WARD: Weather-Aware Road Surface Condition Monitoring Dataset

Soňa Nesnídalová[1,2], Tomáš Kerepecký[1], and Adam Novozámský[1(✉)]

[1] The Czech Academy of Sciences, Institute of Information Theory and Automation, Prague, Czechia
novozamsky@utia.cas.cz

[2] Czech Technical University in Prague, Faculty of Nuclear Sciences and Physical Engineering, Prague, Czechia

Abstract. Road surface condition (RSC) monitoring is essential for enhancing vehicle safety and accident prevention. This study investigates the application of computer vision techniques for real-time sensing of road surface conditions. We introduce a novel dataset named WARD (Weather-Aware Road Dataset), a comprehensive collection of almost 55 000 images collected in real-world driving scenarios across diverse seasonal and weather conditions, designed to advance RSC detection, now available for download. We thoroughly evaluate state-of-the-art computer vision models, specifically MobileNet and EfficientNet, on both the WARD and publicly available RoadSaW datasets, providing insights into their classification performance. MobileNet exhibited superior classification and inference speed results, processing images at up to 30 fps on an affordable GPU. To improve real-time efficiency, we employ temporal smoothing through moving window aggregation. Our findings validate the potential of non-contact, camera-based RSC monitoring, showcasing its practicality and cost-effectiveness compared to other sensors.

1 Introduction

Harsh weather, especially heavy rain or snow, adversely affects driver capabilities, vehicle handling, and the condition of road infrastructure. Several studies have shown that these conditions negatively influence the roadway capacity, speed, and traffic density and increase the risk of traffic accidents [1].

The road friction coefficient, which measures the traction between tires and the pavement, is one of the most critical elements for understanding and analyzing driving safety. Low-friction surfaces significantly increase the risk of accidents by reducing brake effectiveness and causing high-side slips. Various factors influence the friction coefficient, with weather conditions such as rainfall, snow, and ice having the most significant impact. Monitoring surface conditions offers valuable information that could be used as a reference to understand vehicle braking behavior, enabling more effective safety interventions.

In this context, road surface condition (RSC) sensing serves as an essential tool for detecting states that affect friction coefficients, providing key insights

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
E. Rodolà et al. (Eds.): ICIAP 2025, LNCS 16168, pp. 164–175, 2026.
https://doi.org/10.1007/978-3-032-10192-1_14

into vehicle braking dynamics [15]. Consequently, both active safety systems and self-driving vehicles could benefit greatly from the capability to predict such conditions in real-time, enhancing safety and performance [18].

Several sensors have been developed for RSC determination in contact and non-contact forms. Contact sensors face limitations due to their short service life, high installation costs, maintenance requirements, and the inability to provide real-time measurements. Non-contact methods, on the other hand, are portable and support real-time data acquisition [12]. Even though many sensors for RSC detection are already available, utilizing a front-facing in-vehicle camera remains the most accessible non-contact option since it often comes as part of basic vehicle hardware [15]. While a front-facing camera is a cost-effective and widely available sensor, camera-based methods inherently depend on robust, labeled image datasets to train machine learning models for road surface condition detection. Consequently, the availability of such datasets becomes a critical factor in advancing this approach, yet there is a notable lack of publicly available computer vision datasets that explicitly include road surface condition labels. In the absence of dedicated RSC datasets, some studies aim to estimate weather conditions from a wider field of view or even the entire image; however, these approaches often capture the overall atmospheric state and may not accurately reflect the actual condition of the road surface itself [10,16,20]. To our knowledge, only two datasets currently provide labeled data for RSC classification. However, even these exhibit some limitations. For the RoadSaW [4] dataset, the lack of diversity in surface textures, environmental conditions, and real-world anomalies arises from the recordings being captured exclusively on a test track. In RSCD [24] dataset, the preprocessing techniques disregard a significant portion of the lane in front of the vehicle, limiting the view to a small cut-out where the tires would presumably pass.

