

Good Practices and A Strong Baseline for Traffic Anomaly Detection

Yuxiang Zhao^{1,2*} Wenhao Wu^{2*} Yue He² Yingying Li² Xiao Tan² Shifeng Chen^{1†}
¹ Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China
² Department of Computer Vision Technology (VIS), Baidu Inc., China

Abstract

The detection of traffic anomalies is a critical component of the intelligent city transportation management system. Previous works have proposed a variety of notable insights and taken a step forward in this field, however, dealing with the complex traffic environment remains a challenge. Moreover, the lack of high-quality data and the complexity of the traffic scene, motivate us to study this problem from a hand-crafted perspective. In this paper, we propose a straightforward and efficient framework that includes pre-processing, a dynamic track module, and post-processing. With video stabilization, background modeling, and vehicle detection, the pro-processing phase aims to generate candidate anomalies. The dynamic tracking module seeks and locates the start time of anomalies by utilizing vehicle motion patterns and spatiotemporal status. Finally, we use the post-processing to fine-tune the temporal boundary of anomalies. Not surprisingly, our proposed framework was ranked 1st in the NVIDIA AI CITY 2021 leaderboard for traffic anomaly detection. Codes will be available.

1. Introduction

With the ongoing expansion of urban traffic system construction, large-scale traffic management systems and accident warning systems have become an essential component of urban infrastructure development. Traffic anomaly detection is a critical task in video understanding, which has attracted widespread attention from both academia and industry. As shown in Figure 1, due to the complexity of the traffic scene and the diversity of camera views, there remain great challenges in traffic anomaly detection.

Generally, anomalies rarely occur as compared to normal activities. Hence, developing intelligent computer vision methods for automatic video anomaly detection is a pressing need. The goal of a practical anomaly detection system is to timely signal an activity that deviates normal patterns



Figure 1. Examples of traffic anomalies. Red rectangles highlight the abnormal instances.

and identifies the time window of the occurring anomaly. Therefore, anomaly detection can be considered as coarse-level video understanding, which filters out anomalies from normal patterns.

Recently, deep neural networks have achieved great success in image recognition [32, 12, 6], video recognition [4, 11], object detection [28, 27, 21], etc. Deep ConvNets come with excellent modeling capacity and are capable of learning discriminative representations from visual data in large-scale supervised datasets (e.g., ImageNet [9], Kinetics-400 [14], Places [43], MS COCO [19]). However, unlike these above tasks, the application of end-to-end ConvNets to traffic anomaly detection is impeded by the following major obstacle: Training deep ConvNets framework in an end-to-end manner usually requires a large volume of training samples to achieve optimal performance. However, the training set of Track4 only consists of 100 videos, which remain limited in both size and diversity.

These challenges motivate us to study traffic anomaly detection from a hand-crafted perspective. In this paper, we develop a customized architecture with human prior knowledge, which provides a conceptually simple, and robust framework for traffic anomaly detection. As shown in Figure 2, our framework consists of pre-processing, dynamic tracking module, and post-processing.

Concretely, the pipeline of pre-processing including vehicle detection, video stabilization, background modeling,

*The first two authors contributed equally to this work. This work was done when Yuxiang Zhao was a research intern at Baidu.

†Corresponding author.

