

# Hybrid Video Anomaly Detection for Autonomous Driving

Daniel Bogdoll<sup>\*1,2</sup>

bogdoll@fzi.de

Jan Imhof<sup>\*1</sup>

uwxdb@student.kit.edu

Tim Joseph<sup>2</sup>

joseph@fzi.de

Svetlana Pavlitska<sup>1,2</sup>

pavlitska@fzi.de

J. Marius Zöllner<sup>1,2</sup>

zoellner@fzi.de

<sup>1</sup> Karlsruhe Institute of Technology (KIT)  
Karlsruhe, Germany

<sup>2</sup> FZI Research Center for Information  
Technology  
Karlsruhe, Germany

---

## Abstract

In autonomous driving, the most challenging scenarios can only be detected within their temporal context. Most video anomaly detection approaches focus either on surveillance or traffic accidents, which are only a subfield of autonomous driving. We present HF<sup>2</sup>-VAD<sub>AD</sub>, a variation of the HF<sup>2</sup>-VAD surveillance video anomaly detection method for autonomous driving. We learn a representation of normality from a vehicle's ego perspective and evaluate pixel-wise anomaly detections in rare and critical scenarios.

## 1 Introduction

In autonomous driving, the detection of anomalies is crucial for ensuring safety and reliability. Video anomaly detection (VAD) focuses on events in video data that deviate from an expected normality. In autonomous driving, the challenges are complicated by factors such as camera movements, ever-changing backgrounds, and rapid changes in vehicle speed.

Many different types of anomalies exist [4, 8, 25], with many approaches trying to detect them [6]. Reconstruction-based methods focus on learning a representation of normality and labeling deviations from this norm [24, 43]. Predictive approaches detect anomalies when future predictions deviate significantly from observations [40, 67]. However, in autonomous driving, most of the anomaly detection methods that are based on raw sensory data focus on atypical objects rather than atypical behaviors. Temporal anomalies are often detected in object-based trajectory data, assuming perfect perception [9, 27, 31, 58]. The most complex temporal anomalies are *anomalous scenarios*, which pose a high risk of collision [8].

Until recently, mostly accident-focused benchmarks were available to evaluate temporal anomalies in the context of autonomous driving [18]. Given a rise in datasets with novel




Figure 1: Exemplary anomaly detection with  $\text{HF}^2\text{-VAD}_{AD}$ . The vehicle in front performs a sudden braking maneuver, as highlighted in yellow in the top-left graph. The graph shows the anomaly scores per frame, peaking while the vehicle performs the maneuver. The right frame shows pixel-wise anomaly scores for all instances of the class *vehicle*, computing anomaly scores only within class-specific bounding boxes predicted by an object detector.

dynamic anomalies [6, 20, 57], we are interested in the applicability of video anomaly detection methods for autonomous driving. Among recent SotA methods, we choose  $\text{HF}^2\text{-VAD}$  as the representative model [41] and adapt it to the field of autonomous driving. For this purpose, we train a  $\text{HF}^2\text{-VAD}_{AD}$  model on simulated data representing normality and evaluate it on the AnoVox benchmark [6] with *anomalous scenarios* in the form of sudden braking maneuvers of a lead vehicle. We also introduce a pixel-wise evaluation to localize anomalies.

## 2 Related Work

There are three typical camera perspectives for anomaly detection methods in the traffic domain. *Birds-Eye-View (BEV)* and *surveillance* perspectives come from external data providers, while an ego or *dashcam* view can be provided by a vehicle itself [18]. Numerous works have focused on specific scenarios, including the detection of atypical activity in urban surveillance, the analysis of irregularities in pedestrian dynamics, and the monitoring of abnormal behavior patterns in crowds [2, 11, 12, 46, 48, 49, 50, 52, 53] as well as street scenes recorded by city cameras that evaluate vehicle behavior from an outside perspective [24, 54].

Especially the VAD subfields traffic anomaly detection [58, 59] and road accident detection [17] have gained much traction. However, only a few methods specifically analyze behavioral anomalies from the perspective of a vehicle’s onboard cameras. For our comparison in Table 1, we have also excluded frame-wise [10, 23, 55, 56] or supervised [10, 22, 32, 56] methods as well as such that focus on domain-level anomalies [8, 21, 29, 53].

In general video anomaly detection, hybrid methods that combine both reconstruction and prediction methods have emerged as the current state-of-the-art. These methods aim to utilize the strengths of both approaches by leveraging temporal dynamics and reconstruction capabilities to improve anomaly detection [28, 37, 39, 41, 54, 61]. Among these,  $\text{HF}^2\text{-VAD}$  by Liu et al. [41], who built upon the work of Park et al. [47], is commonly used in the literature as a benchmark for the state of the art in video anomaly detection [50] and often serves as the foundation for further advancements [56]. It detects unusual events that deviate from normal patterns for surveillance tasks.

