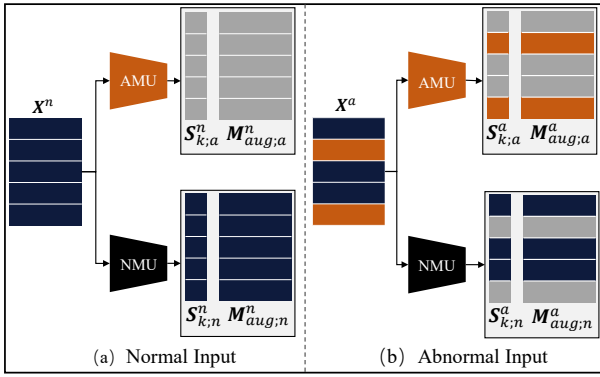


Figure 3: The dual memory units store normal and anomaly instances in different memories. (a) shows that normal video feature goes through dual memory units and obtains relevant output scores and augment features. (b) is the case with abnormal video feature input.

is insufficient for distinguishing hard samples, and we use learnable normal and abnormal memory banks to store corresponding templates $\mathbf{M} \in \mathbb{R}^{M \times D}$, where M is the memory bank number, D is the output feature dimension. We apply query and read operations to encourage correct template learning, and by querying memory banks \mathbf{M} we can assess the similarities between features:

$$\mathbf{S} = \sigma\left(\frac{\mathbf{X}\mathbf{M}^T}{\sqrt{D}}\right), \quad \mathbf{M}_{aug} = \mathbf{S}\mathbf{M}, \quad (3)$$

where σ is sigmoid activation, and $\mathbf{S} \in \mathbb{R}^{N \times M}$ is the query score. Then the memory augmentation feature generated by a read operation is represented as \mathbf{M}_{aug} . \mathbf{S} indicates whether snippets are similar to memory banks. We apply topK selection along the second dimension to determine which snippets are most similar to memory banks:

$$\mathbf{S}_{k;ind}, \mathbf{S}_{k;sc} = \text{topK}(\mathbf{S}, K), \quad \mathbf{S}_{k;i} = \frac{\sum_{j=1}^K S_{k;sc}(i, j)}{K}. \quad (4)$$

In equation (4), $K \ll M$, $\mathbf{S}_{k;ind}$ means the indices of proposed K snippets to be stored, and $\mathbf{S}_{k;sc} \in \mathbb{R}^{N \times K}$ are their query scores. $\mathbf{S}_k \in \mathbb{R}^{N \times 1}$ indicates the highest scores of memory matching.

Figure 3 shows different input features across the dual memory units. While sending normal video embedding $\mathbf{X}^n \in \mathbb{R}^{N \times D}$ into the abnormal memory banks, $\mathbf{S}_{k;a}^n$ and $\mathbf{M}_{aug;a}^n$ are generated by querying and reading operations. Similarly, we can also obtain $\mathbf{S}_{k;n}^n$ and $\mathbf{M}_{aug;n}^n$ as video embedding goes through normal memory banks. The anomaly query scores $\mathbf{S}_{k;a}^n$ are constrained to $\mathbf{0} \in \mathbb{R}^N$ because of the normal input. And the normal query score $\mathbf{S}_{k;n}^n$ should be $\mathbf{1} \in \mathbb{R}^N$.

Likewise, when the abnormal video embedding input to abnormal memory banks, we obtain $\mathbf{S}_{k;a}^a$ and $\mathbf{M}_{aug;a}^a$. Because of the sparseness nature of the anomaly, only several snippets will be activated by anomaly memory banks. We can also obtain $\mathbf{S}_{k;n}^a$ and $\mathbf{M}_{aug;n}^a$ from normal memory

banks. Different from normal input, the anomaly query score $\mathbf{S}_{k;a}^n$ is not all $\mathbf{1}$ because of the weak label of anomaly. Same with MIL, there is at least one element equals 1 in $\mathbf{S}_{k;a}^a$ and $\mathbf{S}_{k;n}^a$.

