Delving into CLIP latent space for Video Anomaly Recognition

Luca Zanella[◇a,**], Benedetta Liberatori[◇a], Willi Menapace[a], Fabio Poiesi[b], Yiming Wang[b], Elisa Ricci[a,b]

[a]*University of Trento, Trento, Italy*
[b]*Fondazione Bruno Kessler, Trento, Italy*

ABSTRACT

We tackle the complex problem of detecting and recognising anomalies in surveillance videos at the frame level, utilising only video-level supervision. We introduce the novel method *AnomalyCLIP*, the first to combine Large Language and Vision (LLV) models, such as CLIP, with multiple instance learning for joint video anomaly detection and classification. Our approach specifically involves manipulating the latent CLIP feature space to identify the normal event subspace, which in turn allows us to effectively learn text-driven directions for abnormal events. When anomalous frames are projected onto these directions, they exhibit a large feature magnitude if they belong to a particular class. We also introduce a computationally efficient Transformer architecture to model short- and long-term temporal dependencies between frames, ultimately producing the final anomaly score and class prediction probabilities. We compare *AnomalyCLIP* against state-of-the-art methods considering three major anomaly detection benchmarks, *i.e.* ShanghaiTech, UCF-Crime, and XD-Violence, and empirically show that it outperforms baselines in recognising video anomalies. Project website and code are available at https://luca-zanella-dvl.github.io/AnomalyCLIP/.

1. Introduction

Video anomaly detection (VAD) is the task of automatically identifying activities that deviate from normal patterns in videos (Suarez and Naval Jr, 2020). VAD has been widely studied by the computer vision and multimedia communities (Bao et al., 2022; Feng et al., 2021b; Mei and Zhang, 2017; Nayak et al., 2021; Sun et al., 2022; Wang et al., 2020; Xu et al., 2019) for several important applications, such as surveillance (Sultani et al., 2018) and industrial monitoring (Roth et al., 2022).

VAD is challenging because data is typically highly imbalanced, *i.e.* normal events are many, whilst abnormal events are rare and sporadic. VAD can be addressed as an out-of-distribution detection problem, *i.e.* one-class classification (OOC) (Liu et al., 2021; Lv et al., 2021; Park et al., 2020; Xu et al., 2019): only visual data corresponding to the normal state is used as training data, and an input test video is classified as normal or abnormal based on its deviation from the learnt

normal state. However, OOC methods can be particularly ineffective in complex real-world applications where normal activities are diverse. An uncommon normal activity may cause a false alarm because it differs from the learnt normal activities. Alternatively, VAD can be addressed with fully-supervised approaches based on frame-level annotations (Bai et al., 2019; Wang et al., 2019). Despite their good performance, they are considered impractical because annotations are costly to produce. Unsupervised approaches can also be used, but their performance in complex settings is not yet satisfactory (Zaheer et al., 2022). For these reasons, the most recent approaches are designed for weakly-supervised learning scenarios (Li et al., 2022a; Sultani et al., 2018; Tian et al., 2021; Wu and Liu, 2021): they exploit video-level supervision.

Whilst existing weakly-supervised VAD methods have shown to be effective in anomaly detection (Li et al., 2022a), they are not designed for recognising anomaly types (*e.g.* shooting vs. explosion). Performing Video Anomaly Recognition (VAR) in addition to VAD, that is not only detecting anomalous events but also recognising the underlying activities, is desirable as it provides more informative and actionable insights. However, addressing VAR in a weakly-supervised setting is highly challenging due to the extreme data imbalance

---

**Corresponding author:
 e-mail: luca.zanella-3@unitn.it (Luca Zanella◇)
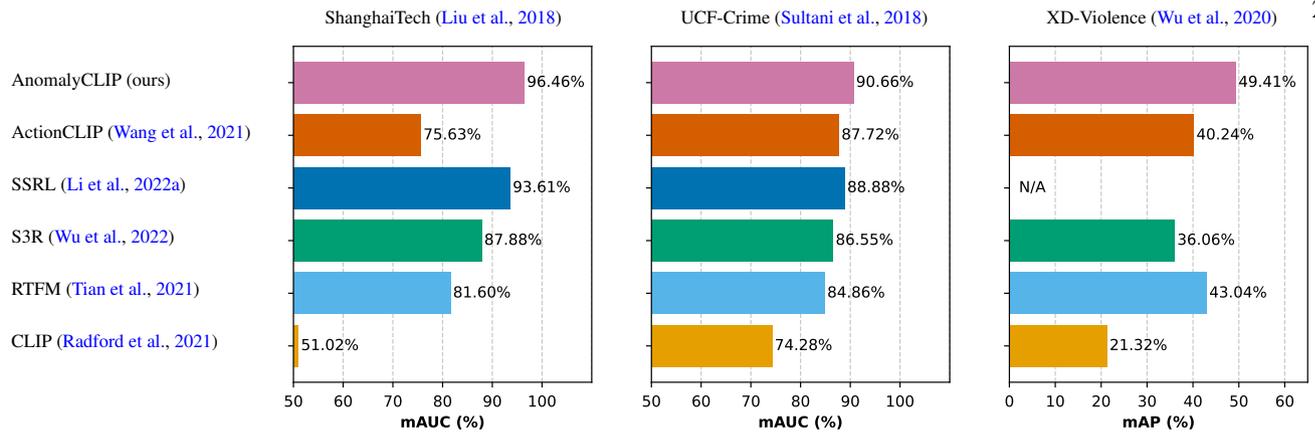 ◇Equal contribution.

Fig. 1: Comparison of various anomaly recognition methods on the ShanghaiTech, UCF-Crime, and XD-Violence datasets in terms of the mean area under the curve (mAUC) of the receiver operating characteristic (ROC) and the mean average precision (mAP) of the precision-recall curve (PRC), which calculate the mean of binary AUC ROC and AP PRC values for all anomalous classes, respectively. A higher mAUC and mAP are crucial for video anomaly recognition as they reflect the model's ability in correctly recognising the correct abnormal class. Notably, our proposed method, AnomalyCLIP, achieves the highest performance on all datasets, surpassing both the state-of-the-art methods on video anomaly detection that are re-purposed for anomaly recognition and CLIP-based video action recognition methods.

and the limited samples representing each anomaly (Sultani et al., 2018).

We have recently experienced the emergence of powerful deep learning models that are trained on massive web-scale datasets (Schuhmann et al., 2021). These models, commonly referred to as Large Language and Vision (LLV) models or foundation models (Radford et al., 2021; Singh et al., 2022), have shown strong generalisation capabilities in several downstream tasks and have become a key ingredient of modern computer vision and multimedia systems. These pre-trained models are publicly available and can be seamlessly integrated into any recognition system. LLV models can also be effectively applied to videos and to supervised action recognition tasks (Wang et al., 2021; Xu et al., 2021).

In this paper, we introduce the first method that jointly addresses VAD and VAR with LLV models. We argue that by leveraging representations derived from LLV models, we can obtain more discriminative features for recognising and classifying abnormal behaviours. However, as supported by our experiments (Fig. 1), a naive application of existing LLV models to VAR-VAD does not suffice due to the imbalance of the training data and the subtle differences between frames of the same video containing and non containing anomalous contents.

Therefore, we propose AnomalyCLIP, a novel solution for VAR based on the CLIP model (Radford et al., 2021), achieving state-of-the-art anomaly recognition performance as shown in Fig. 1.

AnomalyCLIP produces video representations that can be mapped to the textual description of the anomalous event. Rather than directly operating on the CLIP feature space, we re-centre it around a normality prototype, as shown in Fig. 2 (a). In this way, the space assumes important semantics: the magnitude of the features indicates the degree of anomaly, while the direction from the origin indicates the anomaly type. To learn the directions that represent the desired anomaly classes, we propose a Selector model that employs prompt learning and a projection operator tailored to our new space to identify the parts in a video that better match the textual description of the

anomaly. This ability is instrumental to address the data imbalance problem. We use the predictions of the Selector model to implement a semantically-guided Multiple Instance Learning (MIL) strategy that aims to widen the gap between the most anomalous segments of anomalous videos and normal ones. Differently from the features typically employed in VAD that are extracted using temporal-aware backbones (Carreira and Zisserman, 2017; Liu et al., 2022), CLIP visual features do not bear any temporal semantics as it operates at the image level. We thus propose a Temporal model, implemented as an Axial Transformer (Ho et al., 2019), which models both short-term relationships between successive frames and long-term dependencies between parts of the video.

As illustrated in Fig.1, we evaluate the proposed approach on three benchmark datasets, ShanghaiTech (Liu et al., 2018), UCF-Crime (Sultani et al., 2018) and XD-Violence (Wu et al., 2020), and empirically show that our method achieves state-of-the-art performance in VAR.

The contributions of our paper are summarised as follows:

- we propose the first method for VAR that is based on LLV models to detect and classify the type of anomalous events;
- we introduce a transformation of the LLV model feature space driven by a normality prototype to effectively learn the prompt directions for anomaly types;
- we propose a novel Selector model that uses semantic information imbued in the transformed LLV feature space as a robust way to perform MIL segment selection and anomaly recognition;
- we design a Temporal model to better aggregate temporal information by modelling both the short-term relationships between neighbouring frames and the long-term dependencies among segments.

2. Related Works

Video Anomaly Detection. Recognising anomalous behaviours in video surveillance streams is a traditional task in computer vision and multimedia analysis. Existing methods

for VAD can be grouped into four main categories based on the level of supervision available during training. The first group includes fully-supervised methods that assume available frame-level annotations in the training set (Bai et al., 2019; Wang et al., 2019). The second group includes weakly-supervised approaches that only require video-level normal/abnormal annotations (Li et al., 2022a,b; Sultani et al., 2018; Tian et al., 2021; Wu and Liu, 2021). The third group includes one-class classification methods that assume the availability of only normal training data (Liu et al., 2021; Lv et al., 2021; Park et al., 2020). The fourth group includes unsupervised models that do not use training data annotations (Narasimhan, 2018; Zaheer et al., 2022).