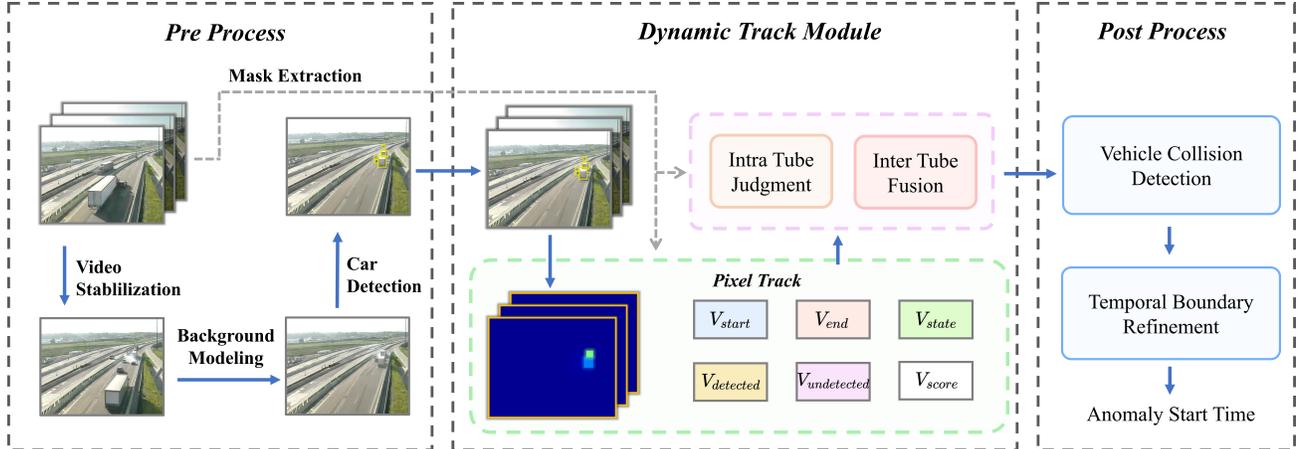


Figure 2. The illustration of hierarchical strategy framework pipeline. This framework includes pre-processing, dynamic tracking module and post-processing.

The pre-processing phase aims to generate the candidate anomalies with video stabilization, background modeling, and vehicle detection. The dynamic tracking module utilizes the motion patterns and spatio-temporal status of vehicles to seek and locate the start time of the anomalies. Finally, we employ the post process to further refine the temporal boundary of anomalies. The experimental results show that our proposed framework is effective and robust in the Track4 test set of the NVIDIA AI CITY 2021 Challenge, which ranks first place in this competition with a 95.24% F1-score and 5.3080 root mean square error, and the final score is 93.55%.

The main contributions are summarized as follows:

- We present a hierarchical architecture with human prior knowledge, which provides a conceptually simple framework for traffic anomaly detection.
- We investigate a series of good practices for temporal localization and anomaly detection (*e.g.*, video stabilization, vehicle collision detection, temporal boundary refinement, *etc.*).
- Results demonstrate that our method outperforms other challengers on the Track 4 test set of the NVIDIA AI CITY 2021 Challenge.

2. Related Work

Anomaly detection is the process of detecting rare or unusual patterns that deviate from the normal behavior, which is called “Outliers” or “Anomalies”. Many studies have explored this field with a variety of methods for a long time. These approaches can be divided into two parts by the types of models, traditional methods and deep learning based methods. Some works use traditional machine learning models like Gaussian mixture models [16], Regression

models [8], histogram-based [41], Dirichlet process mixture models (DPMM) [25], Bayesian network-based models [2] to seek anomalies. With the great development of computer vision leveraged on deep learning, autoencoder based networks and well-designed loss becomes the key point of anomaly forecasting task [23, 15, 26, 36, 37, 5, 34, 40, 42, 20, 24, 7, 35]. Sultani *et al.* [33] construct a new large-scale dataset, called UCF-Crime, which contains real-world anomaly videos.

Traffic anomaly detection is a more fine-grained anomaly detection, which includes multiple kinds of violations of regulations such as driving in the wrong direction, illegal parking, *etc.*

In the past years of NVIDIA AI CITY Challenges, unsupervised traffic anomaly detection methods have become the mainstream in real traffic accident scenarios and promote the development of intelligent transportation. Wei *et al.* [38] first remove the moving vehicles using MOG2 while keeping the stopped vehicles as part of the background, then perform multi-scale detection and classification to detect anomalies. Shine *et al.* [30] show an unsupervised method to tackle this problem including a background extraction stage, an anomaly detection that identifies the stalled vehicles in the background, and a final anomaly confirmation module. Bai *et al.* [1] present a novel spatial-temporal information matrix, which transforms the analysis of a strip trajectory into an analysis of the spatial position. Shine *et al.* [31] conduct background subtraction using GMM and utilize YOLO detector to select anomaly candidates, then detect anomalies by transfer learning without using training data [10]. Li *et al.* [17] present a multi-granularity tracking approach, which combines a box-level branch and a pixel-level branch to analyze the candidate abnormal vehicles at different granularity levels. They ranked first in NVIDIA AI CITY 2020.