Method	Year	High Resolution	Required Sub-Methods						Anomaly Detection	
			ROI	(BEV) Mask	Bbox	Flow	Odometry	Tracking		Motion Pred.
AnoPred [10, 11, 12]	2019	—		×						Learned Normality
CC-Detector [9]	2019	—	×	×						
SSC-TAD [13]	2022	—			×	×			×	Learned Normality
FOL [14, 15]	2023	✓			×	×	×	×		Learned Normality
AnoPred <sub>Mask</sub> [10, 11]	2023	—			×	×				Learned Normality
AccPred [16]	2023	✓		×	×				×	Metric-based
HF <sup>2</sup> -VAD <sub>AD</sub> [17]	2024	✓			×	×				Learned Normality

Table 1: Behavioral, pixel-wise anomaly detection methods for vehicle camera data.

The idea of predicting future frames for the purpose of anomaly detection became popular around 2019, where it was demonstrated both for surveillance tasks [40] and autonomous driving [0, 59]. Liu et al. combined future frame predictions with an optical flow [40] in **AnoPred**, and **CC-Detector** by Bolte et al. focused on semantically distinguishing relevant classes from the rest of an image [9]. Fang et al. proposed self-supervised consistency learning for traffic accident detection (**SSC-TAD**) [13], where the key idea is to detect deviations of appearance, motion, and context relations. For this, they also utilize an explicit graph representation of the scenes. Future object localization (**FOL**) by Yao et al., on the other hand, focused on predicting only the bounding boxes of relevant classes in order to compensate for the motion of the ego vehicle, which makes the prediction of whole frames more complex [59]. Later, Fang et al. built upon their work but removed the need for an odometry estimation for the ego vehicle. Recently, Yao et al. extended their own work and also provided updated **AnoPred<sub>Mask</sub>** version of Liu’s approach by adding masks for relevant objects, similar to Bolte’s approach [60]. Finally, the new DeepAccident dataset by Wang et al. proposed **AccPred**, an accident detection method based on future trajectories and their distances to each other [57]. As shown in Table 1, most of the methods require many additional sub-methods to function, making them rather complex and only recently began to work for high-resolution input images.

### 3 Methodology

As shown in Section 2 and Table 1, existing video anomaly detection methods for autonomous driving are either overly complex or incapable of dealing with high-resolution input data. Based on HF<sup>2</sup>-VAD [40], our adaption HF<sup>2</sup>-VAD<sub>AD</sub> is capable of utilizing high-resolution video streams, learns a representation of normality in the domain of autonomous driving and only requires two sub-methods. Thus, it is the least complex method for video anomaly detection in the field of autonomous driving for high-resolution input.

**HF<sup>2</sup>-VAD.** Here, we provide a brief overview of the hybrid framework HF<sup>2</sup>-VAD [40] we build upon and refer to the original publication for further details. The framework combines Multi-Level Memory modules in an Autoencoder with Skip Connections (ML-MemAE-SC) for optical flow reconstructions and a Conditional Variational Autoencoder (CVAE) for future frame predictions. The framework is trained on normal data and uses discrepancies in the reconstructions and predictions to identify anomalies. Experiments were conducted with well-established surveillance datasets.

The *ML-MemAE-SC* uses an autoencoder structure with multiple memory modules at different levels. This structure improves the model’s ability to learn and retain typical data patterns from the training set and distinguish them from anomalies. The integration of skip connections between the encoder and decoder helps to preserve important data that might




Figure 2: HF<sup>2</sup>-VAD<sub>AD</sub>: Adaptation and extension of HF<sup>2</sup>-VAD for autonomous driving. Optical flows  $y_{1:t}$  and bounding boxes  $x_{1:t}$  for relevant objects are generated for each input image. ML-MemAE-SC reconstructs the optical flows  $\hat{y}_{1:t}$  with memory-modules  $M$  to better reconstruct only normal patterns. The CVAE predicts a future frame  $\hat{x}_{t+1}$ . On this basis, image-wise and localized pixel-wise anomaly scores are generated. Adapted from [40].

otherwise be lost during the encoding process. This part focuses on the accurate reconstruction of normal data. The training of the ML-MemAE-SC includes a composite loss function that summarizes reconstruction and entropy losses. The reconstruction loss is calculated based on the Euclidean distance.

The CVAE plays a central role in predicting future images by using conditioned historical data to simulate probable future scenarios. It determines the latent variable  $z_{t+1}$  by considering both the current image  $x_t$  and the optical flow  $y_{t,t+1}$  between successive images. It works on the premise that a latent variable  $z$ , representing the probabilistic state of the following image, can be derived from the immediate past to subsequently influence the generation of  $\hat{x}_{t+1}$ . It uses the encoded optical flow and the current image to estimate the posterior distribution  $p(z_{t+1} | x_t, y_{t,t+1})$ . A sample from this distribution, which represents the probabilistic state of the next image, is then drawn for image generation. During training, this latent variable  $z$  is combined with the conditions derived from the encoder  $E$  and processed by the decoder  $D$  to create the future image  $\hat{x}_{t+1}$ .

