

Learning Deep Representations of Appearance and Motion for Anomalous Event Detection

Dan Xu¹
danxuhk@gmail.com

Elisa Ricci^{2,3}
eliricci@fbk.eu

Yan Yan^{1,4}, Jingkuan Song¹
yan@disi.unitn.it, jingkuan.song@unitn.it

Nicu Sebe¹
sebe@disi.unitn.it

¹DISI, University of Trento,
Trento, Italy

²Fondazione Bruno Kessler (FBK),
Trento, Italy

³University of Perugia,
Perugia, Italy

⁴ADSC, UIUC Singapore,
Singapore

Abstract

We present a novel unsupervised deep learning framework for anomalous event detection in complex video scenes. While most existing works merely use hand-crafted appearance and motion features, we propose Appearance and Motion DeepNet (AMDN) which utilizes deep neural networks to automatically learn feature representations. To exploit the complementary information of both appearance and motion patterns, we introduce a novel double fusion framework, combining both the benefits of traditional early fusion and late fusion strategies. Specifically, stacked denoising autoencoders are proposed to separately learn both appearance and motion features as well as a joint representation (*early fusion*). Based on the learned representations, multiple one-class SVM models are used to predict the anomaly scores of each input, which are then integrated with a *late fusion* strategy for final anomaly detection. We evaluate the proposed method on two publicly available video surveillance datasets, showing competitive performance with respect to state of the art approaches.

1 Introduction

A fundamental challenge in intelligent video surveillance is to automatically detect abnormal events in long video streams. This problem has attracted considerable attentions from both academia and industry in recent years [1, 16, 17, 21]. Video anomaly detection is also important as it is related to other interesting topics in computer vision, such as dominant behavior detection [23], visual saliency [57] and interestingness prediction [8]. A typical approach to tackle the anomaly detection task is to learn a model which describes normal activities in the video scene and then discovers unusual events by examining patterns which distinctly diverge from the model. However, the complexity of scenes and the deceptive nature of abnormal behaviours make anomaly detection still a very challenging task.

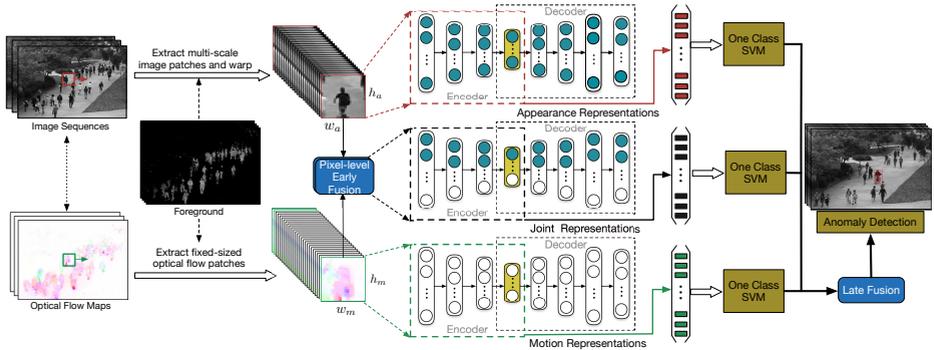


Figure 1: Overview of the proposed AMDN method for anomalous event detection.

Among previous works, several anomaly detection approaches are based on analyzing individual moving objects in the scene. Tracking is usually an initial step for this class of methods. By using accurate tracking algorithms, trajectory extraction can be carried out to further perform trajectory clustering analysis [6, 20] or design representative features [28] to model typical activities and subsequently discover anomalies. In [63], trajectories which are spatially close and have similar motion patterns are identified and used for detecting unusual events. Vaswani *et al.* [29] propose a “shape activity” model to describe moving objects and detect anomalies. However, as tracking performance significantly degrades in the presence of several occluded targets, tracking-based methods are not suitable for analyzing complex and crowded scenes.

To overcome the aforementioned limitations, researchers address the problem by learning spatio/temporal activity patterns in a local or a global context from either 2D image cells or 3D video volumes [2, 9, 8, 18, 21, 22, 24]. This category of methods builds the models based on hand-crafted features extracted from low-level appearance and motion cues, such as color, texture and optical flow. Commonly used low-level features include histogram of oriented gradients (HOG), 3D spatio-temporal gradient, histogram of optical flow (HOF), among others. Cong *et al.* [9] employ multi-scale histograms of optical flow and a sparse coding model and use the reconstruction error as a metric for outlier detection. Mehran *et al.* [18] propose a “social force” model based on optical flow features to represent crowd activity patterns and identify anomalous activities. In [2] co-occurrence statistics of spatio-temporal events are employed in combination with Markov Random Fields (MRFs) to discover unusual activities. Kratz *et al.* [22] introduce a HMMs-based approach for detecting abnormal events by analyzing the motion variation of local space-time volumes. Kim *et al.* [10] propose a method based on local optical flow features and MRFs to spot unusual events. Mahadevan *et al.* [17] introduce a mixtures of dynamic textures model to jointly employ appearance and motion features. However, the adoption of hand-crafted features is a clear limitation of previous methods, as it implies enforcing some a priori knowledge which, in case of complex video surveillance scene, is very difficult to define.