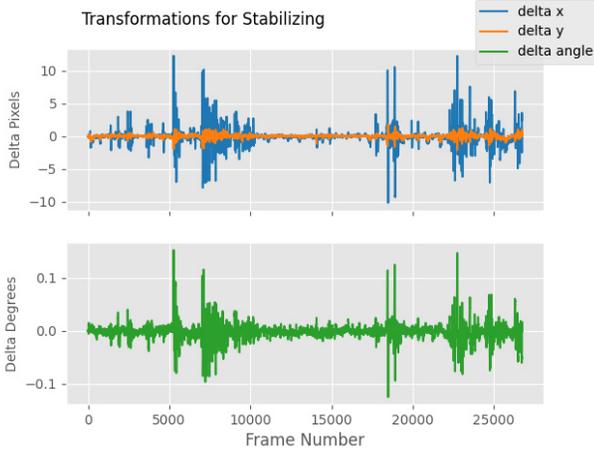


Figure 3. Visualization of video stabilization.



(a) Forward background modeling.



(b) Backward background modeling

Figure 4. Examples of background modeling.

In this paper, we propose a simple and effective framework for traffic anomaly detection in track 4 of the NVIDIA AI CITY 2021 Challenge. We obtain a 0.9524 F1-score with a detection time error of just 5.3080 seconds and our results rank first place on the track 4 test set among all the participant teams.

3. Method

In this section, we first present the pipeline of pre-processing, which includes video stabilization, background modeling, vehicle detection, and mask generation. Then we illustrate our dynamic track module to locate the start time of the anomaly. Finally, we depict our post-processing pipeline which aims to further refine the temporal boundary of anomalies.

3.1. Pre-Processing

Video Stabilization. In a real scenario, the traffic camera under adverse conditions usually results in non-precise motion and occurrence of shaking, which may compromise the stability of the obtained videos, and directly influence the performance of the modules (*e.g.*, detection, background

modeling) in our framework.

To overcome such a problem, we perform digital video stabilization (DVS) to correct camera motion oscillations that occur in the acquisition process, through software techniques, without the use of specific hardware, to enhance visual quality and improve final applications, such as vehicle detection and tracking of vehicles. The whole pipeline of DVS could be roughly divided into two stages: The first stage is to estimate the camera movements. Then the second stage is to correct and smooth the camera motion.

Considering the moving vehicles may affect the quality of motion detection, we apply a robust video stabilization method to compensate the undesired movements of the camera. To be specific, we utilize the combination of feature point matching based on GFTT [29] and calculate sparse optical flow to generate frame-to-frame transformations. A hybrid filter was used in the video stabilization process to smooth the trajectory of transformations. Figure 3 shows the variety of δ_x , δ_y and δ_{angle} in transformation matrix. By using these parameters, we can filter the shacking heavily videos that exceed accumulated threshold δ_t and average threshold δ_{avg} to imply video stabilization.

Background Modeling. It is common knowledge that vehicles involved in a traffic accident or anomaly will often come to a complete stop naturally. As a result, the traffic anomaly detection task is transformed into static vehicle detection. This allows us to effectively detect anomalies of the stopped vehicle on the lane, even when the videos are of poor quality.

Various works in the field of object tracking have attempted to distinguish the foreground and background. Due to the poor video sources, these methods require a robust and adaptive background representation. To dynamically model backgrounds, we use the background modeling approach based on the mixture of Gaussians (MOG) [44]. We perform an ablation analysis on the duration of background modeling to make a trade-off between effectiveness and efficiency. We compare the difference between the forward and backward background modeling results in Figure 4, which shows that backward background modeling results will make stopped vehicles clearer, and we use the forward background modeling approach as an auxiliary method to get a more accurate start time of the anomaly.

Vehicle Detection. For object detection, there are two mainstream methods as follows: one-stage detectors and two-stage detectors. Two-stage detectors are more accurate and require a higher computational time. On the contrary, one-stage detectors benefited from straightforward architectures, which are faster and simpler but might potentially drag down the performance. To obtain robust detection results, in this paper, we build a two-stream vehicle



Figure 5. Visualization of vehicle detection. Left is the results of Faster RCNN and right is results of Cascade RCNN. Red circles highlight the false detection.

detection module that consists of two two-stage detectors: Faster R-CNN [28] and Cascade R-CNN [3]. Concretely, we employ the Faster R-CNN with SENet-152 [13] as our main detector for vehicle detection. The Cascade R-CNN with CBResnet-200 [22], is used as an auxiliary detector to reduce the fail of vehicle detection. Following the common practice, we apply the Feature Pyramid Network (FPN) [18] to build high-level semantic feature maps at all scales. Figure 5 shows the results of our detector.

