

Synthetic Temporal Anomaly Guided End-to-End Video Anomaly Detection

Marcella Astrid^{1,2}, Muhammad Zaigham Zaheer^{1,2}, Seung-Ik Lee^{1,2}

¹University of Science and Technology, ²Electronics and Telecommunications Research Institute,
 Daejeon, South Korea

{marcella.astrid, mzz}@ust.ac.kr, the_silee@etri.re.kr

Abstract

Due to the limited availability of anomaly examples, video anomaly detection is often seen as one-class classification (OCC) problem. A popular way to tackle this problem is by utilizing an autoencoder (AE) trained only on normal data. At test time, the AE is then expected to reconstruct the normal input well while reconstructing the anomalies poorly. However, several studies show that, even with normal data only training, AEs can often start reconstructing anomalies as well which depletes their anomaly detection performance. To mitigate this, we propose a temporal pseudo anomaly synthesizer that generates fake-anomalies using only normal data. An AE is then trained to maximize the reconstruction loss on pseudo anomalies while minimizing this loss on normal data. This way, the AE is encouraged to produce distinguishable reconstructions for normal and anomalous frames. Extensive experiments and analysis on three challenging video anomaly datasets demonstrate the effectiveness of our approach to improve the basic AEs in achieving superiority against several existing state-of-the-art models.

1. Introduction

Anomaly detection in video sequences has recently attracted significant attention because of its importance in surveillance systems [37, 41, 18, 4, 16, 21, 19, 1, 46]. As anomalous events occur rarely in real-life situations and collecting plenty of anomalous examples can be cumbersome, this task is extremely challenging. Therefore, anomaly detection is often seen as one-class classification (OCC) problem in which only normal data is used to train a novelty detection model [7, 18, 4, 35, 39, 17, 14]. At test time, the events that do not conform to the learned representations are considered anomalous.

One common way to tackle the OCC problem is by using a deep autoencoder (AE) [8, 45, 39, 4, 21, 29, 20, 35, 7]. By training to minimize reconstruction error on normal data, the model is encouraged to extract the features represent-

ing normal data in its latent space. This way, at test time, the network is expected to poorly reconstruct the anomalous cases. However, as previously observed by several researchers [47, 26, 39], AEs can also often successfully reconstruct anomalous examples. In such cases, the reconstruction loss between normal and anomalous data may not be discriminative enough to successfully identify the anomalies.

Recently, a new addition to the field of OCC is the idea of utilizing pseudo anomalies generated from the normal training data. For example, Zaheer *et al.* [39] fuse two random images from normal data to generate appearance anomalies and use them to train an image classifier. However, this work needs old and new states of the reconstructor and trains in a two-phase scheme. Furthermore, the approach is limited to appearance and does not consider any temporal information for anomaly detection. On the contrary, our work proposes a simple yet highly effective temporal pseudo anomaly synthesizer to assist the training of an AE in an end-to-end fashion without any bells and whistles. For each pseudo anomaly example as input, our AE model is trained to produce high reconstruction loss. This helps in limiting the capability of AE to reconstruct anomalies at test time.

We owe the inspiration of our approach to the intuition that detecting fast or suddenly changing motion is significantly important and closely related to detecting anomalies. For example, it is a common observation that animals often associate strong motion with dangerous situations [6]. The case is similar for humans as well. For instance, people running unusually may indicate life-threatening situations nearby, such as fires [11, 5] or natural disasters [9, 15]. Moreover, fights or robberies may also be characterized by sudden strong motions. Some other examples may include riding bikes or vehicles on pedestrian sidewalks, over-speeding vehicles, etc. [2, 36, 27]. Therefore, we hypothesize that most of the anomalous events can be characterized by the motion depicted.

