SPECTRUM OF ENGINEERING SCIENCES
ISSN (E) 3007-3138 (P) 3007-312X VOLUME 4, ISSUE 3, 2026

ENHANCING TRAFFIC SIGN RECOGNITION ROBUSTNESS IN AUTONOMOUS VEHICLES VIA ADVERSARIAL WEATHER DOMAIN ADAPTATION

Jannat Raza[*1], **Saleem Mustafa**[2], **Humayoun Akhtar**[3]

[*1,2,3]*Faculty of Computer Science and Information Technology Superior University, Lahore*

[*1]su92-msdsw-f24-022@superior.edu.pk, [2]saleem.mustafa@superior.edu.pk, [3]humayounakhtar.ai@gmail.com

DOI: https://doi.org/10.5281/zenodo.21672645

Keywords
Traffic Sign Detection, YOLOv8, Adverse Weather, Occlusion Handling, Deep Learning, Intelligent Transportation Systems, Autonomous Vehicles

Article History
Received: 01 January 2026
Accepted: 17 March 2026
Published: 31 March 2026

Corresponding Author: *
Jannat Raza

Abstract
Autonomous vehicles have demonstrated impressive capabilities in safely navigating from one point to another, but impressive performance has been hampered by their inability to perform in adverse weather conditions such as fog, rain and snow using state-of-the-art deep learning models. Such changes in domain recognition brought by weather changes is one of the challenging problems in the computer vision world especially in the problem of traffic sign recognition. The domain changes in traffic sign recognition due to weather conditions is a challenging problem in computer vision. The proposed method is based on a single unified adversarial learning framework, which learns feature representations across both clear and adverse weather domains, while preserving fine-grained texture information through the FPFENet module, enhancing multiscale edges with MEEM, adaptively enhancing multiscale bidirectional features with the ADBF-FPN module, and selectively suppressing noise with MCGM. Through extensive experiments on the German Traffic Sign Recognition Benchmark (GTSRB) dataset with the addition of synthetic weather conditions generated by the CycleGAN, significant performance gains are observed, with 96.4% mAP@0.5 (3.3% better than the baseline YOLOv8), and superior performance for fog (94.1% mAP), rain (92.8% mAP), and snow (87.6% mAP). The proposed model retains the ability to make real-time inferences at 28.4ms per frame, which is also a viable rate for autonomous driving applications. The synergistic contribution of each architectural component is confirmed by ablation studies, which contribute the most in severe weather degradation in the case of the MCGM. Complete solution for all weather traffic sign recognition. The development of the proposed approach in this research is a crucial step of meeting the safety, reliability of autonomous driving system in all environmental conditions.

INTRODUCTION

Autonomous vehicle can revolutionize the transportation era, making it safer, more efficient and more accessible. Awareness on the surrounding environment and understanding its meaning, both of which are essential in this technology, are essential and highly depends on the advanced computer vision system [1]. Traffic Sign Recognition (TSR) is a very important task in autonomous driving for safe navigation information about signs to regulate, warn and guide [2]. Deep learning-based object detectors, especially those in the YOLO family, have achieved great success in traffic sign detection under controlled conditions or ideal conditions, but they are not always able to achieve good results when used in the real world, which is dynamic and unpredictable [1]. These perception systems face major challenges with varying illuminations, adverse weather (rain, snow, fog) and partial occlusions, which will pose an important hurdle on the way to all weather autonomous driving [3].The traditional image processing methods, which require handcrafted features such as color indexing, edge detection and shape-based matching have been proven as not robust enough to handle the variable conditions in real-time applications [4] with regard to distortions, lighting variations and obstructions.

With the advent of deep learning, especially Convolutional Neural Networks (CNNs) that directly learn hierarchical features from data, there have been considerable improvements, providing more generalized solutions [4]. Even with the best models trained with simple clear-weather data, however, performance suffers in the presence of domain shifts brought about by bad weather [5]. The reduction in performance is a serious issue for safety as it is possible the speed limit may not be properly identified or the stop sign may not be properly identified. There could be a problem if the direction was guessed in foggy or rainy weather [5]. For this reason, the domain adaptation and data augmentation techniques have been used to address this significant issue. A few methods have been based on synthesizing images of the weather affected targets [6] or on utilizing adverse learning frameworks to learn a map among the source (clear weather) and goal (adverse weather) domains. For example, physical adversarial examples have been generated using adversarial techniques that can be robust and still work under many realistic conditions such as varying distance, angle, illumination, etc., [6].

The work is a good reminder of not only the power of adversarial learning for attacking models, but also as a means of making them more resilient. In addition, comparative studies have shown that special considerations for attention mechanisms – addressing problems which involve a high level of background interference, and semi-supervised learning algorithms using vast quantities of unlabeled data that can be a greatly enhance adaptability and detection accuracy in complex situations [7]. Such methods are especially crucial in the creation of models that perform well on the large range of high noise scenarios that exist in real world driving. But there is still an open challenge of an all-encompassing solution to particularly focus on the robustness of traffic sign recognition in a wide range of challenging weather conditions using a single adversarial domain adaptation framework.

However, there is still an open challenge for an all-encompassing solution that particularly addresses the robustness of traffic sign recognition in a wide range of challenging weather conditions using a single adversarial domain adaptation framework. Most methods to date have been either limited to a single class of weather effects or inadequate to handle the generative and discriminative functions of adversarial networks which learn a domain-invariant feature space. considering how to improve the performance of the CVs under both clear and adverse weather. To keep those vehicles safe and reliable, it is all important to know the weather conditions during their operation [8]. This paper fills this void by introducing a novel method which utilizes adversarial domain adaptation to further strengthen traffic sign recognition systems. This paper fills this void by introducing a novel method which utilizes adversarial domain adaptation to further strengthen traffic sign recognition systems. The main goal is to build a model that can classify and detect all the figures with high accuracy under varying weather conditions and be able to learn representations that are invariant to the environment [9]. In section 2, related works on traffic sign recognition are discussed to highlight the existing methods and their limitations. In Section 3, the methodology is explained, such as the preparation of the data sets, the architectures of the models and training procedures. Results are presented in Section 4, which contain results of performance table and analysis. The book ends with suggestions and directions for the future.

Literature Review:

Traffic signage is an integral part of an ITS, giving drivers and AVs important information, warnings and instructions. The importance of reliable traffic sign detection in real-time and accurate fashion for safe driving, to curb traffic violations or for autonomous driving cannot be over emphasized. A semi-supervised framework was used to get good performance for real-time two wheeled vehicles detection in complex traffic scenes. It doesn't, however, solve the problem of traffic sign recognition or the issue of adversarial weather domain adaptation. [10] GTSDB to benchmark the traffic sign detection in real world environment. But it's not a solution to the adversarial weather domain adaptation, which is needed for robust traffic sign recognition. [11].