To train the dual memory units, we define a dual memory loss consists of four binary cross-entropy (BCE) losses:

$$L_{dm} = \text{BCE}(\mathbf{S}_{k;n}^n, \mathbf{y}_n^n) + \text{BCE}(\mathbf{S}_{k;a}^n, \mathbf{y}_a^n) + \text{BCE}(\mathbf{S}_{k;n}^a, \mathbf{y}_n^a) + \text{BCE}(\mathbf{S}_{k;a}^a, \mathbf{y}_a^a), \quad (5)$$

where $\mathbf{y}_n^n = \mathbf{1} \in \mathbb{R}^N$ and $\mathbf{y}_a^n = \mathbf{0} \in \mathbb{R}^N$. $S_{k;n;k}^a, S_{k;a;k}^a \in \mathbb{R}$ are means of $\mathbf{S}_{k;n}^a, \mathbf{S}_{k;a}^a$ topK result along the first dimension. The labels $\mathbf{y}_n^a, \mathbf{y}_a^a$ are equal to 1. The first two terms of L_{dm} limit the normal memory unit to learn normal patterns, and the last two terms ensure the abnormal pattern learning for abnormal memory unit. With L_{dm} , memory banks store homologous patterns.

With the query scores $\mathbf{S}_{k;n}^n, \mathbf{S}_{k;a}^n$ and $\mathbf{S}_{k;a}^a$, we separate the abnormal and normal feature embedding by triplet loss.

$$L_{trip} = \text{Triplet}(f_a, f_p, f_n), \\ f_a = \text{topK}(\mathbf{S}_{k;n}^n; \mathbf{X}^n), \\ f_p = \text{topK}(\mathbf{S}_{k;n}^a; \mathbf{X}^a), \\ f_n = \text{topK}(\mathbf{S}_{k;a}^a; \mathbf{X}^a), \quad (6)$$

where f_a is the average feature of the interest snippets from \mathbf{X}^n , which indicates the primary part of normal video. The snippets are selected by the topK index of $\mathbf{S}_{k;n}^n$. Similarly, f_p means the normal part in an abnormal video, and f_n means the abnormal part in the abnormal video.

Normal Date Uncertainty Learning

As pointed out by (Tian et al. 2021; Lee et al. 2021), the magnitude of the normal feature is lower than that of anomalies, i.e., $\mathbb{E}[\|x^a\|_2] \geq \mathbb{E}[\|x^n\|_2]$, x^a and x^n are the snippet features of anomaly and normality respectively. Presumably, magnitude feature learning is also appropriate for the augmentation features \mathbf{M}_{aug} of memory output. Hence, the magnitude separability of augmentation features between abnormal and normal is defined as $d(a, n) = \|m^a\|_2 - \|m^n\|_2$. Moreover, effective representation learning of normality should be able to capture both the core and cover variance introduced by noise. We propose to learn the distribution of normality, and then anomaly is out of this distribution (OOD) (Riveram et al. 2022). This idea has been extensively explored in unsupervised approaches, but rarely used in WS-VAD. Similar to unsupervised anomaly detection methods, we propose to use the Gaussian distribution to constraint the latent normal representation z^n , i.e.:

$$p(z_i^n | x_i^n) = \mathcal{N}(z_i^n; \mu_i; \sigma_i^2 I), \quad (7)$$

where μ is the normal feature's mean, and σ is the variance (uncertainty). Both of them can be predicted by an encoder network. To ensure the normal data sampling operation is differentiable, we follow (Chang et al. 2020) and use the re-parameterization to tackle the latent with noise. The DUL features z^n is generated by random noise $\varepsilon \sim \mathcal{N}(0, 1)$:

$$z_i^n = \mu_i^n + \sigma_i^n \varepsilon, \quad (8)$$

Equation (8) represents the latent space of normal data. However, the sample z is impressionable to σ when σ is too large, and it may be far away from the real features μ . To suppress the instability of σ , a Kullback-Leibler diversity(KLD) regularization term is necessary:

$$\begin{aligned} L_{kl} &= \text{KL}(\mathcal{N}(z^n | \mu^n, \sigma^n) || \mathcal{N}(0, 1)) \\ &= -\frac{1}{2D} \sum_{i=1}^D (1 + \log(\sigma_i^n)^2 - (\mu_i^n)^2 - (\sigma_i^n)^2), \end{aligned} \quad (9)$$