To this end, we propose a temporal pseudo anomaly synthesizer that injects synthesized anomaly examples into the

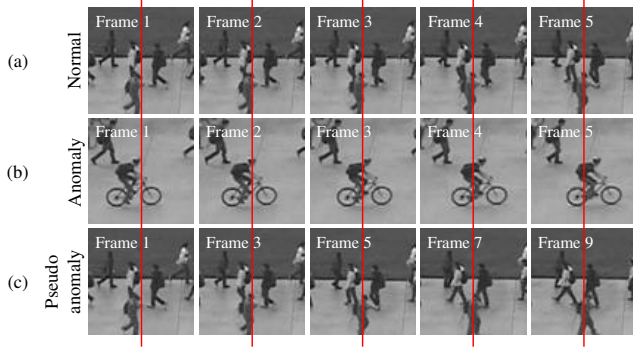


Figure 1. Visualization of normal, anomalous, and pseudo anomalous frames. Given red lines as reference, (a) shows a normal movement pattern in which humans are walking in a usual pace. (b) shows an anomalous movement, i.e. the anomalous object almost completely crosses the reference line within few frames. (c) shows the output of our pseudo anomaly synthesizer generated from normal frames which mimics anomalous movements.

training of an AE. To simulate anomalous movements from normal data, we arbitrarily skip few frames to generate pseudo anomaly sequences as shown in Fig. 1. The overall training is then carried out to minimize the reconstruction loss of normal data while maximizing the reconstruction loss of synthesized anomaly data. Note that, unlike existing motion tracking based anomaly detection methods [3, 32, 43, 24], our approach does not extract any handpicked motion information. Rather, by complementing a temporal pseudo anomaly synthesizer with a deep AE, we harness the power of deep learning to detect a variety of anomalous activities inside videos. Our extensive experiments and analysis demonstrate the superior capability of our approach in three challenging anomaly detection datasets, i.e. Ped2 [16], Avenue [19], and ShanghaiTech [21].

In summary, the contributions of this paper are as follows: 1) We propose to train one-class classifiers with the assistance of temporal pseudo anomaly synthesizer to produce highly discriminative normal and anomalous reconstructions. 2) Extensive experiments demonstrate the superiority of our method compared to a wide range of existing state-of-the-art (SOTA) works [12, 25, 22, 19, 44, 33, 8, 45, 20, 21, 34, 39, 1, 4, 18, 7, 29, 10] on three benchmark datasets.

2. Related Work

Since it is not easy to obtain anomaly examples [31], a popular approach for anomaly detection is one-class classification (OCC) in which the training is carried out using only normal data. Several works train a one-class classifier using the features extracted through object trackers [3, 24, 32, 43]. However, such handpicked features can of-

ten limit the generalization capability of the network on different kinds of activities.

With the recent popularity of deep learning, several researchers [8, 45, 39, 4, 28, 21, 29, 20, 35, 7] utilize autoencoder (AE) based networks to learn normal data representations. An AE is trained only on normal data for reconstruction [7, 29, 8, 45, 35, 39] or prediction [29, 18] tasks. At test time, the model is expected to produce poor reconstructions of anomalous data, which correspond to high anomaly scores. Some AE based approaches use appearance information only [35, 39] while several others use both appearance and motion information [45, 4, 8, 28, 20]. However, AEs can often reconstruct anomalous data as well [47, 26, 39]. Consequently, normal and anomalous data become less distinguishable.

Several attempts have been made to limit the reconstruction capability of an AE on anomalous data. Memory-based networks [7, 29] employ a memory mechanism over the latent space between the encoder and the decoder of an AE. The network is restricted to use only the memorized normalcy definitions which limits its capability to reconstruct anomalous data. However, such networks are highly dependent on the memory size and a small-sized memory may also limit their normal data reconstruction capability. In our approach, we also attempt to limit the reconstruction capability of an AE on anomalous inputs. However, instead of using a memory network, we utilize a pseudo anomaly synthesizer to generate fake anomaly examples and encourage the AE to produce high reconstruction loss on these examples.

The most related to our work is OGNet [39], which fuses normal images to generate appearance anomalies and train the network using both normal and fake anomaly examples. Moreover, OGNet requires a two-phase adversarial training scheme in which a discriminator is trained based on two previously frozen generator models. On the other hand, our approach of synthesizing anomalies is substantially different. We propose to utilize temporal information rather than appearance to synthesize anomalies. Moreover, our method is end-to-end trainable and complements the conventional training of an AE.