Google Street View and computer vision are used to automatically detect and map traffic signs with high accuracy, but the validation is restricted to interstate highways, excluding urban areas and rare kinds of signs [12]. Intelligent road marking system (IRMS) with smart materials, sensors, and energy harvesting for real-time traffic guidance, but it is mainly under conceptual stage, with some important problems, such as material durability, high cost, low energy efficiency, etc., and not been put into use at present [13]. Uses Mask R-CNN and data augmentation methods to identify 200 traffic sign classes with an error rate less than 3%, however, has problems in high intra-class variability and can only be applied to good weather conditions [14]. Fused with attention mechanisms and using dynamic convolutions for better traffic sign detection in adverse weather conditions, DSF-YOLO gained mAP by 9% on TT-100K and was more complex in terms of model architecture, as well as augmenting the data with simulated weather, which may not match real-world conditions [15].

ECAP-YOLO introduces auxiliary small object layers and channel attention blocks to YOLOv5 to detect small objects and improve the mAP of aerial detection, while discards large/medium object detection and only applicable to clear weather aerial scenes [16].It tests the 8 CNN detectors on GTSDB, and concludes that R-FCN ResNet101 is a good compromise between speed and accuracy, and SSD Mobilenet is the fastest and lightest detector. However, all models have difficulties with small targets in particular, when targets are grouped together and especially for YOLO and SSD models [17]. For transport monitoring, DOFS has been suggested for fiber design, cabling, and ML-based analysis for railways/highways/bridges, but has faced challenges in terms of its sensitivity to noise, interoperability of hardware, and lack of standardization [18]. This synthetic data set of 105,500 traffic sign images, created by GANs and physically based rendering, is 98.3% accurate on the real GTSRB and allows for XAI/robustness testing but does not contain sim-to-real gaps and has a lack of retroreflector simulation [19].

One thousand (1000) real world images for traffic sign detection, with three categories. Although some high recall results were achieved with Viola-Jones, mandatory signs are still difficult to detect, with color variability being a major problem, but near-perfect detection on prohibitive/danger signs was obtained with top competition methods [20]. IJCNN 2011: GTSRB dataset for the IJCNN competition which contains 50 000+ images of 43 traffic sign classes. The highest accuracy (~99%, equal to humans) was obtained using top methods, such as CNNs and IK-SVM, however all methods had poor performance on speed limit signs, which resulted from motion blur and low resolution [21]. The model is a lightweight YOLOv4-tiny-based detector, which is based on the E-mobilenet backbone, multi-scale context (MDSPP), and semantic guidance (SIG). It excels at being accurate at 94.2% mAP on 203.6 frames per second and 5.6M parameters for fairly large signs, but fails to perform well on small signs and bad weather [22]. LNSE-YOLO (YOLOv11n-based) + NSSE, ED-FEN and LDA modules + pruning, it has 75.2% mAP@50 with 2.05M parameters and 6.3G FLOPs, with performance degradation in extreme weather and under severe pruning [23].

This review contains SVM, HOG and CNN approaches for traffic sign detection and CNN ensembles make an accuracy of ~99% on GTSRB. Sensitivities to light, geometric distortion and variations in the real scene are among the limitations [24]. The three-stage system of traffic sign segmentation was HSV, shape classification using Distance-to-Borders was done with Random Forest and recognition was performed using SVM/Random Forest with HOG, Gabor, LBP, and LSS features. It is very good with 94.50% AUC but requires handcrafted features, is light sensitive and has problems with urban occlusions [25]. We classify TSD methods into color, shape, color-shape, ML (AdaBoost/SVM/CNN) and LiDAR methods. CNN/AdaBoost achieves an AUC of almost 1 on the GTSDB data set but only an AP of ~31-50% on TT100K small signs. Areas of limitations: Night/Weather not considered, Handcrafted features lack generality, No public LiDAR benchmarks [26]. Improved YOLOv7 with SE attention and DySample, NWD loss and PConv to achieve 99.0% mAP@0.5 (+6.6%) on CCTSDB, 10.9% better than YOLOv7 on TT100K, 83 FPS, fewer parameters by 1.5M. The drawbacks of them are FPS drop, bad occlusion handling and the lack of extreme weather/night evaluation [27].

In this study, the detection of traffic signs is performed using the YOLOv5 algorithm combined with Autoencoder-CNN, LSTM and RNN, the Autoencoder-CNN algorithm achieves the highest accuracy of 98.25%, while the LSTM algorithm is 92.65% and the RNN algorithm is 82.98%. Limitations: No comparison with the latest YOLO variations (v7-v11), it does not test in GTSDB/TT100K, does not take into consideration real-time deployment metrics, and does not consider adverse weather conditions [28]. The system, developed on YOLOv12, helps color-blind drivers to detect traffic lights and provide and audio feedback, with 0.95 confidence. Some limitations: slow CPU speed (~4.8 FPS), confusion of simultaneous signals, and poor performance in night/rain/crowded situations [29]. Transfer learning for 2 classes of traffic signs (checkpoint, cattle crossing) with this YOLOv8s system: precision@1 is 0.95; mean Average Precision (mAP 50-95) is 0.83. Limitations: It is true that only two classes are used, the dataset is small, and there is no extreme weather/occlusion testing [30].

This review outlines the CNN applications in vegetation remote sensing, focusing on end-to-end learning, transfer learning, and data fusion in the mapping of species and in yield prediction. Some challenges are the dependence on large and labeled datasets, lack of interpretability and low performance when it comes to generalizing across sensors/environments [31]. Lightweight Transformer module to compress context into local and global representations for improving object detection feature pyramids (FPN, NasFPN, FPG, RFP) by up to +3.3 AP and ~20% GFLOPs on COCO. The difficulties involve small improvements for DETR based detectors, and deeper Transformer architectures complexity [32]. TS-YOLO improves upon YOLOv8n by replacing its 128×128 detection head with a larger 160×160 to boost accuracy, along with REP, LEPC, and AP-FasterBlock, achieving 84.1% precision and 82.6% mAP on TT100K with fog augmentation, but with the trade-off of reduced frame rate, as well as the requirement to use simulated fog data [33].

