TRAFFIC SIGN DETECTION USING YOLOv8

SUNMATHI K R

Team members:

SRINIVASAN S K, SANTHOSH D, REGALYA SRI M, DR PRIYA K

E-Mail: santhosh10230@gmail.com, sksrinivasan231@gmail.com,

sunmathiramasamy46@gmail.com, regalyasri8@gmail.com, priya.k@kprcas.ac.in

Department of B. Sc. Computer Science with Artificial Intelligence and Data Science

KPR College of Arts Science and Research

Coimbatore, Tamil Nadu, India.

Abstract

Traffic sign detection is an important computer vision task for intelligent transportation systems, advanced driver assistance systems, and autonomous vehicles. Manual identification of traffic signs is time-consuming and may become unreliable when signs appear at different scales, under varying illumination, or within complex road scenes. This chapter presents a deep learning based traffic sign detection system developed using the YOLOv8 object detection model. The proposed system processes road images and identifies traffic signs by generating bounding boxes, class labels, and confidence scores. The model was trained to recognize fifteen traffic-related classes, including Green Light, Red Light, Speed Limit 10, Speed Limit 20, Speed Limit 30, Speed Limit 40, Speed Limit 50, Speed Limit 60, Speed Limit 70, Speed Limit 80, Speed Limit 90, Speed Limit 100, Speed Limit 110, Speed Limit 120, and Stop. The experimental prediction set contained 638 images. Model performance was analyzed using confusion matrices, precision-recall curves, precision-confidence curves, recall-confidence curves, and F1-confidence curves. The precision-recall analysis produced an overall mAP@0.5 of 0.976, while the precision-confidence curve reached 1.00 at a confidence threshold of approximately 0.975. The F1-confidence curve reported a maximum overall F1 score of 0.96 at a confidence value of approximately 0.416. The results demonstrate that the trained YOLOv8 model can detect and classify the targeted traffic signs with high confidence. The system can support applications in road safety, traffic monitoring, driver assistance, and intelligent transportation systems, while future improvements can focus on real-time deployment, challenging environmental conditions, and edge-device implementation and practical deployment support.

Keywords

Traffic Sign Detection, YOLOv8, Object Detection, Deep Learning, Computer Vision, Traffic Monitoring, Intelligent Transportation Systems, Precision, Recall, mAP.

1. Introduction

Traffic signs are essential components of road transportation because they communicate speed restrictions, traffic regulations, warnings, and information to road users. Correct interpretation of these signs contributes directly to road safety. With the development of intelligent transportation systems and autonomous vehicles, automatic traffic sign detection has become an important application of computer vision.

Traditional traffic sign recognition systems often depend on manually designed visual features such as colour, shape, edges, and geometric patterns. Such approaches can perform well in controlled conditions but may become less reliable when signs are small, partially occluded, blurred, rotated, or affected by changing illumination. Deep learning approaches overcome many of these limitations by learning useful features directly from training images.

YOLO (You Only Look Once) is a one-stage object detection family designed to identify objects and their locations efficiently. YOLOv8 is a modern implementation in the Ultralytics ecosystem and is suitable for custom object detection tasks. In this project, YOLOv8 was trained to identify fifteen traffic-related classes. The system produces a bounding box, class label, and confidence score for each detected object.

The completed project was evaluated using multiple diagnostic plots and prediction outputs. The generated results include normalized and raw confusion matrices, precision-recall analysis, confidence-based precision and recall curves, an F1-confidence curve, and sample detection images. These results provide a detailed view of how the model performs across individual traffic-sign categories.

2. Problem Statement

Identifying traffic signs automatically from road images is challenging because traffic signs can differ in size, appearance, illumination, position, and background context. Similar-looking speed-limit signs can also be difficult to distinguish from one another. A practical detection system therefore needs to identify multiple traffic-sign classes accurately while providing confidence information and precise object locations.

The problem addressed in this project is the development of an automated computer-vision system capable of detecting and classifying traffic-related signs from images using a YOLOv8 object detection model.

3. Objectives

• To develop an automated traffic sign detection system using YOLOv8.

• To detect multiple traffic-sign classes from road images.

• To generate bounding boxes around detected signs.

• To assign an appropriate class label and confidence score to each detection.

• To evaluate the trained model using confusion matrices and confidence-based performance curves.

• To analyze precision, recall, F1 score, and mAP@0.5 from the generated evaluation results.

• To demonstrate the practical usefulness of YOLOv8 for traffic monitoring and intelligent transportation applications.

4. Background Study / Literature Review

Traffic sign detection has progressed from colour-, shape-, and feature-based computer vision methods to deep learning based object detection. Modern detectors learn visual representations directly from annotated images, which makes them more adaptable to complex road scenes.

