












Enhancing Highway Safety: Accident Detection on the A9 Test Stretch Using Roadside Sensors

Walter Zimmer  *  Ross Greer ^{†‡}  Xingcheng Zhou *  Rui Song * [§] 
Marc Pavel *  Daniel Lehmborg *  Ahmed Ghita * [¶] 
Akshay Gopalkrishnan [†]  Mohan Trivedi [†]  Alois Knoll * 

<https://tum-traffic-dataset.github.io/tumtraf-a>




Figure 1: Our dataset provides a visual representation of highway accidents with 3D box annotations, track IDs, instance masks, and vehicle trajectories. The accidents were captured using roadside cameras on the A9 Test Bed for Autonomous Driving. The scenes include crashes, overturned vehicles, and some instances of cars catching fire.

Abstract: Road traffic injuries are the leading cause of death for people aged 5-29, resulting in about 1.19 million deaths each year. To reduce these fatalities, it is essential to address human errors like speeding, drunk driving, and distractions. Additionally, faster accident detection and quicker medical response can help save lives. We propose an accident detection framework that combines a rule-based approach with a learning-based one. We introduce a dataset of real-world highway accidents, featuring high-speed crash sequences. It includes 294,924 labeled 2D boxes, 93,012 labeled 3D boxes, and track IDs across 48,144 frames, captured at 10 Hz using four roadside cameras and LiDAR sensors. The dataset covers ten object classes and is released in the OpenLABEL format. Our experiments and analysis demonstrate the reliability of our method.

Keywords: Traffic Safety, Accident Detection, Dataset, Roadside Infrastructure, ITS

*Chair of Artificial Intelligence and Robotics (AIR), Department of Computer Engineering (CE), School of Computation, Information and Technology (CIT), Technical University of Munich (TUM).

 Corresponding author: walter.zimmer@cs.tum.edu

[†]Laboratory for Intelligent and Safe Automobiles (LISA) at the Uni. of California San Diego (UCSD).

[‡]University of California Merced (UCM).

[§]Fraunhofer Institute for Transportation and Infrastructure Systems (IVI).

[¶]SETLabs Research GmbH.

1 Introduction

Data on rare events, such as long-tail events, is needed for improving machine learning models used in robotics, including perception, planning, and control [1]. In autonomous driving, however, gathering this data is both expensive and challenging [2–4]. Events like accidents and near-misses pose serious risks to human life and are difficult to capture naturally [5,6]. Quick accident detection is critical because the time between a crash and medical response can determine survival. Automated accident detection can shorten this response time and help save lives. Our dataset contributes to safer autonomous systems in several ways. First, it enhances learning for key perception tasks like detection [7–17], tracking, segmentation [18] or occupancy prediction [19]. Second, its design enables cooperative perception, where roadside sensors capture the same scene from different angles to reduce blind spots. Finally, these sensors help create digital twins of traffic environments [20], extending visibility beyond a single vehicle’s perspective. This broader view can improve early warnings and safer decision-making in risky traffic situations [21].

Our key contributions are the following:

- We propose an accident detection framework capable of identifying and analyzing crashes and near-misses in real-time in various traffic scenarios.
- We introduce a specialized highway accident dataset designed to capture rare and high-risk events, featuring 294,924 2D annotations and 93,012 3D annotations.
- Our experiments and analysis demonstrate that our approach achieves top performance on this dataset.
- We publicly release the dataset, detection framework, and development tools on our project website.

2 Related Work

Accurate accident detection is crucial for improving road safety and emergency response times. However, existing methods have not been extensively tested on real traffic data from designated test sites. Real-world accident datasets [22] are rare and often lack sufficient data to train deep learning models effectively [23]. In this section, we categorize accident detection approaches into different groups based on their methodologies and data sources.

2.1 Rule-Based Accident Detection

Rule-based methods [24] rely on predefined traffic rules and heuristics to identify accidents. These approaches typically use sensor readings, such as abrupt changes in speed or sudden braking, to infer collisions. While effective in controlled environments, rule-based systems struggle with unseen accident scenarios and complex traffic settings.

2.2 Learning-Based Accident Detection

Machine learning models, particularly deep learning approaches, have been increasingly used for accident detection. These methods use large datasets to learn patterns associated with crashes [25, 26]. Supervised learning techniques train models on labeled accident and non-accident data, whereas unsupervised approaches [27] detect anomalies that may indicate collisions. Despite their effectiveness, these models require extensive labeled data, which is often difficult to obtain.

