

Detecting Contextual Anomalies by Discovering Consistent Spatial Regions

Zhengye Yang
Rensselaer Polytechnic Institute
Department of ECSE, Troy, NY, USA
yangz15@rpi.edu

Richard J. Radke
Rensselaer Polytechnic Institute
Department of ECSE, Troy, NY, USA
rjradke@ecse.rpi.edu

Abstract

We describe a method for modeling spatial context to enable video anomaly detection. The main idea is to discover regions that share similar object-level activities by clustering joint object attributes using Gaussian mixture models. We demonstrate that this straightforward approach, using orders of magnitude fewer parameters than competing models, achieves state-of-the-art performance in the challenging spatial-context-dependent Street Scene dataset. As a side benefit, the high-resolution discovered regions learned by the model also provide explainable normalcy maps for human operators without the need for any pre-trained segmentation model.

1. Introduction

Determining whether an event observed in a video stream is anomalous often depends on its spatial context. For example, pedestrians are expected on sidewalks and cyclists in bike lanes, but not vice versa. In real-world video anomaly detection, the appearance and motion of objects alone are insufficient to classify anomalies; we should learn the “normal” behavior of objects at different locations in the video based on long-term observation.

A major trend in modern video anomaly detection (VAD) algorithms is to train convolutional auto-encoders using a reconstruction proxy task [8, 10, 20]. Future frame prediction [13, 14, 17, 20] can be considered a special form of reconstruction that also addresses temporal regularity. These reconstruction-based methods have achieved great success on widely adopted VAD benchmark datasets, such as UCSD Ped2 [33], CUHK Avenue [16], and ShanghaiTech [14]. However, we observe a large performance drop when such algorithms are applied to datasets in which spatial context plays a critical role in making the anomaly decision, such as the relatively new Street Scene benchmark dataset [23]. One reason is that convolutional auto-encoders are poorly suited to learning spatial context due to their shift-invariant nature [9].

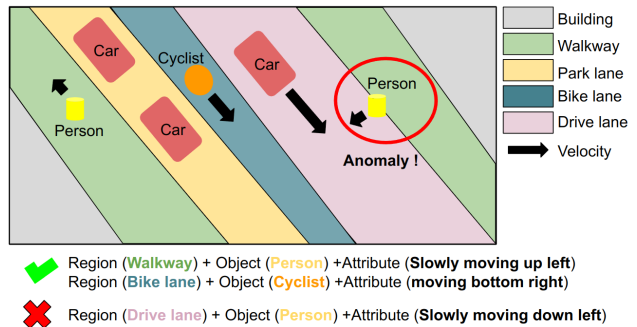


Figure 1. A toy example of building spatial context. Using the discovered semantically meaningful regions, spatial-context-dependent anomalies can be easily detected.

In contrast, high-performing VAD algorithms on the Street Scene dataset (e.g., [23, 25]) usually separate videos into hundreds of rectangular subregions, each with a dedicated normalcy model. Hundreds of subregions not only result in hundreds of normalcy models but also reduce model robustness due to the small number of available object samples per region. However, a typical surveillance camera does not capture hundreds of small regions with different normalcy patterns; meaningful regions should be automatically discovered from long-term observations.

Fig. 1 illustrates a toy experiment mimicking a typical street scenario. Suppose we can extract representations like “pedestrians walk or jog on the sidewalk,” “cyclists ride within the bike lane,” and “cars move quickly in the traffic lane in the correct direction” from training videos. Then, given a new event represented in the same “region-object-attribute” form, we should be able to achieve an almost perfect anomaly detection result. Leveraging object categories and their corresponding attributes is already widespread in the VAD community [19, 24, 28]. However, automatically extracting consistent spatial regions based on these behaviors and using them as the basis for anomaly detection is relatively unexplored territory.

In this paper, we describe a natural approach to discovering spatial regions that have similar expected object motion

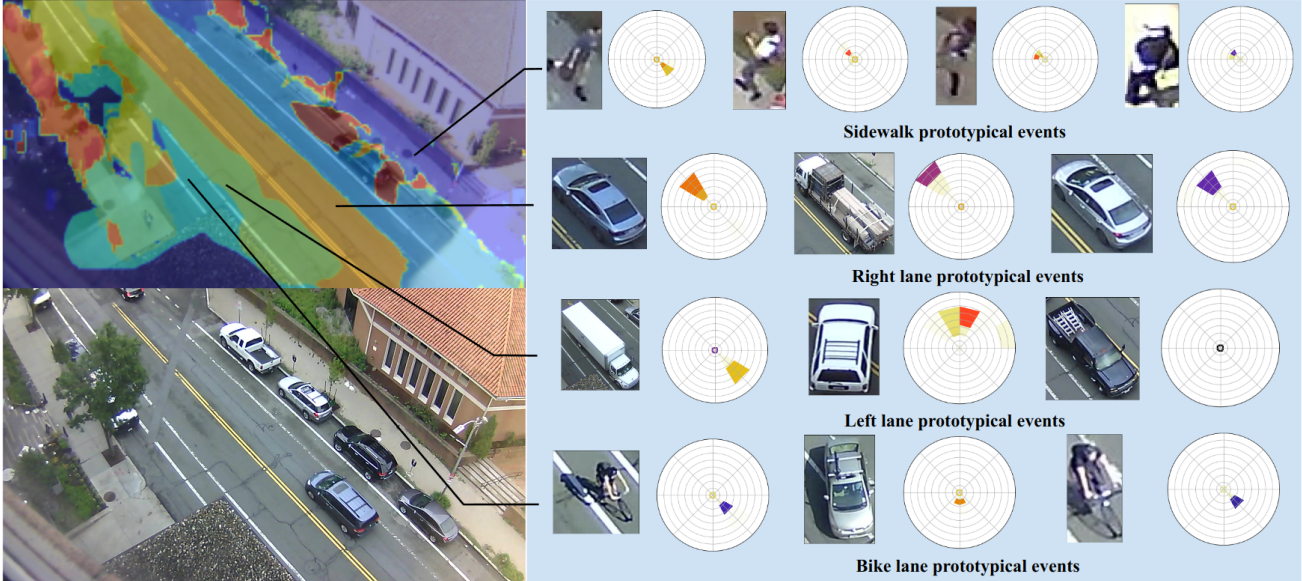


Figure 2. The proposed method discovers regions that have similar object and motion (top left image), illustrated here in the Street Scene dataset (lower left image) [23]. Each region is characterized by a learned mixture of prototypical events in appearance/motion feature space. On the right side we show motion descriptors in magnitude/angle space for learned modes of the mixture model in several regions, along with corresponding objects.