To address the lack of comprehensive, publicly available datasets and to advance real-time camera-based RSC detection, we make the following contributions: (1) We introduce a novel dataset named WARD (Weather-Aware Road Dataset), consisting of almost 55,000 images collected during real-world driving under diverse seasonal and weather conditions, which is available for download. (2) We comprehensively evaluate state-of-the-art computer vision models, focusing on MobileNet and EfficientNet, using both WARD and RoadSaW dataset. This analysis offers valuable insights into their effectiveness in road surface classification tasks. (3) We employ effective data preprocessing techniques, including region-of-interest extraction and bird's-eye view transformations, and implement temporal smoothing through moving window aggregation to enhance real-time performance. (4) We benchmark the models on their accuracy and inference speed, demonstrating the suitability of MobileNet for high frame rate (30 fps) processing in real-time applications.

2 Computer Vision Datasets for RSC Detection

Only a few autonomous driving datasets provide RSC labels. Benchmark datasets such as KITTI [9] and CityScapes [5], widely used for object detec-

tion, primarily focus on urban driving scenarios and lack detailed annotations for weather or surface conditions. However, datasets like BDD100K [23], RobotCar [2], Ithaca365 [7], and ACDC [19] contain certain weather-related annotations for recorded drives. Previous studies, such as [8] or [13], manually labeled these public benchmarks and combined them with self-recorded footage for the purposes of RSC detection. A different approach to creating such dataset was taken in [1,14,17] in the form of roadside camera footage processing instead of using in-vehicle cameras. Other strategies include creating dataset from YouTube videos in [18] or utilizing vehicle response during driving to assign friction to recorded images in [22]. To our knowledge, only two publicly available datasets, RoadSaW [4] and RSCD [24], include explicit surface condition labels. RoadSaW was created with the objective of developing RSC detection techniques. It provides detailed annotations for road surface types (*asphalt*, *concrete* and *cobblestone*) and conditions (*dry*, *damp*, *wet*, and *very wet*). The ground truth provided in terms of wetness estimation is determined by the Mobile Advanced Road Weather Information Sensor MARWIS based on infrared spectroscopy. The recordings included in this dataset were taken over a span of 10 days using cameras mounted on vehicles (a truck or a car) and the reference sensor. The wet conditions were simulated using sprinklers, and the final wetness classes were based on water film height thresholds established for each surface type. This controlled setting allowed the authors to create a balanced dataset with images split nearly equally into each of the classes. The downside of this approach is that data collected exclusively from test track recordings lacks diversity in surface textures, environmental conditions, lighting variations, and real-world anomalies, which limits the dataset's robustness and applicability to varied, real-world scenarios. The dataset can also be merged with a compatible balanced snow coverage dataset, RoadSC [3], which provides three types of snow: *fresh fallen*, *fully packed*, and *partially covered*. The authors in [4] state that the dataset is the first to provide accurate water film height measurements combined with road surface types.

RSCD offers manually created annotations for road surface friction (*dry*, *wet*, *water*, *fresh snow*, *melted snow*, *ice*), unevenness (*smooth*, *slight uneven*, *severe uneven*), and material (*asphalt*, *concrete*, *mud*, *gravel*) across 1 million image samples, gathered in Beijing between January and July 2022, spanning approximately 700 km of roads, captured by a camera mounted on the vehicle's bonnet. These are provided in the form of 360×240 pixel road patches cut from the original images covering the area through which, the authors claim, the tires would pass. The preprocessing techniques significantly limit the view to a small cut-out, disregarding most of the lane in front of the vehicle. Table 1 outlines the aspects that could be considered when selecting datasets for RSC detection. Our self-recorded dataset (WARD), created specifically for this study, is also included for comparison and is described in detail in the following section.

As will be further discussed in Sect. 4, RoadSaW was selected for initial experiments in training our networks from the described public datasets due to its balanced design and highly accurate labeling. RSCD provides only small

Table 1. Comparison of public computer vision datasets employed in automotive tasks with regards to the RSC detection task.

Dataset	BDD100K	KITTI	RobotCar	CityScapes	Ithaca365	ACDC	RoadSaW	RSCD	WARD
Images	12 000 000	14 999	N/A	5 000 (+2 000)	>680 000	8 012	20 244	1 000 000	54 952
Sequences	100 000	22	>100	50	40	N/A	>250	N/A	1 854
Diverse RSC	✓	N/A	✓	×	✓	✓	✓	✓	✓
Diverse Locations	✓	✓	×	✓	×	✓	×	✓	✓

image patches focused on the area, where the vehicle's tires would pass. In contrast, the trapezoidal region of interest used in RoadSaW captures the entire area the vehicle will traverse, offering a more comprehensive representation of road conditions. The snow coverage set provided by the authors has three snow class labels *fresh fallen*, *fully packed* and *partially covered* which were merged into one snow class. In addition, the snow data had to be subsampled from existing images to maintain the balance of the dataset.