In order to enhance the discrimination of normal and anomalous data, some researchers [26, 38, 41, 40, 46, 42, 37] propose to deviate from the fundamental definition of OCC by using real anomaly examples during training. Our method on the other hand utilizes only normal data to synthesize pseudo anomaly examples for training, thus following the conventional OCC protocol.

3. Methodology

In this section, we present our proposed STEAL (Synthetic TEmporal AnomaLy guided end-to-end video anomaly detection) network. Due to the unavailability of




Figure 2. Our approach trains an autoencoder using normal as well as pseudo anomaly sequences. Pseudo anomalies are synthesized using normal data only. The quantity of pseudo anomaly examples is regulated by the probability p . The overall configuration is trained to reduce the reconstruction loss for normal inputs while increasing it for pseudo anomaly inputs.

anomalous examples during training, most of the AE based anomaly detection approaches often fail to discriminate anomalies from normal data at test time. Therefore, we propose the utilization of fake anomaly examples, that are generated by our pseudo anomaly synthesizer using only normal training videos, to enhance the performance of AEs.

3.1. Architecture

Our overall architecture is shown in Fig. 2. We train a conventional AE as our baseline, which takes a sequence of normal frames as input and produces its reconstruction as output. To complement the baseline training, we propose a pseudo anomaly synthesizer that generates fake anomaly examples. These examples are then used for training with a probability p . This way, we limit the reconstruction capability of the AE by forcing it to increase the reconstruction loss on these fake anomaly examples. Finally, the anomaly score is computed using the frame-level reconstruction loss. Each component of our architecture is discussed next:

3.1.1 Autoencoder

In order to capture robust representations, autoencoders (AE) are often designed to take multi-frame inputs [7, 29, 8, 45]. Therefore, we set our AE model to take X as input of size $T \times C \times H \times W$, where T , C , H , and W are the number of frames in the input sequence, number of channels, height of frames, and width of frames, respectively. The reconstruction \hat{X} is then given as:

$$\hat{X} = \mathcal{D}(\mathcal{E}(X)), \quad (1)$$

where \mathcal{E} and \mathcal{D} are the encoder and the decoder of our model, respectively.

Conventionally, an AE is trained to reconstruct using only normal data in the training set. At test time, it is expected to well reconstruct the normal data while poorly reconstruct the anomalous examples. However, this is not always the case. AEs can often “generalize” too well and reconstruct anomalous examples as well [47, 26, 39]. As there are no anomaly examples utilized in the training time, it can be difficult to train a reconstruction based anomaly detection model [39]. To mitigate this problem, we propose a pseudo anomaly synthesizer that can provide fake anomaly examples. These examples are then used to restrain the generation capability of the AE and encourage it to produce high reconstruction loss on any kind of anomalous inputs.

During training, a sequence of video frames X is given to the network as:

$$X = \begin{cases} X_N & \text{with probability } 1 - p, \\ X_P & \text{with probability } p, \end{cases} \quad (2)$$

where X_P is a sequence of frames generated using our proposed pseudo anomaly synthesizer, X_N is a sequence of frames from normal training data, and p is the probability that defines the ratio of pseudo anomaly examples used. Note that the pseudo anomalies are only introduced during training. At test time, we simply input the original sequence of video frames to AE.

3.1.2 Temporal Pseudo Anomaly Synthesizer

In our proposed approach, we follow the OCC protocols by utilizing only normal data to carry out the training. Therefore, in order to generate temporal pseudo anomalies, we utilize the normal training videos only. Similar to the common practices in training a conventional AE [8, 29],

we extract sequences of frames X_N from a training video $V_i = \{I_1, I_2, \dots, I_{K_i}\}$ of length K_i frames by randomly selecting a frame index n from V_i and then taking a fixed number of consecutive T frames thereafter, as described:

$$\begin{aligned} X_N &= (I_n, I_{n+1}, \dots, I_{n+T-1}) \\ &= (I_{n+t})_{0 \leq t < T, n+T-1 \leq K_i}. \end{aligned} \quad (3)$$