2.3 Trajectory-Based Accident Detection

Trajectory-based methods [28] analyze the motion patterns of vehicles to detect unusual behavior that may lead to accidents. These techniques use sensor data to track vehicles in 2D or 3D and identify sudden deviations, abrupt stops, or erratic lane changes. While useful for predicting potential accidents, trajectory-based methods may struggle with false positives and ID switches in highly occluded traffic scenarios.

2.4 Vision-Based Accident Detection

Vision-based approaches [29] utilize camera footage to detect accidents by identifying visual cues such as vehicle deformations, smoke, or fire. These methods often employ convolutional neural networks (CNNs) to classify accident and non-accident scenarios. Although vision-based detection can provide valuable real-time insights, it is sensitive to lighting conditions, occlusions, and camera angles.

2.5 Video-Based Accident Detection

Video-based methods [30] extend image-based approaches by analyzing temporal information in video streams. These techniques apply deep learning models such as recurrent neural networks (RNNs) or transformers to track accident progression over time. Video-based detection enhances accuracy by considering motion dynamics but requires substantial computational resources and well-annotated datasets.

2.6 Accident Detection Using Synthetic Data

To overcome the scarcity of real-world accident data, synthetic datasets have been developed for training and evaluating accident detection models. One notable example is DeepAccident [31], which contains 691 synthetic accident scenarios generated in the CARLA simulator [32]. These accidents are based on crash reports from the National Highway Traffic Safety Administration (NHTSA) and include labeled data from four vehicles and one roadside infrastructure camera. However, a key limitation of DeepAccident is its reliance on simulated environments, which may not fully capture the complexities of real-world crashes. Addressing the sim-to-real gap is essential to improve the generalization of perception models trained on synthetic data.

2.7 Accident Detection Using Real Data

Real-world accident detection datasets are limited due to the difficulty of capturing and annotating crash events. Some existing datasets contain dashcam footage, roadside surveillance videos, or event logs from connected vehicles. However, many of these datasets lack sufficient labeled data to train deep learning models effectively. Furthermore, variations in road conditions, weather, and camera perspectives introduce additional challenges in real-world accident detection.

In summary, while significant progress has been made in accident detection, the scarcity of real-world datasets remains a major challenge. Related approaches may be used to detect both accidents and general traffic anomalies [33,34]. Existing methods often struggle with generalization due to limited real-world training data or the reliance on synthetic environments. To address this gap, our accident detection framework integrates both rule-based and learning-based approaches to reliably identify accidents in 3D. By leveraging real-world roadside sensor data and combining explicit traffic rules with deep learning models, our framework enhances detection accuracy, improves robustness across diverse traffic conditions, and enables faster emergency response times.

3 Accident Detection Framework

We propose an automatic accident detection framework that operates in real-time using roadside cameras. Our method combines both rule-based and learning-based approaches to improve accuracy and reliability. The framework first applies a rule-based approach to detect potential accidents based on vehicle trajectories. If an accident is identified, a learning-based model further verifies the event to minimize false positives.

3.1 Rule-Based Accident Detection

The rule-based approach relies on analyzing vehicle trajectories to detect unusual behavior indicative of an accident. Predefined thresholds, such as sudden stops, erratic lane changes, or prolonged stationary positions on active lanes, are used to classify accidents in real-time. This method provides immediate accident classifications for each vehicle in the current frame, enabling rapid initial detection. The rule-based approach analyzes lane positions, distance measurements, and object speeds to detect potential accidents based on predefined maneuver rules:

$$velocity_i \geq \frac{15 \text{ km/h}}{3.6} \quad (1)$$

$$velocity_i > velocity_lead_i \quad (2)$$

$$velocity_i \geq velocity_j \quad \forall \quad i < j \leq N \quad (3)$$

$$distance_lead_i \geq distance_threshold \quad (4)$$

$$distance_lead_i < \left(\frac{velocity_i - velocity_lead_i}{30} \right)^2 \quad (5)$$

$$ttc_leading_i \leq ttc_threshold \quad (6)$$

If all six rules are met at the same time, the traffic participant is classified as being involved in an accident.