2.1 WARD Dataset

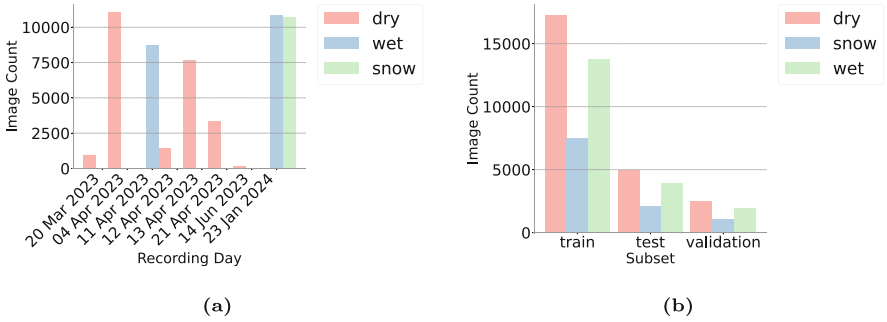

Fig. 1. Image count statistics for WARD. (a) Count of images extracted from the drives on the recorded dates on the x-axis, labeled with the assigned RSC classes. (b) Count of images belonging to the different classes in each of the subsets of the dataset.

To address the limitations of other existing RSC datasets, we introduce WARD, a novel publicly available dataset, to support advancements in road surface condition detection. Our dataset consists of 66 recordings captured by ZED 2 stereo camera mounted behind the rear-view mirror inside the test vehicle with a resolution of 1920×1080 pixels. The camera is factory-calibrated, and its calibration parameters are accessible through the official SDK. The choice of RSC classes for this task was primarily motivated by the objective of recognizing road conditions based on surface state rather than surface type. As a result, we defined three target classes (wet, dry, and snow) regardless of the underlying surface material,

allowing the models to focus on the core aspect of the task we aim to investigate. The images were acquired during multiple seasons of the year, providing variability in lighting and weather conditions. Figure 1a shows the recording dates and the number of images captured throughout the various drives within them, while the color of the bars indicates the class to which they belong in terms of RSC. The featured locations span across multiple smaller towns in Czechia, including some connecting roads, as well as various locations within the city of Prague and its suburban areas. This introduces a great deal of variability in vehicle speed, type of road surface, and road conditions in terms of damage or other imperfections. Since this dataset was created in actual driving conditions (unlike RoadSaW), it also contains samples with obstacles on the road, parts of vehicles, people blocking small sections of the image, shade, or traces of the sidewalk (see these examples in Fig. 2).

Fig. 2. Samples of real-world driving scenarios in WARD featuring speedbumps, traces of sidewalks, shadows, partially visible vehicles, and pedestrian crossings.

From the recordings provided at a rate of 15 fps, the individual frames were extracted and hand-labeled according to the RSC, disregarding the instances where the condition could not be determined with high confidence. The annotator had access to the complete driving scene, with the images arranged in the same chronological order as they appeared in the original videos. In addition to this knowledge, the annotator could also consult the individuals who were present during the data collection on the recording dates, as they were able to provide further context on the weather conditions present, if such information was needed. Since some locations were within cities and towns, frames where the car was still or barely moving were also eliminated, preventing the final dataset from containing multiple images from one spot. After this process, a trapezoidal region of interest was extracted from the labeled images of the road scenes. The trapezoid dimensions and location were determined by visual inspection of the images. It was positioned to approximately match the width of the vehicle's lane, with its height covering a significant portion of the road ahead. Care was taken to avoid making it too large, as this could introduce unwanted noise, such as other vehicles, pedestrians, or, depending on the driving direction, the roadside. In addition to cropping the images, homography mapping was used to obtain a bird's eye view of the road surface (see Fig. 3).