In testing stage, calculation of the variance can be skipped, and we only need the mean to generate robust prediction. Furthermore, as mentioned, anomalies are OOD samples, and we need to enlarge the margin between anomalies and normalities for robustness. The anomaly features are also fed to the mean-encoder to get the embedding μ^a in the same space. We use the magnitude distance loss to separate them in the latent normal space:

$$L_{dis} = \max(0, d - (\|\mu_k^a\|_2^2 - \|z_k^n\|_2^2)), \quad (10)$$

where μ_k^a are the average features of proposed snippets from \mathbf{M}_{aug} . The snippets are selected by topK index operation from $\mathbf{S}_{k;a}^a$. Same with the above operation, z_k^n is obtained from $\mathbf{S}_{k;n}^n$ and \mathbf{M}_{aug} .

Finally, we fuse the DUL features with \mathbf{X} and feed it to the classifier.

Network training and Testing

Training: Pairs of videos with the same amount of normal and abnormal videos are divided into a batch and inputted to our model. Then we generate one score for each snippet and use the BCE loss to discriminate the normality and anomaly with five auxiliary losses:

$$L = L_{cls} + \lambda_1 L_{dm} + \lambda_2 L_{trip} + \lambda_3 L_{kl} + \lambda_4 L_{dis}, \quad (11)$$

where L_{cls} is BCE loss with the supervision of video-level labels, $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are hyperparameters.

Testing: We feed testing videos to our UR-DMU network. In the DUL module, we only use the mean-encoder network to obtain feature embeddings. Finally, the snippet labels are assigned to frame labels.

Experiments

Datasets and Evaluation Metrics

Dataset. Similar to previous works, we conduct experiments on two large video anomaly detection datasets: UCF-Crime (Sultani, Chen, and Shah 2018), and XD-Violence (Wu et al. 2020).

UCF-Crime is a large-scale surveillance dataset that includes 1900 videos with 13 types of anomalous events, e.g., explosion, arrest, and road accidents. The training set with video-level labels contains 800 normal and 810 abnormal videos, while the testing set with frame-level labels has 140 normal and 150 abnormal videos.

XD-Violence is a multi-modal dataset with videos and audios, and it includes different backgrounds: movies, games, live scenes, and etc. There are 4754 videos in the datasets,

	Method	Feature(RGB)	AUC(%)
Un	Hasan et al.(2016)	N/A	50.60
	Lu et al.(2013)	N/A	65.51
	GCL(2022)	ResNext	74.20
Weakly	Sultani et al.(2018)	C3D	75.41
	Sultani et al.(2018)	I3D	76.21
	MA(2019)	PWC	79.00
	GCN (2019)	TSN	82.12
	IBL (2019)	C3D	78.66
	HL-Net (2020)	I3D	82.44
	CLAWS (2020)	C3D	83.03
	MIST (2021)	I3D	82.30
	RTFM (2021)	I3D	84.30
	CRFD (2021)	I3D	84.89
	MSL (2022)	I3D	85.30
	MSL (2022)	VideoSwin	85.62
		Ours	I3D

Table 1: Comparison of the frame-level performance AUC on UCF-Crime testing set with existing approaches.

which contains 3954 training videos with video-level labels and 800 testing videos with frame-level labels.

Evaluation Metrics. We evaluate WS-VAD performance with the area under curve (AUC) of the frame-level receiver operating characteristics (ROC) for UCF-Crime. Meanwhile, the AUC of the frame-level precision-recall curve (AP) is used for XD-Violence. False alarm rate (FAR) and anomaly subset with only abnormal data are also applied in ablation experiments (Wu and Liu 2021; Feng, Hong, and Zheng 2021).