Mask Generation. In general, traffic anomalies usually happen on vehicles driving the main road. Hence, we need to filter the static vehicles on the side roads and parking lots. A practicable solution is to segment out hypothetical abnormal mask regions automatically. To do so, image segmentation is an intuitive approach to distinguish the hypothetical abnormal area. However, due to the complexity of the road scene and limited training videos, it is hard to learn a robust segmentation model.

Inspired by the motion-based mask extraction method in [17]. Then, we suggest an enhanced trajectory-based mask solution to produce finer masks. To obtain the vehicle’s trajectory, we use the Track1 dataset to train the multi-object tracking algorithm DeepSORT [39]. We cluster these trajectories into the primary and secondary parts by calculating the angle of moving direction, which is different from the previous process. The bounding box of the detected vehicle will be used to expand the trajectory-based mask if the absolute value of $\arctan \theta$ is in the primary part. The camera orientation is taken into account in this way, resulting in a more reasonable traffic road mask. Finally, we combine the two masks above to get the final result. These two masks can be used in conjunction, and Figure 6 demonstrates some results of the abnormal mask.

3.2. Dynamic Tracking Module

Pixel Tracking. Inspired by [17], we adopt pixel-level information to filter out suspicious events as much as pos-

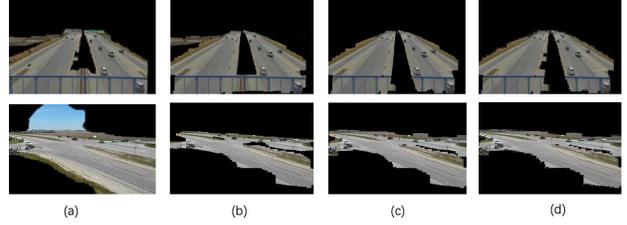


Figure 6. Examples of the abnormal mask. From the left to the right: motion-based mask, original trajectory-based mask, our enhanced trajectory-based mask, and final fused mask. Our trajectory-based mask effectively reduces false auxiliary road recall, and the abnormal area adjacent to the main road is reserved.



Figure 7. Visualization of Pixel track.

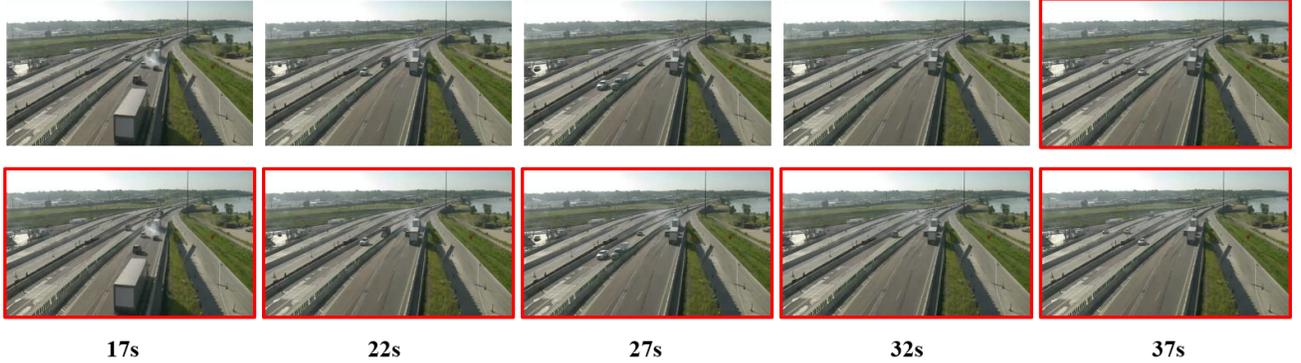
sible. Six spatial-temporal matrices, including $V_{undetected}$, $V_{detected}$, V_{score} , V_{state} , V_{start} and V_{end} , are combined to record anomaly car detection results iteratively. As shown in Figure 2, the pixel region and time stamp of the detection will be continuously accumulated on the six matrixes when a car appears on the background frame. If the peak temporal span exceeded the time threshold, the suspicious results will be thrown out. In addition, we jointly develop two mechanisms to optimize the anomaly detection result as follows: 1) When an abnormal object is observed, we will utilize the IOU algorithm to compare the current object position with the detection result in the next frame. When the intersection of the two regions is greater than 0.5, we update the start timestamp of an anomaly. 2) We develop a backtrack approach based on the spatial-temporal correlation similarity since certain cars do not stop immediately in an accident. When IoU is less than the threshold, the PSNR and color histogram features are extracted for the non-overlapped bounding boxes to calculate box similarity.