After preprocessing the recordings, the resulting images were split into train, test, and validation sets where the training set accounted for 70%, test 20%, and validation 10% of the data. However, this task was challenging because the

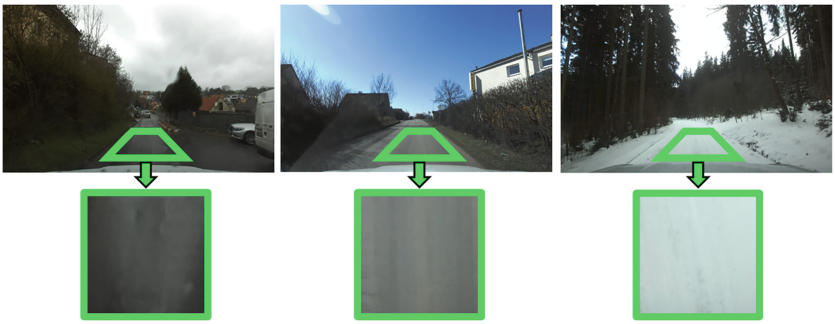

Fig. 3. Example of preprocessing of the recorded data.

recordings varied in length, and the classes were unevenly distributed, making stratified splitting unfeasible. Therefore, all recordings were further split into shorter sequences of approximately 30 frames with a 20-image gap (corresponding to almost 20 m when driving at the lower speed of 50 km/h) between each pair. Since the images show only the region of the road surface in front of the vehicle, disregarding the surroundings, this gap is assumed to be wide enough for the road to display different characteristics in terms of water and snow distribution on the surface, as well as different properties of the road surface itself and slightly changed lighting conditions. This allowed us to divide the data into a train set of 38 655, validation with 5 358, and a test set consisting of 10 939 images. The class counts in these subsets are visualized in Fig. 1b.

3 Methodology

To validate the effectiveness of our created WARD dataset and explore its applicability for real-time RSC detection, we developed a targeted methodology leveraging lightweight deep learning models. We chose to focus on recognizing surface states rather than surface types, defining a three-class classification task: *wet*, *dry*, and *snow*. By disregarding surface type, the models are designed to concentrate on detecting the road surface conditions most relevant to safety, ensuring a practical and efficient solution for real-time applications.

The selection of the CNN model was guided by the importance of inference speed and compact model size for this specific application. Both chosen network structures, *MobileNetV3-Large* [11] and *EfficientNetB0* [21], are created for mobile devices, making them a suitable choice for the target task. Additionally, both models have proven to be very effective in various computer vision problems. EfficientNet typically achieves higher accuracy in classification tasks but at the expense of slightly slower inference speed. Both CNN architectures are implemented in the PyTorch library.

Training. No models with pre-trained weights suitable for this application were found during this research. Therefore, a decision was made to initialize

the weights from the ImageNet-1K dataset [6]. Subsequently, both architectures have the classifier adjusted for the three-class classification task, and the output features of the CNN backbone were set to use dropout. In addition to this regularization technique, data augmentation is also utilized with random rotation in the range of $[-5°, +5°]$, horizontal flipping, brightness alterations, and a subtle added Gaussian blur. Different combinations of these alterations and their magnitudes were also examined, providing less satisfactory results. Both networks were trained using SGD with momentum of 0.9 and a decrease in the learning rate was implemented to be triggered on the validation loss plateau. The loss function applied for training both of these networks was the weighted cross-entropy loss defined as:

$$L = -\sum_{m \in B} \sum_{i=1}^{C} w_i y_i^{(m)} \log(\hat{y}_i^{(m)}), \tag{1}$$

where w_i are class weights, y is the target, $\hat{y}$ output of classifier, C the number of classes and B the input batch. Due to the unbalanced nature of our dataset, the weights were calculated according to the following equation:

$$w_i = \frac{\text{total number of samples}}{\text{number of classes} \times \text{number of samples in class } i} \tag{2}$$

and for the balanced RoadSaW, they were kept at the value of one. Monitoring the validation loss, early stopping was also used to prevent the architectures from overfitting. This combination of regularization techniques showed the best results not only on the training but also on the validation set based on multiple trainings and evaluations.