On the other hand, pseudo anomalies X_P are synthesized by introducing a skip frame parameter s to Eq. (3) as:

$$\begin{aligned} X_P &= (I_n, I_{n+s}, \dots, I_{n+(T-1)s}) \\ &= (I_{n+ts})_{0 \leq t < T, n+(T-1)s \leq K_i, s > 1}. \end{aligned} \quad (4)$$

The skip frame parameter s controls the number of frames we skip to generate temporal pseudo anomaly examples. An example pseudo anomaly sequence with $s = 2$ is visualized in Fig. 1(c).

3.2. Training

In order to learn normal representations, a conventional AE is trained on X_N by minimizing reconstruction loss between the input frame I_{n+t} and its reconstruction \hat{I}_{n+t} as follows:

$$L_N = \frac{1}{T \times C \times H \times W} \sum_{t=0}^{T-1} \left\| \hat{I}_{n+t} - I_{n+t} \right\|_F^2, \quad (5)$$

where $\|\cdot\|_F$ means Frobenius norm.

For X_P generated by our temporal pseudo anomaly synthesizer, the loss can be similarly defined as:

$$L_P = -\frac{1}{T \times C \times H \times W} \sum_{t=0}^{T-1} \left\| \hat{I}_{n+ts} - I_{n+ts} \right\|_F^2. \quad (6)$$

Note the negative sign in Eq. (6), which is introduced to increase the reconstruction loss of pseudo anomaly examples. This helps in limiting the reconstruction capability of our AE on anomalous inputs.

Then, the overall loss L for training takes the form:

$$L = \begin{cases} L_N & \text{if } X = X_N, \\ L_P & \text{if } X = X_P. \end{cases} \quad (7)$$

3.3. Anomaly Score

At test time, concurrent to the existing approaches [7, 29, 18, 8, 45, 39, 41], we predict anomaly scores at frame level. Moreover, we compute these scores by utilizing the reconstruction based Peak Signal to Noise Ratio (PSNR). According to Mathieu *et al.* [23], PSNR is often a better assessment of image quality than the reconstruction loss itself. Recently, it has also been utilized in anomaly detection [29, 18] where PSNR between an input frame and its

reconstruction is used to calculate the anomaly score. In our approach, we compute PSNR \mathcal{P}_t as:

$$\mathcal{P}_t = 10 \log_{10} \frac{[M_{\hat{I}_t}]^2}{\frac{1}{R} \left\| \hat{I}_t - I_t \right\|_F^2}, \quad (8)$$

where R is the total number of pixels in \hat{I}_t , t is the frame index, and $M_{\hat{I}_t}$ is the maximum possible value of \hat{I}_t .

Then, following [29, 18], we normalize the PSNR value to the range of $[0, 1]$ by a min-max normalization over all the frames in a test video V_i as follows:

$$\mathcal{Q}_t = \frac{\mathcal{P}_t - \min_t(\mathcal{P}_t)}{\max_t(\mathcal{P}_t) - \min_t(\mathcal{P}_t)}, \quad (9)$$

where t is the frame index of V_i . In Eq. (9), a higher \mathcal{Q}_t represents lower reconstruction loss compared to the other frames in V_i and vice versa. Therefore, we calculate the final anomaly score \mathcal{A}_t as:

$$\mathcal{A}_t = 1 - \mathcal{Q}_t. \quad (10)$$

4. Experiments

4.1. Datasets

We evaluate our approach on three widely popular video anomaly detection datasets, i.e. Ped2 [16], Avenue [19], and ShanghaiTech [21]. We utilize the standard division of the datasets in which training splits consist of only normal videos. Whereas, every video in each of the test sets contains one or more anomalous portions.

Ped2. This dataset consists of 16 training and 12 test videos [16]. Pedestrians dominate most of the normal frames whereas anomalies include bikes, carts, or skateboards.