Amongst these types of methods, weakly-supervised approaches have gained higher popularity, as they typically yield good results while limiting the annotation effort. Sultani et al. (2018) were the first to formulate weakly-supervised VAD as a multiple-instance learning (MIL) task, dividing each video into short segments that form a set, known as *bag*. Bags generated from abnormal videos are called positive bags, and those generated from normal videos negative bags. Since this pioneering work, MIL has become a paradigm for VAD and several subsequent works have proposed to refine the associated ranking model to more robustly predict anomaly scores. For example, Tian et al. (2021) proposed a Robust Temporal Feature Magnitude (RTFM) loss that is applied to a deep network consisting of a pyramid of dilated convolutions and a self-attention mechanism to model both short-term and long-term relationships between video snippets close in time and events in the whole video. Wu et al. (2022) introduced Self-Supervised Sparse Representation Learning, an approach that combines dictionary-based representation with self-supervised learning techniques to identify abnormal events. Chen et al. (2022) introduced Magnitude-Contrastive Glance-and-Focus Network, a neural network that uses a feature amplification mechanism and a magnitude contrastive loss to enhance the importance of feature discriminative for anomalies. Motivated by the fact that anomalies can occur at any location and at any scale of the video, Li et al. (2022a) proposed Scale-Aware Spatio-Temporal Relation Learning (SSRL), an approach that extends RTFM by not only learning short-term and long-term temporal relationships but also learning multi-scale region-aware features. While SSRL achieves state-of-the-art results in common VAD benchmarks, its high computational complexity limits its applicability. To the best of our knowledge no previous works have explored foundation models (Radford et al., 2021) for VAD, as we propose in this work.

Large Language and Vision models. The emergence of novel large multimodal neural networks (Radford et al., 2021; Schuhmann et al., 2021, 2022; Singh et al., 2022), which can learn joint visual-text embedding spaces, has enabled unprecedented results in several image and video understanding tasks. Current LLV models adopt modality-specific encoders and are trained via contrastive techniques to align the data representations from different modalities (Jia et al., 2021; Radford et al., 2021). Despite their simplicity, these methods have been shown to achieve impressive zero-shot generalisation capabilities. While earlier approaches such as CLIP (Radford et al., 2021) operate on images, LLV models have recently and successfully been extended to the video domains. VideoCLIP (Xu et al., 2021) is an example of this and it is designed to align video and textual representations by contrasting temporally overlapping video-text pairs with mined hard negatives. VideoCLIP can achieve strong zero-shot performance in several video understanding tasks. ActionCLIP (Wang et al., 2021) models action recognition as a video-text matching problem rather than a classical 1-out-of-N majority vote task. Similarly to ours, their method uses the feature space of CLIP to learn semantically-aware representations of videos. However, a direct exploitation of the CLIP feature space fails in capturing information on anomalous events for which a specific adaptation, proposed in this work, is necessary. In addition, action recognition methods often fall short in weakly-supervised VAD tasks due to data imbalance between normal and abnormal events, coupled with the need for frame-level evaluation at test time, despite only having video-level supervision. To the best of our knowledge, no prior work has specifically utilised LLV models to tackle the VAD problem.

3. Proposed approach

Weakly-supervised VAD is the task of learning to classify each frame in a video as either normal or anomalous using a dataset of tuples in the form $(\mathbf{V}, y)$, where $\mathbf{V}$ is a video and y a binary label indicating whether the video contains an anomaly in any of its frames. With respect to VAD, in VAR we introduce the additional task of recognising the *type* of the detected anomaly in each frame. Therefore, VAR considers a dataset of tuples $(\mathbf{V}, c)$, where c indicates the type of anomaly in the video ($c = \emptyset$ means no anomaly is present, thus being *Normal*). In the following, we omit the subscripts for the purpose of readability.

To address the video-level supervision and the imbalance between normal videos and abnormal ones in VAD, the Multiple Instance Learning (MIL) framework (Sultani et al., 2018) is widely used. MIL models each video as a bag of segments $\mathbf{V} = [\mathbf{S}_1, ..., \mathbf{S}_S] \in \mathbb{R}^{S \times F \times D}$, where S is the number of segments, F is the number of frames in each segment, and D is the number of features associated to each frame. Each segment can be seen as $\mathbf{S} = [\boldsymbol{x}_1, ..., \boldsymbol{x}_F] \in \mathbb{R}^{F \times D}$ where $\boldsymbol{x} \in \mathbb{R}^D$ is the feature corresponding to each frame. MIL computes a likelihood of each frame being anomalous, selects the most anomalous ones based on it, and maximises the difference in the predicted likelihood between the normal frames and the ones selected as the most anomalous.

In this paper, we propose to leverage the CLIP model (Radford et al., 2021) to address VAR and show that:

i) the alignment between the visual and textual modalities in the CLIP feature space can be used as an effective likelihood estimator for anomalies; ii) such estimator, not only can detect anomalous occurrences, but also their types; iii) such estimator is effective only when adopting our proposed CLIP space re-centring transformation (see Fig. 2 (a)). Our method is composed of two models as shown in Fig. 2 (b): a *Selector model* and a *Temporal model*. The Selector model $\mathcal{S}$ produces the likelihood that each frame belongs to an anomalous
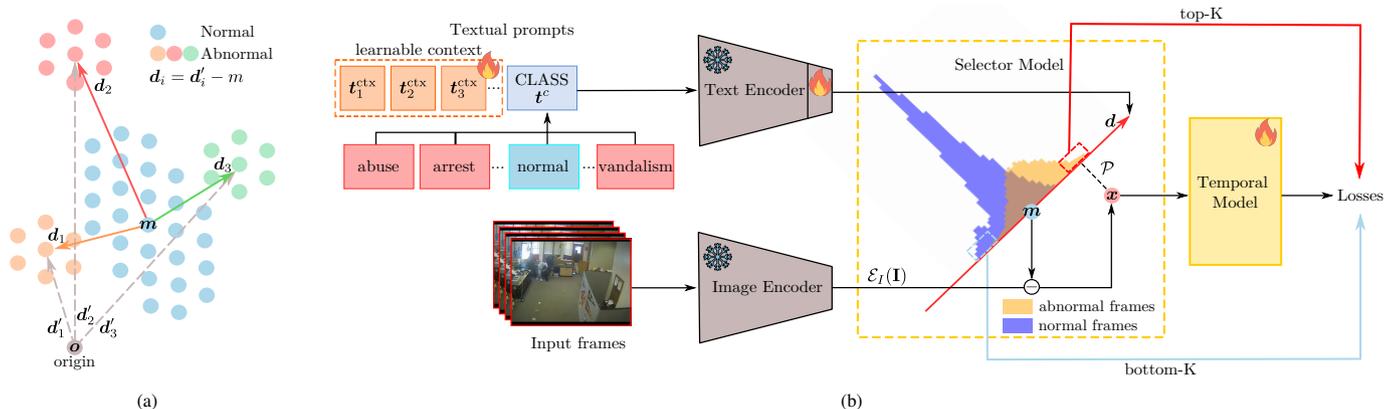
Fig. 2: (a) Illustration of the CLIP space and the effects of the re-centring transformation with features of normal. When the space is not re-centred around the normality prototype m, directions d' are similar, making it difficult to discern anomaly types, and feature magnitude is not linked to the degree of anomaly, making it difficult to identify anomalous events. When re-centred, the distribution of the magnitudes of features projected on each d identifies the degree of detected anomaly of the corresponding type. (b) Illustration of our proposed framework. The Selector model learns directions d using CoOp (Zhou et al., 2022), and uses them to identify the likelihood of each feature x to represent an occurrence of the corresponding anomalous class. MIL selection of the top-K and bottom-K abnormal segments is performed by considering the distribution of likelihoods along the corresponding direction. A Temporal model performs temporal aggregation of the features to produce the final prediction.

class $\mathcal{S}(x) \in \mathbb{R}^C$, where C is the number of anomalous classes. We exploit the vision-text alignment in the CLIP feature space and the CoOp prompt learning approach (Zhou et al., 2022) to estimate this likelihood. The Temporal model $\mathcal{T}$ assigns a binary likelihood to each frame of a video indicating whether the frame is anomalous or normal. Unlike $\mathcal{S}$, $\mathcal{T}$ exploits temporal information to improve predictions and we implement it with a Transformer network (Ho et al., 2019). The predictions from $\mathcal{S}$ and $\mathcal{T}$ are then aggregated to produce a distribution indicating the probability of a frame being normal or abnormal, and which abnormal class it belongs to. We train our model using a combination of MIL and regularisation losses. Importantly, as $\mathcal{T}$ is randomly initialised, the likelihood scores are less reliable, thus we always use the likelihoods produced by $\mathcal{S}$ to perform segment selection in MIL.

We describe the proposed Selector model and Temporal model in detail in Sec. 3.1 and Sec. 3.2, respectively. In Sec. 3.3, we show how we aggregate the predictions of both models for estimating the final probability distribution. Finally, we describe the training and inference in Sec. 3.4.

3.1. Selector model

It is crucial for VAD and VAR to reliably distinguish anomalous and normal frames in anomalous videos given only video-level weak supervision. Motivated by the recent findings in applying LLV models to video action recognition tasks (Wang et al., 2021; Xu et al., 2021), we propose a novel likelihood estimator, encapsulated by our Selector model, that combines the CLIP (Radford et al., 2021) feature space and the CoOp (Zhou et al., 2022) prompt learning approach to learn a set of directions in this space that identify each type of anomaly and their likelihood.

Our main intuition (see Fig. 2 (a)) is that the CLIP feature space presents an underlying structure where the set of CLIP features extracted for each frame in the dataset forms a space that is clustered around a central point which we call the normality prototype. Consequently, the difference between a fea-

ture and the normal prototype determines important characteristics: the magnitude of the distance reflects the likelihood of it being abnormal, while its direction indicates the type of anomaly. Such important characteristics would not be exploited by a naive application of the CLIP feature space to VAR (see Table 9). Unleashing the potential of this space in detecting anomalies thus requires a re-centring transformation, a main contribution of this work.