and appearance, and make anomaly decisions based on what is “normal” for a given region. With high-quality region discovery, we show that training a simple Gaussian mixture model suffices to detect spatial-context-dependent anomalies. This has the side benefits of greatly reducing the number of models required for learning and inference, as well as providing better interpretability to the user.

Our approach achieves state-of-the-art performance on the Street Scene dataset using straightforward Gaussian mixture models to discover a small number of activity regions specified by the user, each defined by a set of prototypical events (modes) in a low-dimensional appearance/motion space. As illustrated in Fig. 2, we automatically discover semantically meaningful regions like traffic lanes, bike lanes, and sidewalks, each of which has typical object appearance and motion. For a new object, we can classify its track as normal or anomalous based on its likelihood according to the learned region model, discovering spatial-context-dependent anomalies like pedestrians in the street or illegally parked cars. The contributions of this paper are threefold:

- We propose a pixel-level clustering method for region discovery and compare it against the common approach of rectangular patches, validating the effectiveness of our strategy in terms of both performance and efficiency.
- We propose a highly effective approach focusing on spatial-context anomalies that achieves state-of-the-art

VAD performance on the Street Scene dataset.

- We demonstrate that the discovered regions provide natural interpretations of the expected normal activities within each region.

2. Related Work

Reconstruction-based Methods. Early VAD approaches were based on dictionary learning (e.g., [16]). The successors to this approach are modern reconstruction-based methods [4, 8, 10, 17, 20, 37, 40] and future-frame-prediction-based methods [13, 14, 17, 20]. In both types of approaches, an auto-encoder is trained to perform reconstruction or prediction, and anomalies are detected based on the reconstruction error of query frames. However, we observe that convolutional neural networks (CNNs), the building blocks shared by many auto-encoders, are naturally unsuitable for learning spatial context due to their shift-invariant properties. Figure 3 illustrates a toy experiment to demonstrate this issue, using the same normal/anomalous classification setup used by most VAD methods. The training data is generated so that red squares only appear in the upper left quadrant and blue circles only appear in the bottom half. Half of the testing data is generated following the same rule, while the other half generates the square and circle positions freely. We see that a convolutional auto-encoder trained from the normal data can also perfectly reconstruct the anomalous frames; spatial context has not been learned. Thus, these types of methods typically perform poorly on

real datasets like Street Scene that contain spatial-context-dependent anomalies.

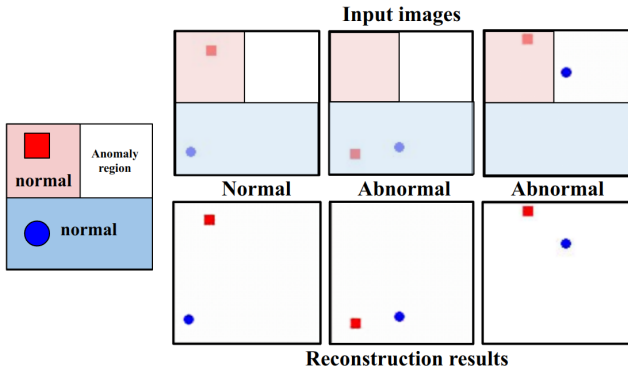


Figure 3. A CNN-based autoencoder fails to learn spatial context due to its shift-invariant nature. The autoencoder can perfectly reconstruct both normal and abnormal samples without understanding what is allowable in each region.

Object-centric Methods. Methods that use pre-trained networks to extract highly discriminative object-centric features often demonstrate superior performance on traditional benchmark datasets [7, 11, 12, 15, 28]. However, their ability to learn spatial context is relatively limited since many benchmark datasets contain only a few anomalous examples that depend on spatial context.

Context Learning in VAD. As VAD algorithms’ performance improves, the state of the art on several benchmarks is basically saturated (e.g., over 99.0 AUC on the UCSD Ped2 dataset). Recent research has pivoted to more nuanced normalcy modeling to detect context-dependent anomalies. Bao *et al.* [1] proposed pixel-level clustering to generate pseudo-labels to train a background segmentation model to help an object-centric auto-encoder predict the next frame. Sun *et al.* [28] proposed a pre-trained segmentation model to create scene-level background features concatenated with motion and appearance vectors to perform scene-aware VAD. However, we must emphasize that spatial context should not be reduced to background appearance (for example, opposing lanes of traffic may have the same background appearance). Further, scene dependence only constrains the normality to be associated with specific cameras or scenes instead of regions within a camera. It is also important to decouple background appearance and spatial context to develop unbiased VAD algorithms.

Yang and Radke [38] proposed a contrastive learning method to learn spatial-temporal context on the frame and patch level, focusing on determining temporal context-dependent anomalies. Singh *et al.*’s EVAL [25] followed a similar strategy to the original Street Scene paper [23], separating a video clip into hundreds of overlapping spatial tubes and storing region-specific exemplar features extracted from a set of attribute networks to detect spatially

dependent anomalies. The follow-on method T-EVAL [26] stores tracklets with corresponding visual features as exemplars and uses nearest neighbor search to detect anomalies. It is worth noting that the search is performed within a 3x3 grid region centered on the query exemplar location to speed up the search process. This can still be considered grid-based modeling. Although there is no dedicated model for each region, when the number of training videos increases, both storage and computational complexity become unmanageable. Our method discovers spatial semantic regions that drastically decrease the number of dedicated models required to cover the scene and achieves better VAD performance compared to these methods.