Implementation Details

We extract the snippet features from pre-trained I3D on Kinetics-400. The same configurations are set for both datasets. In the training stage, each video is divided into $N = 200$ snippets, and the learning rate is 0.0001. We train the model with batch size of 64 and iterator time of 3000, and set the $M^a = 60$ and $M^n = 60$ for UCF-Crime and XD-Violence. The hyperparameters of $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ are set as $(0.1, 0.1, 0.001, 0.0001)$ to balance the total loss. All the operations of topK are equal to $K = \lfloor \frac{M}{16} \rfloor + 1$. In equation (10), $d = 100$.

Results on UCF-Crime

On UCF-Crime dataset, we compare the AUC scores with existing state-of-the-art (SOTA) VAD methods in Table 1. All the methods use RGB features, and our method is based on 10-crop augmentation I3D features. As the results show, without the prior information of anomaly, unsupervised methods are generally lower in rank than weakly supervised methods. When compared with weakly supervised approaches, our method outperforms all the previous methods and achieves the SOTA AUC score of 86.97%. There is a 1.35% improvement for our UR-DMU model when compared to the previous SOTA method MSL (85.62% AUC) (Li, Liu, and Jiao 2022), which is a Transformer-based method with self-training. Compared with the MIL-based method RTFM (84.30% AUC) (Tian et al. 2021), we achieve a considerable improvement of 2.67%.

Method	Feature	AP(%)
Hasan et al.(2016)	RGB	30.77
Sultani et al.(2018)	RGB	73.20
HL-Net (2020)	RGB	73.67
HL-Net (2020)	RGB+Audio	78.64
CRFD (2021)	RGB	75.90
RTFM(2021)	RGB	77.81
MSL(2022)	RGB	78.28
MSL(2022)	RGB	78.59
Pang et al.(2021)	RGB+Audio	81.69
ACF(2022)	RGB+Audio	80.13
Ours	RGB	81.66
Ours	RGB+Audio	81.77

Table 2: Comparison of frame-level performance AP on XD-Violence validation set.

Memory	Dataset					
	UCF(%)			XD(%)		
	AUC	AP	FAR	AUC	AP	FAR
Single A-Mem	83.72	31.46	1.32	88.07	77.59	1.38
Single N-Mem	85.19	35.23	0.98	93.69	80.40	0.48
Dual Mem	86.97	35.59	1.05	94.02	81.66	0.65

Table 3: Comparison of the results of different memory units.

Results on XD-Violence

The AP scores in XD-Violence are displayed with SOTA VAD methods in Table 2. We use the 5-crop I3D RGB and features and VGGish audio features provided by (Wu et al. 2020) for experiments, and achieve 81.66% performance with only RGB features. Compared with the MIL-based method RTFM (77.81% AP), we achieve a 3.85% improvement. Multi-modal methods are also showed in Table 2. With only RGB features, our method outperforms the ACF (Wei et al. 2022), which utilizes the attention module to fuse the video and audio modal. We also obtain the same performances as (Pang et al. 2021b). To evaluate the robustness in multi-modal, we utilize feature concatenation to fuse video and audio features, We can also achieve the best performances with an 81.77% AP score. The results of both datasets indicate our UR-DMU is effective in detecting abnormal events with disturbances.

Ablation Study

Effectiveness of dual memory units. To verify the effect of the different memory units, we split dual memory units into two single memory units. As shown in Table 3, the single abnormal memory unit lacks normal patterns to discriminate what is normality and receives a higher FAR in normality detection. It also indicates that abnormal patterns are not easy to learn on weakly supervised data. The single normal memory unit structure can obtain better results than abnormal memory with rich normal data. The dual memory units get the best performances because of storing both memory patterns. To explore a proper number of the memory banks,