At test time, anomaly detection is performed frame-wise. The frame-wise anomaly score  $S_f$  is a fusion of the flow reconstruction error  $S_r$  and the error in the prediction of future images  $S_p$ , standardized by their respective means  $\mu$  and standard deviations  $\sigma$  from the training samples:

$$S_f = w_r \cdot \frac{S_r - \mu_r}{\sigma_r} + w_p \cdot \frac{S_p - \mu_p}{\sigma_p} \quad (1)$$

**HF<sup>2</sup>-VAD<sub>AD</sub>.** As shown in Figure 2, we first adapt HF<sup>2</sup>-VAD to autonomous driving by generating dense, pixel-wise anomaly scores for the whole frame rather than focusing on the content of unlocalized, detected bounding boxes alone, as can be seen in Figure 1. We learn flow reconstructions and flow-guided frame prediction from the perspective of an ego-vehicle rather than from the perspective of a static surveillance camera. Our training dataset was collected with the CARLA simulation engine [45]. Rather than collecting data with the default autopilot, we employ Roach [46] as the driving agent.

To reflect the complexity of real-world traffic, the training dataset features four different cities, each with its own weather and traffic dynamics. The training data includes 38 scenarios with over 30,000 high-resolution frames with a resolution of 600 x 800 pixels. Crucially, the training dataset contains only normal driving scenes so that the model can internalize






Figure 3: **Sequential Anomaly Detection:** The first image shows the normal driving mode of the lead vehicle, while the last two images show an anomaly with the pixel-wise ground truth overlaid in red. Other regularly driving traffic participants are also present. As visible in the anomaly maps, HF<sup>2</sup>-VAD<sub>AD</sub> successfully detects the unknown maneuver.

a baseline of normal traffic behavior, which is crucial for detecting deviations that indicate anomalies in real-world driving situations. Other normally behaving agents, including lead vehicles, are included in the training data.

Flows are generated by FlowNet2.0 [60] to ensure that both static and motion-based features contribute to the training of the model. Bounding boxes are detected with an off-the-shelf object detection model [45] trained on CARLA data. We first train the ML-MemAE-SC with the extracted flows from the training dataset, then train the CVAE model with the reconstructed flows from the ML-MemAE-SC and the corresponding frames, and finally fine-tune the whole framework. During fine-tuning, the weights of the two models are adjusted together to improve their coordination and the overall accuracy of anomaly detection.

We train the model specifically on foreground objects, i.e., *vehicles*, extracted from these images with bounding boxes, and construct spatiotemporal cubes (STC) for each object, where the dimensions of these cubes are standardized to 32x32. The use of bounding boxes for each vehicle in the scene allows the model to focus specifically on regions within the image where temporal anomalies in the form of unexpected maneuvers can occur.

In addition to frame-wise anomaly detection, we have integrated a pixel-wise anomaly detection. The pixel-wise anomaly score, denoted as  $S_{\text{pixel}}$ , is computed by combining the MSE values from optical flow and image prediction, both of which are normalized to enhance the reliability of the anomaly detection process:

$$S_{\text{pixel}} = w_{r,p} \cdot m_r + w_{p,p} \cdot m_p$$

Here,  $m_r$  and  $m_p$  represent the aggregated robustly scaled MSE values for optical flow and image prediction, respectively. To enhance the robustness and comparability of these MSE values across different datasets, a robust scaling technique is applied. This scaling is important as it mitigates the influence of outliers and non-standard data distributions, thereby improving the reliability of the anomaly detection process. The robust scaling of each MSE value  $x$  is defined as follows:

$$x_{\text{scaled}} = \frac{x - \tilde{x}}{\text{IQR}(X)}$$

Where  $x$  is the original MSE value,  $\tilde{x}$  is the median of the MSE values across the dataset, and  $\text{IQR}(X)$  is the interquartile range, defined as the difference between the 75th and 25th percentiles ( $Q_{75}(X)$  and  $Q_{25}(X)$ ). After robust scaling, the mean value is calculated for each pixel across the RGB channels:




Figure 4: Distribution of braking behaviors of other traffic participants in training (blue) and test (red) data. High-intensity sudden braking scenarios, labeled as anomalous, only occur in the evaluation dataset.

$$m_{\text{aggregated}} = \frac{1}{N} \sum_{i=1}^N x_{\text{scaled},i}$$

Where  $N$  is the number of RGB channels. The final pixel-wise anomaly score  $S_{\text{pixel}}$  is computed by combining the aggregated MSE values from optical flow and image prediction using the weights  $w_{r,p}$  and  $w_{p,p}$ .