Explainability in VAD. Explainability in VAD is a desirable property addressed by prior work in various ways [26, 29, 35]. Szymanowicz *et al.* [29] used a pre-trained action recognition module to provide action labels for each detected anomaly bounding box as a way to explain the anomalous events. Wu *et al.* and EVAL [26] both explain anomalies using high-level attribute features. Recently, several approaches have adopted Large Language Models (LLMs) and Multimodal Large Language Models (MLLMs) for the VAD task [2, 18, 39]. While both LLMs and MLLMs show promise in reasoning and explainability, their success in VAD tasks is due more to the huge overlap between commonsense knowledge encoded in those models and the definition of abnormality in benchmark datasets. Moreover, such commonsense knowledge might introduce harmful bias in determining anomalies using appearance, background bias, or lack of context for the specific area that the camera covers. Our proposed method can provide statistical patterns of appearance and motion that could be further passed to LLMs or MLLMs for natural-language-based explainable VAD.

Region discovery. An older line of work investigates region discovery strategies in video [3, 31, 32, 42]. For example, [3] uses a pixel-level foreground ratio within a time window to find similar activity regions through correlation. Wang *et al.* [32] uses a thermal-diffusion based approach on the optical flow field with clustering to create semantic regions. However, this region clustering is performed at the pixel level and can mainly handle collectively moving dense particles or large objects with a unified motion. Our region discovery method is based on object-level activity and has no requirement about objects’ collective motion or size.

3. Method

As stated in the thought experiment, our goal is to build a model that approximates the “region-object-attribute” idea. To achieve this, our method has three key stages: object representation, region discovery, and anomaly detection, as illustrated in Figures 4 and 5. We rely on the object representation to discover semantically meaningful regions and

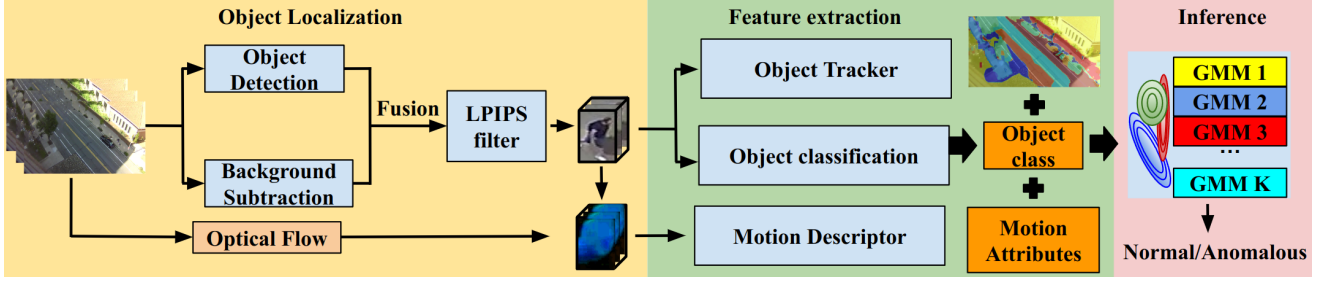


Figure 4. Overview of the proposed system inference pipeline. Extracted object level features are assigned to different GMM models based on region discovery results.

to train normalcy models. To define the problem, suppose we have a collection of N training video clips from the same stationary camera, $\{\mathbf{v}_1, \dots, \mathbf{v}_N\} \in \mathbb{R}^{T \times H \times W \times C}$, where T, H, W, C denote the temporal window size, height, width and color channel respectively. We will create a $H \times W$ map corresponding to K spatial regions, each of which contains a similar set of activities. The learned model in each region then forms the basis for spatial-context-dependent anomaly detection.

3.1. Object representation

Many object-centric VAD methods [1, 7, 15, 24, 28] use pretrained networks to extract various visual features. We also follow this approach, forming tracklets and extracting object attributes from detected object regions.

Object region proposal. We first feed each video clip \mathbf{x} into a pre-trained YOLOv8 model [30] to obtain object bounding boxes. Due to occlusions, varying lighting conditions, and the top-down perspective in the Street Scene dataset, solely relying on bounding boxes from object detection results in many missed detections. To capture as many bounding boxes as possible, we merge the detection results with foreground objects detected using a simple Gaussian mixture background model.

False region removal. Using the fused detection strategy above provides excellent bounding box proposals that include most of the ground-truth objects. However, both object detection and foreground subtraction introduce a significant number of region proposals that only cover pieces of background regions due to false detections. This problem can be greatly mitigated by using LPIPS [41] to compare each region proposal with the learned background, retaining proposals with LPIPS loss greater than a threshold α .

Object class descriptor. Although region proposals from object detection contain category information, region proposals from background subtraction do not. For each detected bounding box, we apply a pre-trained CLIP vision encoder (ViT-B32) [21] and use the CLS token from its output as an appearance feature vector $\mathbf{f}_{app} \in \mathbb{R}^{512}$. We perform zero-shot classification with a set of pre-selected ex-

pected classes (person, bicycle, car, motorcycle) with the text prompt templates used for Imagenet classification in [21]. We use one-hot encoding to represent the object class information $\mathbf{f}_{obj} \in \mathbb{R}^4$ as a compact and highly abstract representation for the following process.