(M_A, M_N)	Dataset					
	UCF(%)			XD(%)		
	AUC	AP	FAR	AUC	AP	FAR
20,20	86.71	31.28	0.82	93.63	81.13	0.39
30,30	86.44	31.29	1.18	93.67	80.46	1.18
40,40	85.09	32.40	1.53	93.07	79.90	0.67
50,50	86.28	33.23	1.63	93.51	80.19	0.74
60,60	86.97	35.59	1.05	94.02	81.66	0.65
70,70	86.47	33.60	0.49	93.06	78.87	1.00
80,80	85.16	34.18	0.96	93.45	79.59	0.31
90,90	86.18	35.14	2.01	93.47	79.87	0.44
100,100	85.23	31.57	1.63	93.92	80.75	0.78

Table 4: Comparison of the results of different memory numbers M_A , and M_N . M_A is the number of abnormal memory banks, and M_N is the number of normal memory banks.

Loss term				Dataset	
L_{dm}	L_{trip}	L_{kl}	L_{dis}	UCF(AUC%)	XD(AP%)
✓	✓			83.46	77.14
✓		✓		82.78	78.05
✓			✓	81.24	76.49
	✓	✓	✓	82.09	78.74
✓		✓	✓	82.78	78.16
✓	✓		✓	83.92	77.57
✓	✓	✓		86.60	80.79
✓	✓	✓	✓	86.97	81.66

Table 5: Comparison of the results of different loss terms.

we set 9 memory cases to search for the best configuration of memory bank numbers. The results are shown in Table 4. M_a and M_n mean the number of abnormal and normal memories. When $M_a = 60, M_n = 60$, our method achieves the best performances on both datasets.

Effectiveness of different loss terms. To explore the impact of four loss terms on our complete model, we report the results of 8 loss combinations in Table 5. The results of the first three cases show that the L_{trip} is most effective to separate the distance between normality and anomaly. From the last five cases, the loss terms of $\{L_{dm}, L_{trip}, L_{kl}\}$ are essential and indispensable for our model. The L_{dis} has a slight effect on detection performance, and control the border between normal and abnormal. The above observations reveal that our loss terms are more suitable for working together.

Effectiveness of proposed modules. To investigate the influences of different modules, we report the results of four combinations of three modules in Table 6, and adopt three additional indicators to comprehensively evaluate our method. The AUC_{sub} and AP_{sub} are separately the AUC score and AP score of the abnormal video subset, which only contains abnormal videos. They are used to confirm the anomaly detection performance for anomaly video. Meanwhile, the FAR is also applied for normal video detection and ensures the performance of normality detection. GL-MHSA module is the feature learning backbone and achieves the lowest performance. When adding the DMU module, all the metrics exhibit great improvements. The increase of AUC/AP scores and the decrease of FAR indicate that the DMU module can store anomaly and normality pat-

Module			Dataset									
			UCF(%)					XD(%)				
GL-MHSA	DMU	NUL	AUC	AP	AUC_{sub}	AP_{sub}	FAR	AUC	AP	AUC_{sub}	AP_{sub}	FAR
✓			82.89	27.07	60.18	28.30	4.86	91.57	75.67	76.17	76.99	0.96
✓	✓		85.53	31.21	70.13	33.79	1.53	93.38	79.80	81.06	81.03	0.62
✓		✓	85.19	33.61	69.74	35.79	2.37	93.34	78.71	78.69	79.44	0.85
✓	✓	✓	86.97	35.59	70.81	37.25	1.05	94.02	81.66	82.36	82.85	0.65

Table 6: Comparison of each module’s performance on UCF-Crime and XD-Violence. AUC_{sub} and AP_{sub} are the scores only using abnormal data. FAR is the false alarm rate with only normal data.

