

Anomaly Detection Using Prediction Error with Spatio-Temporal Convolutional LSTM

Hanh T. M. Tran*, David Hogg

Abstract—In this paper, we propose a novel method for video anomaly detection motivated by an existing architecture for sequence-to-sequence prediction and reconstruction using a spatio-temporal convolutional Long Short-Term Memory (convLSTM). As in previous work on anomaly detection, anomalies arise as spatially localised failures in reconstruction or prediction. In experiments with five benchmark datasets, we show that using prediction gives superior performance to using reconstruction. We also compare performance with different length input/output sequences. Overall, our results using prediction are comparable with the state of the art on the benchmark datasets.

Index Terms—Convolutional LSTM; convolutional autoencoder; prediction error; reconstruction error; anomaly detection.



1. Introduction

AUTOMATIC detecting abnormal events in video has been widely studied in recent years due to its broad range of applications, including wide-area surveillance and health monitoring. This problem is different from event detection where the event is clearly defined, since an anomaly is by definition unknown in advance and may arise from unfamiliar activities or activities in unfamiliar contexts.

The standard approach to anomaly detection has been to learn spatio-temporal models of normal activity using hand-crafted features [1]–[6] or deep feature representations [7], [8]. An abnormality is detected when spatio-temporal patterns are observed that do not conform to the model of normality. Many different low-level features using dense optical flow (e.g., histograms [5], MHOF [3]) and other patterns of spatio-temporal gradient [4] have been used in the past. A model of normality is learned using these features extracted from training data and then used to determine numerical abnormality scores in test data. The model may be of several different kinds, including probabilistic models (e.g, mixture of probabilistic PCA [1], mixture of dynamic texture [2]), domain based (e.g, one-class SVM [6]), sparse coding [3] and Sparse Combination Learning (SCL) [4]. All of these methods have been used for anomaly detection and localization within the image frame.

Recently, deep learning architectures have been successfully applied in many computer vision tasks including video anomaly detection. A key advantage of deep learning methods is that they can learn feature rep-

resentations directly from training data without prior definition. For example, this can be done in an unsupervised manner using auto-encoders [7]–[9]. A stacked de-noising autoencoder can be used to learn appearance and motion features for anomaly detection [8]. A Winner-take-all sparsity constraint combined within the autoencoder has been shown to produce flow-features that are more discriminative for a one class SVM [7] that is trained separately on the compressed representations learnt by the autoencoder.

End-to-end deep learning approaches have also been proposed for anomaly detection [10]–[12], [14], [15]. The convolutional autoencoder (convolutional AE) can be used to learn a model of normality from video, then reconstruction error [10]–[12] or prediction error [14], [15] provide a local measure for anomaly detection. A Generative Adversarial Network (GAN) can be employed to generate a normal distribution over some datasets [11], [14] by jointly optimising with a discriminator that competes to distinguish what is a real normal sample from what is a generated one. The motion dynamic can be learnt using a multi-channel approach, fusing appearance and motion information, and a cross-channel task, forcing the generator to transform raw-pixel data into motion information and vice versa [11] or using FlowNet combined with U-net for a single frame prediction [14]. The combination of a convolutional autoencoder and U-net has also been used to build two-stream network with a shared encoder in which one decoder is for a single frame reconstruction and one is for translating an image to optical flow [15]. Another Spatio-Temporal autoencoder has been proposed for video anomaly detection [16]. The results show that applying 3D convolution in the encoder and 3D deconvolution in the decoder helps to enhance the capability of extracting motion patterns over the temporal dimension.

A memory module is proposed into the AE to address these problems [17], [18]. The encoder inputs a

Hanh T. M. Tran is with the University of Danang - University of Science and Technology, Vietnam (e-mail: hanhtran@dut.udn.vn).

David Hogg is with the School of Computing, University of Leeds, United Kingdom.

*Corresponding author: Hanh T. M. Tran (e-mail: hanhtran@dut.udn.vn)
Manuscript received February 04, 2022; revised April 24, 2022; accepted May 17, 2022.