Motion descriptor. We compute the Farneback optical flow [5] of the whole clip $\mathbf{v}_{flow} \in \mathbb{R}^{(T-1) \times H \times W \times 2}$ to compute motion attributes. For each detected bounding box, we use Histogram-of-Flow (HOF) to summarize the motion \mathbf{v}_{flow} within the box. We use 12 uniformly spaced orientation bins and 1 background ratio bin. Pixels whose magnitude are below a certain threshold (1.5 in our experiment) are considered as background pixels and aggregated in a single bin. The average speed is calculated for each orientation bin. The motion attributes \mathbf{f}_{mot} for the training/inference normalcy model are the dominant orientation and corresponding speed of the given tracklet resulting in $\mathbf{f}_{mot} \in \mathbb{R}^2$.

Tracklet formulation. To incorporate temporal information, we use the DeepSORT algorithm [34] to associate bounding box identities across time with the extracted CLIP embedding \mathbf{f}_{app} . Since the full trajectory lengths might vary, we create tracklets $\mathbf{l} \in \mathbb{R}^{2 \times t_w}$ with a fixed length window t_w . Thus for each detected object we formulate a combined feature $\mathbf{O}_t^i = [\mathbf{f}_{obj}^i, \mathbf{f}_{mot}^i, \mathbf{l}^i]_t$ at each time stamp t , which we use as the basis for the following process.

3.2. Region Discovery Strategies

In frame-based methods, each spatial location has an equal amount of training data. This is not the case when considering the alternative object-centric route. The region representation for object-centric methods requires more careful design to both reduce the number of separate models and improve detection performance in regions where a limited number of samples is available.

Grid Region Separation. The most straightforward way to introduce spatial context in object-centric VAD is to split the image into a non-overlapping rectangular grid, design a model for each rectangle, and make a frame-level determination about each object based on the rectangle containing the center point of its detected bounding box.

In our experiment, we found that using small patches (e.g., 40×40 as used in [23]) results in a large number of patches that don't contain a single example; thus no model can be trained for those regions. We chose an 80×80 rectangle size as our baseline in our comparisons below. This large patch setting still generates 144 separate models.

Region Discovery. Our key innovation is to cluster regions with similar appearance/motion patterns, resulting in a substantial decrease in the number of models to be trained. Our clustering approach is agnostic to the type of feature representation.

To create a high-resolution (around 1 million pixel) region map, we need to first discover what object-level events happen at each pixel location. For each object tracklet moving through the pixel, we determine its dominant motion direction and quantize its speed into 4 log-scale-spaced magnitude bins. Thus for each object, we acquire its class information, dominant motion direction, and speed. We initialize an empty heatmap $\mathbf{M} \in \mathbb{R}^{H \times W \times D}$ with the same size as the original frame and D channels corresponding to the given attributes, which we will incrementally update using the training data.

To provide enough information to create a high-resolution map, we center a Gaussian kernel around the tracklet pixel x_c, y_c , using the σ parameter to control the spatial influence of each moving object on surrounding pixels (and zeroing out the contribution outside the object bounding box). We then add this to the incrementally constructed heatmap \mathbf{M} in (Eqn. 1).

$$\mathbf{G}(x, y) = \exp\left(-\frac{(x - x_c)^2 + (y - y_c)^2}{2\sigma^2}\right)$$

$$\mathbf{1}_{\text{box}}(x, y) = \begin{cases} 1 & \text{if } x_{tl} \leq x \leq x_{br} \text{ and } y_{tl} \leq y \leq y_{br}, \\ 0 & \text{otherwise,} \end{cases}$$

$$\mathbf{M}(x, y, d) \leftarrow \mathbf{M}(x, y, d) + \mathbf{1}_{\text{box}}(x, y) \cdot \mathbf{G}(x, y) \quad (1)$$

After constructing \mathbf{M} , we need to discover the regions that share similar activities. To capture the relationship between attribute channels, we use a Gaussian Mixture model with full covariance (Eqn. 2) to learn K modes from \mathbf{M} . In Section 4 we verify the importance of using full covariance compared to other alternatives. The learned modes can be used as the labels for each discovered region, as shown in Fig. 5.

$$p(\mathbf{f}) = \sum_{i=1}^K \pi_i \cdot p_i(\mathbf{f}) \quad \text{with} \quad p_i(\mathbf{f}) \sim \mathcal{N}(\mu_i, \Sigma_i) \quad (2)$$

3.3. Regional Model

We use another Gaussian Mixture Model (GMM) as the basis for representing the normality distribution of the features in each region. Different regions may exhibit more




Figure 5. The region discovery process. We form a heatmap over all training objects that belong to a certain location. After collecting the histogram for each region, we cluster and form semantic regions.

or less appearance/motion diversity in the objects that pass through them, resulting in variability in the number of components required for the GMM to adequately represent the underlying data distribution. In our experiments, we use the Bayesian Information Criteria (BIC) score to guide the selection of the number of Gaussian modes for each region.

In order to select the optimal K in region discovery, we propose a simple evaluation metric based on symmetric KL divergence. A good region discovery method should result in regions that contain different normalcy distributions. Thus, each region's $p_i(f)$ normalcy model should assign a low likelihood to normal features from other regions. For all region pairs $i, j \in K, i \neq j$, we use the average symmetric KL-divergence μ_{KL}^K (Eqn. 3) of sample scores to measure the region separation quality.

$$D_{\text{sym}}(P_i, P_j) = D_{\text{KL}}(P_i \| P_j) + D_{\text{KL}}(P_j \| P_i)$$

$$\mu_{KL}^K = \frac{1}{K(K-1)} \sum_{i \neq j} D_{\text{sym}}(P_i, P_j) \quad (3)$$

In contrast to rectangular grids, in which each region may only contain a few objects' features as training samples, using our larger discovered regions drastically increases the number of training samples for each regional model.

Learned Regional Normalcy. As a byproduct of our algorithm, we can use the cluster results and the learned mixtures from GMMs to visualize the prototypical normal activities within each region. For example, we can find the object in the training data whose feature is nearest to the mean of each Gaussian in a region's mixture, as in the right hand side of Figure 2.