Moving Window Aggregation. During video data analysis, it was observed that class transitions tend to be less abrupt and more continuous. Thus, moving-window aggregation of predictions in time was utilized to improve the model's performance. This approach involves synthesizing the classification results from a set of neighboring frames and smoothing the final results. Every new prediction is given as a majority class in the current window, which in this case was used at the length of five outputs. This window size would examine a video segment that spans approximately one-sixth of a second, assuming that the camera provides videos at 30 frames per second (fps). It was set small enough to preserve some temporal resolution of the data and to ensure that any rapid transitions in weather conditions can be detected without significant delay.

4 Results

This section presents the results of our experiments. An initial exploration of pre-training the models on publicly available RSC detection datasets, followed by fine-tuning on our dataset, yielded suboptimal performance and is not further discussed. Therefore, we focused on fine-tuning the selected architectures

pre-trained on ImageNet on WARD. The outcomes of our experiments are summarized in Table 2. Both tested architectures performed well on the RoadSaW dataset, achieving almost 96% accuracy across all subsets, with MobileNet delivering slightly better classification metrics. Given the challenges described earlier, both models were trained from scratch on WARD, and the results are also detailed in Table 2. Once again, the classification metrics confirmed the effectiveness of both models for the RSC detection task, closely aligning with the previously observed performance trends. With this data, we observed a validation accuracy and F1 score above 96% using MobileNet, whereas EfficientNet delivered just below this mark.

Table 2. Performance metrics of MobileNet and EfficientNet architectures trained and tested on the RoadSaW and WARD datasets. The table presents accuracy, precision, recall, and F1 scores across training, validation, and test subsets.

	Dataset	Subset	Accuracy	Precision	Recall	F1
MobileNet	RoadSaW	Train	98.08	98.08	98.08	**98.08**
		Test	96.92	97.10	96.92	97.00
		Validation	95.88	95.89	95.88	**95.88**
	WARD	Train	98.24	98.20	98.19	**98.19**
		Test	96.15	95.65	95.02	95.32
		Validation	96.60	96.06	96.02	**96.04**
EfficientNet	RoadSaW	Train	97.64	97.64	97.64	97.64
		Test	96.42	96.54	96.42	96.48
		Validation	94.86	94.88	94.86	94.87
	WARD	Train	97.41	97.51	97.33	97.42
		Test	94.89	94.49	94.14	94.31
		Validation	95.88	95.61	95.48	95.54

Misclassifications were analyzed through confusion matrices, as illustrated on the corresponding test sets in Fig. 4. On RoadSaW, incorrect classifications occurred primarily between the wet and dry classes, which is expected since the snow coverage subset lacked slush or melting snow samples, making these classes easier to differentiate. However, both trained models seem less successful in distinguishing wet and snowy conditions on WARD data, where the classes contain these instances. This makes discriminating the two classes significantly more challenging, as most of the misclassified samples were of the surface covered with slushy, melting snow. Upon closer examination of the incorrect predictions, the misclassifications between wet and dry classes observed on WARD dataset were primarily caused by imperfections on the road surface, which the model likely misinterpreted as water. For RoadSaW, where such surface imperfections were rare or nonexistent, errors due to motion blur, which the model may have

perceived as reflections typically associated with a wet surface, were more common.

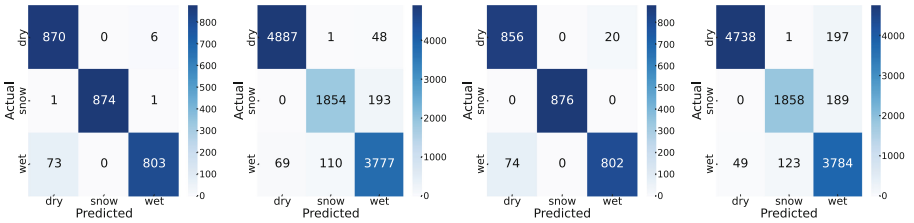

(a) MobileNet on RoadSaW **(b)** MobileNet on WARD **(c)** EfficientNet on RoadSaW **(d)** EfficientNet on WARD

Fig. 4. Confusion matrices for test subsets across different models and datasets.