Avenue. This dataset contains 16 training and 21 test videos [19]. Anomalies include abnormal objects such as bikes and abnormal actions of humans such as unusual walking directions, running, or throwing stuff.

ShanghaiTech. This is by far the largest one-class anomaly detection dataset [21]. It consists of 330 training and 107 test videos. The dataset is recorded at 13 different locations having complex lighting conditions and camera angles. In total, the test videos contain 130 anomalous events including running, riding bicycle, and fighting.

4.2. Experimental Setup

Evaluation criteria. To evaluate our approach, we follow the widely popular frame-level area under the ROC curve (AUC) metric [12, 25, 22, 19, 44, 33, 8, 45, 20, 21, 34, 39, 1, 4, 18, 7, 29, 10]. The ROC curve is obtained by varying the threshold of anomaly scores to plot false and true positive rates across the whole test set of each dataset. Higher AUC values represent more accurate results.

Parameters and implementation details. We adopt a generative architecture recently proposed by Gong *et al.* [7] as our baseline that takes an input sequence X (Eq. (1)) of size $16 \times 1 \times 256 \times 256$ and produces its reconstruction of the same size. All 16 frames are used for computing the reconstruction loss during training (Eq. (5) - (6)). At test time, following [7], only 9th frame out of the 16 frames is considered for anomaly score calculation (Eq. (8) - (10)). However, differently from the original implementation, we remove the memory network and utilize only the autoencoder part. Furthermore, we add Tanh output layer to have the output range of $[-1, 1]$.

The implementation of STEAL Net and the baseline is done in PyTorch [30]. The training is carried out using Adam [13] with a learning rate of 10^{-4} and the batch size is set to 4. The skip frame parameter s in Eq. (4) is set to $\{2, 3, 4, 5\}$, which means s can be randomly selected as 2, 3, 4, 5 each time we generate a pseudo anomaly sequence. The probability p for pseudo anomaly in Eq. (2) is set to 0.01. The baseline in our results refer to the model trained without pseudo anomalies, i.e. the probability p in Eq. (2) is set to 0.

4.3. Quantitative Results

Table 1 shows the AUC comparisons of our proposed STEAL Net on Ped2 [16], Avenue [19], and ShanghaiTech [21] datasets. Our approach yields better performance than the baseline on all three datasets. Specifically, our approach demonstrates an absolute gain of 5.9%, 5.6%, and 2.4% AUC on Ped2, Avenue, and ShanghaiTech datasets respectively. This shows that our approach successfully improves the capability of the baselines on several diverse datasets.

Comparing the performance with existing approaches, we achieve superior performances on all three datasets. Although MNAD-Prediction [29] achieves better performance on Avenue compared to our method, its superiority may be attributed to different base architecture as well as training settings, i.e., prediction task. Nevertheless, compared with MNAD-Reconstruction, which has a similar base architecture as ours and performs the same task of reconstruction, our model yields noticeably better performance, thus demonstrating the superiority of our approach.

Overall, the superior performance of STEAL Net validates our hypothesis that several different kinds of anomalies may be characterized by the motions present within. Therefore, by utilizing a temporal pseudo anomaly synthesizer, we unleash the potential of a deep autoencoder for more accurate anomaly detection.

4.4. Qualitative Results

Input images and the comparisons of their reconstructions produced by STEAL Net and the baseline are shown in the first three columns of Fig. 3. In addition, reconstruc-

Table 1. AUC performance comparison of our approach with several SOTA methods on Ped2, Avenue (Ave), and ShanghaiTech (Sh). Best performances in each dataset are highlighted as bold whereas second best are highlighted as underlined.