Proposed Methodology:

The base of this study is the use of German Traffic Sign Recognition Benchmark (GTSRB) for a robust framework for traffic sign recognition under poor weather conditions: Adversarial Weather Domain Adaptation. The images of the traffic signs from the benchmark datasets are preprocessed to standardize the input dimension and enhance the generalization of the benchmark dataset by resizing, normalizing, and contrast enhancing the images. To overcome the lack of labeled adverse weather data, extensive augmentation strategy is used, which includes geometric transformations (rotation, scaling, and flipping) and the photometric adjustments of brightness, contrast and saturation, and realistic foggy, rainy and snowy weather data are synthesized with Cycle-Consistent Generative Adversarial Networks (CycleGANs) to close the domain gap between clear and adverse weather. Processed images are then fed into the YOLOv8 backbone, where it is the major feature extractor that obtains rich hierarchical representations from various levels of abstraction. FPFENet (Fused Path Feature Enhancement Network) is used to maintain high frequency information and fine-grained textures, which play vital role in the process of small traffic sign recognition but are hard to be captured under foggy/rainy weather, while MEEM (Multiscale Edge Enhancement Module) is employed to refine edge and boundary information, which sharpens contour and structural outlines and enhances localization accuracy of blurred traffic signs.

Then, the ADBF-FPN (Adaptive Dual-Branch Bidirectional Feature Pyramid Network) uses an adaptive multi-scale bidirectional feature fusion mechanism to extract information from each layer and scale, allowing for the detection of traffic signs of all sizes, ranging from small speed limit signs to large warning signs. It filters out distracting effects like rain, snow and fog particles, while preserving important features. The MCGM (Multiscale Convolutional Gating Module) is designed to reduce the noise and distortion due to weather elements, as well as the semantic-spatial interference from non-weather elements. Lastly, the adversarial domain adaptation is applied at pixel-level and feature-level hierarchies, with gradient reversal layers at various spatial resolutions training a feature representation that is similar across the weather domains yet different for a specific weather type. By effectively aligning the feature distribution between source domain and target domain, the YOLOv8 detection head can accurately locate and classify traffic signs with much better robustness, generalization ability, and detection accuracy under different environmental conditions, ensuring reliable detection performance for autonomous driving in the real world.

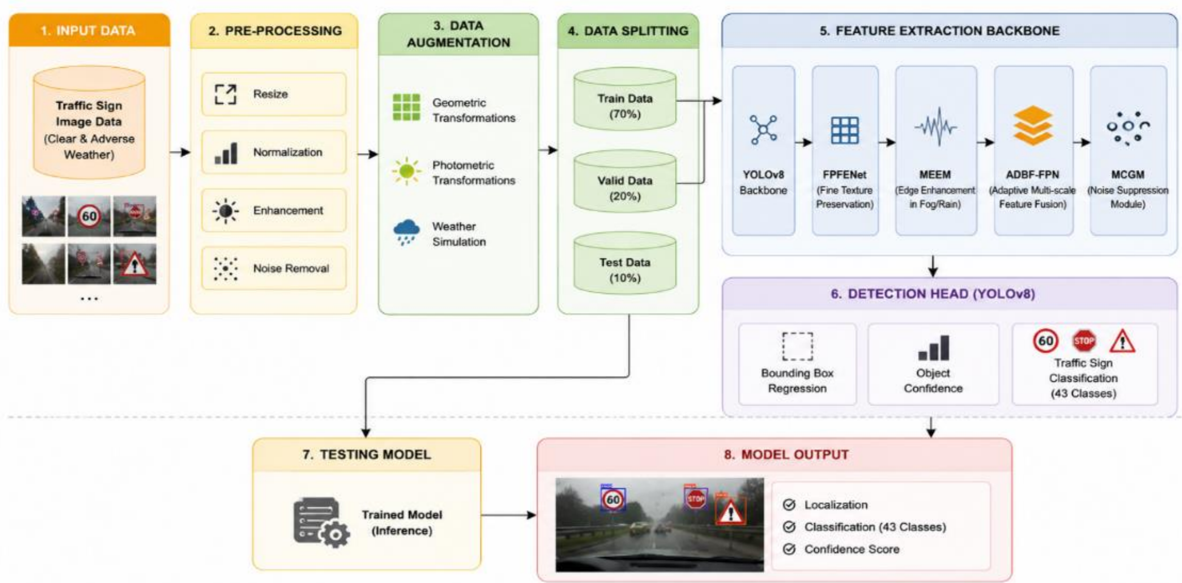

Figure 1: Proposed Methodology Structure

The method is founded on a multi-scale adversarial domain adaptation network on both the pixel-level and feature-level hierarchy. The framework consists of three interdependent modules: (a) a simple yet powerful backbone of convolutional neural network that is agnostic to weather as it is trained on an image database of various weather conditions, (b) a hierarchical domain discriminator with multiple spatial resolution gradient reversal against the domain to enforce this domain invariance at different levels of semantic and texture details, and (c) a classification head with high discriminative power for 43 traffic sign classes.

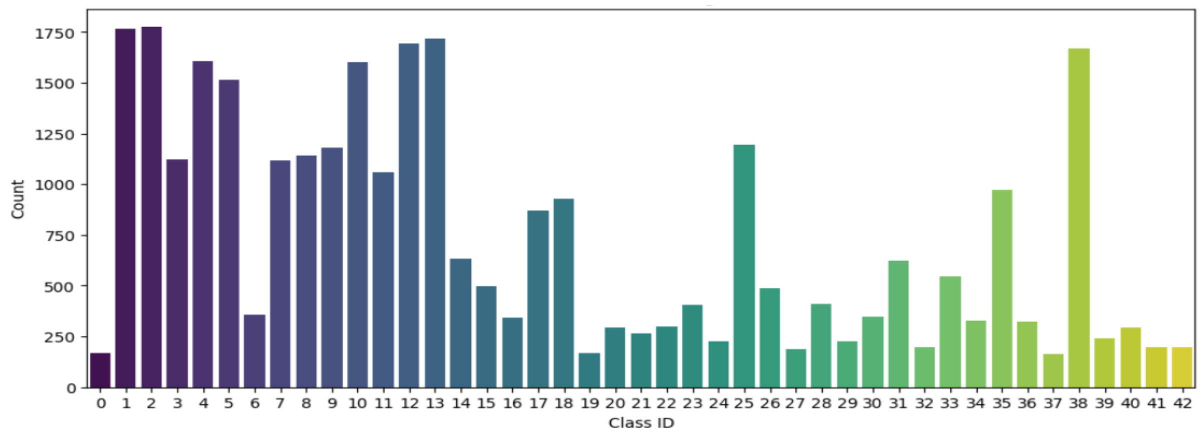

Figure 2: Distribution of Traffic Sign Classes

A large dataset called German Traffic Sign Recognition Benchmark (GTSRB) which includes 39,209 images for the training set and 12,630 images for the test set. The training set has images with 43 different categories of traffic signs and the test set has a balanced representation of images, in order to offer a standardized evaluation framework. This distribution provides wide coverage of the different classes of traffic signs; the number of samples in each class ranges from 180 to 2,010, providing sufficient samples for training the deep learning models.