Following this intuition, we define the normal prototype m as the average feature extracted by the CLIP image encoder $\mathcal{E}_I$ on all N frames $\mathbf{I}$ contained in videos labelled as normal in the dataset:

$$m = \frac{1}{N} \sum_{j=1}^{N} \mathcal{E}_I(\mathbf{I}_j). \tag{1}$$

For each frame $\mathbf{I}$ in the dataset, we produce frame features x by subtracting the normality prototype from the CLIP encoded feature, i.e., $x = \mathcal{E}_I(\mathbf{I}) - m$.

We then exploit the visual-text aligned CLIP feature space and learn the textual prompt embeddings whose directions are used to indicate the anomalous classes. In particular, we employ the prompt learning CoOp method (Zhou et al., 2022) which we find ideal to find such directions as empirically demonstrated by our experiments (see Sec. 4.3).

Given a class c and the textual description of the corresponding label t^c expressed as a sequence of token embeddings, we consider a sequence of learnable context vectors t^{ctx} and derive the corresponding direction for the class $d_c \in \mathbb{R}^D$ as:

$$d_c = \mathcal{E}_T([t^{\text{ctx}}, t^c]) - m, \tag{2}$$

where $\mathcal{E}_T$ indicates the CLIP text encoder. The use of the textual description acts as a prior for the learned direction to match the corresponding type of anomaly, while the context vectors are jointly optimised during training as part of the parameters of $\mathcal{S}$ in order to enable the refinement of the direction. A different direction is learned for each class.

The learned directions serve as the base for our Selector

model $\mathcal{S}$. As shown in Fig. 2(b), the magnitude of the projection of frame feature $\boldsymbol{x}$ on direction $\boldsymbol{d}_c$ indicates the likelihood of the anomalous class c:

$$\mathcal{S}(\boldsymbol{x}) = [\mathcal{P}(\boldsymbol{x}, \boldsymbol{d}_1), ..., \mathcal{P}(\boldsymbol{x}, \boldsymbol{d}_C)] \in \mathbb{R}^C, \qquad (3)$$

where $\mathcal{P}$ indicates our projection operation. However, simply projecting the feature vector on the direction would make the magnitude of the projection susceptible to scale, where anomalous features of one class can potentially have a different magnitude from features of another anomalous class. To mitigate this issue, we perform a batch normalisation (Ioffe and Szegedy, 2015) after the projection which produces a distribution of projected features with zero mean and unitary variance:

$$\mathcal{P}(\boldsymbol{x}, \boldsymbol{d}_i) = \text{BN}\left(\frac{\boldsymbol{x} \cdot \boldsymbol{d}_i}{\|\boldsymbol{d}_i\|}\right), \qquad (4)$$

where BN indicates batch normalisation without affine transformation. As such, we expect within a batch the dominant normal features to be close to the origin and the abnormal features to be at the right side tail of the distribution.

The definition of likelihood can be extended to segments by summing the likelihoods of each frame:

$$\mathcal{S}(\mathbf{S}) = \sum_{i=1}^{F} \mathcal{S}(\boldsymbol{x}_i) \in \mathbb{R}^C \qquad (5)$$

3.2. Temporal Model

The Selector model only learns an initial *time-independent* separation between anomalous and normal frames as the CLIP model operates at the image frame level. However, the temporal information is an important piece of information for VAR that we can exploit. We thus propose the Temporal model $\mathcal{T}$ to model the relationships among frames in both short-term and long-term, to enrich the visual features and to produce the predictions that indicate the likelihood of whether a frame is anomalous:

$$\mathcal{T}(\mathbf{V}) \in \mathbb{R}^{S \times F}. \qquad (6)$$

We use a Transformer architecture to capture the short-term temporal dependencies between frames in a segment and the long-term temporal dependencies between all segments in a video, motivated by their success in relevant sequence modelling tasks (Vaswani et al., 2017). As all the video segments of $\mathbf{V}$ are received as the input, the large number of segments S and frames F increases the computational requirements for self attention. To reduce this cost, we implement $\mathcal{T}$ as an Axial Transformer (Ho et al., 2019) that computes attention separately for the two axes corresponding to the segments and the features in each segment. As suggested by experiments in Sec. 4.3, Axial Transformer is also less prone to over-fitting, a likely case in VAR, as compared to standard Transformer. We terminate the model with a sigmoid activation so that the output likelihood can also be interpreted as a probability.

3.3. Predictions Aggregation

We combine the predictions from $\mathcal{S}$ and $\mathcal{T}$ to obtain the final output: the probabilities indicating whether a frame is normal or anomalous ($p_N(\boldsymbol{x})$ and $p_A(\boldsymbol{x})$) and the probability that a frame presents an anomaly of a certain class ($p_{A,c}(\boldsymbol{x})$).

Given an input frame feature $\boldsymbol{x}$, we define its probability of being anomalous $p_A(\boldsymbol{x})$ as its corresponding output from the Temporal model $\mathcal{T}$. The probability of the frame being normal is $p_N(\boldsymbol{x}) = 1 - p_A(\boldsymbol{x})$. To obtain the probability distribution of the frame to present an anomaly of a specific class $p_{A,c}(\boldsymbol{x})$, we employ the predictions of the Selector model that can be seen as the conditional distribution over the anomalous classes $p_{c|A}(\boldsymbol{x}) = \text{softmax}(\mathcal{S}(\boldsymbol{x}))$. From the definition of conditional probability it follows that $p_{A,c}(\boldsymbol{x}) = p_A(\boldsymbol{x}) * p_{c|A}(\boldsymbol{x})$.

3.4. Training

We train the model following the MIL framework. Specifically, MIL considers a batch with an equal number of normal and anomalous videos, uses the predicted likelihoods to identify the top-K most abnormal segments in anomalous videos, and imposes separation from the other, normal ones (Sultani et al., 2018). Due to the higher capacity of $\mathcal{T}$ with respect to $\mathcal{S}$ and its initial random initialisation, $\mathcal{T}$ can not directly perform this selection since the predicted likelihoods would be excessively noisy. Instead, we use the likelihood predictions from $\mathcal{S}$ to perform MIL segment selection.

Our framework is trained end-to-end using losses on anomalous videos, losses on normal videos, and regularization losses, which we describe in the following.

Given an anomalous video $\mathbf{V}$ of class c, we define the set of top-K most anomalous segments $\mathcal{V}^+ = \{\mathbf{S}_1^+, ..., \mathbf{S}_K^+\}$ and, symmetrically, of bottom-K least anomalous segments $\mathcal{V}^- = \{\mathbf{S}_1^-, ..., \mathbf{S}_K^-\}$ according to the likelihood assigned by the frame-level model $\mathcal{S}$ on the direction corresponding to class c. We consider all frames in $\mathcal{V}^+$ and maximise the likelihood of the corresponding class being predicted by $\mathcal{S}$ by minimising the loss $\mathcal{L}_A^{\text{DIR}}$:

$$\mathcal{L}_A^{\text{DIR}} = -\frac{\sum_{i=1}^{K} \mathcal{S}(\mathbf{S}_i^+)_c}{KF}, \qquad (7)$$

where the likelihood tensor is indexed using the class c. To provide gradients to the temporal model, we also maximise $p_{A,c}(\boldsymbol{x})$ for each frame contained in the segments using cross entropy:

$$\mathcal{L}_{A^+} = -\frac{\sum_{i=1}^{K} \sum_{j=1}^{F} \log(p_{A,c}(\mathbf{S}_{i,j}^+))}{KF}. \qquad (8)$$

Distinguishing normal and anomalous frames in anomalous videos is a challenging problem in VAR due to the appearance similarity between frames of the same video. To foster a better separation between these frames, we additionally consider $\mathcal{V}^-$ and maximise $p_N(\boldsymbol{x})$ for each frame in the segments using cross entropy:

$$\mathcal{L}_{A^-} = -\frac{\sum_{i=1}^{K} \sum_{j=1}^{F} \log(p_N(\mathbf{S}_{i,j}^-))}{KF}, \qquad (9)$$

To leverage the information in normal videos, for each segment $\mathbf{S}_i$ in normal video $\mathbf{V}$, we minimise the likelihood predicted by the Selector model:

$$\mathcal{L}_N^{\text{DIR}} = \frac{\sum_{i=1}^{S} \sum_{c=1}^{C} \mathcal{S}(\mathbf{S}_i)_c}{SFC}. \qquad (10)$$

Following the VAD literature (Feng et al., 2021a; Sultani et al., 2018; Tian et al., 2021) we also require the model to maximise the probability of each frame in its top-K most abnormal segments $\mathcal{V}^+ = \{\mathbf{S}_1^+, ..., \mathbf{S}_K^+\}$ to be normal :

$$\mathcal{L}_{N^+} = -\frac{\sum_{i=1}^{K} \sum_{j=1}^{F} \log(p_N(\mathbf{S}_{i,j}^+))}{KF}. \quad (11)$$

We regularise training with two additional losses (Sultani et al., 2018) on all frames of anomalous videos only. One is a sparsity loss on the predicted scores and encourages the minimal amount of frames to be predicted as abnormal:

$$\mathcal{L}_{\text{spa}} = \frac{\sum_{i=1}^{S} \sum_{j=1}^{F} p_A(\mathbf{V}_{i,j})}{SF} \quad (12)$$

The other is a smoothness term that regularises the predictions along the temporal dimension:

$$\mathcal{L}_{\text{smo}} = \sum_{i=2}^{SF} (p_A(\mathbf{V}_i) - p_A(\mathbf{V}_{i-1})), \quad (13)$$

where indexing is performed on the flattened sequence of frames in the video.