3.2 Learning-Based Accident Detection

Once a potential accident is flagged by the rule-based method, a learning-based approach refines the prediction. We train a YOLOv8 model [35] on our custom accident dataset to detect accidents directly from camera images. Here we use the CVAT annotation tool [36] and annotate accident events with 2D bounding boxes. We split the full dataset into a training (80%), validation (10%), and test (10%) set. It contains 3,725 image frames in total: 2,896 for training, 410 for validation, and 419 for testing. An accident event is visible in 59.48% of images. Table 1 provides an overview of the dataset composition.

Dataset	Train	Validation	Test	Total
Accident frames	1,799	224	292	2,216
Non-accident frames	1,097	186	127	1,509
Total	2,896	410	419	3,725

Table 1: Total number of image frames in the custom dataset. 59.48% of images capture an accident event.

The model filters detections based on a confidence threshold of 0.8. To further reduce false positives, an accident must be detected in at least three consecutive frames before confirmation. Additionally, we aggregate accident detections from multiple roadside cameras to enhance detection accuracy and reduce occlusion effects.

4 Evaluation

4.1 Dataset

Our dataset features 48,144 labeled camera and LiDAR frames with 294,924 2D and 93,012 3D box annotations, tracking labels and trajectory information, and classification of ten different instance types, including cars, trucks, buses, trailers, vans, pedestrians, motorcycles, bicycles, emergency vehicles, and others. The data is captured from roadside cameras and LiDARs during day and nighttime and labeled with the *3D Bounding Box Annotation Toolbox* (3D BAT) [37]. Featured accidents include instances of high-speed lane changes with failure to notice stopped traffic, overturning of vehicles during collision, vehicles catching fire, a variety of emergency response vehicles, and more. Some example cases are illustrated in Fig. 1.

Our dataset can be used as ground truth for the development and verification of AI-based detectors, tracking, fusion algorithms, trajectory prediction, and to understand and analyze the occurrence and the after-effects of naturally occurring high-speed crash incidents and other accidents on the autobahn.

The dataset, as well as its preceding datasets [38–42], are available for academia and industry and include a development kit repository for ease of use [43].

4.2 Experiments

To evaluate our framework, we recorded camera images and processed fused perception results over 128 days, storing the data in rosbag files. The automatic accident analysis was performed on 12,290 15-minute video segments, identifying 831,969 unique vehicles. Our framework detected: 3,748 standing vehicles in active driving lanes, 138 standing vehicles in shoulder lanes, 120 breakdown events, and one accident.

Moreover, we evaluate the accuracy and runtime performance of both accident detection models on our dataset. The rule-based approach needs 10.41 ms per frame (95.05 FPS) on an NVIDIA RTX 3090 GPU. The total processing runtime for a 15-minute rosbag file with 22,500 ROS messages recorded at 25 FPS is 234.25 seconds. This includes the calculation of the lane ID, the distance calculation between all vehicles, and the scenario classification.

Our experiments (see Table 2) show that the learning-based accident detection method achieved a precision of 0.8 and a perfect recall of 1.0 on the test data, ensuring high accuracy in identifying accidents. Additionally, the detection process takes 10.41 ms for the rule-based approach and 16.13 ms (using TensorRT acceleration) for the learning-based approach, making it efficient for real-time applications.

Approach	Accuracy			Runtime [ms]	
	Precision	Recall	F1-Score	AP	
Rule-based Approach	1.000	0.500	0.667	0.97	10.41
Learning-based Approach	0.800	1.000	0.889	0.75	16.13

Table 2: Evaluation of the accident detection framework on the test set.

Qualitative results of our accident detection framework are shown in Fig. 2. We train object detection models using image resolutions of 1920 px, 1280 px, and 960 px, determining that 1280 px yields the best performance.



Figure 2: Visual comparison of the accident detection results (purple) on the test set of our accident dataset. Labeled accidents are displayed in turquoise. This test sequence shows a car accident scenario with a yellow car and a white van. We train object detection models with different input image resolutions (1920 px, 1280 px, and 960 px) and find out that an image resolution of 1280 px works best.

5 Conclusion

Traffic accidents remain a leading cause of fatalities, emphasizing the need for rapid detection to enhance emergency response times. In this work, we introduce a real-world highway accident dataset and propose a hybrid accident detection framework that integrates rule-based and learning-based approaches. Our method demonstrates strong performance in real-time accident detection using roadside infrastructure sensors, effectively identifying crashes and near-miss events. Experimental results confirm the reliability and efficiency of our approach, highlighting its potential to improve traffic safety and emergency response. Future research should focus on enhancing model robustness, expanding predictive capabilities, and bridging the gap between synthetic and real-world accident data.