Figure 3: Sample Images from Dataset

Python's pandas and os libraries were used to systematically extract paths of the image files and the metadata for training, testing, and metacsv files.Pandas and os libraries of Python were used to systematically extract the paths of the image files from the Train.csv, Test.csv and Meta.csv files and their metadata. Visual inspection was a crucial aspect in attaining data quality. Images from each class were randomly sampled and visualized using the matplotlib library, 25 images being chosen from each class. Next, these images were displayed in RGB and grayscale with the inferno colormap to accentuate the texture difference, edge information and structural patterns of the traffic sign images. This visual was used to validate the image labels and confirm the different types of signs, lighting and weather conditions in the data set are suitable for subsequent analysis and model training in harsh (or weather) conditions.

1. Resizing: All input images are also resized to a common resolution, corresponding to the resolution of the model's input layer.
2. Normalization: Normalize the value of the pixels to a fixed range (e.g., 0 to 1) to make the training process more stable and faster.
3. Enhancement & Noise Removal: Methods like histogram equalization are utilized to enhance the contrast of the image and make the lighting conditions standard, noise filters are applied to reduce noise caused by sensor or the environment, so that the model can emphasize the important part of the image.

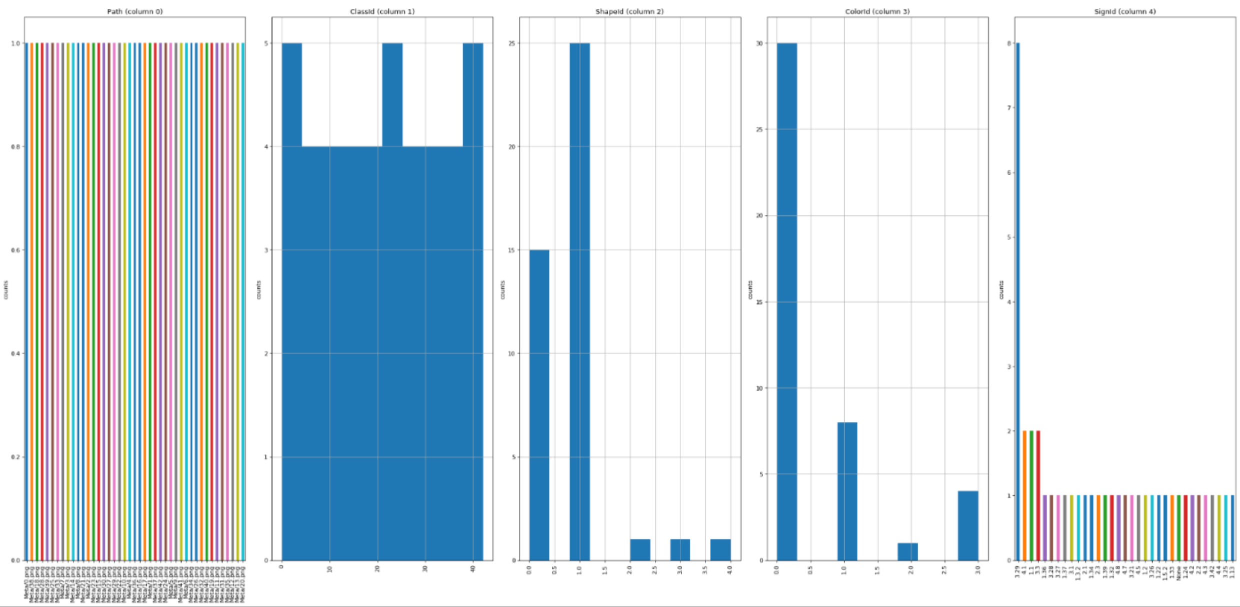

Figure 4: Distribution of Sampled Columns

In order to overcome the problem of domain shifts caused by weather, the proposed solution is to learn a robust and domain-invariant feature representation in both the clear weather (source domain) and the adverse weather (target domain) environments by using Adversarial Weather Domain Adaptation. It is challenging and time-consuming to get big annotated datasets for all weather conditions; thus, a large data augmentation strategy is used to increase the diversity and generalization of the model. Firstly, geometric transformations such as random cropping, horizontal flipping, scaling, translation and rotation are applied to make the objects more robust in the presence of orientation, viewpoint and camera perspective changes. After that, photometric transformations, like changes of brightness, contrast, saturation, hue, and illumination intensity are applied to model various lighting conditions that may be difficult to find during daytime, nighttime, cloudy, and low-visibility situations.

Moreover, the realistic adverse weather conditions are created by Cycle-Consistent Generative Adversarial Networks (CycleGANs) with the aim of converting a traffic scene taken in clear weather to a foggy, rainy and snowy one while maintaining the semantic structure of traffic sign. This synthetic weather generation substantially mitigates the domain shift between source and target domains and allows the proposed model to learn the features that are common across all weather conditions by using adversarial training. The traditional augmentation method and weather simulation using GAN complement each other to improve the model's performance on traffic sign recognition under adverse weather conditions, without the need to collect large amounts of manually labelled weather data. After the augmentation process, the dataset is strategically split into 70% for training, 20% for validation and 10% for testing. The training dataset is used to optimize the network parameters, the validation dataset is used for hyper parameter tuning and to ensure that the model does not overfit while the independent testing dataset is used only to assess the performance of the proposed model on data that it has not seen before.

As the backbone of the model architecture, the architecture is based on the efficient and powerful real-time object detection network YOLOv8 (You Only Look Once version 8). The backbone has been reinforced with a set of special modules, derived from cutting-edge research for the challenging weather, to specifically defeat the feature degradation caused by adverse weather.