Recently, deep learning architectures have been successfully used to tackle various computer vision tasks, such as image classification [13], object detection [6] and activity recognition [27]. However, these works are mainly based on Convolutional Neural Networks and consider a supervised learning scenario. Unsupervised deep learning approaches based on autoencoder networks [60] have also been investigated to address important tasks such as object tracking [62] and face alignment [68]. The key of the success is that, using deep architectures, rich and discriminative features can be learned via multi-layer nonlinear trans-

formations. Therefore, it is reasonable to expect that detecting unusual events in videos can also benefit from deep learning models.

Following this intuition, in this paper we propose a novel Appearance and Motion Deep-Net (AMDN) framework for discovering anomalous activities in complex video surveillance scenes. Opposite to previous works [8, 21, 24], instead of using hand-crafted features to model activity patterns, we propose to learn discriminative feature representations of both appearance and motion patterns in a fully unsupervised manner. A novel approach based on stacked denoising autoencoders (SDAE) [31] is introduced to achieve this goal. An overview of the proposed AMDN is shown in Fig. 1. Low-level visual information including still image patches and dynamic motion fields represented with optical flow is used as input of two separate networks, to first learn appearance and motion features, respectively. To further investigate the correlations between appearance and motion, early fusion is performed by combining image pixels with their corresponding optical flow to learn a joint representation. Finally, for abnormal event prediction, a late fusion strategy is introduced to combine the anomaly scores predicted by multiple one-class SVM classifiers, each corresponding to one of the three learned feature representations. The benefits of the proposed *double fusion* framework (*i.e.* combining both early fusion and late fusion strategies) are confirmed by our extensive experimental evaluation, conducted on two publicly available datasets.

In summary, the main contributions of this paper are: i) As far as we know, we are the first to introduce an unsupervised deep learning framework to automatically construct discriminative representations for video anomaly detection. ii) We propose a new approach to learn appearance and motion features as well as their correlations. Deep learning methods for combining multiple modalities have been investigated in previous works [14, 27]. However, to our knowledge, this is the first work where multimodal deep learning is applied to anomalous event detection. iii) A double fusion scheme is proposed to combine appearance and motion features for discovering unusual activities. It is worth noting that the advantages of combining early and late fusion approaches have been investigated in previous works [14]. However, Lan *et al.* [24] do not consider a deep learning framework. iv) The proposed method is validated on challenging anomaly detection datasets and we obtain very competitive performance compared with the state-of-the-art.

2 AMDN for Abnormal Event Detection

The proposed AMDN framework for detecting anomalous activities is based on two main building blocks (Fig. 1). First, SDAE are used to learn appearance and motion representations of visual data, as well as a joint representation capturing the correlation between appearance and motion features (Sec. 2.1). In the second phase (Sec. 2.2), to detect anomalous events, we propose to train three separate one-class SVMs [29] based on the three different types of learned feature representations. Once the one-class SVM models are learned, given a test sample corresponding to an image patch, three anomaly scores are computed and combined. The combination of the one-class SVM scores is obtained with a novel late fusion scheme. In the following we describe the proposed approach in details.

2.1 Learning Deep Appearance and Motion Representations

In this subsection we present the proposed AMDN for learning deep representations of appearance and motion. In the following we first introduce denoising autoencoders and then

describe the details of the structure and the learning approach of the proposed AMDN.

2.1.1 Denoising Autoencoders

A Denoising Auto-Encoder (DAE) is a one-hidden-layer neural network which is trained to reconstruct a data point \mathbf{x}_i from its (partially) corrupted version $\tilde{\mathbf{x}}_i$ [60]. Typical corrupted inputs are obtained by drawing samples from a conditional distribution $p(\mathbf{x}|\tilde{\mathbf{x}})$ (e.g. common choices for corrupting samples are additive Gaussian white noise or salt-pepper noise). A DAE neural network can be divided into two parts: encoder and decoder, with a single shared hidden layer. These two parts actually attempt to learn two mapping functions, denoted as $f_e(\mathbf{W}, \mathbf{b})$ and $f_d(\mathbf{W}', \mathbf{b}')$, where \mathbf{W}, \mathbf{b} denote the weights and the bias term of the encoder part, and \mathbf{W}', \mathbf{b}' refer to the corresponding parameters of the decoder. For a corrupted input $\tilde{\mathbf{x}}_i$, a compressed hidden layer representation \mathbf{h}_i can be obtained through $\mathbf{h}_i = f_e(\tilde{\mathbf{x}}_i | \mathbf{W}, \mathbf{b}) = \sigma(\mathbf{W}\tilde{\mathbf{x}}_i + \mathbf{b})$. Then, the decoder tries to recover the original input \mathbf{x}_i from \mathbf{h}_i computing $\hat{\mathbf{x}}_i = f_d(\mathbf{h}_i | \mathbf{W}', \mathbf{b}') = s(\mathbf{W}'\mathbf{h}_i + \mathbf{b}')$. The function $\sigma(\cdot)$ and $s(\cdot)$ are activation functions, which are typically nonlinear transformations such as the sigmoid. Using this encoder/decoder structure, the network can learn a more stable and robust feature representations of the input.