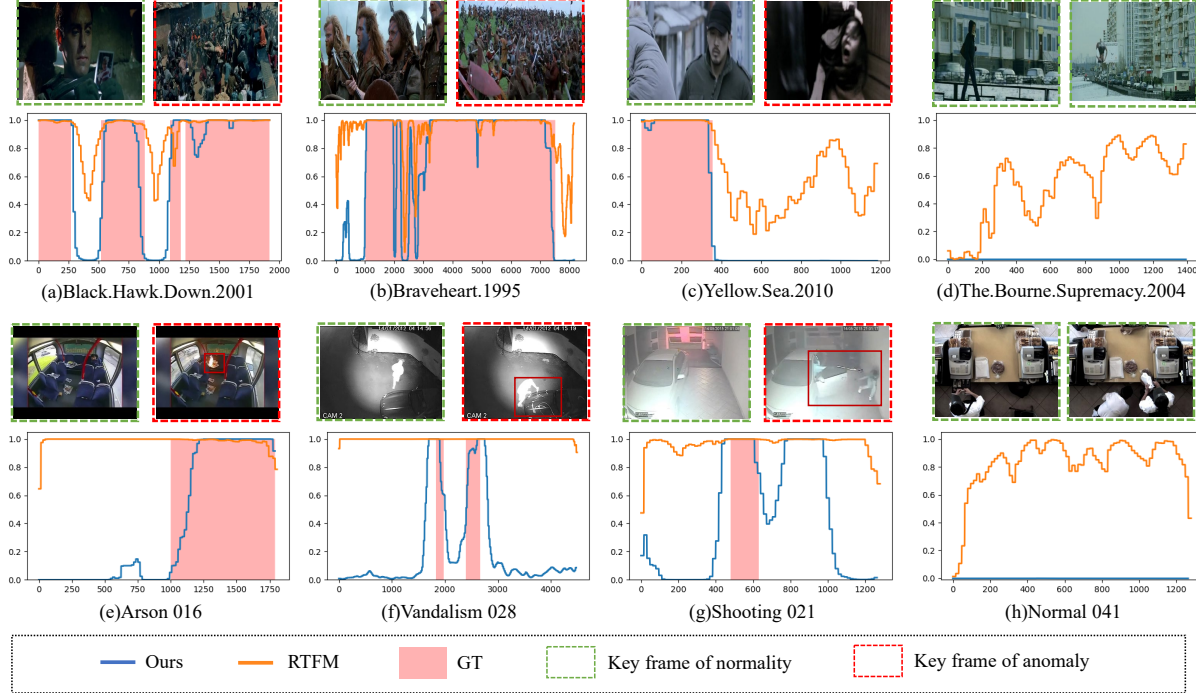


Figure 4: Qualitative results of anomaly detection performances on UCF-Crime and XD-Violence. The visualizations of (a)~(d) are from XD-Violence, and (e)~(h) are from UCF-Crime.

terns. When NUL module is added, each metric has large improvements. But it can not map all the normal patterns into Gaussian space without memory banks auxiliary. Finally, the full model achieves the best performances on two datasets.

Qualitative Results

To further evaluate our method, the qualitative results are shown on two datasets in Figure 4. We compare our method (blue line) with RTFM (orange line) on eight videos to exhibit the frame-level abnormal predictions. The pink area is the ground truth area of the anomaly. We also show the keyframes and mark the abnormal/normal frames to describe what happened in the video, the green rectangles indicate the normal frame and the red rectangles mean the abnormal frame. As exhibited in the figure, our method not only produces a precise detection area but also reduces the FAR on the normal part. RTFM is based on feature magnitude and is effective for scenes that are smooth and change slowly. Results of Figure 4 (f) and (g) show that RTFM fails to make

correct detection. Our method achieves a lower FAR for normal video (d) and (h).

Conclusion

In this work, we propose a UR-DMU model to efficiently separate normality and abnormality for WS-VAD. To obtain more expressive embeddings, we introduce the GL-MHSA module to capture long and short temporal dependencies. Then, we present DMU to store abnormal and normal prototypes under the video-level label, and the DMU loss is devised to encourage memory units learning and separate them. Meanwhile, the DUL scheme is proposed to learn latent normal space from rich normal data with noise perturbation. The predicted anomaly windows and scores of our approach are still relatively coarse and unstable because of using snippet-level features. In the future, we plan to explore more fine-grained locations in both temporal windows and image space with contrastive learning.