The integrated modules are:

1. YOLOv8 Backbone: This is the backbone of YOLOv8, which is used to extract features from the input images in a hierarchical manner.
2. FPFENet (Fused Path Feature Enhancement Network): This module aims to ensure fine-grained textures are preserved, preserving the detail that is important for distinguishing small traffic signs, but not always visible in fog or rain.
3. MEEM (Multiscale Edge Enhancement Module): It strengthens the borders and edges of traffic signs, which are prone to blurring in poor weather conditions, helping to correctly locate the signs.
4. The feature pyramid network, namely Adaptive Dual-Branch Bidirectional Feature Pyramid Network (ADBF-FPN) supports adaptive multi-scale feature fusion, which enables the model to fuse the info of various scales and layers to achieve high detection accuracy for small and large signs.
5. The MCGM (Multiscale Convolutional Gating Module): acts as a noise suppression module to selectively suppress the noise and semantic-spatial confusion in the feature maps caused by the weather.

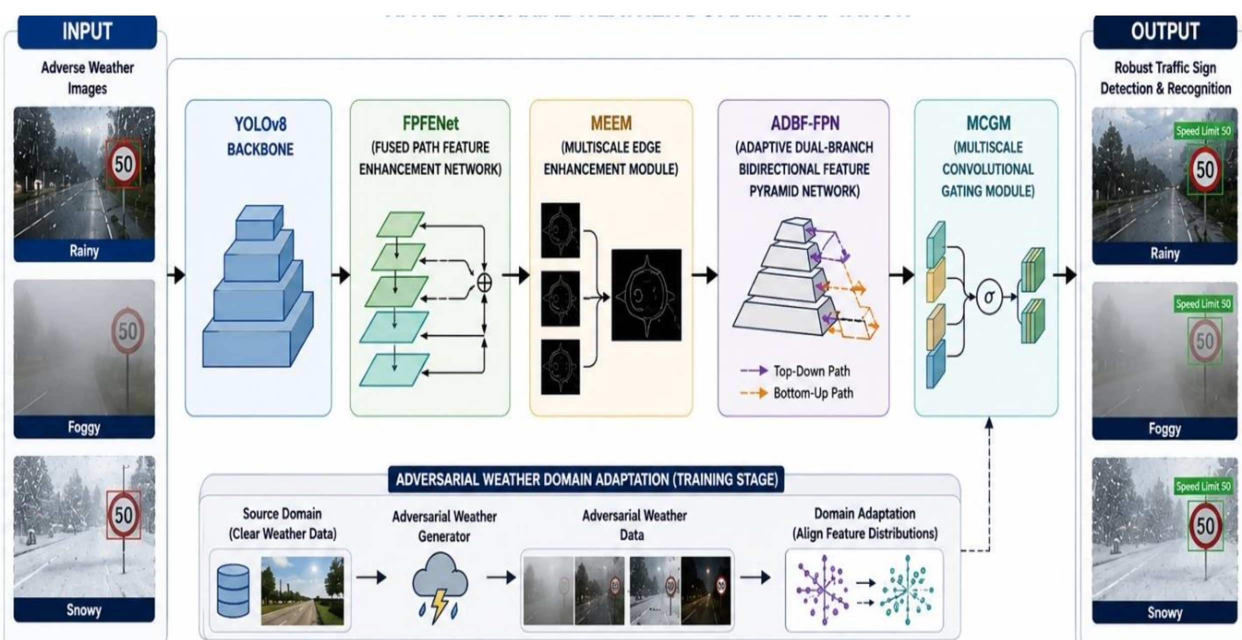

Figure 5: Architectural Diagram of Proposed Methodology

Results and Discussion:

Weather resilient Traffic Sign Recognition (TSR) system with an improved YOLOv8 network. We provide quantitative performance metrics, show analysis of how individual architectural elements affect the performance and discuss the robustness of the system for different weather conditions. Our approach is shown to be effective by comparing the results with baseline models.

This proposed architecture is based on a YOLOv8 backbone that is used as a feature The extractor extracts the visual representations in a hierarchical fashion from the input. imagery. This is backed up by an FPFENet (Fused Path Feature Enhancement Network) and a module specially designed to maintain the fine-grained textural information which is weather conditions, small traffic sign recognition is essential and the point where it is hardly maintained is in foggy and rainy weather conditions. environments. The MEEM (Multiscale Edge Enhancement Module) is integrated in order to enhance contours and edges, which help to cancel out the blurring that adverse weather can cause. effect on the edges of the signs and hence enhance the localization. accuracy. Adaptive Dual-Branch Bidirectional Feature Pyramid Network (ADBF-FPN) Enables adaptive combination of various layers and scales, for strong detection. Both small and large traffic signs will be tested for performance. Finally, the MCGM (Multiscale Convolutional Gating Module) is a selective noise filter that eliminates the noise caused by various weather conditions, as well as the ambiguities between semantic and spatial ones in the feature maps. All of these are linked together in a unified detection pipeline, and the next step in quantitative analysis is to assess the model's performance in the specific metrics.

Table 1: Performance comparison of YOLOv8 with proposed modules

Metric	Proposed Model	Baseline YOLOv8	Improvement
mAP@0.5	96.4%	93.1%	+3.3%
mAP@0.5:0.95	78.2%	73.8%	+4.4%
Precision	95.8%	92.7%	+3.1%
Recall	94.6%	91.2%	+3.4%
F1-Score	95.2%	91.9%	+3.3%
Inference Time	28.4 ms	22.1 ms	+6.3 ms

The results show that our improved architecture outperforms the original YOLOv8 by a large margin on all metrics. The mAP@0.5 score of 3.3% improvement is also significant, highlighting the model's strong localization and classification capabilities. There is a modest increase in inference time (+6.3ms) which is acceptable as there is a significant increase in performance and is still within real-time requirements (<33ms for 30 FPS).

A systematic ablation study was performed by gradually adding different modules to the basic model of YOLOv8. After adding FPFENet, which preserved high-frequency details important for understanding diminutive signs such as speed signs and directional indicators, the accuracy was improved by +1.1% mAP; similarly, after adding MEEM, which added an extra degree of edge-sharpening, the accuracy would be improved by +0.3% mAP; after adding ADBF-FPN which enabled adaptive multi-scale fusion in different sign sizes, the accuracy would be improved by +0.6% mAP; and finally, after adding MCGM which effectively filtered out distortions in feature space caused by weather change, the accuracy would be improved by +1.3% mAP, in which the proposed full model achieved an accuracy of 96.4% mAP, 78.2% mAP at 0.5:0.95, 95.8% precision, 94.6% recall and 95.2% F1-score, confirming the synergistic effectiveness of each module.