6 Future Work

While our framework demonstrates strong performance in real-time accident detection, several areas remain for future improvement. One key challenge is bridging the gap between synthetic and real-world data. Future research should explore domain adaptation techniques and self-supervised learning to enhance model generalization across different environments. Another important direction is improving model robustness under varying traffic conditions, such as heavy congestion, adverse weather, and nighttime scenarios. This could involve integrating multi-modal sensor data, such as LiDAR and radar, to complement camera-based detection.

We also plan to extend our rule-based approach to detect lateral collisions and more complex accident scenarios, such as multi-vehicle crashes and near-miss events. Additionally, refining our learning-based model with larger and more diverse datasets could further reduce false positives and improve real-time decision-making.

In the future, we plan to use Vision-Language Models (VLMs) [44] to enhance the accuracy and reliability of accident detection. Finally, expanding our framework to incorporate predictive accident detection, identifying high-risk situations before a collision occurs, could significantly enhance traffic safety. This may also involve trajectory forecasting.

By addressing these challenges, we aim to develop a more reliable and comprehensive accident detection system that can be deployed in real-world autonomous driving and intelligent transportation applications.

7 Acknowledgments

This research was supported by the Federal Ministry of Education and Research in Germany as part of the AUTOftech.agil project [45] (Grant Number: 01IS22088U). We thank our collaborators and institutions for their valuable contributions and support. Additionally, we appreciate the efforts of the research community in advancing accident detection and traffic safety technologies.

References

- [1] F. Christianos, P. Karkus, B. Ivanovic, S. V. Albrecht, and M. Pavone, “Planning with occluded traffic agents using bi-level variational occlusion models,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5558–5565, IEEE, 2023.
- [2] A. Ghita, B. Antoniussen, W. Zimmer, R. Greer, C. Creß, A. Møgelmoose, M. Trivedi, and A. C. Knoll, “Activeanno3d-an active learning framework for multi-modal 3d object detection,” in *35th IEEE Intelligent Vehicles Symposium (IV) 2024*, 2024.
- [3] N. Kulkarni, A. Rangesh, J. Buck, J. Feltracco, M. Trivedi, N. Deo, R. Greer, S. Sarraf, and S. Sathyanarayana, “Create a large-scale video driving dataset with detailed attributes using amazon sagemaker ground truth,” 2021.
- [4] T. Fingscheidt, H. Gottschalk, and S. Houben, *Deep neural networks and data for automated driving: Robustness, uncertainty quantification, and insights towards safety*. Springer Nature, 2022.
- [5] H. Wang, X. Zhang, Z. Li, J. Li, K. Wang, Z. Lei, and R. Haibing, “Ips300+: a challenging multi-modal data sets for intersection perception system,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2539–2545, IEEE, 2022.
- [6] R. Greer, B. Antoniussen, M. V. Andersen, A. Møgelmoose, and M. M. Trivedi, “The why, when, and how to use active learning in large-data-driven 3d object detection for safe autonomous driving: An empirical exploration,” *arXiv preprint arXiv:2401.16634*, 2024.
- [7] M. Liu, E. Yurtsever, M. Brede, J. Meng, W. Zimmer, X. Zhou, B. L. Zagar, Y. Cui, and A. Knoll, “Graphrelate3d: Context-dependent 3d object detection with inter-object relationship graphs,” *arXiv preprint arXiv:2405.06782*, 2024.
- [8] S. Carta, M. Castrillón-Santana, M. Marras, S. Mohamed, A. S. Podda, R. Saia, M. Sau, and W. Zimmer, “Roadsense3d: A framework for roadside monocular 3d object detection,” in *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, pp. 452–459, 2024.
- [9] W. Zimmer, J. Birkner, M. Brucker, H. T. Nguyen, S. Petrovski, B. Wang, and A. C. Knoll, “Infradet3d: Multi-modal 3d object detection based on roadside infrastructure camera and lidar sensors,” in *2023 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2023.
- [10] W. Zimmer, J. Wu, X. Zhou, and A. C. Knoll, “Real-time and robust 3d object detection with roadside lidars,” in *Proc. of Int. Scientific Conf. on Mobility and Transport: Mobility Innovations for Growing Megacities*, pp. 199–219, Springer, 2023.
- [11] W. Zimmer, E. Ercelik, X. Zhou, X. J. D. Ortiz, and A. Knoll, “A survey of robust 3d object detection methods in point clouds,” *arXiv preprint arXiv:2204.00106*, 2022.