We jointly train the Selector and Temporal models using as final training objective:

$$\mathcal{L} = \mathcal{L}_A^{\text{DIR}} + \mathcal{L}_{A^+} + \mathcal{L}_{A^-} + \mathcal{L}_N^{\text{DIR}} + \mathcal{L}_{N^+} + \lambda_1 \mathcal{L}_{\text{spa}} + \lambda_2 \mathcal{L}_{\text{smo}}. \quad (14)$$

4. Experiments

In this section, we validate our method against a range of baselines taken from state-of-the-art VAD and action recognition methods which we adapt to the VAR task. After introducing the metrics for the novel VAR task, we perform evaluation on three datasets and perform comparison in both the VAD and VAR tasks. An extensive ablation study is performed to justify our main design choices. Sec 4.1 describes our experiment setup in terms of datasets and evaluation protocols. We then present and discuss the results in comparison against state-of-the-art methods in Sec 4.2 and the ablation study in Sec 4.3.

4.1. Experiment Setup

Datasets. We perform our study using three widely-used VAD datasets, *i.e.*, ShanghaiTech (Liu et al., 2018), UCF-Crime (Sultani et al., 2018), and XD-Violence (Wu et al., 2020). *ShanghaiTech* consists of 437 videos, recorded from multiple surveillance cameras in a university campus. A total of 130 abnormal events of 17 anomaly classes are captured in 13 different scenes. We adopt the dataset in the configuration of Zhong et al. (2019) which adapts it to the weakly-supervised setting by organising it into 238 training videos and 199 testing videos. *UCF-Crime* is a large-scale dataset of real-world surveillance videos, containing 1900 long untrimmed videos that cover 13 real-world anomalies with significant impacts on public safety. The training set consists of 800 normal and 810 anomalous videos and the testing set includes the remaining 150 normal and 140 anomalous videos. *XD-Violence* is a large-scale violence detection dataset comprising 4754 untrimmed videos with audio signals and weak labels, divided into a training set of 3954 videos and a test set of 800 videos. With a total duration of 217 hours, the dataset covers various scenarios and captures 6 categories of anomalies. Notably, each violent video may have multiple labels, ranging from 1 to 3. To accommodate our training setup, where only one anomaly type per video is considered, we select the subset of 4463 videos containing at most one anomaly.

Performance Metrics. We perform evaluation in terms of both VAD and VAR. Following previous works, we measure the performance regarding VAD using the area under the curve (AUC) of the frame-level receiver operating characteristics (ROC) as it is agnostic to thresholding for the detection task. A larger frame-level AUC means a better performance in classifying between normal and anomalous events. To measure the VAR performance, we extend the AUC metric to the multi-classification scenario. For each anomalous class, we measure the AUC by considering the anomalous frames of the class as positive and all other frames as negatives. Successively, the mean AUC (mAUC) is computed over all the anomalous classes. Similarly, for the XD-Violence dataset, we follow the established evaluation protocol (Wu et al., 2020) and present VAD results using the average precision (AP) of the precision-recall curve (PRC), while for VAR results we report the mean AP (mAP), which is calculated by averaging the binary AP values across all anomalous classes.

Implementation details. At training time, each video is divided into S non-overlapping blocks. From each block, a random start-index is sampled from which segments of F consecutive frames are considered. If the raw video has length smaller than $S \times F$, we adopt loop padding and repeat the video from the start until the minimum length of $S \times F$ is reached. Each mini-batch of size B used for training is composed of $B/2$ normal clips and $B/2$ anomalous clips. This is a simple but effective way to balance the mini-batch formation, which otherwise will contain mainly normal clips. At inference, to handle videos covering arbitrary temporal windows, we first divide each video $\mathbf{V}$ into S non-overlapping blocks, where each block contains frames whose number is a multiple of F, i.e., $J \times F$, where J depends on the length of $\mathbf{V}^\circ$. We process $\mathbf{V}$ with J inferences to classify all frames in the video. At each j^{th} inference, we extract the j^{th} consecutive F frames from each block, forming segments with a total of $S \times F$ that span the whole video. We then feed the segments into our approach so that our Temporal model can reason the long-term temporal relationships among segments.

For a fair comparison with previous works in VAD (Tian et al., 2021; Wu et al., 2022; Li et al., 2022a), we use $K = 3$ for the MIL selection of the top-K and bottom-K abnormal segments, $S = 32$ number of segments, $F = 16$ frames per segment and $B = 64$ batch size. Please refer to Appendix A for more implementation details and Appendix B for more details on hyper-parameters.

$^\circ$We perform loop padding to ensure that each video is of length $J \times S \times F$

Table 1: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAD and VAR on ShanghaiTech.

Supervision	Method	Features	VAD	VAR	AUC(%)	mAUC(%)
One-class	MNAD (Park et al., 2020)		✓		70.50	
	MPN (Lv et al., 2021)		✓		73.80	
	HF²VAD (Liu et al., 2021)		✓		76.20	
	Zaheer et al. (2022)	ResNext	✓		79.62	
Unsupervised	Zaheer et al. (2022)	ResNext	✓		78.93	
Zero-shot	CLIP (Radford et al., 2021)	ViT-B/16		✓	49.17	51.02
Weakly-supervised	Sultani et al. (2018)	C3D-RGB	✓		86.30	
	IBL (Zhang et al., 2019)	C3D-RGB	✓		82.50	
	Zaheer et al. (2022)	ResNext	✓		86.21	
	GCN (Zhong et al., 2019)	TSN-RGB	✓		84.44	
	MIST (Feng et al., 2021a)	I3D-RGB	✓		94.83	
	Wu et al. (2020)	I3D-RGB	✓			
	CLAWS (Zaheer et al., 2020)	C3D-RGB	✓		89.67	
	RTFM (Tian et al., 2021)	I3D-RGB	✓		97.21	81.60
	Wu and Liu (2021)	I3D-RGB	✓		97.48	
	MSL (Li et al., 2022b)	I3D-RGB	✓		96.08	
	MSL (Li et al., 2022b)	VideoSwin-RGB	✓		97.32	
	S3R (Wu et al., 2022)	I3D-RGB	✓		97.48	87.88
	MGFN (Chen et al., 2022)	I3D-RGB	✓			
	MGFN (Chen et al., 2022)	VideoSwin-RGB	✓			
	SSRL (Li et al., 2022a)	I3D-RGB	✓		97.98	93.61
	ActionCLIP (Wang et al., 2021)	ViT-B/16		✓	96.36	75.63
	AnomalyCLIP (ours)	ViT-B/16	✓	✓	**98.07**	**96.46**

Table 2: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAD and VAR on UCF-Crime.

Supervision	Method	Features	VAD	VAR	AUC(%)	mAUC(%)
One-class	SVM Baseline (Sultani et al., 2018)		✓		50.00	
	SSV (Sohrab et al., 2018)		✓		58.50	
	BODS (Wang and Cherian, 2019)	I3D-RGB	✓		68.26	
	GODS (Wang and Cherian, 2019)	I3D-RGB	✓		70.46	
	Zaheer et al. (2022)	ResNext	✓		74.20	
Un-supervised	Zaheer et al. (2022)	ResNext	✓		71.04	
Zero-shot	CLIP (Radford et al., 2021)	ViT-B/16		✓	58.63	74.28
Weakly-supervised	Sultani et al. (2018)	C3D-RGB	✓		75.41	
	Sultani et al. (2018)	I3D-RGB	✓		77.92	
	IBL (Zhang et al., 2019)	C3D-RGB	✓		78.66	
	Zaheer et al. (2022)	ResNext	✓		79.84	
	GCN (Zhong et al., 2019)	TSN-RGB	✓		82.12	
	MIST (Feng et al., 2021a)	I3D-RGB	✓		82.30	
	Wu et al. (2020)	I3D-RGB	✓		82.44	
	CLAWS (Zaheer et al., 2020)	C3D-RGB	✓		83.03	
	RTFM (Tian et al., 2021)	VideoSwin-RGB	✓		83.31	
	RTFM (Tian et al., 2021)	I3D-RGB	✓		84.03	84.86
	Wu and Liu (2021)	I3D-RGB	✓		84.89	
	MSL (Li et al., 2022b)	I3D-RGB	✓		85.30	
	MSL (Li et al., 2022b)	VideoSwin-RGB	✓		85.62	
	S3R (Wu et al., 2022)	I3D-RGB	✓		85.99	86.55
	MGFN (Chen et al., 2022)	VideoSwin-RGB	✓		86.67	
	MGFN (Chen et al., 2022)	I3D-RGB	✓		86.98	
	SSRL (Li et al., 2022a)	I3D-RGB	✓		87.43	88.88
	ActionCLIP (Wang et al., 2021)	ViT-B/16		✓	82.30	87.72
	AnomalyCLIP (ours)	ViT-B/16	✓	✓	86.36	**90.66**

Table 3: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAD and VAR on XD-Violence.