Given a training set $\mathcal{T} = \{\mathbf{x}_i\}_{i=1}^N$, a DAE learns its parameters $(\mathbf{W}, \mathbf{W}', \mathbf{b}, \mathbf{b}')$ by solving the following regularized least square optimization problem:

$$\min_{\mathbf{W}, \mathbf{W}', \mathbf{b}, \mathbf{b}'} \sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2 + \lambda (\|\mathbf{W}\|_F^2 + \|\mathbf{W}'\|_F^2) \quad (1)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The first term represents the average reconstruction error, while the weight penalty term is introduced for regularization. The parameter λ balances the importance of the two terms. Typically, sparsity constraints are imposed on the output of the hidden units to discover meaningful representations from the data [62]. If we let μ_j be the target sparsity level and $\hat{\mu}_j = \frac{1}{N} \sum_{i=1}^N \mathbf{h}_i^j$ be the average activation values all over all training samples for the j -th unit, an extra penalty term based on cross-entropy, $\varphi(\boldsymbol{\mu} | \hat{\boldsymbol{\mu}}) = -\sum_{j=1}^H [\mu_j \log(\hat{\mu}_j) + (1 - \mu_j) \log(1 - \hat{\mu}_j)]$, can be added to (1) to learn a sparse representation. Here, H is the number of hidden units. The optimization problem (1) has a non-convex objective function and gradient descent can be used to compute a local optima.

2.1.2 AMDN Structure

The proposed AMDN structure consists of three SDAE pipelines (Fig. 1) corresponding to different types of low-level inputs. The three SDAE networks learn appearance and motion features as well as a joint representation of them. We show the basic structures of the proposed SDAE networks in Fig. 2 (a) and (b). Each SDAE consists of two parts: encoder and decoder. For the encoder part, we use an over-complete set of filters in the first layer to capture a representative information from the data. Then, the number of neurons is reduced by half in the next layer until reaching the ‘‘bottleneck’’ hidden layer. The decoder part has a symmetric structure with respect to the encoder part. We now describe the proposed three feature learning pipelines in details.

Appearance representation. This SDAE aims at learning mid-level appearance representations from the original image pixels. To capture rich appearance attributes, a multi-scale sliding-window approach with a stride d is used to extract dense image patches, which are then warped into equal size $w_a \times h_a \times c_a$, where w_a, h_a are the width and height of each

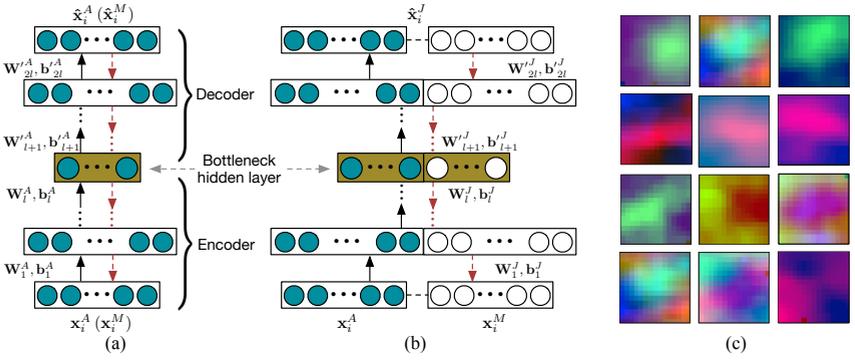


Figure 2: (a) The structure of the appearance and motion representation learning pipelines. (b) The structure of the joint representation learning pipeline. (c) Some examples of the weight filters extracted from the first layer of the joint representation learning pipeline.

patch and c_a is the number of the channels ($c_a = 1$ for gray images). The warped patches $\mathbf{x}_i^A \in \mathbb{R}^{w_a \times h_a \times c_a}$ are used for training. All the patches are linearly normalized into a range $[0, 1]$. We stack 4 encoding layers with $v_a \times w_a \times h_a \times c_a$ neurons in the first layer, where $v_a > 1$ is an amplification factor for constructing an over-complete set of filters.

Motion representation. The motion information is computed with optical flow. We use a sliding window approach with windows of fixed size $w_m \times h_m \times c_m$ ($c_m = 2$ for optical flow magnitude along x and y axes), to generate dense optical flow patches $\mathbf{x}_i^m \in \mathbb{R}^{w_m \times h_m \times c_m}$ for motion representation learning. Similar to the appearance feature pipeline, the patches are normalized into $[0,1]$ within each channel and 4 encoding layers are used. The number of neurons of the first layer is set to $v_m \times w_m \times h_m \times c_m$.