The resulting pixel-wise scores are then used to localize anomalies within the image. A score of 0 is set for all pixels outside the bounding boxes, as shown in Figure 1. In contrast, the frame-wise approach of the original HF<sup>2</sup>-VAD uses simple summation and normalization based on the training statistics. Specifically, the sum of MSE losses across all dimensions is calculated for both optical flow and image predictions, and these summed values are then normalized using the mean and standard deviation from the training data. In the presence of detections, the authors added up all  $32 \times 32$  anomaly patches and chose the largest value as the frame-wise anomaly score. We report these frame-wise results as AUROC<sup>◦</sup> in Section 4.

In contrast, our approach calculates a pixel-wise anomaly score by normalizing and aggregating mean-squared errors of optical flow and frame predictions, then scaling these scores to bounding boxes within each frame. Pixel-wise scores allow for a localization of anomalies for a more granular analysis, compared to a single aggregated score per frame.

## 4 Evaluation

For the evaluation, we used the AnVoX benchmark [8] with *anomalous scenarios* in the form of sudden braking maneuvers performed by lead vehicles, as visible in Figure 3. These are common and dangerous situations in road traffic [83] and are not included in our training data, as Figure 4 shows. In the 13 evaluation scenarios, the ego-vehicle follows a lead vehicle along a longer route under different weather and traffic conditions. At certain intervals, the lead vehicle performs a sudden braking maneuver. Regularly behaving traffic participants are also present in the evaluation dataset. Based on the available pixel-wise ground truth annotations of sudden braking scenarios in AnVoX, as shown in Figure 3, we also generated bounding-box and frame-wise annotations for the evaluation. In AnVoX, the anomalies are annotated starting with the unexpected sudden braking of the leading vehicle and ending with its complete stop.

In order to evaluate the performance of HF<sup>2</sup>-VAD<sub>AD</sub>, we provide both quantitative and qualitative insights. Two metrics are used for evaluation: the already used Area Under the

Model	Domain	Weather	Finetuned	AUROC <sup>◊</sup> ↑	FPR <sub>95</sub> ↓	IoU ↑
<i>Bounding Boxes: Detection</i>						
HF <sup>2</sup> -VAD <sub>AD</sub>	All	All	✓	66.98	<b>2.58</b>	48.10
HF <sup>2</sup> -VAD <sub>AD</sub>	City	Sunshine	✓	68.04	2.48	48.09
HF <sup>2</sup> -VAD <sub>AD</sub>	City	Rain	✓	<u>69.74</u>	<u>2.68</u>	<u>51.30</u>
HF <sup>2</sup> -VAD <sub>AD</sub>	Highway	Sunshine	✓	57.22	3.15	<b>60.71</b>
HF <sup>2</sup> -VAD <sub>AD</sub>	Highway	Rain	✓	<b>76.60</b>	<b>1.34</b>	41.09
<i>Bounding Boxes: Ground Truth</i>						
ML-MemAE-SC	All	All	—	67.49	7.80	—
HF <sup>2</sup> -VAD <sub>AD</sub>	All	All	—	<u>81.01</u>	3.72	—
HF <sup>2</sup> -VAD <sub>AD</sub>	All	All	✓	79.09	<b>2.83</b>	—
HF <sup>2</sup> -VAD <sub>AD</sub>	City	Sunshine	✓	80.74	3.10	—
HF <sup>2</sup> -VAD <sub>AD</sub>	City	Rain	✓	74.08	<u>3.05</u>	—
HF <sup>2</sup> -VAD <sub>AD</sub>	Highway	Sunshine	✓	74.76	3.19	—
HF <sup>2</sup> -VAD <sub>AD</sub>	Highway	Rain	✓	<b>91.74</b>	99.32	—

Table 2: Evaluation and ablation studies of HF<sup>2</sup>-VAD<sub>AD</sub> and ML-MemAE-SC under different scenarios, with **best** and second-best results highlighted.

Receiver Operation Characteristic (AUROC<sup>◊</sup>) for frame-wise evaluations and the False Positive Rate at 95% True Positive Rate (FPR<sub>95</sub>) for pixel-wise evaluation. The AUROC metric, widely used in the VAD literature [19, 41, 42, 44], serves as a measure of VAD accuracy by varying the threshold for anomaly assessment, with higher AUROC values indicating better detection performance. Adding the FPR<sub>95</sub> [13, 26] metric provides insights into the false-positive rates if nearly all anomalies are correctly detected, with lower values indicating better detection performance. Since our detected bounding boxes do not fully match the existing ground truth, we only evaluate the FPR<sub>95</sub> based on the overlapping area and also provide the IoU to demonstrate the quality of the bounding boxes.