Digital Object Identifier 10.31130/ud-jst.2022.289E

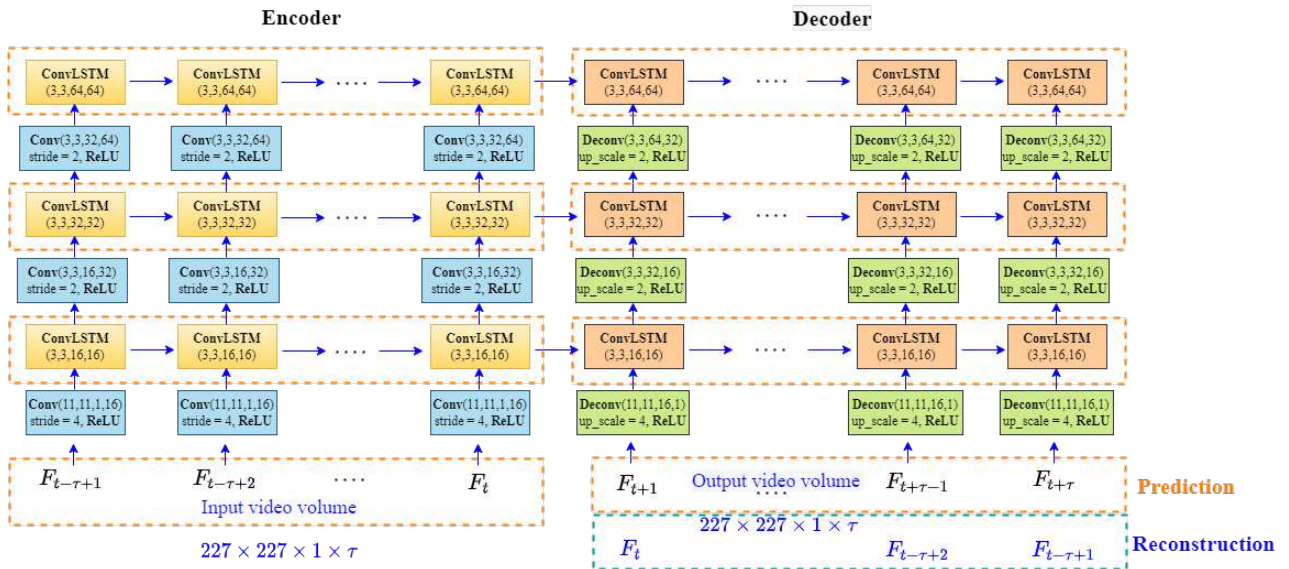


Fig. 1: The encoding-decoding structure used for future prediction or reconstruction with video volumes of τ frames.

normal video frame and extracts feature maps. The encoding features are then used to retrieve prototypical normal patterns in the memory items and to update the memory. Then the feature maps and aggregated memory items are fed into the decoder for reconstructing the input video frame or predicting the next frame. Using cosine similarity and the softmax function for matching probability between incoming encoding features and memory items, the global memory can be read and written to. Since normal patterns in training and testing sets may be different, the memory items are updated during training and testing time, with the use of a predefined threshold to prevent updating on anomaly patterns [18]. However, it is impossible to find an optimal threshold to distinguish between normal and abnormal patterns under various scenarios. Meta-learning methodology is introduced into a Dynamic Prototype Unit (DPU) to learn prototypes for encoding normal dynamics and to enable the fast adaption capacity to a new scene with only a few training frames [19]. As in previous work [18], the DPU inputs the encoding feature maps, which are outputs of the encoder part of U-net, to generate a pool of dynamic prototypes. However it is trained in a fully differential attention manner in which attention mapping functions are implemented as fully connected layers and updated using gradient descent style. After training the AE backbone using only frame prediction loss, the DPU module is trained with the meta-training phase using frame pairs sampled from videos of diverse scenes. In the testing phase, in order to adapt the model to a new scene, the first few frames of the sequence in this scene are used to construct K-shot input-output frame pairs. The results show that the DPU is more memory-efficient than the memory module in previous work [17], [18].