Moving Window Aggregation. Employing the moving window aggregation has shown to be effective in improving the classification results on the test sets of both datasets as can be seen in Table 3. Even with a small window of 5 predictions, the trained MobileNet architecture achieved test accuracy and an F1 score of almost 98% on RoadSaW. The impact of errors caused by delays from aggregation is more pronounced in our dataset due to its class imbalance, which leads to overlapping classes within the evaluated batches. Nonetheless, all the evaluated metrics show slight improvements, indicating that window aggregation could effectively improve the RSC detection results.

Table 3. Results of both models on the test sets using moving window aggregation.

	Dataset	Aggregation	Accuracy	F1
MobileNet	RoadSaW	×	96.92	97.00
		✓	97.72	97.78
	WARD	×	96.15	95.32
		✓	96.54	95.62
EfficientNet	RoadSaW	×	96.42	96.48
		✓	97.49	97.57
	WARD	×	94.89	94.31
		✓	95.93	95.18

Prediction Speed. For this particular application, the prediction speed of the chosen models is also taken into account. The inference times reported in Table 4 are given as a mean of 200 predictions on GPU NVIDIA GeForce RTX 3050 Ti and a CPU AMD Ryzen 7 5800U. With the ability to process an image in 23 ms

on GPU, MobileNet shows promise for real-time applications, including videos running at 30 frames per second in this context. With slower inference, even on GPU, EfficientNet is still very effective and can provide results in real-time on both CPU and GPU with 15 fps. Although both architectures provide very fast inference, only MobileNet has the potential to be used at a 30 fps rate.

Table 4. Inference speed on GPU and CPU for both networks given as a mean of 200 prediction times on WARD.

Architecture	Inference Time [ms]	
	GPU	CPU
EfficientNet	35.53	55.54
MobileNet	23.20	40.00

5 Conclusion

This study addressed the issue of weather-aware road surface condition monitoring, an important factor in enhancing vehicle safety and accident prevention. We introduced the WARD dataset, a comprehensive collection of real-world driving scenarios across diverse seasonal and weather conditions designed to advance RSC detection using computer vision. The dataset consists of 66 recordings during eight recording dates, split into short 30-frame drives and reduced to road surface cut-outs located in front of the vehicle transformed into a top view. The splitting of the recordings was a necessary step in the creation process, as uneven class distribution made it impossible to perform a structured split for training, validation, and testing. The WARD dataset is publicly available to support further advancements in RSC sensing research.

Through comprehensive experiments on both the RoadSaW dataset and the newly created WARD dataset, MobileNet and EfficientNet demonstrated strong performance. MobileNet excelled in both accuracy and inference speed, making it particularly suitable for real-time applications. Additional evaluation was provided for both models, employing window aggregation of predictions to smooth out the final results. This approach is in line with the gradual nature of changes in road surface conditions in real-world driving scenarios and was shown to improve the accuracy of both models.

Our findings validate the potential of non-contact, camera-based RSC monitoring, showcasing its practicality and cost-effectiveness compared to traditional methods. Integrating innovative preprocessing and temporal smoothing techniques further underscores the feasibility of real-time applications in modern vehicles.

Among the shortcomings of our dataset are its restriction to a single country and the use of ROI settings specifically fine-tuned to our setup, which may limit

its generalizability. Future work may include expanding data collection across countries, optimizing ROI dimensions, or exploring quantization to improve robustness, transferability, and efficiency for deployment on resource-constrained systems.

The WARD dataset was specifically prepared for the research community to support the training and testing of their models and is publicly available for scientific purposes on Kaggle - https://adamnovozamsky.com/ward.

Acknowledgments. This study was funded by the Czech Science Foundation, grant no GA24-10069S.

Disclosure of Interests. The authors have no competing interests to declare relevant to this article's content.