3.4. Inference

During inference, we extract object tracklets and features just as in the training process. We assign objects to their enclosing region (either a rectangular box as in comparison methods or the non-rectangular regions discovered by our method) based on the first center position of their tracklets and evaluate using the corresponding model shown in Figure 4. For GMM models, we use the negative log-likelihood (NLL) as the indicator of anomaly. To compute a frame-level anomaly score, we simply select the highest NLL score from tracklets within each frame.

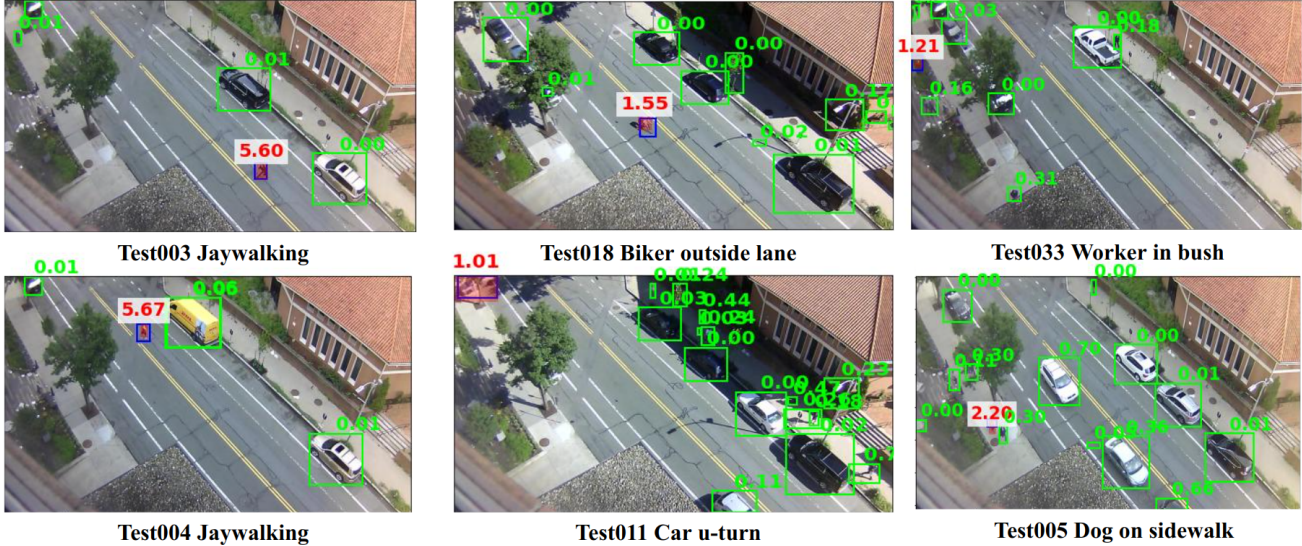


Figure 6. Examples of anomaly detection results. Green boxes indicate normal objects, blue boxes indicate ground truth anomalies, and red indicates detected anomalies. Anomaly scores for detected objects are shown above the bounding boxes.

4. Experimental Results

4.1. Dataset and Evaluation Metrics

Datasets. We mainly evaluate our method on the Street Scene dataset [23] since it is one of the few publicly available datasets to address spatial-context anomalies. This contains 46 training and 35 testing videos under a single-scene setting with 720×1280 resolution. The training and testing videos contain 56,847 and 146,410 frames respectively. The normal videos contain typical events like pedestrians walking and vehicles/cyclists moving properly in their lanes. However, the testing videos contain 17 types of spatial-context anomalies such as jaywalking, cars outside the correct lane, and cyclists on the sidewalk. We also report our performance on the widely adopted Ped2 [33] and ShanghaiTech [14] datasets. Since ShanghaiTech contains 13 different scenes, our method needs to create region maps for each camera and train the models separately.

Evaluation criteria. We use both the standard frame level Area-Under-Curve (AUC) [33] and the Region/Track-based Detection Criteria (RBDC/TBDC) proposed in [23]. For detailed definition, we refer readers to [22]. We report Micro-AUC as the frame level evaluation metric that concatenates all test videos and computes the AUC score. RBDC and TBDC measure the AUC of the region/track matching rate across false positive rates ranging from 0 to 1. While [23] highlights the shortcomings of using the frame-level AUC metric (e.g., its inability to account for multiple false positives or multiple anomalies within a frame), RBDC and TBDC also have limitations. Most studies that employ RBDC/TBDC metrics in their evaluations adhere to the original thresholds: a region match is defined as having at least 0.1 intersection-over-union between the predicted

and ground truth regions, and a track match is considered successful if at least 10% of the anomalous trajectory is detected.

4.2. Implementation details

We provide implementation details in this section, for more detailed information, please refer to supplementary.

Object detection fusion. The detection results are created from both a pre-trained YOLOv8-x [30] model and background subtraction. To obtain foreground objects, we initialize a Mixture of Gaussians background model [27] with variance threshold set to 16 at every individual video clip.

Acquiring tracklets. We set the DeepSORT algorithm with a minimum of 3 frames associated for each track. t_w is set to 3 for Street Scene and ShanghaiTech and 1 for UCSD Ped2.

Region discovery and normalcy model settings. The GMMs for region discovery and for normalcy modelling are set to use full covariance. The maximum component per region of BIC search is set to 20. We search the parameter K over the range [2, 16]. We apply Gaussian filtering to smooth the anomaly score.