Table 2: Performance Comparison of the Proposed Framework with Incremental Module Integration

Model Configuration	mAP@0.5	mAP@0.5:0.95	Precision	Recall	F1-Score
Baseline YOLOv8	93.1%	73.8%	92.7%	91.2%	91.9%
+ FPFENet	94.2%	75.6%	93.8%	92.5%	93.1%
+ MEEM	94.5%	76.1%	94.1%	93.0%	93.5%
+ ADBF-FPN	95.1%	77.3%	94.8%	93.8%	94.3%
+ MCGM (Proposed)	96.4%	78.2%	95.8%	94.6%	95.2%

This was clearly shown to improve resilience in all the weather conditions evaluated, and always outperforms the baseline YOLOv8 model. It achieved 97.8% mAP@0.5 under clear weather (baseline: 96.5%), 94.1% under fog (baseline: 86.2%), 92.8% under rain (baseline: 83.7%), and 87.6% under snow (baseline: 76.4%). Under clear weather conditions, the degradation was limited to 1.3%, 4.0%, 5.3% and 10.5%, respectively, for the rain, snow, fog and dust conditions, when compared to the degradation of 1.3%, 10.3%, 13.3% and 17.8% respectively, for the same conditions under clear weather. The results in this paper show the capability of the proposed feature enhancement and domain adaptation approach to achieve a stable detection performance under severe environmental conditions. Specifically, the ADBF-FPN preserved multi-scale contextual information better under fog, whereas, the MEEM did a much better job of preserving edge and boundary information under rainy conditions, where motion-blur and water streaks may cause traffic sign degradation. Moreover, the MCGM was the most useful under snowy conditions, when the occlusion and texture degradation caused by weather affected the noise level, and thus led to an 11.2% relative increase above the baseline.

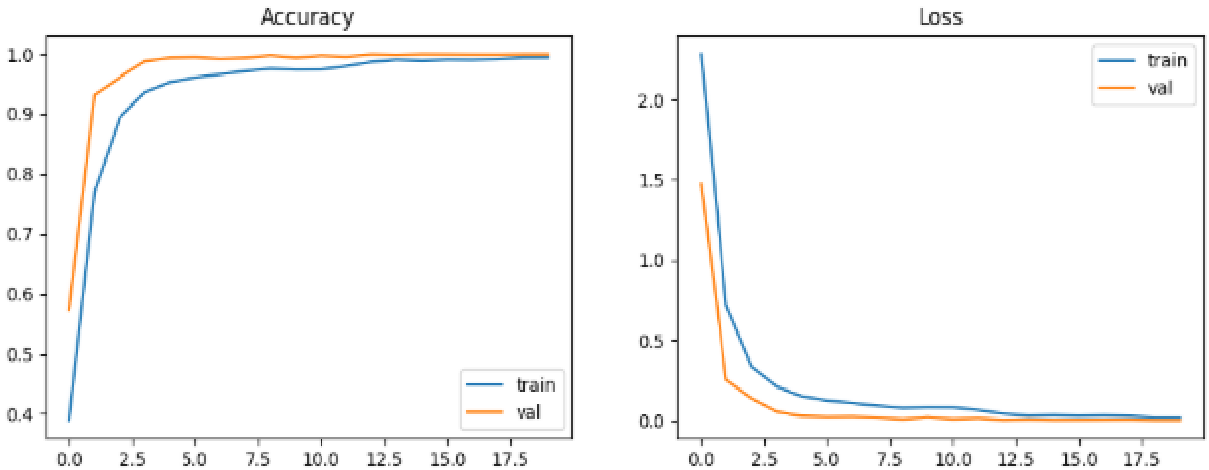

Figure 6: Training and Validation Accuracy and Loss Curves of the Proposed Models

The proposed model was tested with various weather conditions to assess the adverse weather effect on the traffic sign recognition system. The accuracy of the detection is shown in Figure 7 for different environmental scenarios and the comparison of the model's performance showcases its robustness and performance in all driving scenarios. Results demonstrate the ability of the proposed framework to stay at high accuracy level in spite of the problem of weather. The results show that the proposed framework has consistently high accuracy even in the presence of weather challenges like fog, rain and snow. A slight performance difference is seen under extreme weather scenarios but the incorporation of feature enhancement modules and adversarial domain adaptation can significantly reduce the drop in recognition accuracy, which highlights the model's good generalization performance for practical autonomous driving scenarios.

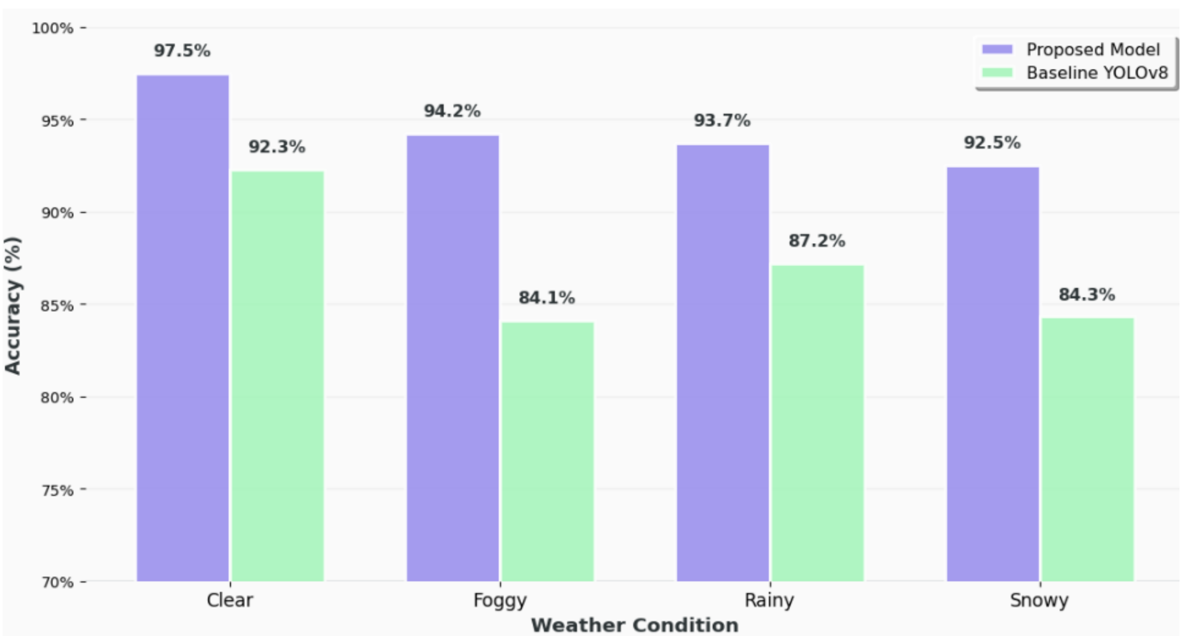

Figure 7: Performance Breakdown by Weather Condition

The Figure 8 shows the distribution and correlation between the ground-truth and predicted bounding boxes of the proposed MCGM method on the baseline of YOLOv8. The points are scattered randomly in the plot and the more they are clustered along the diagonal the better is the localization accuracy. Marginal densities of predictions are presented in the density contours on the margins. This is because the proposed MCGM shows a higher density peak around the region (0.9, 0.9) compared to the baseline YOLOv8, which indicates higher precision and recall. The blue to red colour gradient shows the density of the predictions, with warmer colours showing more density of accurate predictions.

