STAN: SPATIO-TEMPORAL ADVERSARIAL NETWORKS FOR ABNORMAL EVENT DETECTION

Sangmin Lee, Hak Gu Kim, and Yong Man Ro*

Image and Video Systems Lab., School of Electrical Engineering, KAIST, South Korea

ABSTRACT

In this paper, we propose a novel abnormal event detection method with spatio-temporal adversarial networks (STAN). We devise a spatio-temporal generator which synthesizes an inter-frame by considering spatio-temporal characteristics with bidirectional ConvLSTM. A proposed spatio-temporal discriminator determines whether an input sequence is realnormal or not with 3D convolutional layers. These two networks are trained in an adversarial way to effectively encode spatio-temporal features of normal patterns. After the learning, the generator and the discriminator can be independently used as detectors, and deviations from the learned normal patterns are detected as abnormalities. Experimental results show that the proposed method achieved competitive performance compared to the state-of-the-art methods. Further, for the interpretation, we visualize the location of abnormal events detected by the proposed networks using a generator loss and discriminator gradients.

Index Terms— Abnormal event detection, adversarial learning, spatio-temporal features, interpretation

1. INTRODUCTION

With a rapid development of storage and video acquisition devices, the video surveillance system is widely used for privacy protection, process monitoring in factories, criminal tracking, etc. Most of the surveillance scenes captured for a long time are meaningless and normal. It is time-consuming and labor-intensive for people to watch long hours of meaningless scenes. To address the problem, automatic abnormal event (*i.e.*, meaningful moment) detection in videos has increasingly attracted attention in image processing and computer vision fields.

There were many existing methods for automatically detecting abnormal events in videos. In [1, 2], trajectory based abnormal event detection methods which utilize the dynamic information of objects were proposed. Hu *et al.* [1] proposed method for learning statistical motion patterns with multiobject tracking. Zhou *et al.* [2] proposed statistical model to detect abnormal behaviors using Kanade-Lucas-Tomasi Feature Tracker (KLT). However, trajectory based methods are not robust to occlusions and crowded scenes since most of the object tracking algorithms are vulnerable to such environments. In [3, 4], handcraft feature based abnormal event detection methods were proposed. Kaltsa *et al.* [3] utilized a descriptor created from Histograms of Oriented Gradients (HOG) and Histograms of Oriented Swarms (HOS). In [4], distributions of spatio-temporal oriented energy were used to model behavior. Although the handcraft feature based methods are more robust to occlusions and complex scenes than trajectory based methods, they require prior knowledge to design proper features for various events.

In recent years, deep learning has attracted attention in many challenging tasks with better performance compared to the traditional methods [5, 6, 7]. In this flow, several methods using deep learning have been proposed for abnormal event detection. Hasan et al. [8] proposed the convolutional autoencoder based method for learning temporal regularity. In [9], the ConvLSTM [10] was employed to consider the spatiotemporal characteristics of event patterns. Recently, generative adversarial networks (GAN) based method [11] was proposed. In [11], one GAN is trained to generate an optical flow map from a frame while the other GAN is trained to generate a frame from an optical flow map. Although the method of [11] achieved good performance by using GAN, its performance could be sensitive to the quality of the estimated optical flow map. It could not be robust to scenes with occlusions in which it is difficult to well estimate the optical flow map.

In this paper, we propose the spatio-temporal adversarial networks (STAN) which learn the spatio-temporal features of normal patterns. Learning the characteristics of abnormal events is very difficult because the definition of such events is not spatially or temporally bounded. Thus, it is reasonable to learn normal patterns and consider deviations from the learned patterns as abnormalities. With the proposed STAN, we directly generate spatio-temporal scene without using the optical flow map. The STAN contains two networks which are the spatio-temporal generator and the spatio-temporal discriminator. The proposed generator consists of three network modules which are a spatial encoder for encoding spatial features of frames, a bidirectional ConvLSTM for encoding temporal feature of the scene, and a spatial decoder for generating an inter-frame. By using the bidirectional configuration, the spatio-temporal features of normal patterns can be encoded

^{*} Corresponding author (ymro@ee.kaist.ac.kr)



Fig. 1. Spatio-temporal adversarial networks at training time.

in forward and backward directions. The proposed discriminator is a 3D convolutional neural network for determining whether an input sequence is real or not.

By the adversarial learning, the generator tries to generate the inter-frame in order to deceive the discriminator while the discriminator tries to discriminate between a real sequence and a fake sequence. Note that the fake sequence includes the generated inter-frame. At testing time, abnormal events are detected by using a pixel-wise loss for the generator and an adversarial loss for the discriminator.