Supervision	Method	Features	VAD	VAR	AP(%)	mAP(%)
Zero-shot	CLIP (Radford et al., 2021)	ViT-B/16		✓	27.21	21.32
Weakly-supervised	Wu et al. (2020)	C3D-RGB	✓		67.19	
	Wu et al. (2020)	I3D-RGB	✓		73.20	
	MSL (Li et al., 2022b)	C3D-RGB	✓		75.53	
	Wu and Liu (2021)	I3D-RGB	✓		75.90	
	RTFM (Tian et al., 2021)	I3D-RGB	✓		77.81	43.04
	MSL (Li et al., 2022b)	I3D-RGB	✓		78.28	
	MSL (Li et al., 2022b)	VideoSwin-RGB	✓		78.58	
	S3R (Wu et al., 2022)	I3D-RGB	✓		**80.26**	36.06
	MGFN (Chen et al., 2022)	I3D-RGB	✓		79.19	
	MGFN (Chen et al., 2022)	VideoSwin-RGB	✓		80.11	
	ActionCLIP (Wang et al., 2021)	ViT-B/16		✓	61.01	40.24
	AnomalyCLIP (ours)	ViT-B/16	✓	✓	78.51	**49.41**

4.2. Evaluation Against Baselines

Regarding VAD, we compare AnomalyCLIP against state-of-the-art methods with different supervision setups, including one-class (Park et al., 2020; Liu et al., 2021; Lv et al., 2021), unsupervised (Zaheer et al., 2022) and weakly-supervised (Li

et al., 2022a; Tian et al., 2021; Wu et al., 2022). As none of the above-mentioned methods address the VAR task, we produce baselines by re-purposing some best-performing VAD methods including RTFM (Tian et al., 2021), S3R (Wu et al., 2022) and SSRL (Li et al., 2022a)°, and CLIP-based baselines (Radford et al., 2021; Wang et al., 2021):

- *Multi-classification with RTFM (Tian et al., 2021), S3R (Wu et al., 2022) and SSRL (Li et al., 2022a) (weakly-supervised).*
  We keep the original pretrained model frozen and add a multi-class classification head that we train to predict the class using a cross entropy objective on the top-K most anomalous segments selected as in the original method. These baselines are weakly-supervised.
- *CLIP (Radford et al., 2021) (zero-shot).* We achieve the classification by soft-maxing of the cosine similarities of the input frame feature x with vectors corresponding to the embedding of the textual prompt *"a video from a CCTV camera of a {class}"* using the pre-trained CLIP model.
- *ActionCLIP (Wang et al., 2021) (weakly-supervised).* We retrain ActionCLIP (Wang et al., 2021) on our datasets by propagating the video-level anomaly labels to each frame of the corresponding video.

Table 1 presents the results on ShanghaiTech (Liu et al., 2018). Although ShanghaiTech is a rather saturated dataset for VAD due to its simplicity in scenarios, AnomalyCLIP scores the state-of-the-art results on both VAD and VAR, with +0.09% and +2.85% in terms of AUC ROC and mAUC ROC, respectively. ActionCLIP (Wang et al., 2021) performs poorly in terms of mAUC, which we attribute to the low proportion of abnormal events in ShanghaiTech that makes the MIL selection strategy of particular importance to avoid incorrect supervisory signals on normal frames of abnormal videos. In contrast, our proposal has a better recognition of the positive instances of abnormal videos, thus achieving better performance even when anomalies are rare. AnomalyCLIP achieves a large improvement of +45.44% in terms of mAUC against zero-shot CLIP, demonstrating that a naive application of a VAR pipeline in the CLIP space does not yield satisfactory results. A revision of this space, implemented as our proposed transformation, is necessary to use it effectively.

Table 2 reports the results on UCF-Crime (Sultani et al., 2018). Our method exhibits the best discrimination of the anomalous classes, achieving the highest mAUC ROC among baselines. Similar to ShanghaiTech, it also achieves an improvement in terms of mAUC against zero-shot CLIP, verifying the importance of our proposed adaptation of the CLIP space. Compared to ActionCLIP (Wang et al., 2021), our Anomaly-CLIP obtains +2.94% in terms of mAUC, highlighting the need for a MIL framework to mitigate mis-assignment of anomalous class labels to normal frames of anomalous videos. It is also worth noting that the higher mAUC obtained by ActionCLIP does not result in a competitive AUC ROC on VAD, which implicates a worse separation between normal and abnormal frames. When compared to the best performing method

---

° We thank authors for making their code and models publicly available

Table 4: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAR on UCF-Crime. The table highlights the top performers, with cells highlighted in red representing first place, cells in orange representing second place, and cells in yellow representing third place.

Method	Class													mAUC
	Abuse	Arrest	Arson	Assault	Burglary	Explosion	Fighting	RoadAcc.	Robbery	Shooting	Shoplifting	Stealing	Vandalism	
RTFM Tian et al. (2021)	79.99	62.57	90.53	82.27	85.53	92.76	85.21	90.31	81.17	82.82	92.56	90.23	87.20	84.86
S3R Wu et al. (2022)	86.38	68.45	92.19	93.55	86.91	93.55	81.69	85.03	82.07	85.32	91.64	94.59	83.82	86.55
SSRL Li et al. (2022a)	95.33	79.26	93.27	91.74	89.06	92.25	87.36	80.24	87.75	84.50	92.31	94.22	88.17	88.88
CLIP zero-shot Radford et al. (2021)	57.37	80.65	93.72	80.83	74.34	90.31	83.54	87.46	70.22	63.99	71.21	45.49	66.45	74.28
ActionCLIP Wang et al. (2021)	91.88	90.47	89.21	86.87	81.31	94.08	83.23	94.34	82.82	70.53	91.60	94.06	89.89	87.72
AnomalyCLIP	75.03	94.56	96.66	94.80	90.08	94.79	88.76	93.30	86.85	87.45	89.47	97.00	89.78	90.66

Table 5: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAR on ShanghaiTech. The table highlights the top performers, with cells highlighted in red representing first place, cells in orange representing second place, and cells in yellow representing third place.

Method	Class															mAUC
	Car	Chasing	Circuit	Fall	Fighting	Jumping	Monocycle	Push	Robbery	Running	Skateboard	Stoop	ThrowingObj.	Vaudeville	Vehicle	
RTFM Tian et al. (2021)	99.70	95.41	99.83	70.19	97.36	89.14	37.99	35.28	67.01	90.59	96.81	64.11	97.93	91.75	90.85	81.60
S3R Wu et al. (2022)	98.71	96.80	99.97	85.63	95.93	69.33	96.82	54.76	61.19	94.43	96.92	75.46	97.63	97.78	96.84	87.88
SSRL Li et al. (2022a)	99.35	97.31	99.95	91.24	96.88	93.07	89.74	90.62	91.81	94.47	97.73	71.81	98.44	96.32	95.49	93.61
CLIP zero-shot Radford et al. (2021)	61.65	77.88	5.95	61.73	79.37	23.68	77.78	63.36	37.71	54.39	76.15	8.47	44.10	65.97	27.08	51.02
ActionCLIP Wang et al. (2021)	98.50	93.86	98.59	16.38	97.45	89.63	98.05	8.14	67.36	78.25	97.10	0.76	97.70	98.65	93.97	75.63
AnomalyCLIP	98.08	96.66	97.97	96.69	98.03	95.48	86.89	97.99	95.00	97.95	97.29	98.62	96.50	96.97	96.79	96.46

Table 6: Results of the state-of-the-art methods and our AnomalyCLIP in terms of VAR on XD-Violence. The table highlights the top performers, with cells highlighted in red representing first place, cells in orange representing second place, and cells in yellow representing third place.

Method	Class						mAP
	Abuse	CarAccident	Explosion	Fighting	Riot	Shooting	
RTFM Tian et al. (2021)	9.25	25.36	53.53	61.73	90.38	18.01	43.04
S3R Wu et al. (2022)	2.63	23.82	45.29	49.88	90.41	4.34	36.06
CLIP zero-shot Radford et al. (2021)	0.32	12.21	22.26	25.25	66.60	1.26	21.32
ActionCLIP Wang et al. (2021)	2.73	25.15	55.28	58.09	87.31	12.87	40.24
AnomalyCLIP	6.10	31.31	68.75	71.44	92.74	26.13	49.41

SSRL (Li et al., 2022a) on VAD, our method obtains an improvement of +1.78% in terms of mAUC on VAR, while being slightly worse with −1.07% in terms of AUC ROC on VAD.

Table 3 shows the results on XD-Violence (Wu et al., 2020). AnomalyCLIP outperforms other state-of-the-art methods on VAR achieving the highest mAP. Compared to the VAD baselines' models, AnomalyCLIP outperforms RTFM (Tian et al., 2021) and demonstrates performance close to S3R (Wu et al., 2022). Please refer to Appendix C for further details on how we obtain results on XD-Violence.

Tables 4, 5, and 6 display the multi-class AUC and AP for each individual abnormal class. The proposed method has a clear advantage when applied to the UCF-Crime and XD-Violence datasets, which are generally considered to be complex benchmarks in anomaly detection. Our method achieves the best mAUC and mAP on average, while it is less advantageous when dealing with anomalies that exhibit slight deviations from normal patterns, such as Shoplifting in UCF-Crime. The advantage of our proposed method is less noticeable when applied to the ShanghaiTech dataset, which captures simple scenes where most methods have achieved a saturated performance.

Fig. 3 presents the qualitative results of our proposed AnomalyCLIP in detecting and recognising anomalies within a set of UCF-Crime, ShanghaiTech, XD-Violence test videos. The model is capable of predicting both the presence of anomalies in test videos and the category of the anomalous event. In video *Normal_Video_246* from UCF-Crime (Row 2, Column 2), it can be seen how some frames have a higher-than-expected proba-bility of being abnormal. It is interesting to note how in the video *RoadAccidents133* from UCF-Crime (Row 1, Column 2) the anomaly score remains high even in the aftermath of the accident. It is also interesting to note that for Normal videos, AnomalyCLIP is able to obtain a relatively low anomaly probability all over the frames, meaning our model has learnt a robust normal representation among Normal videos. Please refer to Appendix E for more results on the test videos. Furthermore, for a more intuitive understanding of the results presented in the paper, we invite readers to access the website https://luca-zanella-dvl.github.io/AnomalyCLIP, where easily accessible qualitative results are available.

4.3. Ablation

In this section, we perform ablations of our method to validate our main design choices with UCF-Crime: the way in which we represent and learn directions, the transformations applied to the CLIP space and the employed way for estimating the likelihood of anomaly, the choice of architecture for the Temporal model, training objectives, and the impact of using features extracted from different backbones.