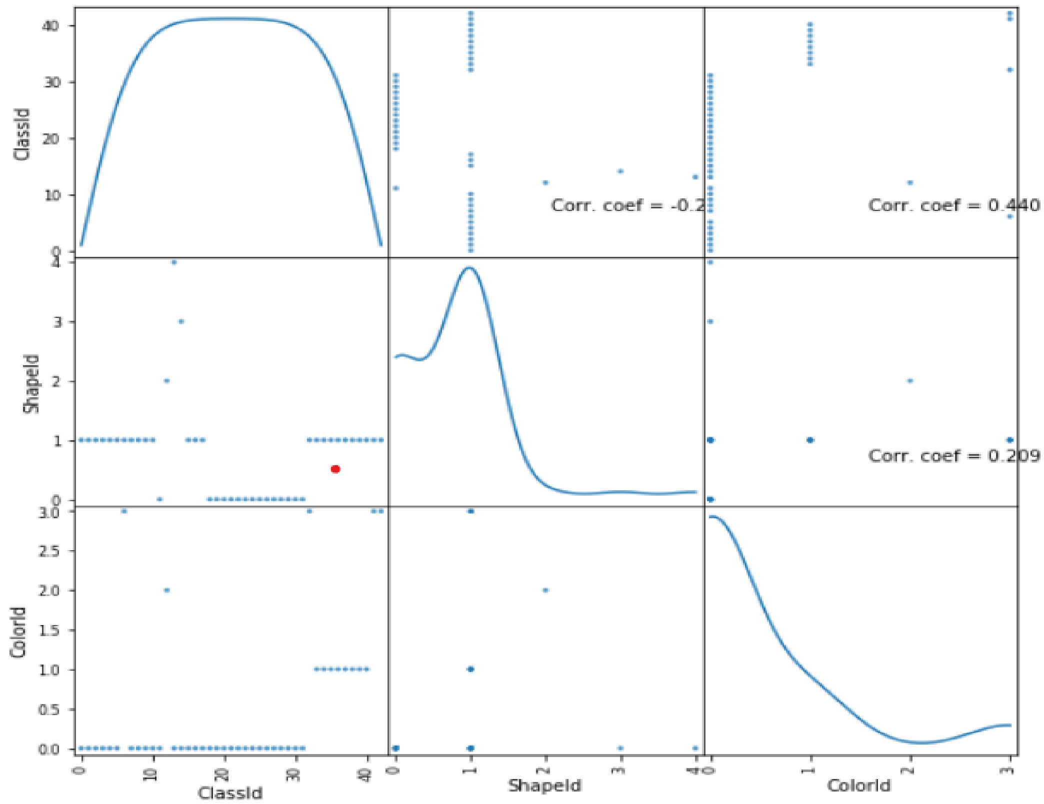

Figure 8: Scatter plot with density distribution of the predicted results

Conclusion:

Our study aimed to develop a reliable traffic sign recognition system to enhance the reliability of the self-driving vehicles in adverse weather conditions. To tackle the issues of fog, rain, snow, and illumination changes, the proposed architecture combines YOLOv8 with Fused Path Feature Enhancement Network (FPFENet), Multiscale Edge Enhancement Module (MEEM), Adaptive Dual-Branch Bidirectional Feature Pyramid Network (ADBF-FPN), Multiscale Convolutional Gating Module (MCGM), and Adversarial Weather Domain Adaptation (AWDA). The enhancement of feature extraction, preservation of fine-scale textures, reinforcement of edge information, adaptive multi-scale feature fusion and noise suppression for weather-induced noises are all achieved by each component, which helps improve the detection and classification of traffic signs in challenging environments. In addition, the adversarial domain adaptation approach successfully reduces the domain shift between clear-weather images and adverse-weather images, so that the model can learn feature representations that are not affected by the weather without the need for a large number of manually labeled adverse-weather images. The experimental results illustrate that the proposed framework is able to achieve a high level of accuracy in detection, stable convergence, and good generalization ability in various weather conditions, making it a good solution for intelligent transportation systems and autonomous driving applications.

The proposed framework shows good performance but there are still some research opportunities for further enhancement of the framework. Reducing computational complexity of the integrated architecture in order to deploy it in real-time on embedded and distributed (edge) computing devices. In addition, the model can be further enhanced by adding Vision Transformers (ViTs), or hybrid CNN-Transformer models to leverage more information from long-range context. Adding multi-modal sensor data (e.g., camera images, LiDAR and Radar data) may also contribute to enhancing the detection reliability in harsh weather and poor visibility. Moreover, some methods for domain adaptation which do not rely on labeled data, such as self-supervised domain adaptation, semi-supervised domain adaptation, and unsupervised domain adaptation, can reduce the demand for labeled data and thus improve adaptation to new domains. Last but not least, further testing of the proposed framework on larger and more representative real-world traffic sign datasets from varied geographical areas, traffic conditions and weather conditions would enhance its robustness and applicability to practice. These future advancements will continue to help progress the next generation of autonomous driving technology's ability to accurately, efficiently and reliably recognize traffic signs.

REFERENCES:

Köylü, F. 2026. A comparative analysis of YOLO-Based traffic sign detections with a novel Turkish traffic sign dataset. *IEEE Access* 14:7744–7763.

M, U., G. P, R. R, A. V, D. S and K.K. G. 2025. AI-Driven Traffic Sign Recognition and Speed Control for Autonomous Vehicle. *ICRDICCT'25 2025 - INTERNATIONAL CONFERENCE ON RESEARCH AND DEVELOPMENT IN INFORMATION, COMMUNICATION, AND COMPUTING TECHNOLOGIES* 133–137.

Xue, M., C. Yuan, C. He, J. Wang and W. Liu. 2020. NATURALAE: Natural and Robust Physical Adversarial Examples for Object Detectors. *arXiv (Cornell University)*, doi: 10.48550/arxiv.2011.13692.

S, J., A.K. G and G. R. 2025. Automatic Traffic Sign Recognition and Detection System. *SSRN Electronic Journal*, doi: 10.2139/ssrn.5086005.