Experimental results show that the proposed method performs competitively compared to the state-of-the-art methods, especially on complex scenes with large objects and frequent occlusions. In addition, we visually interpret the proposed networks by locating detected abnormal events through both the generator and the discriminator.

The remainder of this paper is organized as follows. Section 2 describes the proposed method in detail. In Section 3, experimental results are shown for qualitative and quantitative evaluation. Finally, conclusions are drawn in Section 4.

2. PROPOSED METHOD

Fig. 1 shows the overall procedure of the proposed STAN at training time. The proposed spatio-temporal generator generates the inter-frame by considering neighboring frames. Then, consecutive frames including the generated inter-frame are considered as a fake sequence while a real sequence contains only real frames. Through the adversarial learning, the spatio-temporal discriminator is trained to determine whether the input sequence is real or fake. The generator is trained to generate the inter-frame that can fool the discriminator. Details are described in following sub-sections.

2.1. Spatio-temporal generator

The role of the proposed generator is to generate a clear frame for the normal scene and a distorted frame for the abnormal scene. In video data, each frame has a high correlation with neighboring frames. We devise the bidirectional ConvLSTM to generate the inter-frame by considering the spatio-temporal characteristics of the neighboring frames. With the bidirectional configuration, the spatio-temporal features of normal patterns are encoded effectively. As a result, scenes with spatio-temporal abnormal patterns can be sensitively detected.

The proposed generator consists of the spatial encoder, the bidirectional ConvLSTM, and the spatial decoder. By the spatial encoder, latent spatial features are encoded for representing visual information of each frame. The spatial features of the previous 5 frames and the later 5 frames are fed to the forward ConvLSTM and the backward ConvLSTM respectively, in order to encode the spatio-temporal features. Then, the hidden states and the cell states resulting from each ConvLSTM module are concatenated in pairs to combine forward and backward features. As shown in Fig. 1, the combined ConvLSTM outputs the latent feature of the inter-frame. Finally, the inter-frame is generated by the spatial decoder.

To generate the desirable inter-frame, a realism loss and a pixel-wise loss are designed for the generator. By minimizing the realism loss, the generator can force the discriminator to fail to classify the fake sequence as fake. Let G_{θ} and D_{ϕ} denote the generator function with parameter θ and the discriminator function with parameter ϕ , respectively. Let $\mathbf{X}_t \in \mathbb{R}^{224 \times 224 \times 1}$ and $\hat{\mathbf{X}}_t \in \mathbb{R}^{224 \times 224 \times 1}$ denote the *t*-th real frame and the *t*-th generated frame by G_{θ} , respectively. The realism loss can be written as

$$\ell_{\text{real}}(\boldsymbol{\theta}; t) = -\log(D_{\boldsymbol{\phi}}(\hat{\mathbf{S}}_t)), \tag{1}$$

where $\hat{\mathbf{S}}_t \in \mathbb{R}^{11 \times 224 \times 224 \times 1}$ is the fake sequence including *t*-th generated frame $\hat{\mathbf{X}}_t$.

In addition, by minimizing the pixel-wise loss between the *t*-th real frame \mathbf{X}_t and the *t*-th generated frame $\hat{\mathbf{X}}_t$, the generated frame resembles the real frame in pixel-level. The pixel-wise loss can be written as

$$\ell_{\text{pixel}}(\boldsymbol{\theta}; t) = \|G_{\boldsymbol{\theta}}(\mathbf{X}_{t-5}, ..., \mathbf{X}_{t-1}, \mathbf{X}_{t+1}, ..., \mathbf{X}_{t+5}) - \mathbf{X}_t\|_2.$$
(2)

Finally, the total loss for the generator is defined as a combination of the two losses, which can be written as

$$\mathcal{L}_{G}(\boldsymbol{\theta}) = \sum_{t \in batch} \ell_{real}(\boldsymbol{\theta}; t) + \lambda \ell_{pixel}(\boldsymbol{\theta}; t), \qquad (3)$$

where λ is a hyper-parameter to balance the realism loss and the pixel-wise loss.

 Table 1. Architectures of the proposed networks.