Another approach to learning regular spatio-temporal patterns is to use a convolutional LSTM [13],

[20]. The motivation is that reconstruction over a longer duration using the memory of the LSTM should capture more complex flow patterns. The convolutional network is used to encode each frame, then feeding these encoding tensors to Convolutional LSTMs to memorize the change of the appearance which corresponds to motion information [20]. Two Deconvolutional Networks (DeconvNet) are used, one for reconstructing past frames and to identify whether an anomaly occurs; and one for reconstructing the current frame. Thus the reconstruction error is an indicator of the change in appearance or motion. The temporal unit in [13], [20] is applied on the final spatial stage, which encodes high level representations. Interleaving RNNs between spatial convolution layers has recently been shown to improve performance on precipitation now-casting [21]. The model can learn temporal information on hierarchical spatial representations from low-level to high-level. In our work, we adopt the same architecture, except that we remain with convolutional LSTMs instead of the complex trajGRU RNN [21]. Our results show a comparable level of performance to the state of the art on benchmark datasets with fewer model parameters than state of the art models. Moreover, using prediction gives better performance than reconstruction. Finally, performance varies as expected with different prediction windows.

2. Architecture

Figure 1 illustrates the encoding-decoding structure for future prediction or reconstruction, motivated by earlier work [21] and adapted for anomaly detection. At each time step, the network takes a video volume of τ video frames $F_{t-\tau+1}, \dots, F_t$, and generates an output volume of the same size, predicting the future $F_{t+1}, \dots, F_{t+\tau}$ or reconstructing the input in reverse order $F_t, \dots, F_{t-\tau+1}$.

2.1. Encoding-decoding model

The structure consists of two networks, an encoding network and a decoding network (Fig. 1). The encoder contains three convolutional layers, each followed by leaky ReLU with negative slope equal to 0.2 [22]. In order to do down-sampling, we use all three convolutional layers with stride. The strided convolution allows the network to learn its own spatial down-sampling. Similarly, three deconvolution layers are used in the decoder to learn its own spatial up-sampling. The goal of temporal encoding is to capture and compress changes due to motion in the input sequence into encoding hidden states that allow the decoder to reconstruct the input or predict the future.

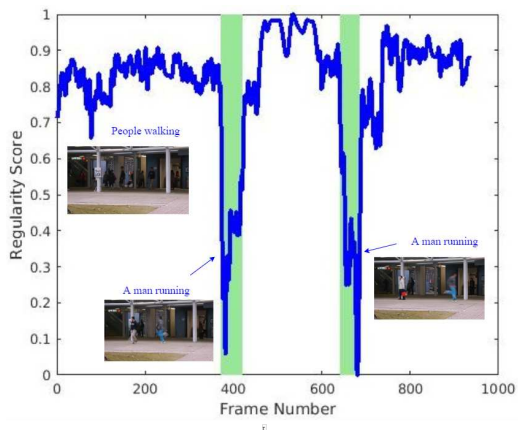


Fig. 2: The regularity score of video sequence of the CUHK Avenue dataset [4]. The score decreases when an anomaly (a running man) appears on the scene.

Spatio-temporal LSTM cells [23] are employed as a temporal encoder/decoder. At each time t , the convolutional LSTM (convLSTM) module receives as input a new video frame after projection in the spatial feature space. This is used together with the memory content and output of the previous step $t - 1$ to compute new memory activations. Interleaving multiple convLSTMs between convolutional layers helps the model learn spatio-temporal dynamic information at different levels. The high level states capture global spatial-temporal representations while the lower level states retain the detail of local spatio-temporal representations. After the last frame is read, the decoding LSTMs take corresponding states from the encoder as their initial states and output an estimate for the target sequence (Fig. 1). The low-level states are combined with the up-sampling outputs as the initial states and inputs of decoding LSTMs helps to aggregate low-level information to the up-scaling data stream. Therefore, the output contains details on both background and object (Fig. 3).