Joint appearance and motion representation. The above-mentioned SDAE learn appearance and motion features separately. Taking into account the correlations between motion and appearance, we propose to couple these two pipelines to learn a joint representation. The network training data $\mathbf{x}_i^J \in \mathbb{R}^{w_j \times h_j \times (c_a + c_m)}$ are obtained through a pixel-level early fusion of the gray image patches and the corresponding optical flow patches.

2.1.3 AMDN Training

We train the AMDN with two steps: pretraining and fine-tuning. The layer-wise pretraining learns one single denoising auto-encoder at a time using (1) with sparsity constraints (Sec.2.1.1). The input is corrupted to learn the mapping function $f_e(\cdot)$, which is then used to produce the representation for the next layer with uncorrupted inputs. By using a greedy layer-wise pretraining, the denoising autoencoders can be stacked to build a multi-layer feed-forward deep neural network, *i.e.* a stacked denoising autoencoder. The network parameters are initialized through pretraining all layers, and then fine-tuning is used to adjust parameters over the whole network.

Fine-tuning treats all the layers of an SDAE as a single model. Given a training set $\mathcal{T}^k = \{\mathbf{x}_i^k\}_{i=1}^{N^k}$ with N^k training samples ($k \in \{A, M, J\}$ corresponds to appearance, motion and joint representation, respectively), the backpropagation algorithm can be used to fine-tune the network. The following objective function is used for fine-tuning the SDAE with $2L + 1$ layers:

$$J(\mathcal{T}^k) = \sum_i^{N^k} \|\mathbf{x}_i^k - \hat{\mathbf{x}}_i^k\|_2^2 + \lambda_F \sum_{i=1}^L (\|\mathbf{W}_i^k\|_F^2 + \|\mathbf{W}_i^k\|_F^2), \quad (2)$$

where λ_F is a user defined parameter. To speed up the convergence during training, stochastic gradient descent (SGD) is employed and the training set is divided into mini-batches with size N_b^k . Fig. 2 (c) shows some of the learned filters in the first layer after fine-tuning for the joint representation learning pipeline.

After fine-tuning the whole network, the learned features representation can be computed to perform video anomaly detection. Theoretically, the output of each layer in an SDAE can be used as a novel learned feature representation. In this work, we choose the output of the ‘‘bottleneck’’ hidden layer to obtain a more compact representation. Let \mathbf{x}_i^k be the i -th input data sample, and $\sigma_l^k(\mathbf{W}_l^k, \mathbf{b}_l^k)$ be the mapping function of the l -th hidden layer of the k -th SDAE pipeline. The learned features, \mathbf{s}_i^k , can be extracted through a forward pass computing $\mathbf{s}_i^k = \sigma_L(\sigma_{L-1}(\cdots \sigma_1(\mathbf{W}_1^k \mathbf{x}_i^k + \mathbf{b}_1^k)))$, where the L -th hidden layer is the ‘‘bottleneck’’ hidden layer.

2.2 Abnormal Event Detection with Deep Representations

We formulate the video anomaly detection problem as a patch-based binary categorization problem, *i.e.* given a test frame we obtain $M_I \times N_I$ patches via sliding window with a stride d and classify each patch as corresponding to a normal or abnormal region. Specifically, given each test patch t we compute three anomaly scores $A^k(\mathbf{s}_t^k)$, $k \in \{A, M, J\}$, using one-class SVM models and the computed features representations \mathbf{s}_t^k . The three scores are then linearly combined to obtain the final anomaly score $\mathcal{A}(\mathbf{s}_t^k) = \sum_{k \in \{A, M, J\}} \alpha^k A^k(\mathbf{s}_t^k)$. A pre-processing of dynamic background subtraction can be carried out to improve the computing efficiency during the test phase as the anomalies exist in the foreground region.

2.2.1 One-class SVM Modeling

One-class SVM is a widely used algorithm for outlier detection, where the main idea is to learn a hypersphere in the feature space and map most of the training data into it. The outliers of the data distribution correspond to point lying outside the hypersphere. Formally, given a set of training samples $\mathcal{S} = \{\mathbf{s}_i^k\}_{i=1}^{N^k}$, the underlying problem of one-class SVM can be formulated as the following quadratic program:

$$\begin{aligned} \min_{\mathbf{w}, \rho} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + \frac{1}{vN^k} \sum_{i=1}^{N^k} \xi_i - \rho \\ \text{s.t.} \quad & \mathbf{w}^T \Phi(\mathbf{s}_i^k) \geq \rho - \xi_i, \xi_i \geq 0. \end{aligned} \quad (3)$$