## 4.1 Quantitative Results

As shown in Table 2, we perform various experiments with varying conditions. Comparing city and highway scenarios, we observe a strong sensitivity of the model to these environments, but no clear trend emerges. For highway settings with bad weather, we observe poor performance. As the pixel-wise anomaly scores are only calculated within bounding boxes and all other pixels are set to 0, there is generally a low false positive rate. However, if more than 5% of the ground truth anomalies are not detected, detecting all anomalous pixels requires a threshold close to zero. This results in the entire image being detected as anomalous, causing the FPR<sub>95</sub> to rise sharply, as shown in Figure 5.

To determine optimal values for the weights, as introduced in Section 3, we perform an empirical study, as shown in Table 3. While we observe best results in the case when we completely neglect image predictions, we decided to set  $w_r, w_{r_p} = 10$  and  $w_p, w_{p_p} = 0.1$ , as the performance metrics were too small to be significant, and we retain the advantages of both modules. The study shows that the model performs significantly worse when image




Figure 5: Comparison of ROC curves for the setting *highway with rain*. The x marks the  $FPR_{95}$  position. On the left, a low  $FPR_{95}$  is achieved with predicted bounding boxes. On the right, the  $FPR_{95}$  rose sharply with ground truth bounding boxes.

$w_r$ and $w_{rp}$	$w_p$ and $w_{pp}$	$FPR_{95} \downarrow$	$AUROC^\circ \uparrow$
10	0.1	<b>2.83</b>	<u>79.09</u>
1	0.1	<u>2.86</u>	78.78
1	0	<b>2.83</b>	<b>79.13</b>
0	1	96.67	69.15
0.1	1	4.05	76.20
0.1	10	95.78	70.70

Table 3: Empirical study to determine the effects of different weights for  $HF^2$ -VAD<sub>AD</sub>.

prediction is weighted higher, which is in contrast to the original implementation of  $HF^2$ -VAD, in which the frame prediction was weighted higher than the flow reconstruction, with the weight for frame prediction set to 1 and the weight for flow reconstruction set to 0.1.

**Ablation Studies.** To better understand the performance of our approach, we perform three types of ablation studies, as shown in Table 2, with the focus on the first three entries in the lower part of the table. First, we only utilize the intermediate outputs of the ML-MemAE-SC to evaluate isolated flow reconstructions.

Second, we evaluate both bounding boxes generated from an off-the-shelf object detection model [45] and ground truth bounding boxes from the evaluation dataset to isolate the effect of model failures. We observe that bounding boxes are best detected in highway settings with good weather. The IoU performance suffers especially from poor detections of distant objects. Since mostly irrelevant vehicles are missing in the detections, we see improved  $FPR_{95}$  values for bounding box detections, but at the cost of reduced AUROC scores.

Lastly, we examine the fine-tuning stage during training. After fine-tuning,  $HF^2$ -VAD<sub>AD</sub> shows an improvement in precision. While the AUROC decreases slightly, the model achieves a better balance in the performance metrics with a significantly lower  $FPR_{95}$ .

## 4.2 Qualitative Results

In Figure 6, we show a series of comparative visualizations demonstrating the performance of  $HF^2$ -VAD<sub>AD</sub> in traffic scenarios. The model is able to predict regular behaviors very well,




Figure 6: On the left, two normal scenarios are shown. Here, the model’s next-frame predictions (a) match the actual next frames (b) closely. On the right, we have shown sudden brake scenarios. Here, significant prediction differences arise, leading to high anomaly values (c).

while atypical ones lead to poor predictions and, thus, high anomaly scores. The pixel-wise anomaly maps serve as a meaningful indicator, with more highlighted areas signaling larger prediction discrepancies. For normal driving patterns, these maps show only small deviations from the ground truth, indicating that the model can predict the subsequent images well. For anomalies, the model often predicts, as it learned, that the vehicle maintains its distance. This leads to positive anomaly detections, as the vehicle is currently anomalous and getting closer due to the braking maneuver. This can be recognized by the increased highlighting in the difference maps.

**Failure Cases.** As shown in Figure 7, there are certain cases which are challenging. The first image shows how the model often incorrectly produces high anomaly scores when a lead vehicle is turning rather than going straight. This may be due to the model primarily learns straight-line driving patterns and the training data does not contain enough examples of lead vehicles turning directly in front of the car. The middle image shows a scene where the object detection model has problems with detecting objects in unfavorable weather conditions, such as heavy rain, so anomaly detection is not possible.

## 5 Conclusion

In summary, our research has shown that HF<sup>2</sup>-VAD, a framework originally developed for detecting anomalies in surveillance systems, can be effectively transferred to autonomous driving. By extending it to dense pixel-wise anomaly maps and creating a dedicated training and test dataset, we have validated the effectiveness of HF<sup>2</sup>-VAD<sub>AD</sub> in the vehicle context, bridging the gap between static monitoring and dynamic vehicle environments.

Our contributions further include the extension to dense anomaly detections, domain-specific weight adaptations, and dealing with multiple bounding boxes separately rather than combining them into one score, using robust scaling. In the future, we aim to further refine the framework to better deal with current failure cases and evaluate more types of anomalies to examine its performance among a wider variety of scenarios, such as wrong-way drivers.