2.2. Input data layer

The input to the model is a video volume consisting of τ consecutive frames. Each frame is extracted from raw video, converted to a gray-scale image and resized to 227×227 . The pixel values are scaled to the

range $[0, 1]$. We stack τ frames in the 4th dimension into video volumes and use them as the input of size $227 \times 227 \times 1 \times \tau$ to the encoder. Following [10], we generate more video sequences by concatenating frames with skipping strides of 1, 2 and 3, thereby simulating faster motion patterns. Although speed can be important in anomaly detection, we still carry out this augmentation to minimise over-fitting and to have a fair comparison with [10], [13]. Unlike [10], we do not stack precomputed optical flow into our input volume, in the expectation that the network can learn the necessary patterns of motion.

3. Training

The weights \mathbf{W}_l and biases b_l of each layer l are learned by minimizing the regularized least squares error:

$$\frac{1}{2N\tau} \sum_{n=1}^N \|\theta_n - \hat{\theta}_n\|_2^2 + \frac{\lambda}{2} \sum_l \|\mathbf{W}_l\|_2^2 \quad (1)$$

where $\hat{\theta}_n$ is the predicted frame sequence (or the reconstructed frame sequence) from the model and θ_n is the target sequence. The first term is the prediction error (or the reconstruction error) and the second term is to regularize the weights. λ is a hyper-parameter used to balance the importance of two terms.

The weights in each convolutional layer are initialized from a zero-mean Gaussian distribution with standard deviation calculated from the number of input channels and the spatial filter size of the layer [24]. This is a robust initialization method that particularly considers the rectifier nonlinearities. We initialize the weights for convLSTM using a zero-mean Gaussian distribution with a fixed standard deviation of 0.01. The biases for all layers are initialized to zero. The input-to-hidden and hidden-to-hidden convolutional filters in the convLSTM cell are the same size.

3.1. Anomalous event detection

The Adam [25] method is used to optimize the error in Eq. 1 with batch size $N = 4$, momentum of 0.9 and 0.999, and weight decay $\lambda = 5 \times 10^{-4}$ [26]. We train a network separately on each dataset so that the model learns the specific normal patterns. An event may be normal in one dataset but abnormal in another. For example, people going towards the turnstile to enter the subway station is normal in the Subway Entrance dataset but abnormal in the Subway Exit dataset. We start training the model with a learning rate of 10^{-4} . After 80 epochs, we stop training and use the model for anomaly detection.

4. Regularity score for anomaly detection

Once the model is trained, the prediction error between each output frame \hat{F}_i and the target frame F_i in the video sequence is computed, then errors of all τ

TABLE 1: Performance comparison with the state of the art.

Method	AUC/EER (%)				
	UCSDPed1	UCSDPed2	CUHK Avenue	Subway Entrance	Subway Exit
Conv-WTA [7]	91.6/ <u>14.8</u>	95/ 9.5	81/26.5	-	-
AMDN [8]	<u>92.1</u> /16	90.8/17.1	-	-	-
GAN [11]	-	93.5/15.6	-	-	-
Conv-AE [10]	81/27.9	90/21.7	70.2/25.1	94.3 /26.0	80.7/9.9
ST-AE [13]	89.9/ 12.5	87.4/ <u>12.0</u>	80.3/ 20.7	84.7/ <u>23.7</u>	<u>94.0</u> / <u>9.5</u>
Past-Current-LSTM [20]	75.5/-	88.1/-	77/-	<u>93.3</u> / <u>-</u>	87.7/ <u>-</u>
STAE-3D [16]	92.3 /15.3	91.2/16.7	77.1/33.8	-	-
FlowNet-Unet-GAN [14]	83.1/-	95.4/-	85.1/-	-	-
Two-streams AE [15]	-	94.1/-	83.3/-	-	-
MemAE [17]	-	96.2/-	86.9/-	-	-
LMN [18]	-	97 /-	88.5/-	-	-
MPD* [19]	83.2/-	95.1/-	84.0/-	-	-
MPD [19]	85.1/-	<u>96.9</u> / <u>-</u>	89.5 / <u>-</u>	-	-
Ours (prediction)	80.8/25.1	92.3/14.4	84.8/ <u>22.4</u>	90.2/ 15.9	95/8