4.3. Quantitative Analysis

Table 1 summarizes our algorithm’s performance on the context-dependent Street Scene dataset. We use a 9×16 grid (i.e., 80×80 pixel regions) as a baseline, denoted “Grid Sep.” in the table. Our method outperforms the previous algorithms in terms of the frame-level AUC and the RBDC criterion. While other methods achieve higher TBDC, our algorithm demonstrates competitive TBDC with 1-2 orders of magnitude fewer regions and compact feature representation. This indicates that our feature design can enable con-

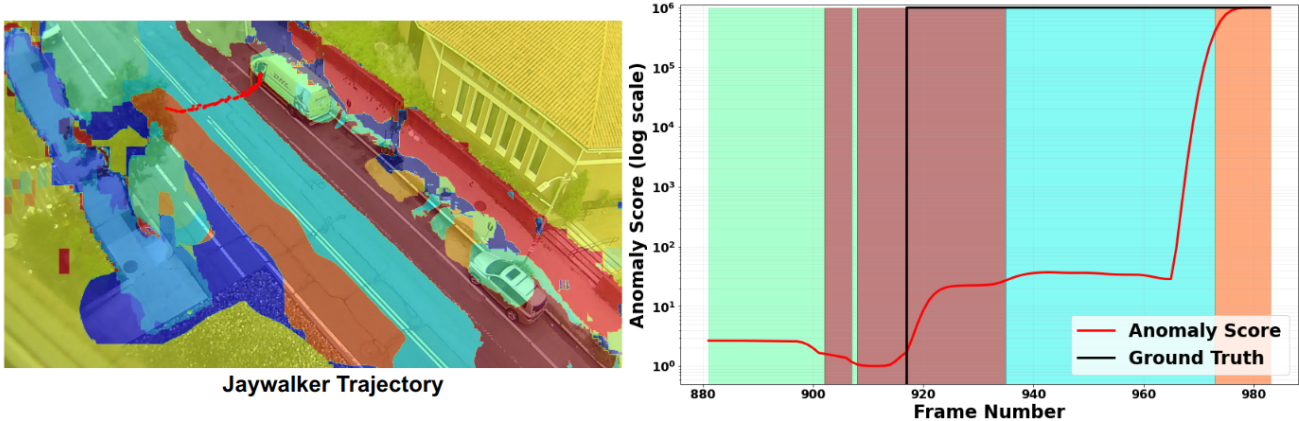


Figure 7. An example of the anomaly detection response to a jaywalker trajectory in **Test004**. The jaywalker exits a car to the bike lane and walks across the traffic lane. **Left**: The red trajectory indicates the jaywalker trajectory captured by our method. **Right**: The anomaly detection scores for this object trajectory. The background colors of the graph correspond to distinct spatial regions automatically obtained from region discovery. The jaywalker motion is normal in the parking region and bike lane (light green/maroon backgrounds) but our method generates a high anomaly score when the jaywalker is in the traffic lane regions (light blue and orange backgrounds).

sistent detection for each anomaly trajectory.

Method	Reg./Dim.	RBDC	TBDC	AUC
Flow [23]	2304/11200	11.0	52.0	51.0
FG [23]	2304/22400	21.0	53.0	61.0
EVAL [25]	900/513	24.3	<u>64.5</u>	-
T-EVAL [26]	- /150	<u>30.9</u>	72.9	-
Grid Sep.	144/6	25.9	56.1	70.2
Region Disc.	12/6	34.0	62.5	67.0

Table 1. Performance comparison on the Street Scene Dataset; Bold indicates the best performance and underline indicates the second best performance. Reg. and Dim. stand for the number of regions and dimension of feature.

Method	Venue	RBDC	TBDC	AUC
AED [7]	<i>PAMI21</i>	69.2	93.2	98.7
HF2VAD [15]	<i>ICCV21</i>	-	-	99.3
SSMTL [6]	<i>CVPR21</i>	72.8	91.2	97.5
EVAL [25]	<i>CVPR22</i>	87.4	95.1	-
HSC [28]	<i>CVPR23</i>	-	-	98.2
T-EVAL [26]	<i>CVPRW24</i>	84.2	<u>97.9</u>	-
Ours	-	<u>84.8</u>	99.1	<u>99.0</u>

Table 2. Performance comparison on the UCSD Ped2 Dataset.

Although the main focus of our proposed method is to detect spatial-context dependent anomalies, we also test our algorithm on the UCSD Ped2 dataset [33] (Table 2) and the ShanghaiTech dataset [14] (Table 3). Our method uses a simple 6-dimensional feature to achieve state-of-the-art performance in video anomaly tracking and second best video anomaly localization in UCSD Ped2, further illustrating the effectiveness and compactness of our proposed method. Our model also shows competitive performance

on the ShanghaiTech dataset.

Method	Venue	RBDC	TBDC	AUC
AED [7]	<i>PAMI21</i>	41.3	78.8	82.7
HF2VAD [15]	<i>ICCV21</i>	-	-	76.1
SSMTL [6]	<i>CVPR21</i>	42.8	83.9	83.4
EVAL [25]	<i>CVPR22</i>	<u>59.2</u>	89.4	76.6
HSC [28]	<i>CVPR23</i>	-	-	83.4
T-EVAL [26]	<i>CVPRW24</i>	59.6	87.6	-
Ours	-	46.8	73.2	81.3

Table 3. Performance comparison on the ShanghaiTech Dataset.

4.4. Qualitative examples

Figure 6 illustrates successful anomaly detection results from our algorithm on the Street Scene dataset, demonstrating its ability to detect difficult examples like a parking attendant ticketing a car and a dog on the sidewalk. The same anomaly threshold is used in all examples of the figure.

We drill down into the anomaly detection criterion in Figure 7, illustrating the response as a pedestrian exits a car into the bike lane and later walks across the traffic lanes. Within the bike lane and parking region (maroon), the anomaly score depends on the jaywalker’s motion attributes. When the jaywalker appears in the traffic lanes, our method immediately generates a high anomaly score.

4.5. Parameter analysis

In this section, we validate our choices for important components and parameters in our proposed method.