where \mathbf{w} is the learned weight vector, ρ is the offset, $\Phi(\cdot)$ is a feature projection function which maps feature vector \mathbf{s}_i^k into a higher dimensional feature space. The user defined parameter $v \in (0, 1]$ regulates the expected fraction of outliers distributed outside the hypersphere. Introducing a nonlinear mapping, the projection function $\Phi(\cdot)$ can be defined implicitly by introducing an associated kernel function $k(\mathbf{s}_i^k, \mathbf{s}_j^k) = \Phi(\mathbf{s}_i^k)^T \Phi(\mathbf{s}_j^k)$ and (3) can be solved in the corresponding dual form [25]. In our experiments we consider a rbf kernel, $k(\mathbf{s}_i^k, \mathbf{s}_j^k) = e^{-\frac{\|\mathbf{s}_i^k - \mathbf{s}_j^k\|^2}{2\sigma^2}}$. Given the optimal \mathbf{w} and ρ obtained by solving (3), an outlier score for a test sample \mathbf{s}_t^k of the k -th SDAE pipeline can be estimated by computing $A^k(\mathbf{s}_t^k) = \rho - \mathbf{w}^T \Phi(\mathbf{s}_t^k)$.

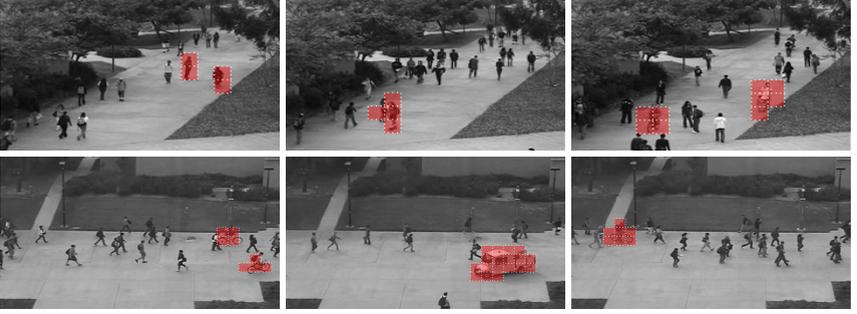


Figure 3: Examples of anomaly detection results on Ped1 (top) and Ped2 (bottom) sequences. Our approach can detect anomalies such as bikes, vehicles and skaters and provides accurate localization.

2.2.2 Late Fusion for Anomaly Detection

An unsupervised late fusion scheme is designed to automatically learn the weight vector $\boldsymbol{\alpha} = [\alpha^A, \alpha^M, \alpha^J]$. The weight learning approach is based on the following optimization:

$$\begin{aligned} \min_{\mathbf{W}_s^k, \alpha^k} \quad & \sum_k \alpha^k \text{tr} \left(\mathbf{W}_s^k \mathbf{S}^k (\mathbf{W}_s^k \mathbf{S}^k)^T \right) + \lambda_s \|\boldsymbol{\alpha}\|_2^2 \\ \text{s.t.} \quad & \alpha^k > 0, \quad \sum_k \alpha^k = 1 \end{aligned} \quad (4)$$

where $\mathbf{S}^k = [\mathbf{s}_1^k, \dots, \mathbf{s}_{N^k}^k]$ is the matrix of training set samples, \mathbf{W}_s^k maps the k -th feature \mathbf{s}_i^k into a new subspace and $\mathbf{W}_s^k \mathbf{S}^k (\mathbf{W}_s^k \mathbf{S}^k)^T$ represents the covariance of k -th feature type in the new subspace. The term $\|\boldsymbol{\alpha}\|_2^2$ is introduced to avoid overfitting and λ_s is a user defined parameter. To solve (4), we first get \mathbf{W}_s^k as the d eigenvectors of $\mathbf{S}^k \mathbf{S}^{kT}$ corresponding to the d -largest eigenvalues. Then, $\boldsymbol{\alpha}$ can be obtained by solving the simplex problem [39]:

$$\min_{\alpha^k > 0, \sum_k \alpha^k = 1} \frac{1}{2} \|\boldsymbol{\alpha} - \mathbf{c}\|_2^2 \quad (5)$$

with $\mathbf{c} = [c^A, c^M, c^J]$, $c^k = -\frac{1}{2\lambda_s} \text{tr} \left(\mathbf{W}_s^k \mathbf{S}^k (\mathbf{W}_s^k \mathbf{S}^k)^T \right)$. Then for each patch t , we identify if it corresponds to an abnormal activity by computing the associated anomaly score $\mathcal{A}(\mathbf{s}_t^k)$ and comparing it with a threshold η , *i.e.* $\mathcal{A}(\mathbf{s}_t^k) \underset{\text{abnormal}}{\overset{\text{normal}}{\leq}} \eta$.

3 Experimental Results

Datasets and Experimental Setup. The proposed method is mainly implemented in Matlab and C++ based on Caffe framework [9]. The code for optical flow calculation is written in C++ and wrapped with Matlab mex for computational efficiency [15]. For one-class SVMs, we use the LIBSVM library (version 3.2) [9]. The experiments are carried out on a PC with a middle-level graphics card (NVIDIA Quadro K4000) and a multi-core 2.1 GHz CPU with 32 GB memory. Two publicly available datasets, the UCSD (Ped1 and Ped2) dataset [17] and the Train dataset [56], are used to evaluate the performance of the proposed approach.