Then, in the following phase, we begin to process the candidate results. First, we filter the true anomaly start time using intra tube judgment. Second, we use intra tube fusion to join tubes from the same vehicle together.

Intra tube judgment. The candidate tube may contain several vehicles, making it difficult for the model to deter-



(a) Video 33 on test set



(b) Video 73 on test set

Figure 8. Example results of vehicle collision detection. For video 33, we backtrack the start time from the stopped time (612 seconds) to the crash time (603 seconds). For video 73, we backtrack the start time from the delayed stopped time (37 seconds) to the crash time (17 seconds).

mine the true anomaly start time. So we conduct intra tube judgment to remove the parts that do not belong to current anomalous vehical. Here we denote the i -th tube as U^i , and the j -th frame of it as U_j^i . Reg^i are regions of tube U^i in chronological order. The start time of tube U^i is denoted by U_{Start}^i , the end time is denoted by U_{End}^i . The regions in Reg^i across the timeline are composited by the set $S(Reg^i)$. We first calculate the similarity between every element in $S(Reg^i)$ and the mean over the interval \bar{U}^i , and store it in the set $S(Sim_j^i)$. We use SSIM as the similarity measure which ranges from (0, 1).

Then we compute mean-subtracted value $S'(Sim_j^i)$ for every element in $S(Sim_j^i)$. We judge whether the region in j -th frame belongs to current anomalous car or not in this manner:

$$\begin{cases} S'(Sim_j^i) > Thresim, & False \\ LowerBound < S'(Sim_j^i) \leq Thresim, & True \end{cases}$$

That is to say, the $S'(Sim_j^i)$ which exceeds the threshold $Thresim$ demonstrates that Reg_j^i is not the bbox of the current anomalous car on the global time. This comparison is

performed in a chronological order until Sim_j^i is lower than $Thresim$ and j -th timestamp is recorded as a potential start of an anomaly. If a successive satisfied duration from j -th timestamp exceeds $\gamma(U_{End}^i - U_{Start}^i)$, then j -th timestamp is assigned as the starting time of the true anomaly. γ is a hyperparameter that controls the tolerance. This process should be continued until all the candidate tubes are judged.

Inter tube fusion. One vehical can appears in two tubes. In order to fuse candidate tubes which represent to the same vehicle, we perform following steps to fuse these tubes. Generally, one vehical appeared in two spatial locations in a picture are highly similar in apperence and bounding box scale. Based on this hypothesis, we compute the similarity between the mean of region sets of two tubes, $Sim(\bar{U}^i, \bar{U}^k)$. If the similarity results exceed a defined threshold, the initial start time and end time are $min(U_{Start}^i, U_{Start}^j)$ and $max(U_{End}^i, U_{End}^j)$ respectively.

3.3. Post Processing

Vehicle Collision Detection. Stalled vehicles and crashes are the most common anomalies in the training set. The anomaly start time for a stalled vehicle is when the vehicle comes to a complete stop. The methods described above are capable of accurately predicting the stalled vehicle’s abnormal start time. The start time for a single-vehicle crash or a multiple-vehicle crash is the moment the first crash happens rather than the time the vehicle comes to a complete stop. As a result, the above method’s prediction would be somewhat delayed. To that end, we devise a method for detecting vehicle collisions. First, we obtain the estimated time when the vehicle comes to a full stop. Then, along the temporal axis, we trace this vehicle backwards.

By analyzing the videos of crashes, we find that there will be obvious changes in the foreground at the moment of collision on the surrounding areas the wrecked car. Meanwhile, the background image of the surrounding areas will be different between the time of the crash and after the vehicle has stopped completely. Based on the above, we can judge the accurate collision time of vehicles. The results is shown in Figure 8.