Region proposal vs. semantic segmentation. To evaluate the effectiveness of our region proposal, we directly compare the performance in the Street Scene dataset with region maps generated from our region proposal and the off-the-shelf semantic segmentation model SegFormer [36] (with

the largest variant B5 pretrained on ADE20K [43]) shown in Table 4. We use the average logit output of all training videos with blocking foreground objects to generate the region map (SegFormer) of a single scene. We also mimic the process of our region proposals to cluster the output logits into the same number of clusters (SegFormer clustering). Since the segmentation model lacks motion cues, there is a significant performance improvement of our region proposals over the off-the-shelf semantic segmentation model. For more details, see the supplementary material.

Region map	RBDC	TBDC	AUC
SegFormer	20.5	51.5	59.3
SegFormer clustering	29.1	46.7	61.3
Ours	34.0	62.5	67.0

Table 4. Performance using different region proposals in the Street Scene Dataset.

Deep features vs. object class for spatial context-dependent anomalies. As we mentioned above, using object classifications instead of deep appearance features can significantly reduce computational load and the number of model parameters. Table 5 compares the performance using full CLIP features [21] vs. the one-hot classification of expected objects in the scene. We see that the performance is actually very similar.

Optimal region discovery for spatial context-dependent anomalies. The performances for different numbers of regions are summarized in Table 6, showing that 12 region model using KL-divergence criteria μ_{KL} achieves both the highest divergence and best performance in our experiments. We also report the divergence in the grid separation scenario; the low divergence further indicates the redundancy of having overly granular regions.

To verify the effectiveness of different attribute features, we also performed an ablation study in Table 7. We can see

Feature type	Feature Dim.	RBDC	TBDC
CLIP	512	33.1	60.1
Object CLS	4	34.0	62.5

Table 5. Performance comparison between using CLIP features and object classification in the Street Scene Dataset.

Number of Regions	μ_{KL}	RBDC	TBDC	AUC
6	9.23	21.2	48.9	61.0
8	10.21	25.9	50.9	64.0
10	11.64	28.9	53.9	66.0
12	14.75	34.0	62.5	67.0
16	13.47	29.9	57.9	67.0
144	4.99	25.9	56.1	70.2

Table 6. Performance using different numbers of discovered regions in Street Scene Dataset.

that all attributes are required for the best performance, and we observe that the different attributes often complement each other.

Obj. Cls.	Ori.	Mag.	RBDC	TBDC	AUC
✓			31.2	45.1	67.9
	✓		29.0	49.8	63.3
		✓	10.7	52.4	62.0
✓		✓	26.2	48.0	66.0
✓	✓		35.6	56.7	73.1
✓	✓	✓	34.0	62.5	67.0

Table 7. Performance using different feature branches in Street Scene dataset. Obj., Cls., Ori. and Mag. stand for object, class, motion orientation, and motion magnitude respectively.

5. Conclusions

We proposed a simple but effective method for detecting spatial-context anomalies using region discovery. There are several limitations of the proposed method as shown in Figure 8. Since our model relies on appearance and motion features within a short time window, its ability to detect long-term trajectory-oriented anomalies like soliciting is limited. This also partially explains why we achieve state-of-the-art performance on the region-based criteria RBDC but not on the track-based criteria TBDC. We identified several opportunities to improve the proposed algorithm in future work. First, the current hard region assignment might magnify the impact of incorrect region clustering. It would be helpful to explore soft region assignment which should alleviate this problem. For more person action oriented datasets like ShanghaiTech, incorporating pose features [11, 24] could further boost performance.




Figure 8. Examples of false alarms and misses. Green boxes indicate normal objects, blue boxes indicate ground truth anomalies, and red indicates detected anomalies.

6. Acknowledgement

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award 22STESE00001-04-00. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