frames are summed up to form the prediction error for a volume as follows:

$$e(t) = \sum_{i=t+1}^{i=t+\tau} \|\hat{F}_i - F_i\|_2 \quad (2)$$

The prediction error then is normalized to compute a regularity score $s(t)$ of a testing volume as follows [10]:

$$s(t) = 1 - \frac{e(t) - \min_{t'} e(t')}{\max_{t'} e(t')} \quad (3)$$

where $\min_{t'} e(t')$ and $\max_{t'} e(t')$ are calculated over the prediction errors of all volumes in the same test video. If the regularity score $s(t)$ is less than a threshold, the corresponding test volume is abnormal.

We also use the same architecture for reconstruction in our experiments. Instead of using the next τ frames as the target sequence, we use the input sequence in reverse order as the target. Replacing the target sequence in Eq. 2, we obtain the reconstruction error and use it for anomaly detection with the reconstruction model.

5. Experiments

Our method is evaluated both quantitatively and qualitatively. We modify and use Caffe [27] for all our experiments. Code and trained models are available at https://github.com/t2mhanh/convLSTM_Prediction_AnomalyDetection.

5.1. Datasets

Our models are trained on five of the most commonly used datasets for anomaly detection: UCSD (UCSDPed1 and UCSDPed2) [2], CUHK Avenue [4], Subway (Entrance and Exit) [5]. The UCSD and CUHK datasets have separate training videos which contain mostly normal events. The first 12 minutes of Subway Entrance and the first 5 minutes of Subway Exit are used for training.

5.2. Anomalous event detection

Two performance metrics are employed for evaluation and comparison with state of the art results: Equal Error Rate (EER) and Area Under the ROC Curve (AUC). The regularity score of each volume determines whether it is normal or abnormal. We follow the intuition that testing video volumes containing normal events generate high regularity scores (Eq. 3) since they are similar to training data. A testing video sequence containing an anomaly gives a lower score. Setting different thresholds on the regularity score, volumes are classified into those that contain an anomaly and those that do not. These predictions are compared with ground-truth to give the equal error rate (EER) and area under the curve (AUC) of the resulting ROC curve (TPR versus FPR) generated by varying an acceptance threshold. Good performance has a low EER and high AUC.

Table 1 shows that the model trained for prediction performs comparably to state of the art results. Performance on UCSDPed1 is relatively poor, whilst for CUHK Avenue, the AUC is better than most methods, except FlowNet-Unet-GAN [14], MemAE [17], LMN [18], MPD [19]. However, MemAE [17], LMN [18] and MPD [19] have more parameters than our models which is shown in Table 3.

TABLE 2: Comparison of AUC/EER with different models. τ is the number of frames in an input sequence and a target sequence.

Method	AUC/EER (%)		
	UCSDPed1	UCSDPed2	CUHK Avenue
Reconstruction	75.6/28.9	87.5/17.1	81.4/26.1
Prediction	$\tau = 2$	78.3/27.1	86.1/21.1
	$\tau = 5$	80.8/25.1	92.3/14.4
	$\tau = 8$	79/26.5	89.6/18.5
			83.2/23.2

Table 2 shows the results when different models are used. In the table, ‘‘Reconstruction’’ is for a model trained for reconstructing a sequence of 5 frames and ‘‘Prediction’’ is for models trained to predict τ frames.

The model trained for future prediction gives better results than the reconstruction model. This may be because prediction will always try to draw back to normality, whereas reconstruction works from pre-sight of an anomalous sequence. The quality comparison between reconstruction and prediction is shown in Fig. 3.