Temporal Boundary Refinement. Forward background modeling will cause the anomaly to be displayed as after-images after a delay, causing a temporal prediction delay. To obtain a more accurate temporal localization, we use reverse background modeling and appearance similarity to refine the anomaly’s start and end times.

4. Experiments

NVIDIA AI CITY 2021 Challenge presents a dataset of 100 videos for training and 150 videos for testing. The dataset includes events that occurred naturally or abnormally in a variety of severe weather conditions, such as raining, snowing, or fogging. Each video is about 15 minutes long and has a resolution of 800×410 pixels. In this section, we first introduce the experimental setup, then we provide the implementation details. Finally, we describe the evaluation metrics and present the results of our method.

4.1. Implementation Details

Vehicle Detection. To train the Faster R-CNN and Cascade R-CNN, we use SGD with momentum 0.9 and weight decay $1e-4$. With an initial learning rate of 0.01 and a mini-batch of 8, the Faster R-CNN model is trained for 50K iterations. The learning rate is reduced by a factor of 10 at iteration 30K and 40K. Besides, the Cascade R-CNN model is trained with an initial learning rate of 0.005 and reduced by a factor of 10 at iteration 40K and 60K. To obtain more data for training, following the official instructions, we use data across different challenge tracks (*i.e.*, Track1 and Track3)

to pre-train our detectors, then we finetune the model with the training videos of Track4. The shorter side of the input images is resized to 800 pixels while keeping the aspect ratio. We use 5 scale anchors of $\{16^2, 32^2, 64^2, 128^2, 256^2\}$ and ground-truth boxes are associated with anchors, which have been assigned to pyramid levels. The NMS is used to filter the final vehicle detection result, which is made up of the union of two models with a threshold of 0.8. Moreover, we train our models with the PaddlePaddle deep learning framework ¹.

Video Stabilization. For distinguishing a shacking video, the accumulated threshold δ_t is set to 17200 for video stabilization, and the average threshold δ_{avg} is set to 0.645.

Road Mask. Following the protocol in [17], we use the same parameters in motion-based mask. In our enhanced trajectory-based mask, the minimum trajectory length n is set to 5, and the minimum distance d of the trajectory is set to 50. Further, threshold of the angle of moving direction to distinguish primary and secondary parts is set to 0.8.

Pixel Tracking. The thresholds for the normal-suspicious and suspicious/abnormal-normal state transitions are all set to three consecutive frames. The time thresholds for suspicious and coarse anomaly candidates have been set to 20s and 30s, respectively. The relaxed constraint satisfaction ratio T_{ratio}^r is 0.6 and the shortest traceback time T_{time} is 30s. The IoU thresholds T_{IoU} and T_{IoU}^r are 0.3 and 0.5, respectively; the PSNR thresholds T_{PSNR} and T_{PSNR}^r are 18 and 20, respectively; and the color histogram thresholds T_{Color} and T_{Color}^r are 0.88 and 0.9, respectively.

Intra Tube Judgment & Inter tube fusion. In Intra tube, the min and max threshold of ssim similarity $Thresh_{similarity}$ is set to 0.25 and 0.6 respectively. The duration time ratio γ is fixed to 0.3. In inter tube, the PSNR threshold between two tube is set to 18.

Vehicle Collision Detection. The threshold of surrounding area is set to 50, the threshold of foreground changes is set to 1000 and the similarity threshold of the background image is 0.9.

4.2. Evaluation Metrics

Evaluation for Track 4 will be based on model anomaly detection performance, measured by the F1-score, and detection time error, measured by RMSE. Specifically, the track 4 score will be computed as:

$$S4 = F1 \times (1 - \text{NRMSE}), \quad (1)$$

¹<https://github.com/paddlepaddle/paddle>

where F1 is the F1-score and NRMSE is the normalized root mean square error (RMSE). The S4 score ranges between 0 and 1, and higher scores are better.

The F1-score is computed as Eq. 2. For the purpose of computing the F1-score, a true-positive (TP) detection will be considered as the predicted anomaly within 10 seconds of the true anomaly (i.e., seconds before or after) that has the highest confidence score. Each predicted anomaly will only be a TP for one true anomaly. A false-positive (FP) is a predicted anomaly that is not a TP for some anomaly. Finally, a false-negative (FN) is a true anomaly that was not predicted.