- [12] W. Zimmer, M. Grabler, and A. Knoll, “Real-time and robust 3d object detection within road-side lidars using domain adaptation,” *arXiv preprint arXiv:2204.00132*, 2022.
- [13] M. P. Philipsen, M. B. Jensen, A. Møgelmoose, T. B. Moeslund, and M. M. Trivedi, “Traffic light detection: A learning algorithm and evaluations on challenging dataset,” in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pp. 2341–2345, IEEE, 2015.
- [14] H. Abualsaud, S. Liu, D. B. Lu, K. Situ, A. Rangesh, and M. M. Trivedi, “Laneaf: Robust multi-lane detection with affinity fields,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7477–7484, 2021.
- [15] W. Zimmer, R. Pranamulia, X. Zhou, M. Liu, and A. C. Knoll, “PointCompress3d – a point cloud compression framework for roadside LiDARs in intelligent transportation systems.”
- [16] M. Liu, E. Yurtsever, M. Brede, J. Meng, W. Zimmer, X. Zhou, B. L. Zagar, Y. Cui, and A. Knoll, “GraphRelate3d: Context-dependent 3d object detection with inter-object relationship graphs.”
- [17] S. Mohamed, W. Zimmer, R. Greer, A. A. Ghita, M. Castrillon-Santana, M. Trivedi, A. Knoll, S. M. Carta, and M. Marras, “Transfer learning from simulated to real scenes for monocular 3d object detection.”
- [18] R. Greer, A. Gopalkrishnan, M. Keskar, and M. M. Trivedi, “Patterns of vehicle lights: Addressing complexities of camera-based vehicle light datasets and metrics,” *Pattern Recognition Letters*, vol. 178, pp. 209–215, 2024.
- [19] R. Song, C. Liang, H. Cao, Z. Yan, W. Zimmer, M. Gross, A. Festag, and A. Knoll, “Collaborative semantic occupancy prediction with hybrid feature fusion in connected automated vehicles,” in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17996–18006. ISSN: 2575-7075.
- [20] P. Kremer, N. Nourani-Vatani, and S. Park, “A digital twin for teleoperation of vehicles in urban environments,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12521–12527, IEEE, 2023.
- [21] R. Greer, N. Deo, A. Rangesh, M. Trivedi, and P. Gunaratne, “Safe control transitions: Machine vision based observable readiness index and data-driven takeover time prediction,” in *27th International Technical Conference on the Enhanced Safety of Vehicles (ESV) National Highway Traffic Safety Administration*, no. 23-0331, 2023.
- [22] Y. Xu, C. Huang, Y. Nan, and S. Lian, “TAD: A large-scale benchmark for traffic accidents detection from video surveillance.”
- [23] M. Liu, E. Yurtsever, J. Fossaert, X. Zhou, W. Zimmer, Y. Cui, B. L. Zagar, and A. C. Knoll, “A survey on autonomous driving datasets: Statistics, annotation quality, and a future outlook,” *IEEE Transactions on Intelligent Vehicles*, 2024.