Representation and Learning of the Directions. In the ablation shown in Table 7, we evaluate the choice of the CoOp (Zhou et al., 2022) framework to learn directions in the CLIP space. When CoOp is removed, we directly learn the directions from randomly initialized points in the CLIP space (Row 1) or make use of fixed engineered prompts of the form *"a video from a CCTV camera of a {class}"* (Row 2). Both choices result in degradation of the results, indicating that text-
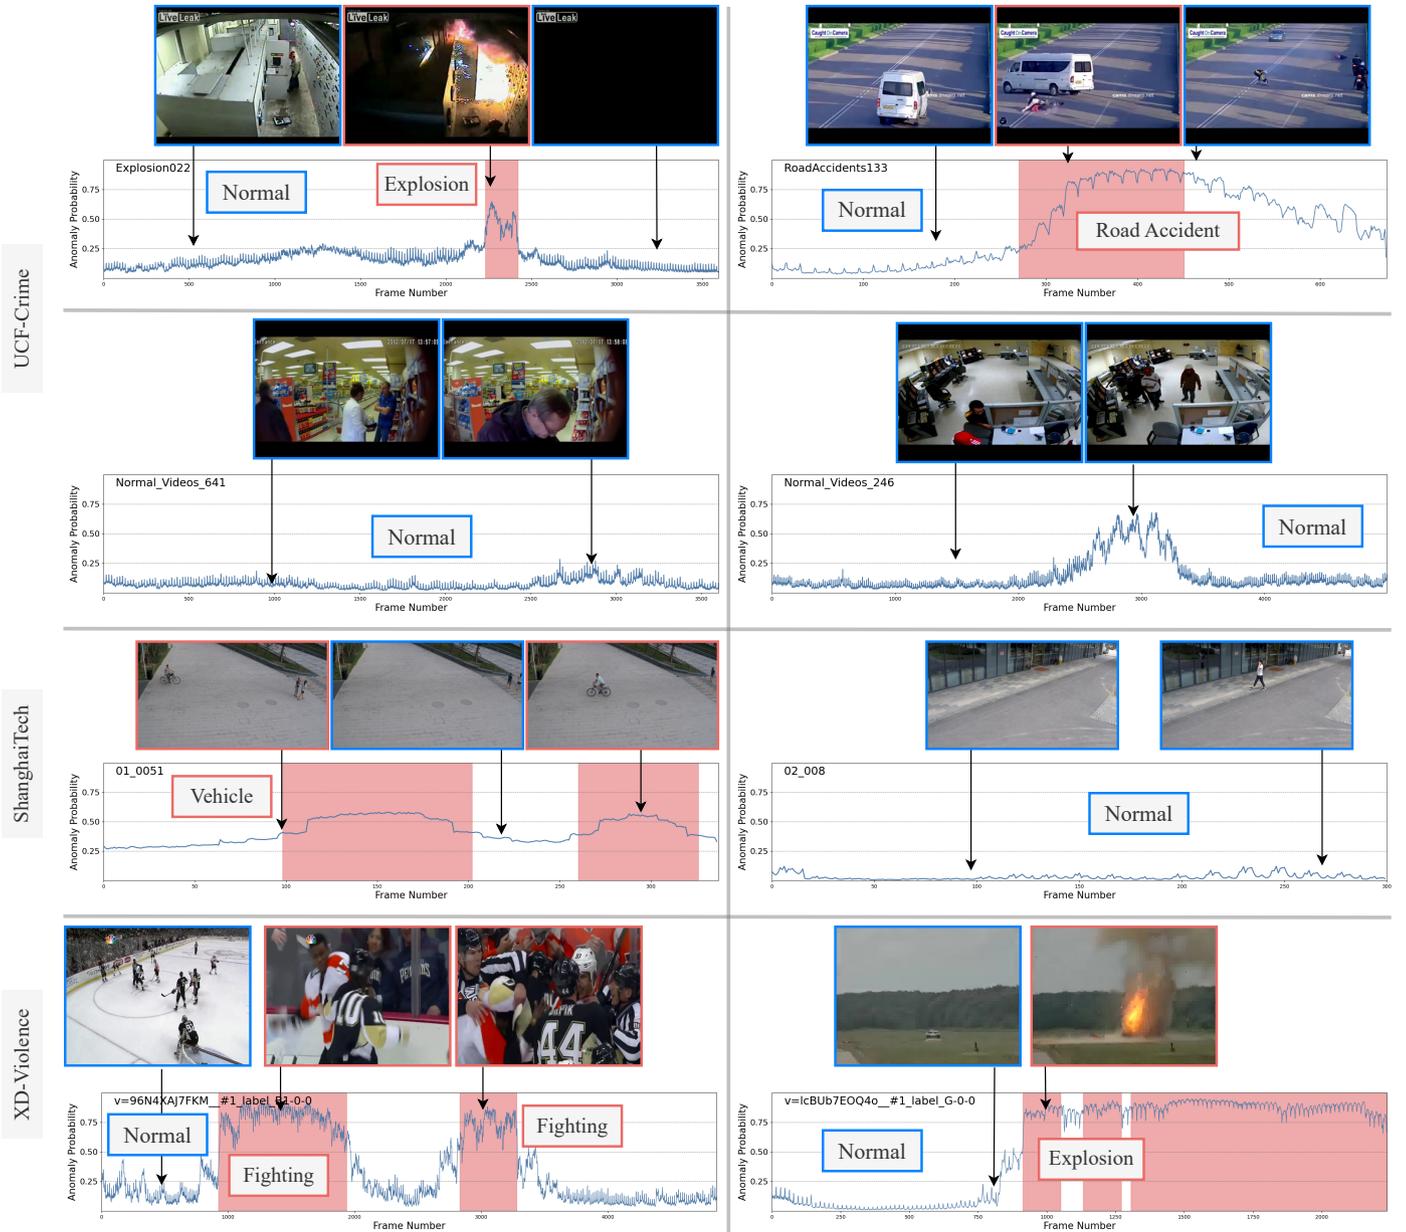
Fig. 3: Qualitative results for VAR on four test videos from UCF-Crime (the top two rows), two test videos from ShanghaiTech (the third row), and two test videos from XD-Violence (the bottom row). For each video, we show at the bottom the predicted probability of each frame being anomalous by our model over the number of frames. We showcase some key frames to reflect the relevance between the predicted anomaly probability and the visual content. The red shaded areas denote the temporal ground-truth of anomalies. We also indicate the predicted anomalous class for detected abnormal frames in the red boxes, while videos without detected anomalies are indicated with blue boxes as Normal.

Table 7: Ablation on representation and learning of the directions of abnormality. 'Finetuning' indicates that the last projection layer is fine-tuned. The final configuration of our model is represented by the row highlighted in grey in the table.

Text encoder	Directions	AUC	mAUC
No	Direct Optimisation	84.98	69.86
Frozen	Engineered Prompts	84.66	81.35
Frozen	CoOp	85.88	87.39
Finetuning	CoOp	**86.36**	**90.66**

Table 8: Comparisons of different architectural choices for the CoOp module. 'Shared' means that all the classes share a unified context, otherwise each class has a specific context. The final configuration of our model is represented by the row highlighted in grey in the table.

Context vectors	Shared	AUC	mAUC
4		86.16	**91.05**
8		**86.36**	90.66
16		85.82	90.65
8	✓	85.97	90.01

Table 9: Ablation of different likelihood estimation methods, feature space transformations and MIL selection. 'Features' indicates the transformation applied to CLIP features. The final configuration of our model is represented by the row highlighted in grey in the table.

Likelihood	Features	MIL Selection	AUC	mAUC
cosine sim.	CLIP	cosine sim.	85.59	83.69
$\mathcal{S}$	CLIP - m	feature magnitude	84.92	89.82
$\mathcal{S}$	CLIP - m	$\mathcal{S}$	**86.36**	**90.66**

Table 10: Comparisons of different architectural choices for the Temporal model. The final configuration of our model is represented by the row highlighted in grey in the table.

Temporal Model	Short-term	Long-term	AUC	mAUC
MLP			74.86	84.46
Transformer	✓		84.69	88.38
Transformer		✓	85.10	89.29
MTN		✓	82.71	87.65
Axial Transformer	✓	✓	**86.36**	**90.66**

guided initialization of the directions and directions finetuning are both necessary. Furthermore, we show that unfreezing the last projection of the text encoder (Row 4) enables a greater freedom in finetuning the discovered directions, yielding the best results.

In the ablation shown in Table 8, we evaluate the architectural choices on the CoOp module to learn directions in the CLIP space. Specifically, we experimented by varying the number of context vectors t^{ctx} used from 4 to 8 to 16, and using shared or class-specific context vectors. Although using 4 context vectors results in a slightly higher mAUC score, we eventually opted to use 8 context vectors because they produce a higher AUC score. Results (Row 2 and 4) show that learning a specific set of context vectors for each class, is more tailored to fine-grained categories, rather than relying on more generic shared context vectors for all classes.

Likelihood Estimation and CLIP Latent Space Transformation. The way in which the extracted CLIP features are transformed and the chosen likelihood estimation method play a crucial role in the quality of segment selection. We evaluate several choices in this procedure in Table 9. Directly using the CLIP space and cosine similarities with the learned directions as likelihood estimators (Row 1) produces the worst VAR results, indicating that the use of the normality prototype m is of high importance in the context of anomaly detection. Second, Row 2 shows that MIL segment selection as a function of the feature magnitude without accounting for the direction is not as effective, given that the large magnitude could be attributed to irrelevant factors.

Temporal Model Architecture. Capturing temporal information is an essential aspect of VAR since it provides insights into the behaviour of objects and scenes over time. Table 10 shows results for different architectures of $\mathcal{T}$ *i.e.* a 3-layer MLP, two Transformer Encoders (Vaswani et al., 2017), the multi-scale temporal network (MTN), designed in RTFM and used in S3R and SSRL, and the employed Axial Transformer. In particu-

lar, one transformer encoder (Row 2) performs self-attention on each independent 16-frame segment, solely modelling short-term dependencies. The other (Row 3) applies self-attention on segment embeddings, which are obtained by averaging 16-frame feature embeddings within each segment, thereby only modelling long-term dependencies. To ensure a fair comparison, both transformers are designed to have a capacity similar to that of the Axial Transformer. The reduced performance of the MLP baseline (Row 1) indicates the necessity of considering temporal information which is not readily available in the extracted CLIP features. The Axial transformer can capture temporal dependencies and outperform the compared architectures.