Wang et al. investigated an improved YOLOv5 network for real-time multi-scale traffic sign detection. Their work addressed the difficulty of detecting signs at different scales and proposed an adaptive feature-pyramid approach together with data augmentation for improved robustness.

Zhang et al. introduced CCTSDB 2021 as a more comprehensive traffic sign detection benchmark. The dataset and benchmark emphasize challenging road conditions and provide a useful basis for evaluating traffic sign detectors.

Recent research has increasingly explored YOLOv8 for traffic-sign detection. Studies published in 2024 proposed improved YOLOv8 architectures for small signs, complex backgrounds, adverse weather, and real-time autonomous driving applications. These works show that feature enhancement, attention mechanisms, and improved loss functions can further improve the performance of YOLO-based traffic sign detectors.

The present project follows this general research direction by using YOLOv8 as the base detector and evaluating its performance on a fifteen-class traffic-sign dataset. Unlike a classification-only system, the proposed approach performs object detection, meaning that it identifies both what the traffic sign is and where it is located.

5. Proposed Approach

The proposed system follows a complete object-detection workflow, beginning with image data and ending with visual detection results and quantitative evaluation.

DATASET

↓

DATA COLLECTION

↓

DATA PREPROCESSING

↓

ANNOTATION

↓

DATASET SPLITTING

↓

YOLOv8 MODEL TRAINING

↓

MODEL VALIDATION

↓

PREDICTION

↓

DETECTION OUTPUT

↓

PERFORMANCE EVALUATION

The workflow begins with the traffic-sign image dataset. Images are prepared and associated with object annotations. The annotated data is then used to train the YOLOv8 detector. After training, unseen images are passed through the model. The model returns predicted bounding boxes, class names, and confidence values. Finally, the prediction results are analyzed using standard object-detection metrics and diagnostic plots.

5.1 Methodology Stages

Collect and organize traffic-sign images.

Prepare the images and corresponding object annotations.

Train the YOLOv8 object detection model.

Validate the model during the training process.

Run inference on the prediction/evaluation images.

- Generate bounding-box based detection outputs.

- Evaluate the model using precision, recall, F1 score, mAP, and confusion matrices.

6. Data Collection

The project uses a traffic-sign image dataset containing fifteen target classes. The class list was obtained directly from the generated YOLOv8 evaluation plots and therefore represents the classes used by the trained model.

No.	Traffic Sign Class
1	Green Light
2	Red Light
3	Speed Limit 10
4	Speed Limit 20
5	Speed Limit 30
6	Speed Limit 40
7	Speed Limit 50
8	Speed Limit 60
9	Speed Limit 70
10	Speed Limit 80
11	Speed Limit 90
12	Speed Limit 100
13	Speed Limit 110

No.	Traffic Sign Class
14	Speed Limit 120
15	Stop

The prediction output generated during the completed project reports 638 predicted images. These images were used to demonstrate the trained model's ability to detect traffic signs in road scenes.

The dataset contains both traffic-light categories and speed-limit/stop categories. This combination allows the model to learn different visual patterns, including circular regulatory signs and rectangular traffic-light objects.

7. Data Processing

7.1 Image Preparation

The images are prepared in a format suitable for YOLOv8 training and inference. Consistent image processing helps the network learn meaningful visual features and maintain stable training behaviour.

7.2 Annotation

Each target traffic sign is represented by a bounding box and an associated class label. These annotations provide the spatial information required by an object detector to learn where traffic signs occur in an image.

7.3 Data Augmentation

Object-detection pipelines can use augmentation operations such as scaling, translation, rotation, and appearance variation to improve robustness. Such transformations help a detector learn from images that differ in position, size, and visual conditions.

7.4 Dataset Organization

The prepared data is organized for model training, validation, and prediction. Training data is used to learn model parameters, validation data is used to monitor performance, and prediction data is used to examine the final model on images not used for parameter learning.

8. YOLOv8 Model

YOLOv8 is a modern one-stage object detector provided through the Ultralytics framework. It is designed to perform object localization and classification efficiently. The model receives an image and produces detections consisting of bounding-box coordinates, class predictions, and confidence values.

8.1 Main Components

- Backbone: extracts hierarchical visual features from the input image.
- Feature-processing layers: combine information at different feature levels so objects of different sizes can be represented.
- Detection head: predicts object locations, classes, and confidence scores.
- Post-processing: filters and organizes detections to provide the final set of bounding boxes.

8.2 Why YOLOv8 Was Selected

YOLOv8 was selected because the project requires simultaneous localization and classification of multiple traffic signs. The model provides a practical balance between detection accuracy and computational efficiency and supports custom datasets through the Ultralytics training pipeline.