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{\text{TP}}{\text{TP} + \frac{1}{2}(\text{FP} + \text{FN})}. \quad (2)$$

We compute the detection time error as the RMSE of the ground truth anomaly time and predicted anomaly time for all TP predictions. In order to eliminate jitter during submissions, normalization will be done using min-max normalization with a minimum value of 0 and a maximum value of 300, which represents a reasonable range of RMSE values for the task. NRMSE is the normalized root mean square error (RMSE) as follow:

$$\text{NRMSE} = \frac{\min(\sqrt{\frac{1}{\text{TP}} \sum_{i=1}^{\text{TP}} (t_i^p - t_i^{gt})^2}, 300)}{300}. \quad (3)$$

4.3. Experiments results

We evaluate our model on NVIDIA AI CITY 2021 Challenge test dataset. The F1-score is 0.9524, and the RMSE is 5.3080, as shown in Table 1. The results demonstrate the effectiveness and robustness of our method. The leaderboard for all teams is shown in Table 2, and we are in **first** place with a score of 0.9355.

Table 1. Our results on Track4 test set.

F1	RMSE	S4 Score
0.9524	5.3080	0.9355

5. Conclusion

In this paper, we propose a simple and effective framework for detecting traffic anomalies (e.g., stalled car, car crash). In Track 4 of the NVIDIA AI CITY 2021 Challenge, we obtain a 0.9524 F1-score with a detection time error of just 5.3080 seconds. Our results rank first place on the Track 4 test set among all the participant teams, which demonstrates that the superiority of our method.

Acknowledgment

This work is supported by Key-Area Research and Development Program of Guangdong Province

Table 2. The final rank and S4 score on NVIDIA AI CITY 2021 Track4.

Rank	Team ID	Team Name	Score
1	76	BaiduVIS&SIAT	0.9355
2	158	BD	0.922
3	92	WHU-IIP	0.9197
4	90	SIS Lab	0.8597
5	153	Titan Mizzou	0.5686
6	48	BUPT-MCPRL2	0.289
7	26	Attention Please!	0.2184
8	154	Alchera	0.1418
9	12	CET	0.1401

(2019B010155003), Shenzhen Science and Technology Innovation Commission (JCYJ20200109114835623), National Natural Science Foundation of China (U1713203), and the Scientific Instrument Developing project of the Chinese Academy of Sciences (YJKYYQ20190028).

References

- [1] Shuai Bai, Zhiqun He, Yu Lei, Wei Wu, Chengkai Zhu, Ming Sun, and Junjie Yan. Traffic anomaly detection via perspective map based on spatial-temporal information matrix. In *CVPR Workshops*, pages 117–124, 2019. 2
- [2] Calum G Blair and Neil M Robertson. Event-driven dynamic platform selection for power-aware real-time anomaly detection in video. In *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 3, pages 54–63. IEEE, 2014. 2
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018. 4
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1
- [5] Ming-Ching Chang, Yi Wei, Nenghui Song, and Siwei Lyu. Video analytics in smart transportation for the aic’18 challenge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 61–68, 2018. 2
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1
- [7] Zhengying Chen, Yonghong Tian, Wei Zeng, and Tiejun Huang. Detecting abnormal behaviors in surveillance videos based on fuzzy clustering and multiple auto-encoders. In *2015 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2015. 2
- [8] Kai-Wen Cheng, Yie-Tarn Chen, and Wen-Hsien Fang. Gaussian process regression-based video anomaly detection and localization with hierarchical feature representation.