| Methods | Ped2 [16] | Ave [19] | Sh [21] |
|----------------------------|--------------|--------------|--------------|
| MPPCA [12] | 69.3% | - | - |
| MPPCA+SFA [12] | 61.3% | - | - |
| Mehran <i>et al.</i> [25] | 55.6% | - | - |
| MDT [22] | 82.9% | - | - |
| Lu <i>et al.</i> [19] | - | 80.9% | - |
| LSHF [44] | 91.0% | - | - |
| Ramachandra and Jones [33] | 88.3% | 72.0% | - |
| AE-Conv2D [8] | 85.0% | 80.0% | 60.9% |
| AE-Conv3D [45] | 91.2% | 77.1% | - |
| AE-ConvLSTM [20] | 88.1% | 77.0% | - |
| TSC [21] | 91.0% | 80.6% | 67.9% |
| StackRNN [21] | 92.2% | 81.7% | 68.0% |
| AbnormalGAN [34] | 93.5% | - | - |
| OGNet [39] | <u>98.1%</u> | - | - |
| LSA [1] | 95.4% | - | 72.5% |
| Cluster AE [4] | 96.5% | 86.0% | 73.3% |
| Frame-Pred [18] | 95.4% | 85.1% | 72.8% |
| MemAE [7] | 94.1% | 83.3% | 71.2% |
| MNAD-Prediction [29] | 97.0% | 88.5% | 70.5% |
| MNAD-Reconstruction [29] | 90.2% | 82.8% | 69.8% |
| TAM-Net [10] | <u>98.1%</u> | 78.3% | - |
| Baseline | 92.5% | 81.5% | 71.3% |
| STEAL Net | 98.4% | <u>87.1%</u> | 73.7% |

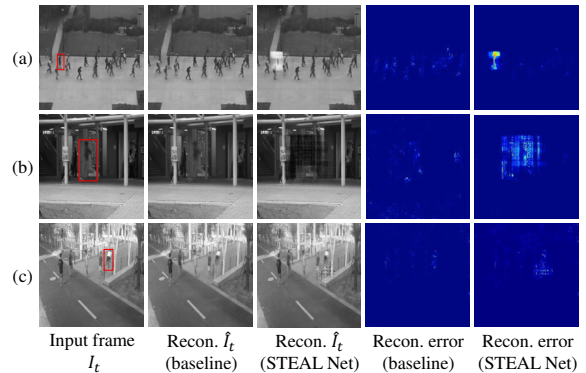


Figure 3. Reconstruction comparison between our approach and the baseline on normal and anomalous frames from Ped2, Avenue, and ShanghaiTech datasets. Our approach specifically attempts to distort the anomalous objects thus producing more discrimination. Anomaly ground truths are marked as red boxes.

tion error heat maps are visualized in the last two columns. These heat maps are generated by computing the squared error of each pixel between the reconstruction and the input frame followed by min-max normalization. As our approach is specifically trained to produce high reconstruction loss on anomalous inputs, it consequently attempts to satisfy this condition by distorting the anomalous regions. On the other hand, concurrent with the reports in several research works [47, 26, 39], the baseline reconstructs the

anomalies well, thus often failing to produce noticeable discrimination.

5. Conclusion

We proposed the utilization of pseudo anomalies, that are generated using only normal data, to assist the training of an autoencoder (AE) for video anomaly detection. In addition to the conventional training of AEs where a network only attempts to minimize reconstruction error on inputs, we further encourage the network to maximize this loss on pseudo anomalies. The effectiveness of our approach in complementing the AEs to achieve superiority against several existing state-of-the-art models is extensively analyzed on three challenging video anomaly datasets.

6. Acknowledgements

This work was supported by the ICT R&D program of MSIT/IITP. [2019-0-01309, Development of AI Technology for Guidance of a Mobile Robot to its Goal with Uncertain Maps in Indoor/Outdoor Environments]. Also, we thank Ki-In Na and Jae-Yeong Lee for the discussions and support in improving our work.