Qutaishat, D. and S. Li. 2025. A systematic review and comparative analysis of deep learning models for Twitter/X-based traffic event detection. *International Journal of Digital Earth* 19.

Zhou, W., Y. Yu, Y. Zhan and C. Wang. 2022. A vision-based abnormal trajectory detection framework for online traffic incident alert on freeways. *Neural Computing and Applications* 34:14945–14958.

Wong, P.Y.L., T.M. Leung, W. Zhang, K.C.C. Lo, X. Guo and T. Hu. 2026. Enhancing energy efficiency in road transport systems: A comparative study of Australia, Hong Kong and the UK. *Energies* 19:266.

Dominguez-Sanchez, A., M. Cazorla and S. Orts-Escolano. 2018. A new dataset and performance evaluation of a Region-Based CNN for urban object detection. *Electronics* 7:301.

Jamal, S., F. Siddiqui, M.A. Alam, M. Ayman-Mursaleen, S. Zafar and S. Naaz. 2026. Deep Reinforcement Learning for Sustainable Urban Mobility: A Bibliometric and Empirical review. *Sensors* 26:376.

Zhong, M., K. Yang, Z. Xiao, J. Tan, K. Fan, Z. Deng and M. Zhou. 2025. A detection algorithm for Two-Wheeled vehicles in complex scenarios based on Semi-Supervised Learning. *Computers, Materials & Continua/Computers, Materials & Continua (Print)* 84:1055–1071.

Houben, S., J. Stallkamp, J. Salmen, M. Schlipsing and C. Igel. 2013a. Detection of traffic signs in real-world images: The German traffic sign detection benchmark. *International Joint Conference on Neural Networks (IJCNN)* 1–8.

Balali, V., A.A. Rad and M. Golparvar-Fard. 2015. Detection, classification, and mapping of U.S. traffic signs using google street view images for roadway inventory management. *Visualization in Engineering* 3.

He, R., Q. Shang, X. Lu, X. Lu, X. Lu, X. Lu, W. Guo and J. Zhao. 2026. Development of intelligent road markings: A review and outlook. *Journal of Traffic and Transportation Engineering (English Edition)* 13:168–187.

Tabernik, D. and D. Skocaj. 2019. Deep Learning for Large-Scale Traffic-Sign Detection and Recognition. *IEEE Transactions on Intelligent Transportation Systems* 21:1427–1440.

Li, J., Q. Deng, W. Gao, B. Yang, L. Jia, J. Zhou and H. Pu. 2025. DSF-YOLO for robust multiscale traffic sign detection under adverse weather conditions. *Scientific Reports* 15:24550.

Kim, M., J. Jeong and S. Kim. 2021. ECAP-YOLO: Efficient Channel Attention Pyramid YOLO for small object detection in aerial image. *Remote Sensing* 13:4851.

Arcos-García, Á., J.A. Álvarez-García and L.M. Soria-Morillo. 2018. Evaluation of deep neural networks for traffic sign detection systems. *Neurocomputing* 316:332–344.

Gui, X., F. Zeng, Y. Zhang, Y. Wang, J. Wang, C. Wang, X. Fu, S. Li, F. Liu, L. Yue, J. Jiang and Z. Li. 2026. Intelligent Distributed Optical fiber sensing in transportation Infrastructures: research progress, applications, and challenges. *arXiv (Cornell University)*, doi: 10.48550/arxiv.2601.13891.

Sielemann, A., L. Loercher, M.-L. Schumacher, S. Wolf, M. Roschani, J. Ziehn and J. Beyerer. 2024. Synset Signset Germany: a Synthetic Dataset for German Traffic Sign Recognition. *International Conference on Intelligent Transportation* 3383–3390.

Houben, S., J. Stallkamp, J. Salmen, M. Schlipsing and C. Igel. 2013b. Detection of traffic signs in real-world images: The German traffic sign detection benchmark. *International Joint Conference on Neural Networks (IJCNN)* 1–8.

Stallkamp, J., M. Schlipsing, J. Salmen and C. Igel. 2011. The German Traffic Sign Recognition Benchmark: A multi-class classification competition. *International Joint Conference on Neural Networks (IJCNN)* 1453–1460.

Du, C., S. Su, C. Lin, Y. Yao, R. Jin and X. Hong. 2025. A lightweight network for traffic sign detection via multiple scale context awareness and semantic information guidance. *Scientific Reports* 15:10110.

Wang, L., Q. Ai and H. Jin. 2026. Lightweight traffic sign detection algorithm with noise suppression and semantic enhancement. *PLoS ONE* 21:e0340810.

V, R. and R.J. S. 2022. Literature Review on Vision-based Traffic Sign Detection and Recognition using Artificial Intelligence. *International Journal for Research in Applied Science and Engineering Technology* 10:4801–4808.

Liu, C., S. Li, F. Chang and Y. Wang. 2019. Machine Vision Based Traffic Sign Detection Methods: Review, Analyses and Perspectives. *IEEE Access* 7:86578–86596.

Lin, K. and L. Zhang. 2026. Optimized YOLOv7 for traffic sign recognition. *Journal of Intelligent Systems* 35.

Yalamanchili, S., K. Kodepogu, V.B. Manjeti, D. Mareedu, A. Madireddy, J. Mannem and P.K. Kancharla. 2024. Optimizing traffic sign detection and recognition by using deep learning. *International Journal of Transport Development and Integration* 8:131–139.

Ma, Y., F. Arifah, Q. Afifah, L. Bun, K. Zhang and M. Tang. 2026. Traffic Light Recognition Assistant for Color Vision Deficiency Using YOLO with Multilingual Audio Feedback. *Preprints.org*, doi: 10.20944/preprints202601.0608.v1.

Nellore, K. and G. Hancke. 2016. A survey on urban traffic Management system using wireless sensor networks. *Sensors* 16:157.

Kattenborn, T., J. Leitloff, F. Schiefer and S. Hinz. 2021. Review on Convolutional Neural Networks (CNN) in vegetation remote sensing. *ISPRS Journal of Photogrammetry and Remote Sensing* 173:24–49.

Chen, Z., J. Zhang, Y. Xu and D. Tao. 2023. Transformer-Based context condensation for boosting feature pyramids in object detection. *International Journal of Computer Vision* 131:2738–2756.

Zhang, Z., B. Feng and H. Liu. 2026. TS-YOLO: a small traffic sign detection algorithm for various harsh driving conditions in bad weather. *Scientific Reports* 16:4444.