Table 11 shows the results for different values of the embedding size and the number of layers. In the final architecture we use 1 layer and an embedding size of 256, for a total of 10.4 M trainable parameters.

Losses. Table 12 illustrates the contribution of the losses on the Selector model's outputs, where we progressively remove the losses from the full training objective. The loss on abnormal videos contributes to improved VAD and VAR results on UCF-Crime. Similarly, using the loss on normal videos improves the results on Shanghaitech and XD-Violence, as can be seen in Tables D.16 and D.17 of Appendix D.

Table 13 similarly shows the contribution of the losses on the aggregated model's output, where we remove each from the complete training objective. We validate that each of the proposed losses promotes performance on both the VAD and VAR tasks.

The bottom-K least anomalous segments $\mathcal{V}^- = \{\mathbf{S}_1^-, ..., \mathbf{S}_K^-\}$ of anomalous videos proved to be beneficial for learning the Temporal Model. Inspired by this, we analyse the impact of incorporating this set of frames into the Selector Model loss by minimising the loss:

$$\mathcal{L}_{A^-}^{\text{DIR}} = \frac{\sum_{i=1}^K \mathcal{S}(\mathbf{S}_i^-)_c}{KF}, \tag{15}$$

Table 11: Comparisons of different architectural choices for the Axial Transformer. The final configuration of our model is represented by the row highlighted in grey in the table.

Embedding size	Number of layers	AUC	mAUC
64	1	82.83	90.10
128	1	84.97	90.53
256	1	**86.36**	**90.66**
512	1	85.51	89.28
256	2	85.89	89.67
256	3	85.15	88.14

Table 12: Ablation of the losses on the Selector model. The final configuration of our model is represented by the row highlighted in grey in the table.

$\mathcal{L}_A^{\mathrm{DIR}}$	$\mathcal{L}_N^{\mathrm{DIR}}$	AUC	mAUC
		85.89	89.34
	✓	85.91	87.26
✓		**86.46**	**90.75**
✓	✓	86.36	90.66

Moreover, instead of using all segments of normal videos in the Selector Model loss, we evaluate the impact of using only the top-K most abnormal segments $\mathcal{V}^+ = \{\mathbf{S}_1^+, ..., \mathbf{S}_K^+\}$ by minimising the likelihood predicted by the Selector Model:

$$\mathcal{L}_{N^+}^{\mathrm{DIR}} = \frac{\sum_{i=1}^K \mathcal{S}(\mathbf{S}_i^+)_c}{KF} \tag{16}$$

In Table 14, we present our findings, which indicate that modifying the loss function in either of two ways cause a degradation of performance. Specifically, our experiments (Row 1) demonstrate that using the bottom-K least abnormal segments is only effective when learning the Temporal Model. This is because if there is no clear separation between the bottom-K and top-K abnormal features, the Selector Model can lead to incorrectly selected bottom-K features that prevent it from learning good directions in the feature space. However, incorporating the bottom-K least abnormal segments becomes beneficial in the Temporal Model, which has a greater capacity. Furthermore, our experiments indicate that using all normal segments (Row 3) provides a more robust estimation of the direction from normal to anomalous compared to using only the top-K most abnormal segments (Row 2).

Feature Representation. The purpose of this ablation study is to determine the most suitable feature space for the proposed method *AnomalyCLIP*. To achieve this, we first investigate whether the space learned by the Selector Model can be applied to the Temporal Model. This C-dimensional space is formed by projecting each frame feature $\boldsymbol{x}$ onto every $\boldsymbol{d}_c$ direction, where C represents the number of anomalous classes. Our results, presented in Table 15, indicate that using only this space leads to sub-optimal model performance (Row 3). This finding highlights the necessity of incorporating the information contained in the original feature space as well. We also experiment with using I3D features for both the Selector Model and the Temporal Model (Row 1), but the results demonstrate that the model using these features performs worse. We attribute this to the

Table 13: Ablation of losses on the aggregated outputs. The final configuration of our model is represented by the row highlighted in grey in the table.

$\mathcal{L}_{A^+}$	$\mathcal{L}_{A^-}$	$\mathcal{L}_{N^+}$	AUC	mAUC
	✓	✓	45.23	69.57
✓		✓	84.50	**90.88**
✓	✓		80.96	86.10
✓	✓	✓	**86.36**	90.66

Table 14: Ablation on the variation of Selector model losses. The final configuration of our model is represented by the row highlighted in grey in the table.

$\mathcal{L}_A^{\mathrm{DIR}}$	$\mathcal{L}_N^{\mathrm{DIR}}$	$\mathcal{L}_{N^+}^{\mathrm{DIR}}$	$\mathcal{L}_{A^-}^{\mathrm{DIR}}$	AUC	mAUC
✓	✓		✓	**86.41**	88.29
✓		✓		86.17	90.53
✓	✓			86.36	**90.66**

fact that I3D features are mapped to a region of space that is not aligned with the text features, unlike the features generated by CLIP's image encoder. For this reason, we also experimented using I3D features for the Temporal Model and features from CLIP's image encoder for the Selector Model (Row 2). The result of this experiment further emphasises that the latent space of CLIP is a more semantic space in which anomalous events of different classes are more separated, which in turn leads to superior discriminative ability in detecting and recognising anomalous events.

5. Conclusions

In this work, we addressed the challenging task of Video Anomaly Recognition that extends the scope of Video Anomaly Detection by further requiring the classification of the anomalous activities. We proposed *AnomalyCLIP*, the first method that leverages LLV models in the context of VAR. Our work shed light on the fact that a naive application of existing LLV models (Radford et al., 2021; Wang et al., 2021) to VAR leads to unsatisfactory performance and we demonstrated that several technical design choices are required to build a multimodal deep network for detecting and classifying abnormal behaviours. We also performed an extensive experimental evaluation showing that *AnomalyCLIP* achieves state-or-the-art VAR results on the benchmark ShanghaiTech (Liu et al., 2018), UCF-Crime (Sultani et al., 2018), and XD-Violence (Wu et al., 2020) datasets. As future work, we plan to extend our method in open-set scenarios to reflect the real-world applications where anomalies are often not pre-defined. We will also investigate the applicability of our method in other multi-modal tasks, e.g., fine-grained classification.

References

Bai, S., He, Z., Lei, Y., Wu, W., Zhu, C., Sun, M., Yan, J., 2019. Traffic anomaly detection via perspective map based on spatial-temporal information matrix, in: CVPR Workshops.

Bao, Q., Liu, F., Liu, Y., Jiao, L., Liu, X., Li, L., 2022. Hierarchical scene normality-binding modeling for anomaly detection in surveillance videos, in: ACM Multimedia.

Carreira, J., Zisserman, A., 2017. Quo vadis, action recognition? a new model and the kinetics dataset, in: CVPR.

Table 15: Comparisons of different features. The final configuration of our model is represented by the row highlighted in grey in the table.

Selector Model	Temporal Model	AUC	mAUC
I3D-RGB	I3D-RGB	65.05	84.24
ViT-B/16	I3D-RGB	78.11	88.26
ViT-B/16	$\mathcal{S}(x)$	84.44	86.78
ViT-B/16	ViT-B/16	**86.36**	**90.66**

Chen, Y., Liu, Z., Zhang, B., Fok, W., Qi, X., Wu, Y.C., 2022. Mgfn: Magnitude-contrastive glance-and-focus network for weakly-supervised video anomaly detection. arXiv .

Feng, J.C., Hong, F.T., Zheng, W.S., 2021a. Mist: Multiple instance self-training framework for video anomaly detection, in: CVPR.

Feng, X., Song, D., Chen, Y., Chen, Z., Ni, J., Chen, H., 2021b. Convolutional transformer based dual discriminator generative adversarial networks for video anomaly detection, in: ACM Multimedia.

Ho, J., Kalchbrenner, N., Weissenborn, D., Salimans, T., 2019. Axial attention in multidimensional transformers. arXiv .

Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: ICML.

Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T., 2021. Scaling up visual and vision-language representation learning with noisy text supervision, in: ICML.

Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al., 2017. The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 .

Li, G., Cai, G., Zeng, X., Zhao, R., 2022a. Scale-aware spatio-temporal relation learning for video anomaly detection, in: ECCV, Springer.

Li, S., Liu, F., Jiao, L., 2022b. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection, in: AAAI.

Liu, W., Luo, W., Lian, D., Gao, S., 2018. Future frame prediction for anomaly detection–a new baseline, in: CVPR.

Liu, Z., Nie, Y., Long, C., Zhang, Q., Li, G., 2021. A hybrid video anomaly detection framework via memory-augmented flow reconstruction and flow-guided frame prediction, in: ICCV.

Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H., 2022. Video swin transformer, in: CVPR.

Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization, in: ICLR.

Lv, H., Chen, C., Cui, Z., Xu, C., Li, Y., Yang, J., 2021. Learning normal dynamics in videos with meta prototype network, in: CVPR.

Mei, T., Zhang, C., 2017. Deep learning for intelligent video analysis, in: ACM Multimedia.

Narasimhan, M.G., 2018. Dynamic video anomaly detection and localization using sparse denoising autoencoders. Multimedia Tools and Applications .

Nayak, R., Pati, U.C., Das, S.K., 2021. A comprehensive review on deep learning-based methods for video anomaly detection. Image and Vision Computing 106.

Park, H., Noh, J., Ham, B., 2020. Learning memory-guided normality for anomaly detection, in: CVPR.

Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al., 2021. Learning transferable visual models from natural language supervision, in: ICML.

Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P., 2022. Towards total recall in industrial anomaly detection, in: CVPR.

Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al., 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. arXiv .

Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A., 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv .

Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., Kiela, D., 2022. Flava: A foundational language and vision alignment model, in: CVPR.

Sohrab, F., Raitoharju, J., Gabbouj, M., Iosifidis, A., 2018. Subspace support vector data description, in: ICPR.

Suarez, J.J.P., Naval Jr, P.C., 2020. A survey on deep learning techniques for video anomaly detection. arXiv .