Generator			Discriminator			
Layer	Filter/ Stride	Output Size $(h \times w \times c)$	Layer	Filter/ Stride	Output Size $(l \times h \times w \times c)$	
Conv1	5×5/ (2, 2)	112×112×16	2D Convil	5×5×5/	7~112~112~22	
Conv2	5×5/ (2, 2)	56×56×32	3D Collv I	(1, 2, 2)	/ ^ 112 × 112 × 32	
Conv3	3×3/ (2, 2)	28×28×64	2D Convil	3×5×5/	5~56~56~64	
Conv4	3×3/ (1, 1)	28×28×128	3D Collv1	(1, 2, 2)	J × J0 × J0 × 04	
Forward	3×3/	20 20 <i>(</i> 1) 20 G)		3×3×3/	3×28×28×128	
ConvLSTM (1, 1)		28×28×64	3D Conv3	(1, 2, 2)		
Backward	3×3/			3×3×3/	1×14×14×256	
ConvLSTM	(1, 1)	28×28×64	3D Conv4	(1, 2, 2)		
Combined	3×3/			1×3×3/	1×7×7×512	
ConvLSTM	(1, 1)	28×28×128	3D Conv5	(1, 2, 2)		
DeConv1	3×3/ (1, 1)	28×28×64		1×3×3/	1×7×7×1	
DeConv2	$3 \times 3/$ (2, 2)	56×56×32	3D Conv5	(1, 1, 1)		
DeConv3	5×5/ (2, 2)	112×112×16				
DeConv4	5×5/ (2, 2)	224×224×1				

2.2. Spatio-temporal discriminator

The proposed spatio-temporal discriminator consists of 3D convolutional layers. The 3D CNN can reliably determine whether the input sequence is real or fake by considering both spatial and temporal characteristics of the scene [12]. The discriminator outputs 7×7 probability map for employing the local adversarial loss [13].

To distinguish the real sequence and the fake sequence, an adversarial loss is designed for the discriminator. By minimizing the loss, the discriminator can discriminate between the real sequence and the fake sequence. The adversarial loss for the discriminator can be written as

$$\mathcal{L}_D(\boldsymbol{\phi}) = \sum_{t \in batch} -\log(1 - D_{\boldsymbol{\phi}}(\hat{\mathbf{S}}_t)) - \log(D_{\boldsymbol{\phi}}(\mathbf{S}_t)), \quad (4)$$

where $\mathbf{S}_t \in \mathbb{R}^{11 \times 224 \times 224 \times 1}$ is the real sequence. The first term, $-\log(1 - D_{\phi}(\hat{\mathbf{S}}_t))$, allows the discriminator to classify the fake sequence as fake. The second term, $-\log(D_{\phi}(\mathbf{S}_t))$, allows the discriminator to classify the real sequence as real.

The proposed discriminator plays two roles. First, during the adversarial learning, the discriminator helps the generator learn the spatio-temporal features of normal patterns. Second, after the adversarial learning, the discriminator alone can detect abnormal events. The discriminator is trained to determine only real sequence containing normal events as real. Thus, the discriminator considers the real sequence containing abnormal events as not real. Consequently, the second term in (4) has a higher value for video sequences including abnormal events and can be used to detect abnormalities. Table 1 shows architecture details of the networks.

2.3. Abnormality score

To quantify detected abnormalities, we devise a novel abnormality score by using the losses of the generator and the discriminator. The generator and the discriminator complement detection results of each other. Thereby, robust abnormality detection can be achieved. The proposed abnormality loss can be defined as

$$\ell_s(t) = \|G_{\boldsymbol{\theta}}(\mathbf{X}_{t-5}, ..., \mathbf{X}_{t-1}, \mathbf{X}_{t+1}, ..., \mathbf{X}_{t+5}) - \mathbf{X}_t\|_2 - \lambda_s \log(D_{\boldsymbol{\phi}}(\mathbf{S}_t)),$$
(5)

where λ_s is a hyper-parameter to balance the generator detection and the discriminator detection.

Then, by normalizing $\ell_s(t)$, abnormality score s(t) at t-th frame can be written as

$$s(t) = \frac{\ell_s(t) - \min_t \ell_s(t)}{\max_t \ell_s(t) - \min_t \ell_s(t)}.$$
(6)

Note that video sequences containing abnormal events have higher abnormality scores.

3. EXPERIMENTS AND RESULTS

3.1. Dataset

For experiments, we used three datasets: UCSD Ped1 [14], UCSD Ped2 [14] and Avenue [15] datasets. The training sets of these datasets contain only normal events while the testing sets contain both normal and abnormal events. The UCSD Ped1 dataset consists of 34 clips for training and 36 clips for testing. The UCSD Ped2 dataset consists of 16 clips for training and 12 clips for testing. The Avenue dataset consists of 16 clips for training and 21 clips for testing. The UCSD Ped1 and Ped2 datasets contain the motion of small objects in a broad area. On the other hand, the Avenue dataset contains complex motion of large objects and frequent occlusions.