References

- [1] Davide Abati, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Latent space autoregression for novelty detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2019. 1, 2, 4, 5
- [2] Lisa Aultman-Hall and M Georgina Kaltenecker. Toronto bicycle commuter safety rates. *Accident Analysis & Prevention*, 31(6):675–686, 1999. 1
- [3] Arslan Basharat, Alexei Gritai, and Mubarak Shah. Learning object motion patterns for anomaly detection and improved object detection. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2
- [4] Yunpeng Chang, Tu Zhigang, Xie Wei, and Yuan Junsong. Clustering driven deep autoencoder for video anomaly detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 4, 5
- [5] Dominic Elliott and Denis Smith. Football stadia disasters in the united kingdom: learning from tragedy? *Industrial & Environmental Crisis Quarterly*, 7(3):205–229, 1993. 1
- [6] Christopher S Evans, Joseph M Macedonia, and Peter Marler. Effects of apparent size and speed on the response of chickens, *gallus gallus*, to computer-generated simulations of aerial predators. *Animal Behaviour*, 46(1):1–11, 1993. 1
- [7] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1705–1714, 2019. 1, 2, 3, 4, 5
- [8] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 733–742, 2016. 1, 2, 3, 4, 5
- [9] Mingli Hu, Zhigeng Fang, Sifeng Liu, Wei Chen, and Baohua Yang. Study on evolution mechanism of fateful stampede accident based on graphical evaluation and review technique. In *2010 IEEE International Conference on Systems, Man and Cybernetics*, pages 2976–2978. IEEE, 2010. 1
- [10] Xiangli Ji, Bairong Li, and Yuesheng Zhu. Tam-net: Temporal enhanced appearance-to-motion generative network for video anomaly detection. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 2, 4, 5
- [11] John P Keating. The myth of panic. *Fire Journal*, 76(3):57–61, 1982. 1
- [12] Jaechul Kim and Kristen Grauman. Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2928. IEEE, 2009. 2, 4, 5
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [14] Yuh-Jye Lee, Yi-Ren Yeh, and Yu-Chiang Frank Wang. Anomaly detection via online over-sampling principal component analysis. *IEEE transactions on knowledge and data engineering*, 25(7):1460–1470, 2012. 1
- [15] Meifang Li, Yongxiang Zhao, Lerong He, Wenxiao Chen, and Xianfeng Xu. The parameter calibration and optimization of social force model for the real-life 2013 ya’an earthquake evacuation in china. *Safety science*, 79:243–253, 2015. 1
- [16] Weixin Li, Vijay Mahadevan, and Nuno Vasconcelos. Anomaly detection and localization in crowded scenes. *IEEE transactions on pattern analysis and machine intelligence*, 36(1):18–32, 2013. 1, 2, 4, 5
- [17] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422. IEEE, 2008. 1
- [18] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection - a new baseline. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6536–6545, 2018. 1, 2, 4, 5
- [19] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal event detection at 150 fps in matlab. In *Proceedings of the IEEE international conference on computer vision*, pages 2720–2727, 2013. 1, 2, 4, 5
- [20] Weixin Luo, Wen Liu, and Shenghua Gao. Remembering history with convolutional lstm for anomaly detection. In *2017 IEEE International Conference on Multimedia and Expo (ICME)*, pages 439–444. IEEE, 2017. 1, 2, 4, 5
- [21] Weixin Luo, Wen Liu, and Shenghua Gao. A revisit of sparse coding based anomaly detection in stacked rnn framework. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 341–349, 2017. 1, 2, 4, 5