References

- [1] Qianye Bao, Fang Liu, Yang Liu, Licheng Jiao, Xu Liu, and Lingling Li. Hierarchical Scene Normality-Binding Modeling for Anomaly Detection in Surveillance Videos. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6103–6112. ACM, 2022. 3, 4
- [2] Rohit Bharadwaj, Hanan Gani, Muzammal Naseer, Fahad Shahbaz Khan, and Salman Khan. VANE-Bench: Video Anomaly Evaluation Benchmark for Conversational LMMs, June 2024. arXiv:2406.10326 [cs]. 3
- [3] Chen Change Loy, Tao Xiang, and Shaogang Gong. Multi-camera activity correlation analysis. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1988–1995, Miami, FL, June 2009. IEEE. 3
- [4] MyeongAh Cho, Taeoh Kim, Woo Jin Kim, Suhwan Cho, and Sangyoung Lee. Unsupervised video anomaly detection via normalizing flows with implicit latent features. *Pattern Recognition*, 129:108703, 2022. 2
- [5] Gunnar Farnesk. Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings 13*, pages 363–370. Springer, 2003. 4
- [6] Mariana-Iuliana Georgescu, Antonio Barbalau, Radu Tudor Ionescu, Fahad Shahbaz Khan, Marius Popescu, and Mubarak Shah. Anomaly Detection in Video via Self-Supervised and Multi-Task Learning. In *CVPR*, pages 12737–12747, 2021. 7
- [7] Mariana Iuliana Georgescu, Radu Ionescu, Fahad Shahbaz Khan, Marius Popescu, and Mubarak Shah. A Background-Agnostic Framework with Adversarial Training for Abnormal Event Detection in Video. *IEEE TPAMI*, pages 4505–4523, 2021. 3, 4, 7
- [8] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton Van Den Hengel. Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection. In *ICCV*, pages 1705–1714, 2019. 1, 2
- [9] Anirudh Goyal and Yoshua Bengio. Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 478(2266):20210068, Oct. 2022. 1
- [10] 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. 1, 2
- [11] Or Hirschorn and Shai Avidan. Normalizing Flows for Human Pose Anomaly Detection. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13499–13508, Paris, France, Oct. 2023. IEEE. 3, 8
- [12] 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 7834–7843, 2019. 3
- [13] Wenrui Liu, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Diversity-measurable anomaly detection. In *CVPR*, pages 12147–12156, 2023. 1, 2
- [14] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection—a new baseline. In *CVPR*, pages 6536–6545, 2018. 1, 2, 6, 7
- [15] 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 *ICCV*, pages 13588–13597, 2021. 3, 4, 7
- [16] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal Event Detection at 150 FPS in MATLAB. In *ICCV*, pages 2720–2727, 2013. 1, 2
- [17] Hui Lv, Chen Chen, Zhen Cui, Chunyan Xu, Yong Li, and Jian Yang. Learning normal dynamics in videos with meta prototype network. In *CVPR*, pages 15425–15434, 2021. 1, 2
- [18] Hui Lv and Qianru Sun. Video Anomaly Detection and Explanation via Large Language Models, Jan. 2024. arXiv:2401.05702 [cs]. 3
- [19] Jakub Micorek, Horst Possegger, Dominik Narnhofer, Horst Bischof, and Mateusz Kozinski. Mulde: Multiscale log-density estimation via denoising score matching for video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18868–18877, 2024. 1
- [20] Hyunjong Park, Jongyoung Noh, and Bumsub Ham. Learning Memory-Guided Normality for Anomaly Detection. In *CVPR*, pages 14360–14369, 2020. 1, 2
- [21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 4, 8
- [22] Bharathkumar Ramachandra, Michael Jones, and Ranga Raju Vatsavai. A Survey of Single-Scene Video Anomaly Detection. *IEEE TPAMI*, 44(5):2293–2312, 2020. 6
- [23] Bharathkumar Ramachandra and Michael J. Jones. Street Scene: A new dataset and evaluation protocol for video anomaly detection. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2558–2567, 2020. 1, 2, 3, 5, 6, 7
- [24] Tal Reiss and Yedid Hoshen. Attribute-based Representations for Accurate and Interpretable Video Anomaly Detection, Dec. 2022. arXiv:2212.00789 [cs]. 1, 4, 8
- [25] Ashish Singh, Michael J Jones, and Erik G Learned-Miller. Eval: Explainable video anomaly localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18717–18726, 2023. 1, 3, 7
- [26] Ashish Singh, Michael J Jones, and Erik G Learned-Miller. Tracklet-based explainable video anomaly localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3992–4001, 2024. 3, 7
- [27] Chris Stauffer and W Eric L Grimson. Adaptive background mixture models for real-time tracking. In *Proceedings. 1999*

- IEEE computer society conference on computer vision and pattern recognition (Cat. No PR00149)*, volume 2, pages 246–252. IEEE, 1999. 6
- [28] Shengyang Sun and Xiaojin Gong. Hierarchical Semantic Contrast for Scene-aware Video Anomaly Detection. In *CVPR*, pages 22846–22856, 2023. 1, 3, 4, 7
- [29] Stanislaw Szymanowicz, James Charles, and Roberto Cipolla. Discrete neural representations for explainable anomaly detection. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1506–1514, Waikoloa, HI, USA, Jan. 2022. IEEE. 3
- [30] Rejin Varghese and Sambath M. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6, 2024. 4, 6
- [31] Qi Wang, Mulin Chen, Feiping Nie, and Xuelong Li. Detecting Coherent Groups in Crowd Scenes by Multiview Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1):46–58, Jan. 2020. 3
- [32] Weiyue Wang, Weiyao Lin, Yuanzhe Chen, Jianxin Wu, Jingdong Wang, and Bin Sheng. Finding Coherent Motions and Semantic Regions in Crowd Scenes: A Diffusion and Clustering Approach. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8689, pages 756–771. Springer International Publishing, Cham, 2014. Series Title: Lecture Notes in Computer Science. 3
- [33] Weixin Li, Vijay Mahadevan, and Nuno Vasconcelos. Anomaly Detection and Localization in Crowded Scenes. *IEEE TPAMI*, 36(1):18–32, Jan. 2014. 1, 6, 7
- [34] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017. 4
- [35] Chongke Wu, Sicong Shao, Cihan Tunc, Pratik Satam, and Salim Hariri. An explainable and efficient deep learning framework for video anomaly detection. *Cluster Computing*, 25(4):2715–2737, Aug. 2022. 3
- [36] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34:12077–12090, 2021. 7
- [37] Zhiwei Yang, Jing Liu, Zhaoyang Wu, Peng Wu, and Xiaotao Liu. Video event restoration based on keyframes for video anomaly detection. In *CVPR*, pages 14592–14601, 2023. 2
- [38] Zhengye Yang and Richard J. Radke. Context-aware Video Anomaly Detection in Long-Term Datasets. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4002–4011, Seattle, WA, USA, June 2024. IEEE. 3
- [39] Huaxin Zhang, Xiaohao Xu, Xiang Wang, Jialong Zuo, Chuchu Han, Xiaonan Huang, Changxin Gao, Yuehuan Wang, and Nong Sang. Holmes-VAD: Towards Unbiased and Explainable Video Anomaly Detection via Multi-modal LLM, June 2024. arXiv:2406.12235 [cs]. 3
- [40] Menghao Zhang, Jingyu Wang, Qi Qi, Haifeng Sun, Zirui Zhuang, Pengfei Ren, Ruilong Ma, and Jianxin Liao. Multi-scale video anomaly detection by multi-grained spatio-temporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17385–17394, 2024. 2
- [41] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595. IEEE, June 2018. 4
- [42] Bolei Zhou, Xiaoou Tang, Hepeng Zhang, and Xiaogang Wang. Measuring Crowd Collectiveness. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1586–1599, Aug. 2014. 3
- [43] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene Parsing through ADE20K Dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, Honolulu, HI, July 2017. IEEE. 8