Reconstruction (error $e(t) = 21.57$) - UCSDPed2 - biker



Prediction (error $e(t) = 41.03$) - UCSDPed2 - biker



Reconstruction (error $e(t) = 30.28$) - UCSDPed1 - car



Prediction (error $e(t) = 66.25$) - UCSDPed1 - car



Reconstruction (error $e(t) = 11$) - CUHK Avenue - running



Prediction (error $e(t) = 59.81$) - CUHK Avenue - running

Fig. 3: Prediction and reconstruction of third frame out of 5 (middle), compared to target frame (left); accumulated per-pixel error over 5 frames as blue-green-red colour map (right). Ground truth anomalies shown as rectangles. Taken from UCSDPed2, UCSDPed1 and CUHK Avenue. Best viewed in color

The number of model parameters for the method against different end-to-end trainable models in the

state of the art are compared in Table 3. We achieve 75 fps for anomaly detection with a GeForce GTX TITAN X, faster than other state of the art methods with the same setting [18].

TABLE 3: Comparison of model complexity and testing speed.

Methods	Parameters (M)	FPS
Conv-AE [10]	8.4	—
ST-AE [13]	1.1	—
STAE-3D [16]	0.5	—
MemAE [17]	6.2	45
LMN [18]	15.0	67
Ours	0.85	75

As can be seen in Fig. 3, the future prediction of a biker becomes worse than the prediction of a pedestrian. The model is trained mostly on video sequences containing pedestrians, the prediction of the biker looks similar to the pedestrian. Here the prediction error is significantly larger than the reconstruction error.

6. Conclusion

We have adapted a state of the art predictive encoder-decoder deep network to detect abnormal events in video. We evaluate detection performance using both sequence prediction and reconstruction, and show that prediction gives superior performance on anomaly detection. For the prediction model, we obtain competitive performance to state of the art methods on five standard datasets. Finally, we evaluate performance across different prediction windows, encompassing varying levels of motion complexity. Our future work includes investigating the fusion of gray-scale images and optical flow on input.

Acknowledgment

This work was supported by The Murata Science Foundation and The University of Danang - University of Science and Technology, code number of Project: T2021-02-05MSF.

References

- [1] Kim J, Grauman K. Observe locally, "infer globally: a space-time MRF for detecting abnormal activities with incremental updates", *2009 IEEE conference on computer vision and pattern recognition*, 2009 Jun 20 (pp. 2921-2928).
- [2] Mahadevan, Vijay, Weixin Li, Viral Bhalodia, and Nuno Vasconcelos, "Anomaly detection in crowded scenes", *In 2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 1975-1981, IEEE, 2010.
- [3] Cong, Yang, Junsong Yuan, and Ji Liu, "Sparse reconstruction cost for abnormal event detection", *In CVPR*, 2011, pp. 3449-3456. IEEE, 2011.
- [4] Lu, Cewu, Jianping Shi, and Jiaya Jia, "Abnormal event detection at 150 fps in matlab," *In Proceedings of the IEEE international conference on computer vision*, pp. 2720-2727, 2013.
- [5] Adam, Amit, Ehud Rivlin, Ilan Shimshoni, and Daviv Reinitz, "Robust real-time unusual event detection using multiple fixed-location monitors," *IEEE transactions on pattern analysis and machine intelligence* 30, no. 3 (2008): 555-560.