References

1. Abdelraouf, A., Abdel-Aty, M., Wu, Y.: Using vision transformers for spatial-context-aware rain and road surface condition detection on freeways. IEEE Trans. Intell. Transp. Syst. **23**(10), 18546–18556 (2022). https://doi.org/10.1109/TITS.2022.3150715
2. Barnes, D., Gadd, M., Murcutt, P., Newman, P., Posner, I.: The oxford radar robotcar dataset: a radar extension to the oxford robotcar dataset. arXiv preprint arXiv: 1909.01300 (2019)
3. Cordes, K., Broszio, H.: Camera-based road snow coverage estimation. In: 2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pp. 4013–4021 (2023). https://doi.org/10.1109/ICCVW60793.2023.00433
4. Cordes, K., Reinders, C., Hindricks, P., Lammers, J., Rosenhahn, B., Broszio, H.: Roadsaw: a large-scale dataset for camera-based road surface and wetness estimation. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 4439–4448 (2022). https://doi.org/10.1109/CVPRW56347.2022.00490
5. Cordts, M., et al.: The cityscapes dataset for semantic urban scene understanding. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 248–255. IEEE (2009). https://doi.org/10.1109/CVPR.2009.5206848
7. Diaz-Ruiz, C.A., et al.: Ithaca365: Dataset and driving perception under repeated and challenging weather conditions (2022)
8. Fink, D., Busch, A., Wielitzka, M., Ortmaier, T.: Resource efficient classification of road conditions through cnn pruning. IFAC-PapersOnLine **53**(2), 13958–13963 (2020). https://doi.org/10.1016/j.ifacol.2020.12.913, 21st IFAC World Congress
9. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. Int. J. Robot. Res. (IJRR) (2013)
10. Gragnaniello, D., Greco, A., Sansone, C., Vento, B.: Real-time joint recognition of weather and ground surface conditions by a multi-task deep network. Eng. Appl. Artif. Intell. **139**, 109543 (2025)

11. Howard, A., et al.: Searching for mobilenetv3 (2019)
12. Ishii, Y., Fuchida, M., Ozeki, K., Nakamura, A.: Investigation on the use of polarized images for frozen road surface recognition. In: 2025 IEEE/SICE International Symposium on System Integration (SII), pp. 1581–1586 (2025). https://doi.org/10.1109/SII59315.2025.10870979
13. Karunasekera, H., et al.: Varying road surface condition estimation in ego and adjacent lanes. In: 2023 IEEE Intelligent Vehicles Symposium (IV), pp. 1–6 (2023). https://doi.org/10.1109/IV55152.2023.10186540
14. Khan, M.N., Ahmed, M.M.: Weather and surface condition detection based on road-side webcams: application of pre-trained convolutional neural network. Int. J. Transp. Sci. Technol. **11**(3), 468–483 (2022). https://doi.org/10.1016/j.ijtst.2021.06.003
15. Ma, Y., Wang, M., Feng, Q., He, Z., Tian, M.: Current non-contact road surface condition detection schemes and technical challenges. Sensors **22**(24) (2022). https://doi.org/10.3390/s22249583
16. Ojala, R., Alamikkotervo, E.: Road surface friction estimation for winter conditions utilising general visual features (2024). https://arxiv.org/abs/2404.16578
17. Ramanna, S., Sengoz, C., Kehler, S., Pham, D.: Near real-time map building with multi-class image set labeling and classification of road conditions using convolutional neural networks. Appl. Artif. Intell. **35**(11), 803–833 (2021). https://doi.org/10.1080/08839514.2021.1935590
18. Roychowdhury, S., Zhao, M., Wallin, A., Ohlsson, N., Jonasson, M.: Machine learning models for road surface and friction estimation using front-camera images. In: 2018 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2018). https://doi.org/10.1109/IJCNN.2018.8489188
19. Sakaridis, C., Dai, D., Van Gool, L.: ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (10 2021)
20. Samo, M., Mafeni Mase, J.M., Figueredo, G.: Deep learning with attention mechanisms for road weather detection. Sensors **23**(2) (2023). https://doi.org/10.3390/s23020798, https://www.mdpi.com/1424-8220/23/2/798
21. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks (2020)
22. Vosahlik, D., Cech, J., Hanis, T., Konopisky, A., Rurtle, T., Svancar, J., Twardzik, T.: Self-supervised learning of camera-based drivable surface friction. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), pp. 2773–2780 (2021). https://doi.org/10.1109/ITSC48978.2021.9564894
23. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning (2020)
24. Zhao, T., He, J., Lv, J., Min, D., Wei, Y.: A comprehensive implementation of road surface classification for vehicle driving assistance: Dataset, models, and deployment. IEEE Trans. Intell. Transp. Syst. **24**(8), 8361–8370 (2023). https://doi.org/10.1109/TITS.2023.3264588