- [22] Vijay Mahadevan, Weixin Li, Viral Bhalodia, and Nuno Vasconcelos. Anomaly detection in crowded scenes. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1975–1981. IEEE, 2010. 2, 4, 5
- [23] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 4
- [24] Gérard Medioni, Isaac Cohen, François Brémont, Somboon Hongeng, and Ramakant Nevatia. Event detection and analysis from video streams. *IEEE Transactions on pattern analysis and machine intelligence*, 23(8):873–889, 2001. 2
- [25] Ramin Mehran, Alexis Oyama, and Mubarak Shah. Abnormal crowd behavior detection using social force model. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 935–942. IEEE, 2009. 2, 4, 5
- [26] Asim Munawar, Phongtharin Vinayavekhin, and Giovanni De Magistris. Limiting the reconstruction capability of generative neural network using negative learning. In *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2017. 1, 2, 3, 5
- [27] AP Muthusamy, M Rajendran, K Ramesh, and P Sivaprakash. A review on road traffic accident and related factors. *International Journal of Applied Engineering Research*, 10(11):28177–28183, 2015. 1
- [28] Trong-Nguyen Nguyen and Jean Meunier. Anomaly detection in video sequence with appearance-motion correspondence. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1273–1283, 2019. 2
- [29] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning memory-guided normality for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14372–14381, 2020. 1, 2, 3, 4, 5
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 5
- [31] Josh Patterson and Adam Gibson. *Deep learning: A practitioner's approach*. "O'Reilly Media, Inc.", 2017. 2
- [32] Claudio Piciarelli, Christian Micheloni, and Gian Luca Foresti. Trajectory-based anomalous event detection. *IEEE Transactions on Circuits and Systems for video Technology*, 18(11):1544–1554, 2008. 2
- [33] Bharathkumar Ramachandra and Michael Jones. Street scene: A new dataset and evaluation protocol for video anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2569–2578, 2020. 2, 4, 5
- [34] Mahdyar Ravanbakhsh, Moin Nabi, Enver Sangineto, Lucio Marcenaro, Carlo Regazzoni, and Nicu Sebe. Abnormal event detection in videos using generative adversarial nets. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 1577–1581. IEEE, 2017. 2, 4, 5
- [35] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3379–3388, 2018. 1, 2
- [36] Neal Sikka, C Vila, M Stratton, Mateen Ghassemi, and Ali Pourmand. Sharing the sidewalk: A case of e-scooter related pedestrian injury. *The American journal of emergency medicine*, 37(9):1807–e5, 2019. 1
- [37] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6479–6488, 2018. 1, 2
- [38] Yuki Yamanaka, Tomoharu Iwata, Hiroshi Takahashi, Masanori Yamada, and Sekitoshi Kanai. Autoencoding binary classifiers for supervised anomaly detection. In *Pacific Rim International Conference on Artificial Intelligence*, pages 647–659. Springer, 2019. 2
- [39] Muhammad Zaigham Zaheer, Jin-ha Lee, Marcella Astrid, and Seung-Ik Lee. Old is gold: Redefining the adversarially learned one-class classifier training paradigm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14183–14193, 2020. 1, 2, 3, 4, 5
- [40] Muhammad Zaigham Zaheer, Jin-Ha Lee, Marcella Astrid, Arif Mahmood, and Seung-Ik Lee. Cleaning label noise with clusters for minimally supervised anomaly detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, June 2020. 2
- [41] Muhammad Zaigham Zaheer, Arif Mahmood, Marcella Astrid, and Seung-Ik Lee. Claws: Clustering assisted weakly supervised learning with normalcy suppression for anomalous event detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 4
- [42] Muhammad Zaigham Zaheer, Arif Mahmood, Hochul Shin, and Seung-Ik Lee. A self-reasoning framework for anomaly detection using video-level labels. *IEEE Signal Processing Letters*, 27:1705–1709, 2020. 2
- [43] Tianzhu Zhang, Hanqing Lu, and Stan Z Li. Learning semantic scene models by object classification and trajectory clustering. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1940–1947. IEEE, 2009. 2
- [44] Ying Zhang, Huchuan Lu, Lihe Zhang, Xiang Ruan, and Shun Sakai. Video anomaly detection based on locality sensitive hashing filters. *Pattern Recognition*, 59:302–311, 2016. 2, 4, 5
- [45] Yiru Zhao, Bing Deng, Chen Shen, Yao Liu, Hongtao Lu, and Xian-Sheng Hua. Spatio-temporal autoencoder for video anomaly detection. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1933–1941, 2017. 1, 2, 3, 4, 5
- [46] Jia-Xing Zhong, Nannan Li, Weijie Kong, Shan Liu, Thomas H Li, and Ge Li. Graph convolutional label noise cleaner: Train a plug-and-play action classifier for anomaly detection. In *Proceedings of the IEEE Conference on Com-*

- puter Vision and Pattern Recognition*, pages 1237–1246, 2019. [1](#), [2](#)
- [47] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning Representations*, 2018. [1](#), [2](#), [3](#), [5](#)