3.2. Implementation details

We used the Adam [16] to optimize the networks with a learning rate of 0.0002 and a batch size of 3. At training time, first, only the generator was trained to minimize the pixel-wise loss. Then, the generator and the discriminator were trained in the adversarial way to minimize $\mathcal{L}_G(\theta)$ and $\mathcal{L}_D(\phi)$ alternately. We set λ as 1 in (3) and λ_s as maximum value ratio divided by 10 of the two terms in (5). At the end of the generator and the discriminator, the tanh and the sigmoid were used as activation functions respectively. For all the other parts, the exponential linear unit (ELU) was used.

	AUC (%)			
Method	UCSD Ped1	UCSD Ped2	Avenue	
MPPCA [17]	59.0	69.3	N / A	
Social force (SF) [18]	67.5	55.6	N / A	
MPPCA + SF [14]	68.8	61.3	N / A	
Detection at 150fps [15]	91.8	N / A	80.9	
AMDN [19]	92.1	90.8	N / A	
Conv-AE [8]	81.0	90.0	70.2	
ConvLSTM-AE [20]	75.5	88.1	77.0	
Unmasking [21]	68.4	82.2	80.6	
Optical flow-GAN [11]	97.4	93.5	N / A	
Stacked RNN [22]	N / A	92.2	81.7	
Proposed method (G_{θ})	81.6	95.9	86.6	
Proposed method (D_{ϕ})	71.6	86.2	81.8	
Proposed method $(G_{ heta} ext{ and } D_{\phi})$	82.1	96.5	87.2	

 Table 2. Frame-level performance of the proposed method.

 Table 3. Event-level performance of the proposed method.

Method	Correct Detection / False Alarm		Precision (%)	
	UCSD Ped1	UCSD Ped2	UCSD Ped1	UCSD Ped2
IT-AE [8]	36/11	12/3	76.6	80.0
Conv-AE [8]	38/6	12/1	86.3	92.3
Proposed method $(G_{ heta} ext{ and } D_{\phi})$	37 / 3	12/0	92.5	100.0

3.3. Performance evaluation

To evaluate the performance of the proposed method, two evaluation metrics were employed. First, we employed the frame-level evaluation [23] based on the area under curves (AUC). Second, we employed event-level evaluation [8] based on the number of detected events.

Table 2 shows frame-level performance compared with handcraft and sparsity based methods [14, 15, 17, 18] and deep learning based methods [8, 11, 19, 20, 21, 22]. The AUC values of the other methods were taken from each paper and [19]. Interestingly, the proposed method even with G_{θ} or D_{ϕ} showed respectably high AUC values on the Avenue. The proposed method with both G_{θ} and D_{ϕ} outperformed the other methods including state-of-the-art methods. On the other hand, the performance on the UCSD Ped1 showed a bit lower than that on the other datasets. We observed that the UCSD Ped1 included several corrupted frames in testing clips as shown in Fig. 2. By the proposed method, these corrupted frames were detected as abnormal events even though such frames were denoted as normal events in ground truth labels.

Table 3 shows event-level performance comparisons with correct detection, false alarm, and precision. The proposed method achieved higher precision results compared to the other methods in event-level evaluation.

3.4. Interpretation of the proposed networks

Fig. 3 shows localization and visualization of abnormal events detected by the proposed method for interpreting the net-



Fig. 2. Examples of the corrupted frames in UCSD Ped1.



Fig. 3. Localization and visualization results of the abnormal events on UCSD Ped1 (first row), UCSD Ped2 (Second row), and Avenue (third row). (a) Real frame, (b) Generated frame, (c) Abnormality visualization by the generator, and (d) Abnormality visualization by the discriminator.

works. Fig. 3(c) shows visualization results obtained from the pixel-wise loss between the real frame and the generated frame. The proposed generator could detect the abnormal event regions such as a car, a bicycle, and a jumping person. Further, we visualized gradients of the discriminator using guided backpropagation [24]. Fig. 3(d) shows gradients of activation at the output of the discriminator with respect to the input real sequences including the abnormal events. We observed that magnitude of the gradients was significantly high around the abnormal event regions.

Through these results, we confirmed that each network performs abnormal event detection with the proper basis. The proposed networks could not only determine whether each frame is abnormal or not, but also detect where abnormal events occur at each frame.