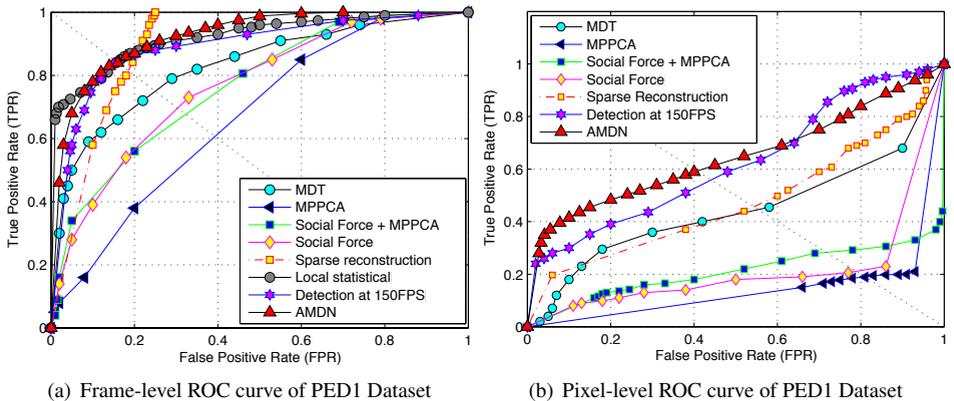


Figure 4: UCSD dataset (Ped1 sequence): comparison of frame-level and pixel-level anomaly detection results with state of the art methods.

Algorithm	Ped1(frame)		Ped1(pixel)		Ped2	
	EER	AUC	EER	AUC	EER	AUC
MPPCA [16]	40%	59.0%	81%	20.5%	30%	69.3%
Social force [17]	31%	67.5%	79%	19.7%	42%	55.6%
Social force+MPPCA [16]	32%	66.8%	71%	21.3%	36%	61.3%
Sparse reconstruction [18]	19%	–	54%	45.3%	–	–
Mixture dynamic texture [19]	25%	81.8%	58%	44.1%	25%	82.9%
Local Statistical Aggregates [20]	16%	92.7%	–	–	–	–
Detection at 150 FPS [21]	15%	91.8%	43%	63.8%	–	–
Joint representation (early fusion)	22%	84.9%	47.1%	57.8%	24%	81.5%
Fusion of appearance and motion pipelines (late fusion)	18%	89.1%	43.6%	62.1%	19%	87.3%
AMDN (double fusion)	16%	92.1%	40.1%	67.2%	17%	90.8%

Table 1: UCSD dataset: comparison in terms of EER (Equal Error Rate) and AUC (Area Under ROC) with the state of the art methods on Ped1 and Ped2.

The **UCSD pedestrian** dataset [17] is a challenging anomaly detection dataset including two subsets: Ped1 and Ped2. The video sequences depict different crowded scenes and anomalies include bicycles, vehicles, skateboarders and wheelchairs. In some frames the anomalies occur at multiple locations. Ped1 has 34 training and 16 test image sequences with about 3,400 anomalous and 5,500 normal frames, and the image resolution is 238×158 pixels. Ped2 has 16 training and 12 test image sequences with about 1,652 anomalous and 346 normal frames. The image resolution is 360×240 pixels.

The **Train** dataset [16] depicts moving people inside a train. This is also a challenging abnormal event detection dataset due to dynamic illumination changes and camera shake problems. The dataset consists of 19218 frames, and the anomalous events are mainly due to unusual movements of people on the train.

Quantitative evaluation. In the first series of experiments we evaluate the performance of the proposed method on the UCSD dataset. For appearance learning, patches are extracted using a sliding window approach at three different scales, *i.e.* 15×15 , 18×18 and 20×20 pixels. This generates more than 50 million image patches, 10 million of which are randomly sampled and warped into the same size ($w_a \times h_a = 15 \times 15$ pixels) for training. For learning the motion representation, the patch size is fixed to $w_m \times h_m = 15 \times 15$ pixels, and 6 million training patches are randomly sampled. In the test phase, we use sliding widow with a size of 15×15 and a stride $d = 15$. The number of neurons of the first layer of the appearance and motion network is both set to 1024, while for the joint pipeline is 2048. Then the encoder