- IEEE Transactions on Image Processing*, 24(12):5288–5301, 2015. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1
- [10] Keval Doshi and Yasin Yilmaz. Fast unsupervised anomaly detection in traffic videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 624–625, 2020. 2
- [11] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016. 1
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [13] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 4
- [14] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1
- [15] Xiaoli Li and Ze-min Cai. Anomaly detection techniques in surveillance videos. In *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 54–59. IEEE, 2016. 2
- [16] Yanshan Li, Weiming Liu, and Qinghua Huang. Traffic anomaly detection based on image descriptor in videos. *Multimedia tools and applications*, 75(5):2487–2505, 2016. 2
- [17] Yingying Li, Jie Wu, Xue Bai, Xipeng Yang, Xiao Tan, Guanbin Li, Shilei Wen, Hongwu Zhang, and Errui Ding. Multi-granularity tracking with modularized components for unsupervised vehicles anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 586–587, 2020. 2, 4, 6
- [18] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 4
- [19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1
- [20] Sophia WTT Liu, Henry YT Ngan, Michael K Ng, and Steven J Simske. Accumulated relative density outlier detection for large scale traffic data. *Electronic Imaging*, 2018(9):239–1, 2018. 2
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 1
- [22] Yudong Liu, Yongtao Wang, Siwei Wang, TingTing Liang, Qijie Zhao, Zhi Tang, and Haibin Ling. Cbnet: A novel composite backbone network architecture for object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11653–11660, 2020. 4
- [23] Jefferson Ryan Medel and Andreas Savakis. Anomaly detection in video using predictive convolutional long short-term memory networks. *arXiv preprint arXiv:1612.00390*, 2016. 2
- [24] Xuan Mo, Vishal Monga, Raja Bala, and Zhigang Fan. Adaptive sparse representations for video anomaly detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 24(4):631–645, 2013. 2
- [25] Henry YT Ngan, Nelson HC Yung, and Anthony GO Yeh. Outlier detection in traffic data based on the dirichlet process mixture model. *IET intelligent transport systems*, 9(7):773–781, 2015. 2
- [26] N Patil and Prabir Kumar Biswas. Global abnormal events detection in surveillance video—a hierarchical approach. In *2016 Sixth International Symposium on Embedded Computing and System Design (ISED)*, pages 217–222. IEEE, 2016. 2
- [27] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015. 1, 4
- [29] Jianbo Shi et al. Good features to track. In *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, pages 593–600. IEEE, 1994. 3
- [30] Linu Shine, Anitha Edison, and CV Jiji. A comparative study of faster r-cnn models for anomaly detection in 2019 ai city challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 306–314, 2019. 2
- [31] Linu Shine et al. Fractional data distillation model for anomaly detection in traffic videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 606–607, 2020. 2
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [33] 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. 2
- [34] Hanlin Tan, Yongping Zhai, Yu Liu, and Maojun Zhang. Fast anomaly detection in traffic surveillance video based on robust sparse optical flow. In *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 1976–1980. IEEE, 2016. 2
- [35] Li-Li Wang, Henry YT Ngan, and Nelson HC Yung. Automatic incident classification for large-scale traffic data by adaptive boosting svm. *Information Sciences*, 467:59–73, 2018. 2

- [36] Tian Wang, Meina Qiao, Aichun Zhu, Yida Niu, Ce Li, and Hichem Snoussi. Abnormal event detection via covariance matrix for optical flow based feature. *Multimedia Tools and Applications*, 77(13):17375–17395, 2018. 2
- [37] Tian Wang and Hichem Snoussi. Detection of abnormal visual events via global optical flow orientation histogram. *IEEE Transactions on Information Forensics and Security*, 9(6):988–998, 2014. 2
- [38] JiaYi Wei, JianFei Zhao, YanYun Zhao, and ZhiCheng Zhao. Unsupervised anomaly detection for traffic surveillance based on background modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 129–136, 2018. 2
- [39] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE international conference on image processing (ICIP)*, pages 3645–3649. IEEE, 2017. 4
- [40] Bosi Yu, Yazhou Liu, and Quansen Sun. A content-adaptively sparse reconstruction method for abnormal events detection with low-rank property. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(4):704–716, 2016. 2
- [41] Ying Zhang, Huchuan Lu, Lihe Zhang, and Xiang Ruan. Combining motion and appearance cues for anomaly detection. *Pattern Recognition*, 51:443–452, 2016. 2
- [42] Zhong Zhang, Xing Mei, and Baihua Xiao. Abnormal event detection via compact low-rank sparse learning. *IEEE Intelligent Systems*, 31(2):29–36, 2015. 2
- [43] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. 1
- [44] Zoran Zivkovic. Improved adaptive gaussian mixture model for background subtraction. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 2, pages 28–31. IEEE, 2004. 3