4. CONCLUSIONS

In this paper, we proposed the novel generative model based abnormal event detection method. To effectively represent the spatio-temporal features of normal patterns, the spatiotemporal generator and the spatio-temporal discriminator were trained in the adversarial way. As a result, the proposed method achieved competitive performance compared to the state-of-the-art methods. In particular, the proposed method far outperformed other methods on the dataset containing complex motion and frequent occlusions. Further, by visualizing the abnormal event regions detected by the proposed method, we could interpret how the STAN determine abnormal event at each scene.

5. REFERENCES

- Weiming Hu, Xuejuan Xiao, Zhouyu Fu, Dan Xie, Tieniu Tan, and Steve Maybank, "A system for learning statistical motion patterns," *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 9, pp. 1450–1464, 2006.
- [2] Shifu Zhou, Wei Shen, Dan Zeng, and Zhijiang Zhang, "Unusual event detection in crowded scenes by trajectory analysis," in *ICASSP*. IEEE, 2015, pp. 1300–1304.
- [3] Vagia Kaltsa, Alexia Briassouli, Ioannis Kompatsiaris, Leontios J Hadjileontiadis, and Michael Gerasimos Strintzis, "Swarm intelligence for detecting interesting events in crowded environments," *IEEE transactions on image processing*, vol. 24, no. 7, pp. 2153–2166, 2015.
- [4] Andrei Zaharescu and Richard Wildes, "Anomalous behaviour detection using spatiotemporal oriented energies, subset inclusion histogram comparison and eventdriven processing," in ECCV. Springer, 2010, pp. 563– 576.
- [5] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012, pp. 1097–1105.
- [6] Dae Hoe Kim, Wissam Baddar, Jinhyeok Jang, and Yong Man Ro, "Multi-objective based spatio-temporal feature representation learning robust to expression intensity variations for facial expression recognition," *IEEE Transactions on Affective Computing*, 2017.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *NIPS*, 2014, pp. 2672–2680.
- [8] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis, "Learning temporal regularity in video sequences," in *CVPR*. IEEE, 2016, pp. 733–742.
- [9] Jefferson Ryan Medel and Andreas Savakis, "Anomaly detection in video using predictive convolutional long short-term memory networks," *arXiv preprint arXiv:1612.00390*, 2016.
- [10] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo, "Convolutional lstm network: A machine learning approach for precipitation nowcasting," in *NIPS*, 2015, pp. 802–810.
- [11] Mahdyar Ravanbakhsh, Moin Nabi, Enver Sangineto, Lucio Marcenaro, Carlo Regazzoni, and Nicu Sebe, "Abnormal event detection in videos using generative adversarial nets," in *ICIP*. IEEE, 2017.

- [12] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*. IEEE, 2015, pp. 4489–4497.
- [13] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, and Russ Webb, "Learning from simulated and unsupervised images through adversarial training," in *CVPR*, 2017, pp. 2242–2251.
- [14] Vijay Mahadevan, Weixin Li, Viral Bhalodia, and Nuno Vasconcelos, "Anomaly detection in crowded scenes," in *CVPR*. IEEE, 2010, pp. 1975–1981.
- [15] Cewu Lu, Jianping Shi, and Jiaya Jia, "Abnormal event detection at 150 fps in matlab," in *ICCV*. IEEE, 2013, pp. 2720–2727.
- [16] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [17] Jaechul Kim and Kristen Grauman, "Observe locally, infer globally: A space-time mrf for detecting abnormal activities with incremental updates," in *CVPR*. IEEE, 2009, pp. 2921–2928.
- [18] Ramin Mehran, Alexis Oyama, and Mubarak Shah, "Abnormal crowd behavior detection using social force model," in *CVPR*. IEEE, 2009, pp. 935–942.
- [19] Dan Xu, Yan Yan, Elisa Ricci, and Nicu Sebe, "Detecting anomalous events in videos by learning deep representations of appearance and motion," *Computer Vision and Image Understanding*, vol. 156, pp. 117–127, 2017.
- [20] Weixin Luo, Wen Liu, and Shenghua Gao, "Remembering history with convolutional lstm for anomaly detection," in *ICME*. IEEE, 2017, pp. 439–444.
- [21] Radu Tudor Ionescu, Sorina Smeureanu, Bogdan Alexe, and Marius Popescu, "Unmasking the abnormal events in video," in *ICCV*. IEEE, 2017, pp. 2914–2922.
- [22] Weixin Luo, Wen Liu, and Shenghua Gao, "A revisit of sparse coding based anomaly detection in stacked rnn framework," in *ICCV*. IEEE, 2017, pp. 341 – 349.
- [23] Weixin Li, Vijay Mahadevan, and Nuno Vasconcelos, "Anomaly detection and localization in crowded scenes," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 1, pp. 18–32, 2014.
- [24] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.