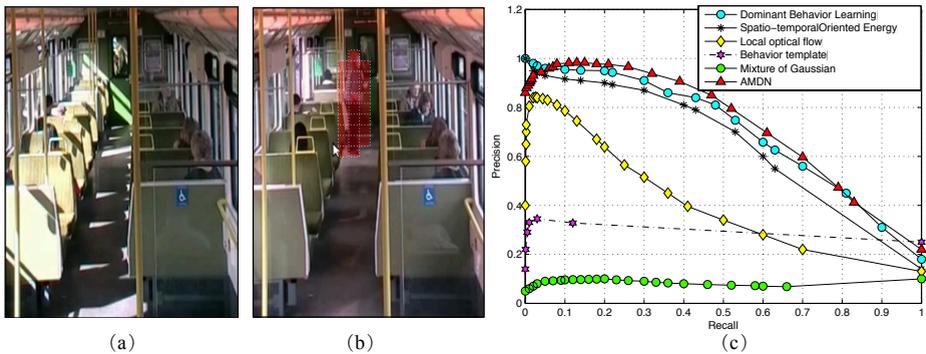


Figure 5: Train dataset: (a) a frame depicting typical activities, (b) an example of a detected anomaly. (c) precision/recall curve.

part can be simply defined as: $1024(2048) \rightarrow 512(1024) \rightarrow 256(512) \rightarrow 128(256)$, and the decoder part is a symmetric structure. For the pre-training of the DAE, the corrupted inputs are produced by adding a Gaussian noise with variance 0.0003. The network training is based on the SGD with a momentum parameter set to 0.9. We use a fixed learning rates $\lambda = 0.01$, $\lambda_F = 0.0001$ and a mini-batch size $N_b = 256$. For one-class SVMs, the parameter ν is tuned with cross validation. The learned late fusion weights $[\alpha^A, \alpha^M, \alpha^J]$ are obtained with $\lambda_s = 0.1$ and for Ped1 and Ped2 are $[0.2, 0.5, 0.3]$ and $[0.2, 0.4, 0.4]$, respectively.

Some examples of anomaly detection results on the UCSD dataset are shown in Fig. 3. To perform a quantitative evaluation, we use both a frame-level ground truth and a pixel-level ground truth. The frame-level ground truth represents whether one or more anomalies occur in a test frame. The pixel-level ground truth is used to evaluate the anomaly localization performance. If the detected anomaly region is more than 40% overlapping with the annotated region, it is a true detection. We carry out a frame-level evaluation on both Ped1 and Ped2. Ped1 provides 10 test image sequences with pixel-level ground truth. The pixel-level evaluation is performed on these test sequences.

Fig. 4 (a) and (b) show the frame-level and pixel-level detection results on Ped1. The ROC curve is produced by varying the threshold parameter η . Table 1 shows a quantitative comparison in terms of Area Under Curve (AUC) and Equal Error Rate (EER) of our method with several state-of-the-art approaches. From the frame-level evaluation, it is evident that our method outperforms most previous methods and that its performance are competitive with the best two baselines [24] and [16]. Moreover, considering pixel-level evaluation, *i.e.* accuracy in anomaly localization, our method outperforms all the competing approaches both in terms of EER and AUC. Table 1 also demonstrate the advantages of the proposed double fusion strategy. Our AMDN guarantees better performance than early fusion and late fusion approaches. Specifically, for early fusion we only consider the learned joint appearance/motion representation and a single one-class SVM. For late fusion we use the two separate appearance and motion pipelines and the proposed fusion scheme but we discard the joint representation pipeline. Interestingly, in this application the late fusion strategy outperforms early fusion.

In the second series of experiments we consider the Train dataset. For parameters, we use the same experimental setting of the UCSD experiments, except from the parameters λ_F and N_b which are set to 0.00001 and 100, respectively. The learned late fusion weights are $[\alpha^A, \alpha^M, \alpha^J] = [0.3, 0.4, 0.3]$. We compare the proposed approach with several methods,

including dominant behavior learning [23], spatio-temporal oriented energies [36], local optical flow [4], behavior templates [10] and mixture of Gaussian. From the precision/recall curve shown in Fig. 5 (c), it is clear that our method outperforms all the baselines.

4 Conclusions

We presented a novel unsupervised learning approach for video anomaly detection based on deep representations. The proposed method is based on multiple stacked autoencoder networks for learning both appearance and motion representations of scene activities. A double fusion scheme is designed to combine the learned feature representations. Extensive experiments on two challenging datasets demonstrate the effectiveness of the proposed approach and show competitive performance with respect to existing methods. Future works include investigating other network architectures, alternative approaches for fusing multimodal data in the context of SDAE, and extending our framework using multi-task learning [34, 35] for detecting anomalies in heterogeneous scenes.

Acknowledgments

This work was partially supported by the MIUR Cluster project Active Ageing at Home, the EC H2020 project ACANTO and by A*STAR Singapore under the Human-Centered Cyber-physical Systems (HCCS) grant. The authors also would like to thank NVIDIA for GPU donation.