Figure 7: **Failure Cases:** Left: The model can produce high pixel-wise anomaly scores for turning vehicles. The encircled area in the frame-wise graph shows the duration of the curve. Middle: A missing bounding box detection during bad weather situations leads to no anomaly detection. Right: The bounding box only captures a part of the lead vehicle, and a truck in the back is completely missed.

## Acknowledgment

This work results from the just better DATA project supported by the German Federal Ministry for Economic Affairs and Climate Action (BMWK), grant number 19A23003H.

## References

- [1] Georgios Aivatoglou, Nikolaos Oikonomou, Georgios Spanos, Kristina Livitckaia, Konstantinos Votis, and Dimitrios Tzovaras. EVENT: Real-time Video Feed Anomaly Detection for Enhanced Security in Autonomous Vehicles. In *Mediterranean Conference on Control and Automation (MED)*, 2023.
- [2] Amnah Aldayri and Waleed Albattah. Taxonomy of Anomaly Detection Techniques in Crowd Scenes. *Sensors*, 22, 2022.
- [3] Shuang Bai, Chao Han, and Shan An. Recognizing Anomalies in Urban Road Scenes Through Analysing Single Images Captured by Cameras on Vehicles. *Sensing and Imaging*, 19, 2018.
- [4] Daniel Bogdoll, Jasmin Breitenstein, Florian Heidecker, Maarten Bieshaar, Bernhard Sick, Tim Fingscheidt, and J. Marius Zollner. Description of corner cases in automated driving: Goals and challenges. In *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2021.
- [5] Daniel Bogdoll, Maximilian Nitsche, and J. Marius Zöllner. Anomaly Detection in Autonomous Driving: A Survey. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022.
- [6] Daniel Bogdoll, Iram Hamdard, Lukas Namgyu Rößler, Felix Geisler, Muhammed Bayram, Felix Wang, Jan Imhof, Miguel de Campos, Anushervon Tabarov, Yitian Yang, Hanno Gottschalk, and J. Marius Zöllner. AnoVox: A Benchmark for Multi-modal Anomaly Detection in Autonomous Driving. In *European Conference on Computer Vision Workshop (ECCVW)*, 2024.

- [7] Jan-Aike Bolte, Andreas Bar, Daniel Lipinski, and Tim Fingscheidt. Towards corner case detection for autonomous driving. In *IEEE Intelligent Vehicles Symposium (IV)*, 2019.
- [8] Jasmin Breitenstein, Jan-Aike Termöhlen, Daniel Lipinski, and Tim Fingscheidt. Systematization of corner cases for visual perception in automated driving. In *IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [9] Antonia Breuer, Quy Le Xuan, Jan-Aike Termöhlen, Silviu Homoceanu, and Tim Fingscheidt. Quo Vadis? Meaningful Multiple Trajectory Hypotheses Prediction in Autonomous Driving. In *IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021.
- [10] Fu-Hsiang Chan, Yu-Ting Chen, Yu Xiang, and Min Sun. Anticipating Accidents in Dashcam Videos. In Shang-Hong Lai, Vincent Lepetit, Ko Nishino, and Yoichi Sato, editors, *Asian Conference on Computer Vision (ACCV)*, 2017.
- [11] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41, 2009.
- [12] S. Chandrakala, K. Deepak, and G. Revathy. Anomaly detection in surveillance videos: a thematic taxonomy of deep models, review and performance analysis. *Artificial Intelligence Review*, 56, 2023.
- [13] Giancarlo Di Biase, Hermann Blum, Roland Siegwart, and Cesar Cadena. Pixel-wise Anomaly Detection in Complex Driving Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [14] Youcef Djenouri, Asma Belhadi, Jerry Chun-Wei Lin, Djamel Djenouri, and Alberto Cano. A Survey on Urban Traffic Anomalies Detection Algorithms. *IEEE Access*, 7, 2019.
- [15] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An Open Urban Driving Simulator. In *Conference on Robot Learning (CoRL)*, 2017.
- [16] Jianwu Fang, Jiahuan Qiao, Jie Bai, Hongkai Yu, and Jianru Xue. Traffic Accident Detection via Self-Supervised Consistency Learning in Driving Scenarios. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 23, 2022.
- [17] Jianwu Fang, Lei-lei Li, Junfei Zhou, Junbin Xiao, Hongkai Yu, Chen Lv, Jianru Xue, and Tat-Seng Chua. Abductive Ego-View Accident Video Understanding for Safe Driving Perception. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [18] Jianwu Fang, Jiahuan Qiao, Jianru Xue, and Zhengguo Li. Vision-Based Traffic Accident Detection and Anticipation: A Survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 34, 2024.
- [19] 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 *International Conference on Computer Vision (ICCV)*, 2019.