Sultani, W., Chen, C., Shah, M., 2018. Real-world anomaly detection in surveillance videos, in: CVPR.

Sun, C., Jia, Y., Wu, Y., 2022. Evidential reasoning for video anomaly detection, in: ACM Multimedia.

Tian, Y., Pang, G., Chen, Y., Singh, R., Verjans, J.W., Carneiro, G., 2021. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning, in: ICCV.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. NeurIPS .

Wang, G., Yuan, X., Zheng, A., Hsu, H.M., Hwang, J.N., 2019. Anomaly candidate identification and starting time estimation of vehicles from traffic videos, in: CVPR workshops.

Wang, J., Cherian, A., 2019. Gods: Generalized one-class discriminative subspaces for anomaly detection, in: ICCV.

Wang, M., Xing, J., Liu, Y., 2021. Actionclip: A new paradigm for video action recognition. arXiv .

Wang, Z., Zou, Y., Zhang, Z., 2020. Cluster attention contrast for video anomaly detection, in: ACM Multimedia.

Wu, J.C., Hsieh, H.Y., Chen, D.J., Fuh, C.S., Liu, T.L., 2022. Self-supervised sparse representation for video anomaly detection, in: ECCV.

Wu, P., Liu, J., 2021. Learning causal temporal relation and feature discrimination for anomaly detection. IEEE Transactions on Image Processing .

Wu, P., Liu, J., Shi, Y., Sun, Y., Shao, F., Wu, Z., Yang, Z., 2020. Not only look, but also listen: Learning multimodal violence detection under weak supervision, in: ECCV.

Xu, H., Ghosh, G., Huang, P.Y., Okhonko, D., Aghajanyan, A., Metze, F., Zettlemoyer, L., Feichtenhofer, C., 2021. Videoclip: Contrastive pre-training for zero-shot video-text understanding, in: EMNLP.

Xu, K., Sun, T., Jiang, X., 2019. Video anomaly detection and localization based on an adaptive intra-frame classification network. IEEE Transactions on Multimedia .

Zaheer, M.Z., Mahmood, A., Astrid, M., Lee, S.I., 2020. Claws: Clustering assisted weakly supervised learning with normalcy suppression for anomalous event detection, in: ECCV.

Zaheer, M.Z., Mahmood, A., Khan, M.H., Segu, M., Yu, F., Lee, S.I., 2022. Generative cooperative learning for unsupervised video anomaly detection, in: CVPR.

Zhang, J., Qing, L., Miao, J., 2019. Temporal convolutional network with complementary inner bag loss for weakly supervised anomaly detection, in: ICIP.

Zhong, J.X., Li, N., Kong, W., Liu, S., Li, T.H., Li, G., 2019. Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection, in: CVPR.

Zhou, K., Yang, J., Loy, C.C., Liu, Z., 2022. Learning to prompt for vision-language models. International Journal of Computer Vision .

In this appendix, we provide further details on the implementation and training of the proposed *AnomalyCLIP*. We also provide more details on how we obtain XD-Violence results. Furthermore, we report supplementary results of the ablation performed on the loss of the *Selector model* to support our design choices. Lastly, we offer additional qualitative results.

Appendix A. Implementation Details

Similarly to CoOp (Zhou et al., 2022), context vectors t^{ctx} are randomly initialised by drawing from a zero-mean Gaussian distribution with standard deviation equal to 0.02. We use the CLIP image encoder (Radford et al., 2021), specifically the ViT-B/16 implementation, without fine-tuning, and apply standard CLIP image augmentations to each frame. As supported by the ablation in Table 11, we employ a one-layer axial transformer (Ho et al., 2019) for the Temporal Model with an embedding size of 256 for UCF-Crime (Sultani et al., 2018) and 128 for XD-Violence (Wu et al., 2020), and a two-layer axial transformer with an embedding size of 256 for ShanghaiTech (Liu et al., 2018). In the case of UCF-Crime and XD-Violence, we use the image features of the CLIP space as input to the Temporal Model. However, for ShanghaiTech, we observe an improvement in performance by incorporating the output of the Selector Model as an additional input to the Temporal Model. This is likely because ShanghaiTech is less challenging than the other two and, as a result, the Selector Model already provides sufficient discriminative features.

Consistent with previous work on VAD Tian et al. (2021); Wu et al. (2022); Chen et al. (2022); Li et al. (2022a), we incorporate a random masking strategy during the selection process operated by $\mathcal{S}$. Specifically, we randomly mask 70% of the segments to prevent the model from repeatedly selecting the same segments. This approach ensures a more diverse and representative selection of segments, thus improving the overall performance.

Appendix B. Training Details

Training was performed using the AdamW optimiser (Loshchilov and Hutter, 2019) with parameters $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-8}$ and weight decay $w = 0.2$. We tuned the learning rate and the number of epochs based on the behaviour of the training loss. Specifically, the learning rate is set to 5×10^{-4}, 10^{-5} and 5×10^{-6} for ShanghaiTech, UCF-Crime, and XD-Violence, respectively, warmed up for 10% of the total training epochs and decayed to zero following a cosine annealing schedule. The number of epochs is set to 50 for UCF-Crime and XD-Violence, while it is set to 100 for ShanghaiTech, due to its smaller size. We set the weight for each loss term to 1 without tuning. Following previous work (Tian et al., 2021; Wu et al., 2022; Li et al., 2022a), we use $\lambda_1 = 8 \times 10^{-3}$ ad $\lambda_2 = 8 \times 10^{-4}$ for sparsity and smoothness regularisation terms, respectively.

Table D.16: Ablation of the losses on the Selector model on ShanghaiTech (Liu et al., 2018) The final configuration of our model is represented by the row highlighted in grey in the table.

$\mathcal{L}_A^{\text{DIR}}$	$\mathcal{L}_N^{\text{DIR}}$	AUC	mAUC
		97.86	95.92
	✓	97.35	96.00
✓		97.95	96.35
✓	✓	**98.07**	**96.46**

Table D.17: Ablation of the losses on the Selector model on XD-Violence (Wu et al., 2020) The final configuration of our model is represented by the row highlighted in grey in the table.

$\mathcal{L}_A^{\text{DIR}}$	$\mathcal{L}_N^{\text{DIR}}$	AP	mAP
		77.45	47.74
	✓	**78.69**	48.03
✓		78.16	49.02
✓	✓	78.51	**49.41**

Appendix C. Reproducibility XD-Violence

As the original implementations of RTFM (Tian et al., 2021) and S3R (Wu et al., 2022) do not provide neither the code nor the trained models for XD-Violence (Wu et al., 2020), we made the necessary adaptations to support XD-Violence based on the information available in the original papers and on the open-source platform Github, and used the 2048-D features extracted after the final average pooling layer of the I3D ResNet50 model, pre-trained on Kinetics400 (Kay et al., 2017). First, we pre-train their models on the entire XD-Violence dataset and save the checkpoint at the training iteration that obtains the highest average precision (AP) on the test set, following their training protocol. Subsequently, we maintain the original pre-trained model frozen and introduce a multiclass classification head. This newly introduced head undergoes training following the methodology outlined in Sec. 4.2.

Appendix D. Ablation

Losses. Tables D.16 and D.17 illustrate the contribution of the losses on the Selector model's outputs, where we progressively remove the losses from the full training objective, on ShanghaiTech (Liu et al., 2018) and XD-Violence (Wu et al., 2020), respectively. Both the losses on anomalous and normal videos contribute to better VAD and VAR results.

Appendix E. Qualitative Results

Fig. E.4 presents additional qualitative results of our proposed AnomalyCLIP in detecting and recognising anomalies within a set of UCF-Crime and ShanghaiTech test videos. The model is capable of predicting both the presence of anomalies in test videos and the category of the anomalous event. Videos *Arson016* (Row 1, Column 1), *Arrest001* (Row 1, Column 2) and *Burglary033* (Row 2, Column 2) serve as good examples of the effectiveness of the proposed method. The anomalies are temporally located, and the ground-truth labels (as indicated in

Fig. E.4: Qualitative results for VAR on twelve test videos from UCF-Crime (the top three rows), ShanghaiTech (the fourth row) and XD-Violence (the bottom row). For each video, we show at the bottom the predicted probability of each frame being anomalous by our model over the number of frames. We showcase some key frames to reflect the relevance between the predicted anomaly probability and the visual content. The red shaded areas denote the temporal ground-truth of anomalies. We also indicate the predicted anomalous class for detected abnormal frames in the red boxes, while videos without detected anomalies are indicated with blue boxes as Normal.

the video name) are correctly identified. However, it is worth noting that in *Arson016* some frames are misjudged as Explosion, which is nevertheless a similar type of anomaly.

One failure case is observed in the sample *Shoplifting039* (Row 2, Column 2), where the proposed method fails to detect the anomaly. The reason for this failure could be attributed to the fact that the annotated anomaly is visually very similar to a normal situation, making it difficult even for humans to understand that a shoplifting is taking place and not an authorised person moving an object. This result underscores the challenge of accurately detecting anomalies in complex and visually similar scenarios. In video *Robbery102* (Row 2, Column 1), the anomaly is correctly located but wrongly classified as Assault, indicating the challenges of VAR.

Videos *Shooting032* (Row 4, Column 2) and *Fighting033* (Row 4, Column 1) are interesting examples that highlight the ability of the proposed method to detect anomalous situations even in the aftermath of the anomaly. In these videos, the anomaly probability remains high even after the anomalous situation annotated in the ground truth has ended, correctly indicating that there is still something anomalous happening.

The videos from ShanghaiTech (Row 5, Columns 1-2) also provide insights into the performance of the proposed method. In the video on the left, the anomaly is correctly classified as a vehicle. However, there is also a false alarm, which represents a failure case. On the right side of the last row, the video shows a monocycle anomaly that is wrongly classified as Running. It is reasonable to assume that the fast movement of the person riding the monocycle could have contributed to this misclassification.