References

- [1] A. Adam, E. Rivlin, I. Shimshoni, and D. Reinitz. Robust real-time unusual event detection using multiple fixed-location monitors. *IEEE Transactions on PAMI*, 30(3): 555–560, 2008.
- [2] Y. Benezeth, P. Jodoin, V. Saligrama, and C. Rosenberger. Abnormal events detection based on spatio-temporal co-occurrences. In *CVPR*, 2009.
- [3] C. Chang and C. Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27, 2011.
- [4] Y. Cong, J. Yuan, and J. Liu. Sparse reconstruction cost for abnormal event detection. In *CVPR*, 2011.
- [5] Z. Fu, W. Hu, and T. Tan. Similarity based vehicle trajectory clustering and anomaly detection. In *ICIP*, 2005.
- [6] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [7] H. Grabner, F. Nater, M. Druey, and L. Van Gool. Visual interestingness in image sequences. In *ACM MM*, 2013.
- [8] T. Hospedales, S. Gong, and T. Xiang. Video behaviour mining using a dynamic topic model. *IJCV*, 98(3):303–323, 2012.

- [9] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM MM*, 2014.
- [10] P. Jodoin, J. Konrad, and V. Saligrama. Modeling background activity for behavior subtraction. In *ICDSC*, 2008.
- [11] J. Kim and K. Grauman. Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates. In *CVPR*, 2009.
- [12] L. Kratz and K. Nishino. Anomaly detection in extremely crowded scenes using spatio-temporal motion pattern models. In *CVPR*, 2009.
- [13] A. Krizhevsky, I. Sutskever, and G.E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [14] Z. Lan, L. Bao, S. Yu, W. Liu, and A. G. Hauptmann. Double fusion for multimedia event detection. In *Advances in Multimedia Modeling*, 2012.
- [15] C. Liu. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [16] C. Lu, J. Shi, and J. Jia. Abnormal event detection at 150 fps in matlab. In *ICCV*, 2013.
- [17] V. Mahadevan, W. Li, V. Bhalodia, and N. Vasconcelos. Anomaly detection in crowded scenes. In *CVPR*, 2010.
- [18] R. Mehran, A. Oyama, and M. Shah. Abnormal crowd behavior detection using social force model. In *CVPR*, 2009.
- [19] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng. Multimodal deep learning. In *ICML*, 2011.
- [20] C. Piciarelli, C. Micheloni, and G. Foresti. Trajectory-based anomalous event detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(11):1544–1554, 2008.
- [21] V. Reddy, C. Sanderson, and B.C. Lovell. Improved anomaly detection in crowded scenes via cell-based analysis of foreground speed, size and texture. In *CVPRW*, 2011.
- [22] E. Ricci, G. Zen, N. Sebe, and S. Messelodi. A prototype learning framework using emd: Application to complex scenes analysis. *IEEE Transactions on PAMI*, 35(3): 513–526, 2013.
- [23] M. Roshtkhari and M. Levine. Online dominant and anomalous behavior detection in videos. In *CVPR*, 2013.
- [24] V. Saligrama and Z. Chen. Video anomaly detection based on local statistical aggregates. In *CVPR*, 2012.
- [25] B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, and R. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.

- [26] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [27] N. Srivastava and R. R. Salakhutdinov. Multimodal learning with deep boltzmann machines. In *NIPS*, 2012.
- [28] C. Stauffer and W.E.L. Grimson. Learning patterns of activity using real-time tracking. *IEEE Transactions on PAMI*, 22(8):747–757, 2000.
- [29] N. Vaswani, A. Roy-Chowdhury, and R. Chellappa. " shape activity": a continuous-state hmm for moving/deforming shapes with application to abnormal activity detection. *IEEE Transactions on Image Processing*, 14(10):1603–1616, 2005.
- [30] P. Vincent, H. Larochelle, Y. Bengio, and P. Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008.
- [31] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *JMLR*, 11:3371–3408, 2010.
- [32] N. Wang and D. Yeung. Learning a deep compact image representation for visual tracking. In *NIPS*, 2013.
- [33] X. Wang, K. Tieu, and E. Grimson. Learning semantic scene models by trajectory analysis. In *ECCV*, 2006.
- [34] Y. Yan, E. Ricci, R. Subramanian, O. Lanz, and N. Sebe. No matter where you are: Flexible graph-guided multi-task learning for multi-view head pose classification under target motion. In *ICCV*, 2013.
- [35] Y. Yan, E. Ricci, R. Subramanian, G. Liu, and N. Sebe. Multi-task linear discriminant analysis for multi-view action recognition. *IEEE Transactions on Image Processing*, 23(12):5599–5611, 2014.
- [36] R. Zaharescu, A. and Wildes. Anomalous behaviour detection using spatiotemporal oriented energies, subset inclusion histogram comparison and event-driven processing. In *ECCV*, 2010.
- [37] Y. Zhai and M. Shah. Visual attention detection in video sequences using spatiotemporal cues. In *ACM MM*, 2006.
- [38] J. Zhang, S. Shan, M. Kan, and X. Chen. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In *ECCV*, 2014.
- [39] G. Zoutendijk. Methods of feasible directions: a study in linear and non-linear programming. 1960.