- [20] Lei Gong, Yu Zhang, Yingqing Xia, Yanyong Zhang, and Jianmin Ji. SDAC: A Multimodal Synthetic Dataset for Anomaly and Corner Case Detection in Autonomous Driving. *AAAI Conference on Artificial Intelligence*, 38, 2024.
- [21] Ruben Grewal, Paolo Tonella, and Andrea Stocco. Predicting Safety Misbehaviours in Autonomous Driving Systems using Uncertainty Quantification. *arXiv:2404.18573*, 2024.
- [22] Feten Hajri and Hajer Fradi. Vision Transformers for Road Accident Detection from Dashboard Cameras. In *IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 2022.
- [23] Sanjay Haresh, Sateesh Kumar, M. Zeeshan Zia, and Quoc-Huy Tran. Towards Anomaly Detection in Dashcam Videos. In *IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [24] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K. Roy-Chowdhury, and Larry S. Davis. Learning Temporal Regularity in Video Sequences. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [25] Florian Heidecker, Jasmin Breitenstein, Kevin Rösch, Jonas Löhdefink, Maarten Bieshaar, Christoph Stiller, Tim Fingscheidt, and Bernhard Sick. An application-driven conceptualization of corner cases for perception in highly automated driving. In *IEEE Intelligent Vehicles Symposium (IV)*, 2021.
- [26] Jens Henriksson, Christian Berger, Markus Borg, Lars Tornberg, Sankar Raman Sathyamoorthy, and Cristofer Englund. Performance Analysis of Out-of-Distribution Detection on Trained Neural Networks. *Information and Software Technology*, 130, 2021.
- [27] Kai-Chieh Hsu, Karen Leung, Yuxiao Chen, Jaime F. Fisac, and Marco Pavone. Interpretable trajectory prediction for autonomous vehicles via counterfactual responsibility. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [28] Heqing Huang, Bing Zhao, Fei Gao, Penghui Chen, Jun Wang, and Amir Hussain. A Novel Unsupervised Video Anomaly Detection Framework Based on Optical Flow Reconstruction and Erased Frame Prediction. *Sensors*, 23, 2023.
- [29] Manzoor Hussain, Nazakat Ali, and Jang-Eui Hong. DeepGuard: A framework for safeguarding autonomous driving systems from inconsistent behaviour. *Automated Software Engineering*, 29, 2021.
- [30] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [31] Ruochen Jiao, Juyang Bai, Xiangguo Liu, Takami Sato, Xiaowei Yuan, Qi Alfred Chen, and Qi Zhu. Learning representation for anomaly detection of vehicle trajectories. *arXiv:2303.05000*, 2023.
- [32] Hoon Kim, Kangwook Lee, Gyeongjo Hwang, and Changho Suh. Crash to Not Crash: Learn to Identify Dangerous Vehicles Using a Simulator. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 2019.

- [33] Hoon Kim, Kangwook Lee, Gyeongjo Hwang, and Changho Suh. Crash to Not Crash: Learn to Identify Dangerous Vehicles Using a Simulator. *AAAI Conference on Artificial Intelligence*, 33, 2019.
- [34] Santhosh Kelathodi Kumaran, Debi Prosad Dogra, and Partha Pratim Roy. Anomaly Detection in Road Traffic Using Visual Surveillance: A Survey. *ACM Computing Surveys*, 53, 2021.
- [35] Sofiane Lagraa, Maxime Cailac, Sean Rivera, Frederic Beck, and Radu State. Real-Time Attack Detection on Robot Cameras: A Self-Driving Car Application. In *IEEE International Conference on Robotic Computing (IRC)*, 2019.
- [36] Trung-Nghia Le, Shintaro Ono, Akihiro Sugimoto, and Hiroshi Kawasaki. Attention R-CNN for Accident Detection. In *IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [37] Qun Li, Rui Yang, Yaying Shen, Ziyi Zhang, and Xianzhong Long. Dual-branch framework with convolutional attentive block for video anomaly detection. In *IEEE Chinese Control Conference (CCC)*, 2023.
- [38] Rongqin Liang, Yuanman Li, Jiantao Zhou, and Xia Li. Text-driven traffic anomaly detection with temporal high-frequency modeling in driving videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.
- [39] Ting Liu, Chengqing Zhang, Xiaodong Niu, and Liming Wang. Spatio-temporal prediction and reconstruction network for video anomaly detection. *Plos One*, 17, 2022.
- [40] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future Frame Prediction for Anomaly Detection - A New Baseline. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [41] Zhian Liu, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. A hybrid video anomaly detection framework via memory-augmented flow reconstruction and flow-guided frame prediction. In *IEEE/CVF International conference on computer vision (ICCV)*, 2021.
- [42] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal Event Detection at 150 FPS in MATLAB. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [43] Weixin Luo, Wen Liu, and Shenghua Gao. Remembering history with convolutional LSTM for anomaly detection. *IEEE International Conference on Multimedia and Expo (ICME)*, 2017.
- [44] Vijay Mahadevan, Weixin Li, Viral Bhalodia, and Nuno Vasconcelos. Anomaly detection in crowded scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [45] MinyiLin. Roboflow Universe: CARLA Computer Vision Project. <https://universe.roboflow.com/minyilin-3nzak/carla-f4116>, 2022. Accessed: 2924-06-07.
- [46] Rashmiranjan Nayak, Umesh Chandra Pati, and Santos Kumar Das. A comprehensive review on deep learning-based methods for video anomaly detection. *Image and Vision Computing*, 106, 2021.