8.3 Model Output

For each accepted detection, the model provides a class label, a confidence score, and a bounding box. For example, the project output identifies a Speed Limit 80 sign with a confidence of 0.95 and a Green Light with a confidence of 0.96 in the displayed prediction image.

9. Training and Evaluation

The YOLOv8 model was trained using the prepared traffic-sign dataset. During training, the network learned visual features associated with the fifteen target classes. After training, inference was performed on the project prediction set and the generated evaluation plots were used to analyze performance.

9.1 Evaluation Metrics

Metric	Meaning
Precision	Proportion of predicted positive detections that are correct.
Recall	Proportion of relevant objects that are successfully detected.
F1 Score	Harmonic mean of precision and recall.
mAP@0.5	Mean Average Precision calculated at an IoU threshold of 0.50.
Confidence	Model score associated with the confidence of a detection.

9.2 Evaluation Figures

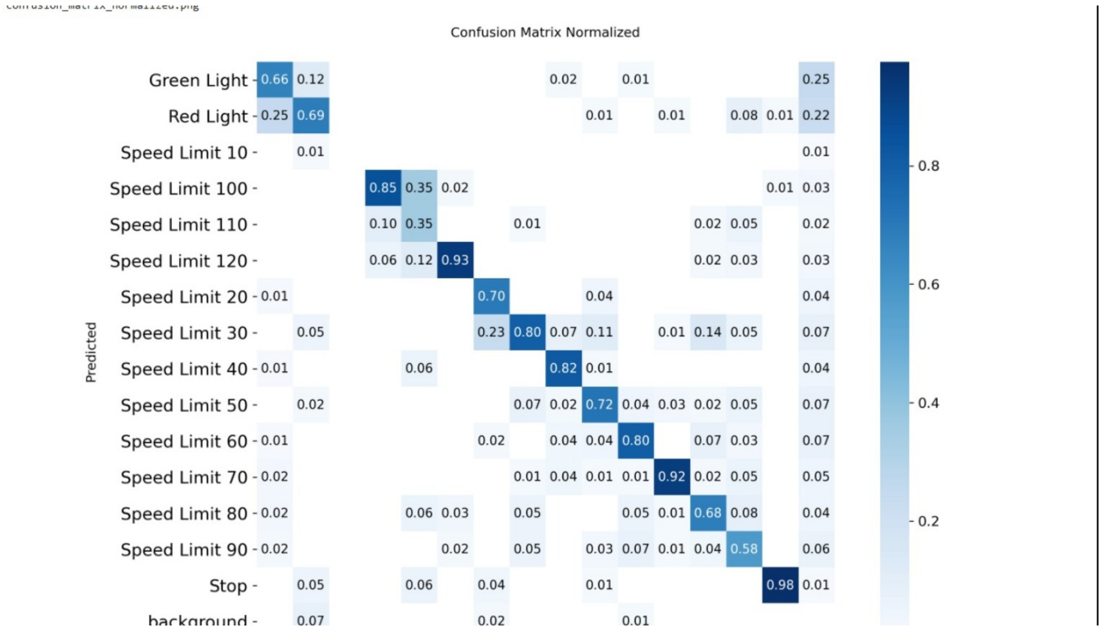

Figure 1. Normalized confusion matrix generated from the trained YOLOv8 model.

The normalized confusion matrix shows the proportion of predictions associated with each class. Strong diagonal values indicate that many examples are correctly associated with their target classes, while off-diagonal values indicate confusion between classes. The plot also shows that some errors occur between visually similar speed-limit classes.