- [24] A. Pirdavani, E. De Pauw, T. Brijs, S. Daniels, M. Magis, T. Bellemans, and G. Wets, “Application of a rule-based approach in real-time crash risk prediction model development using loop detector data,” vol. 16, no. 8, pp. 786–791. Publisher: Taylor & Francis .eprint: <https://doi.org/10.1080/15389588.2015.1017572>.
- [25] A. Zahid, T. Qasim, N. Bhatti, and M. Zia, “A data-driven approach for road accident detection in surveillance videos,” vol. 83, no. 6, pp. 17217–17231.
- [26] V. A. Adewopo and N. Elsayed, “Smart city transportation: Deep learning ensemble approach for traffic accident detection,” vol. 12, pp. 59134–59147. Conference Name: IEEE Access.
- [27] Y. Yao, X. Wang, M. Xu, Z. Pu, Y. Wang, E. Atkins, and D. J. Crandall, “DoTA: Unsupervised detection of traffic anomaly in driving videos,” vol. 45, no. 1, pp. 444–459. Conf. Name: IEEE Trans. on Pattern Analysis and Machine Intelligence.
- [28] D. Yang, Y. Wu, F. Sun, J. Chen, D. Zhai, and C. Fu, “Freeway accident detection and classification based on the multi-vehicle trajectory data and deep learning model,” vol. 130, p. 103303.
- [29] J. Fang, J. Qiao, J. Xue, and Z. Li, “Vision-based traffic accident detection and anticipation: A survey,” vol. 34, no. 4, pp. 1983–1999. Conference Name: IEEE Transactions on Circuits and Systems for Video Technology.
- [30] B. Maaloul, A. Taleb-Ahmed, S. Niar, N. Harb, and C. Valderrama, “Adaptive video-based algorithm for accident detection on highways,” in *2017 12th IEEE International Symposium on Industrial Embedded Systems (SIES)*, pp. 1–6. ISSN: 2150-3117.
- [31] T. Wang, S. Kim, J. Wenxuan, E. Xie, C. Ge, J. Chen, Z. Li, and P. Luo, “Deepaccident: A motion and accident prediction benchmark for v2x autonomous driving,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, pp. 5599–5606, 2024.
- [32] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, “CARLA: An open urban driving simulator,” in *Proceedings of the 1st Annual Conference on Robot Learning*, pp. 1–16, PMLR. ISSN: 2640-3498.
- [33] R. Greer and M. Trivedi, “Towards explainable, safe autonomous driving with language embeddings for novelty identification and active learning: Framework and experimental analysis with real-world data sets,” *arXiv preprint arXiv:2402.07320*, 2024.
- [34] R. Greer, S. Desai, L. Rakla, A. Gopalkrishnan, A. Alofi, and M. Trivedi, “Pedestrian behavior maps for safety advisories: Champ framework and real-world data analysis,” in *2023 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1–8, IEEE, 2023.
- [35] G. Jocher, A. Chaurasia, and J. Qiu, “Ultralytics yolov8,” 2023.
- [36] CVAT.ai Corporation, “Computer vision annotation tool (CVAT).” original-date: 2018-06-29T14:02:45Z.

- [37] W. Zimmer, A. Rangesh, and M. Trivedi, “3d bat: A semi-automatic, web-based 3d annotation toolbox for full-surround, multi-modal data streams,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1816–1821, IEEE, 2019.
- [38] C. Creß, W. Zimmer, L. Strand, M. Fortkord, S. Dai, V. Lakshminarasimhan, and A. Knoll, “A9-dataset: Multi-sensor infrastructure-based dataset for mobility research,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*, pp. 965–970.
- [39] W. Zimmer, C. Creß, H. T. Nguyen, and A. C. Knoll, “TUMTraf intersection dataset: All you need for urban 3d camera-LiDAR roadside perception,” in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1030–1037. ISSN: 2153-0017.
- [40] C. Creß, W. Zimmer, N. Purschke, B. N. Doan, S. Kirchner, V. Lakshminarasimhan, L. Strand, and A. C. Knoll, “Tumtraf event: Calibration and fusion resulting in a dataset for roadside event-based and rgb cameras,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [41] W. Zimmer, G. A. Wardana, S. Sritharan, X. Zhou, R. Song, and A. C. Knoll, “Tumtraf v2x cooperative perception dataset,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22668–22677, 2024.
- [42] X. Zhou, D. Fu, W. Zimmer, M. Liu, V. Lakshminarasimhan, L. Strand, and A. C. Knoll, “WARM-3d: A weakly-supervised sim2real domain adaptation framework for roadside monocular 3d object detection.”
- [43] W. Zimmer, C. Creß, X. Zhou, and A. Knoll, “TUM traffic dataset development kit.” <https://github.com/tum-traffic-dataset/tum-traffic-dataset-dev-kit>.
- [44] X. Zhou, M. Liu, E. Yurtsever, B. L. Zagar, W. Zimmer, H. Cao, and A. C. Knoll, “Vision language models in autonomous driving: A survey and outlook,” pp. 1–20. Conference Name: IEEE Transactions on Intelligent Vehicles.
- [45] R. van Kempen, “AUTOtech.agil : Architecture and technologies for orchestrating automotive agility.” Conference Name: 32. Aachen Colloquium Sustainable Mobility.