- [6] Wang, Siqi, En Zhu, Jianping Yin, and Fatih Porikli, "Anomaly detection in crowded scenes by SL-HOF descriptor and foreground classification," *In 2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 3398-3403, IEEE, 2016.
- [7] Tran, Hanh TM, and David Hogg, "Anomaly detection using a convolutional winner-take-all autoencoder," *In Proceedings of the British Machine Vision Conference 2017*, British Machine Vision Association, 2017.
- [8] Xu, Dan, Elisa Ricci, Yan Yan, Jingkuan Song, and Nicu Sebe, "Learning deep representations of appearance and motion for anomalous event detection," *arXiv preprint arXiv:1510.01553* (2015).
- [9] Tran, Thi Minh Hanh. "Anomaly Detection in Video." *PhD diss.*, University of Leeds, 2018.
- [10] Hasan, Mahmudul, 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*, pp. 733-742, 2016.
- [11] Ravanbakhsh, Mahdyar, Enver Sangineto, Moin Nabi, and Nicu Sebe, "Training adversarial discriminators for cross-channel abnormal event detection in crowds," *In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1896-1904, IEEE, 2019.
- [12] Zhao, Bin, Li Fei-Fei, and Eric P. Xing, "Online detection of unusual events in videos via dynamic sparse coding," *In CVPR 2011*, pp. 3313-3320. IEEE, 2011.
- [13] Chong, Yong Shean, and Yong Haur Tay, "Abnormal event detection in videos using spatiotemporal autoencoder," *In International symposium on neural networks*, pp. 189-196. Springer, Cham, 2017.
- [14] Liu, Wen, 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*, pp. 6536-6545, 2018.
- [15] Nguyen, Trong-Nguyen, and Jean Meunier, "Anomaly detection in video sequence with appearance-motion correspondence," *In Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1273-1283, 2019.
- [16] Zhao, Yiru, 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*, pp. 1933-1941, 2017.
- [17] Gong, Dong, 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/CVF International Conference on Computer Vision*, pp. 1705-1714, 2019.
- [18] Park, Hyunjong, 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*, pp. 14372-14381, 2020.
- [19] Lv, Hui, Chen Chen, Zhen Cui, Chunyan Xu, Yong Li, and Jian Yang, "Learning normal dynamics in videos with meta prototype network," *In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15425-15434, 2021.
- [20] Luo, Weixin, Wen Liu, and Shenghua Gao, "Remembering history with convolutional lstm for anomaly detection," *In 2017 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 439-444. IEEE, 2017.
- [21] Shi, Xingjian, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo, "Deep learning for precipitation nowcasting: A benchmark and a new model," *Advances in neural information processing systems*, 30 (2017).
- [22] Maas, Andrew L., Awni Y. Hannun, and Andrew Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," *In Proc. icml*, vol. 30, no. 1, p. 3, 2013.
- [23] Shi, Xingjian, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," *Advances in neural information processing systems*, 28 (2015).
- [24] He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," *In Proceedings of the IEEE international conference on computer vision*, pp. 1026-1034, 2015.
- [25] Kingma, Diederik P., and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980* (2014).
- [26] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, 25 (2012).
- [27] Jia, Yangqing, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell, "Caffe: Convolutional architecture for fast feature embedding," *In Proceedings of the 22nd ACM international conference on Multimedia*, pp. 675-678, 2014.



Hanh T. M. Tran is currently a Lecturer with the Department of Electronics and Telecommunications, the University of Danang - University of Science and Technology, Vietnam, where she joined since 2009. She received the B.Eng. and M.Eng. degrees in Electronics and Telecommunications from the University of Danang - University of Science and Technology in 2008 and 2011, respectively. She obtained the Ph.D. degree from the University of Leeds,

United Kingdom, in 2018. She was a Visiting Researcher with the Arizona State University, Arizona, USA, in 2012. Her main research interests include image/video processing, machine learning, deep learning, anomaly detection, object detection and recognition.



David Hogg is Professor of Artificial Intelligence at the University of Leeds. His research is on artificial intelligence and particularly in computer vision. He has been Pro-Vice-Chancellor for Research and Innovation at the University of Leeds, visiting professor at the MIT Media Lab, chair of the EPSRC ICT Strategic Advisory Team, and chair of the Academic Advisory Group of the Worldwide Universities Network. He is a Fellow of the European Association for

Artificial Intelligence (EurAI), a Distinguished Fellow of the British Machine Vision Association, and a Fellow of the International Association for Pattern Recognition. He is Director of the UKRI Centre for Doctoral Training in Artificial Intelligence for Medical Diagnosis and Care at the University of Leeds