Acknowledgments

This work is supported by the National Key Research and Development Program of China under Grant No.2020YFB1805601, National Natural Science Foundation of China (NSFC No. 62272184), and CCF-Tencent Open Research Fund (CCF-Tencent RAGR20220120). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

References

- Carreira, J.; and Zisserman, A. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*, 4727–4733.
- Chang, J.; Lan, Z.; Cheng, C.; and Wei, Y. 2020. Data Uncertainty Learning in Face Recognition. In *CVPR*, 5709–5718.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Feng, J.; Hong, F.; and Zheng, W. 2021. MIST: Multiple Instance Self-Training Framework for Video Anomaly Detection. In *CVPR*, 14009–14018.
- Gong, D.; Liu, L.; Le, V.; Saha, B.; Mansour, M. R.; Venkatesh, S.; and Hengel, A. V. D. 2019. Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection. In *ICCV*, 1705–1714.
- Ilse, M.; Tomczak, J. M.; and Welling, M. 2018. Attention-based deep multiple instance learning. In *ICML*, 2132–2141.
- Lee, P.; Uh, Y.; and Byun, H. 2020. Background Suppression Network for Weakly-Supervised Temporal Action Localization. In *AAAI*, 11320–11327.
- Lee, P.; Wang, J.; Lu, Y.; and Byun, H. 2021. Weakly-supervised Temporal Action Localization by Uncertainty Modeling. In *AAAI*, 1854–1862.
- Li, S.; Liu, F.; and Jiao, L. 2022. Self-Training Multi-Sequence Learning with Transformer for Weakly Supervised Video Anomaly Detection. In *AAAI*, 1395–1403.
- Li, X.; Dai, Y.; Ge, Y.; Liu, J.; Shan, Y.; and Duan, L. 2022. Uncertainty Modeling for Out-of-Distribution Generalization. *arXiv preprint arXiv:2202.03958*.
- Liu, W.; Luo, W.; Lian, D.; and Gao, S. 2018. Future Frame Prediction for Anomaly Detection - A New Baseline. In *CVPR*, 6536–6545.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021a. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *ICCV*, 9992–10002.
- Liu, Z.; Nie, Y.; Long, C.; Zhang, Q.; and Li, G. 2021b. A Hybrid Video Anomaly Detection Framework via Memory-Augmented Flow Reconstruction and Flow-Guided Frame Prediction. In *ICCV*, 13568–13577.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2021c. Video Swin Transformer. *arXiv preprint arXiv:2106.13230*.
- Lu, C.; Shi, J.; and Jia, J. 2013. Abnormal Event Detection at 150 FPS in MATLAB. In *ICCV*, 2720–2727.
- Luo, W.; Zhang, T.; Yang, W.; Liu, J.; Mei, T.; Wu, F.; and Zhang, Y. 2021. Action Unit Memory Network for Weakly Supervised Temporal Action Localization. In *CVPR*, 9969–9979.
- Nguyen, P.; Liu, T.; Prasad, G.; and Han, B. 2018. Weakly Supervised Action Localization by Sparse Temporal Pooling Network. In *CVPR*, 6752–6761.
- Pang, G.; Shen, C.; Cao, L.; and Hengel, A. V. D. 2021a. Deep Learning for Anomaly Detection: A Review. *ACM Computing Surveys*, 54(2): 1–38.
- Pang, G.; Yan, C.; Shen, C.; Hengel, A. V. D.; and Bai, X. 2020. Self-Trained Deep Ordinal Regression for End-to-End Video Anomaly Detection. In *CVPR*, 12170–12179.
- Pang, W.; He, Q.; Hu, Y.; and Li, Y. 2021b. Violence Detection in Videos Based on Fusing Visual and Audio Information. In *ICASSP*, 2260–2264.
- Riveram, A. R.; Khan, A.; Bekkouch, I. E. I.; and Sheikh, T. S. 2022. Anomaly Detection Based on Zero-Shot Outlier Synthesis and Hierarchical Feature Distillation. *IEEE Transactions on Neural Networks and Learning Systems*, 33(1): 281–291.
- Stephen, K.; and Menon, V. 2020. Learning Memory-guided Normality for Anomaly Detection. In *CVPR*, 14360–14369.
- Sultani, W.; Chen, C.; and Shah, M. 2018. Real-World Anomaly Detection in Surveillance Videos. In *CVPR*, 6479–6488.
- Tian, Y.; Pang, G.; Chen, Y.; Singh, R.; Verjans, J. W.; and Carneiro, G. 2021. Weakly-supervised Video Anomaly Detection with Robust Temporal Feature Magnitude Learning. In *ICCV*, 4955–4966.
- Tran, D.; Bourdev, L. D.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning Spatiotemporal Features with 3D Convolutional Networks. In *ICCV*, 4489–4497.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NIPS*, 5998–6008.
- Wang, X.; Girshick, R. B.; Gupta, A.; and He, K. 2018a. Non-Local Neural Networks. In *CVPR*, 7794–7803.
- Wang, X.; and Gupta, A. 2018. Videos as Space-Time Region Graphs. In *ECCV*, 413–431.
- Wang, X.; Yan, Y.; Tang, P.; Bai, X.; and Liu, W. 2018b. Revisiting multiple instance neural networks. *Pattern Recognit*, 74: 15–24.
- Wei, D.; Liu, C.; Liu, Y.; Liu, J.; Zhu, X.; and Zeng, X. 2022. Look, Listen and Pay More Attention: Fusing Multi-Modal Information for Video Violence Detection. In *ICASSP*, 1980–1984.
- Wu, J.; Zhang, W.; Li, G.; Wu, H.; Tan, X.; Li, Y.; Ding, E.; and Lin, L. 2021. Weakly-Supervised Spatio-Temporal Anomaly Detection in Surveillance Video. In *IJCAI*, 1172–1178.
- Wu, P.; and Liu, J. 2021. Learning Causal Temporal Relation and Feature Discrimination for Anomaly Detection. *IEEE Transactions on Image Processing*, 30: 3512–3527.