- [47] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning Memory-Guided Normality for Anomaly Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [48] N. Patil and Prabir Kumar Biswas. A survey of video datasets for anomaly detection in automated surveillance. In *International Symposium on Embedded Computing and System Design (ISED)*, 2016.
- [49] O. P. Popoola and Kejun Wang. Video-Based Abnormal Human Behavior Recognition—A Review. *IEEE Transactions on Systems, Man, and Cybernetics*, 42, 2012.
- [50] Bharathkumar Ramachandra, Michael J. Jones, and Ranga Raju Vatsavai. A Survey of Single-Scene Video Anomaly Detection. *arXiv:2004.05993*, 2020.
- [51] Nicolae-Catalin Ristea, Florinel-Alin Croitoru, Radu Tudor Ionescu, Marius Popescu, Fahad Shahbaz Khan, and Mubarak Shah. Self-Distilled Masked Auto-Encoders are Efficient Video Anomaly Detectors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [52] K. K. Santhosh, D. P. Dogra, and P. P. Roy. Anomaly Detection in Road Traffic Using Visual Surveillance: A Survey. *ACM Computing Surveys*, 53, 2020.
- [53] Andrea Stocco, Michael Weiss, Marco Calzana, and Paolo Tonella. Misbehaviour Prediction for Autonomous Driving Systems. In *IEEE/ACM International Conference on Software Engineering (ICSE)*, 2020.
- [54] Yao Tang, Lin Zhao, Shanshan Zhang, Chen Gong, Guangyu Li, and Jian Yang. Integrating prediction and reconstruction for anomaly detection. *Pattern Recognition Letters*, 129, 2020.
- [55] Tung Minh Tran, Tu N. Vu, Nguyen D. Vo, Tam V. Nguyen, and Khang Nguyen. Anomaly Analysis in Images and Videos: A Comprehensive Review. *ACM Computing Surveys*, 55, 2022.
- [56] Weilin Wan, Weizhong Zhang, and Cheng Jin. Pose-Motion Video Anomaly Detection via Memory-Augmented Reconstruction and Conditional Variational Prediction. In *IEEE International Conference on Multimedia and Expo (ICME)*, 2023.
- [57] Tianqi Wang, Sukmin Kim, Ji Wenxuan, Enze Xie, Chongjian Ge, Junsong Chen, Zhenguo Li, and Ping Luo. DeepAccident: A Motion and Accident Prediction Benchmark for V2X Autonomous Driving. *AAAI Conference on Artificial Intelligence*, 38, 2024.
- [58] Julian Wiederer, Arij Bouazizi, Marco Troina, Ulrich Kressel, and Vasileios Belagianis. Anomaly Detection in Multi-Agent Trajectories for Automated Driving. In *Conference on Robot Learning (CoRL)*, 2021.
- [59] Yu Yao, Mingze Xu, Yuchen Wang, David J. Crandall, and Ella M. Atkins. Unsupervised Traffic Accident Detection in First-Person Videos. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

- [60] Yu Yao, Xizi Wang, Mingze Xu, Zelin Pu, Yuchen Wang, Ella Atkins, and David J. Crandall. DoTA: Unsupervised Detection of Traffic Anomaly in Driving Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 2023.
- [61] Muchao Ye, Xiaojiang Peng, Weihao Gan, Wei Wu, and Yu Qiao. AnoPCN: Video Anomaly Detection via Deep Predictive Coding Network. *ACM International Conference on Multimedia*, 2019.
- [62] Guang Yu, Siqi Wang, Zhiping Cai, En Zhu, Chuanfu Xu, Jianping Yin, and Marius Kloft. Cloze Test Helps: Effective Video Anomaly Detection via Learning to Complete Video Events. In *ACM International Conference on Multimedia*, 2020.
- [63] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [64] Yuxiang Zhao, Wenhao Wu, Yue He, Yingying Li, Xiao Tan, and Shifeng Chen. Good Practices and a Strong Baseline for Traffic Anomaly Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2021.
- [65] Zhili Zhou, Xiaohua Dong, Zhetao Li, Keping Yu, Chun Ding, and Yimin Yang. Spatio-Temporal Feature Encoding for Traffic Accident Detection in VANET Environment. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 23, 2022.