- Wu, P.; Liu, J.; Shi, Y.; Shao, F.; Wu, Z.; and Yang, Z. 2020. Not only Look, But Also Listen: Learning Multimodal Violence Detection Under Weak Supervision. In *ECCV*, 322–339.
- Yang, J.; Li, C.; Zhang, P.; Dai, X.; Xiao, B.; Yuan, L.; and Gao, J. 2021a. Focal Self-attention for Local-Global Interactions in Vision Transformers. *arXiv preprint arXiv:2107.00641*.
- Yang, Z.; Qin, J.; and Huang, D. 2022. ACGNet: Action Complement Graph Network for Weakly-Supervised Temporal Action Localization. In *AAAI*, 3090–3098.
- Yang, Z.; Zhang, T.; Bozchalooi, I. S.; and Darve, E. 2021b. Memory-Augmented Generative Adversarial Networks for Anomaly Detection. *IEEE Transactions on Neural Networks and Learning Systems*, 33(6): 2324–2334.
- Zaheer, M. Z.; Mahmood, A.; Astrid, M.; and Lee, S.-I. 2020. CLAWS: Clustering Assisted Weakly Supervised Learning with Normalcy Suppression for Anomalous Event Detection. In *ECCV*, 358–376.
- Zhang, J.; Qing, L.; and Miao, J. 2019. Temporal Convolutional Network with Complementary Inner Bag Loss for Weakly Supervised Anomaly Detection. In *ICIP*, 4030–4034.
- Zhong, J.; Li, N.; Kong, W.; Liu, S.; Li, T. H.; and Li, G. 2019. Graph Convolutional Label Noise Cleaner: Train a Plug-And-Play Action Classifier for Anomaly Detection. In *CVPR*, 1237–1246.
- Zhu, Y.; and Newsam, S. D. 2019. Motion-Aware Feature for Improved Video Anomaly Detection. In *BMVC*, 270.
- Zou, Y.; Yu, Z.; Liu, X.; and Kumar, J., B.V.K. Vijaya adn Wang. 2019. Confidence Regularized Self-Training. In *ICCV*, 5981–5990.