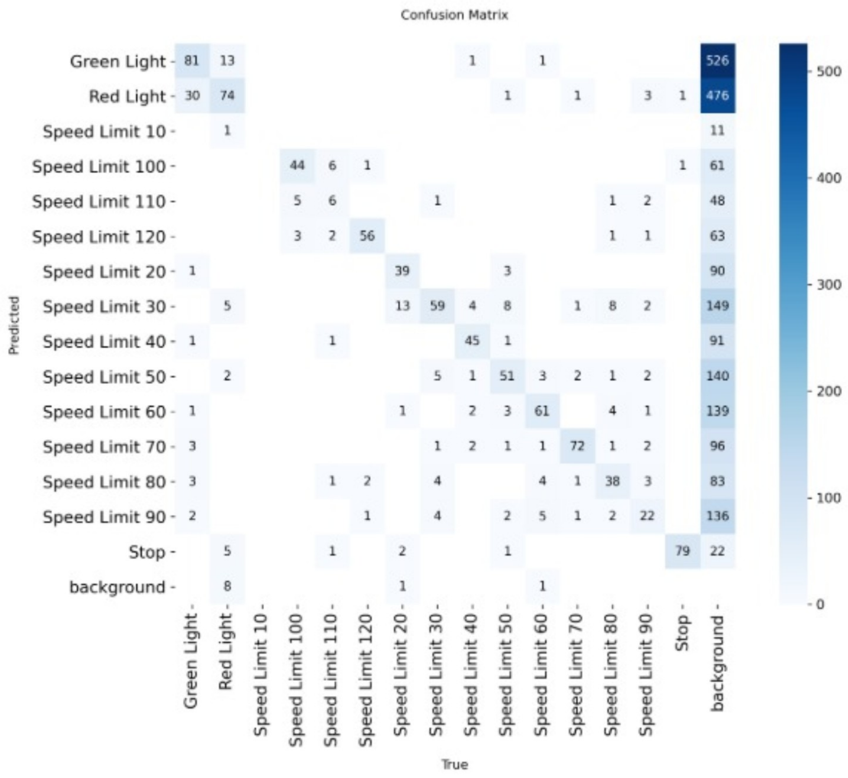

Figure 2. Raw confusion matrix showing the number of predictions for each class.

The raw confusion matrix provides the absolute counts underlying the normalized view. It is useful for identifying which classes have more examples and where misclassifications occur. The background column/row also indicates cases where detections and ground-truth objects do not correspond perfectly.

10. Model Prediction and Output

After training, the model was used to perform inference on road images. The output is visualized using coloured bounding boxes together with the predicted class and confidence value. This makes the model output easy to interpret and provides direct evidence of object localization.

Figure 3. Sample YOLOv8 prediction showing Speed Limit 80 and Green Light detections.

In the sample prediction, the model identifies a Speed Limit 80 sign with a confidence score of 0.95 and a Green Light with a confidence score of 0.96. The bounding boxes are placed around the detected objects, demonstrating that the system performs localization as well as classification. The figure also reports 638 predicted images for the completed prediction run.

11. Results and Discussion

The generated evaluation plots show strong overall performance of the trained traffic-sign detector. The project result was reported as 97.5% accuracy. In addition, the detailed YOLOv8 evaluation plots provide metric-specific values that should be reported separately rather than treating accuracy and mAP as the same measurement.

11.1 Precision-Recall Analysis

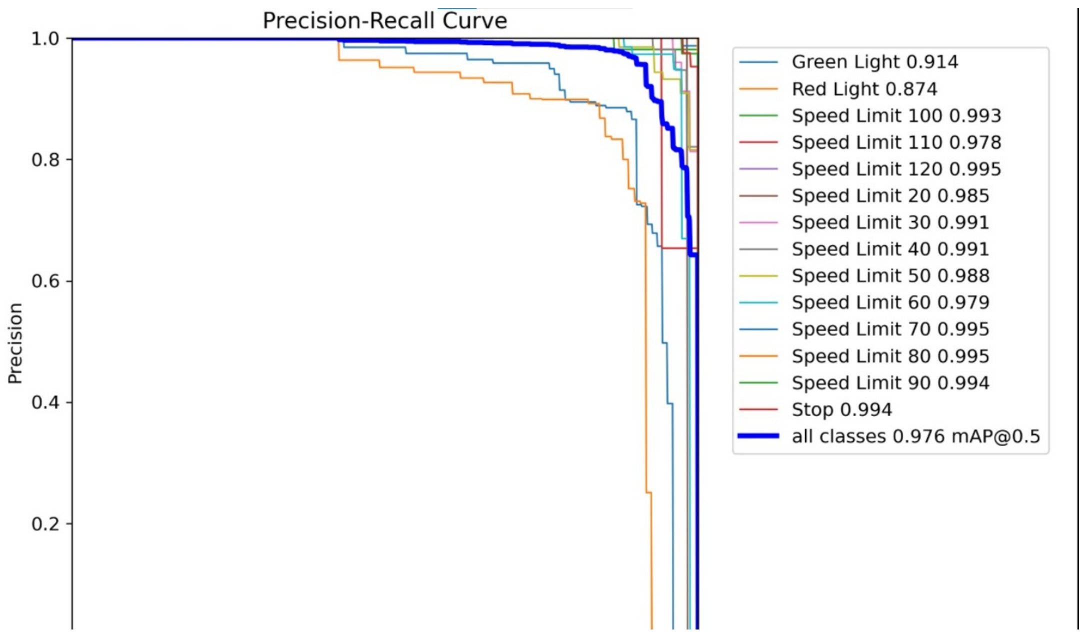

Figure 4. Precision-Recall curve for all traffic-sign classes.

The precision-recall curve shows the relationship between precision and recall for each class. The bold overall curve reports an mAP@0.5 of 0.976. Several classes have average precision values above 0.99, while Green Light and Red Light have lower values than many of the speed-limit categories. This indicates that traffic-light categories are relatively more challenging for the model than several speed-limit classes.

11.2 F1-Confidence Analysis

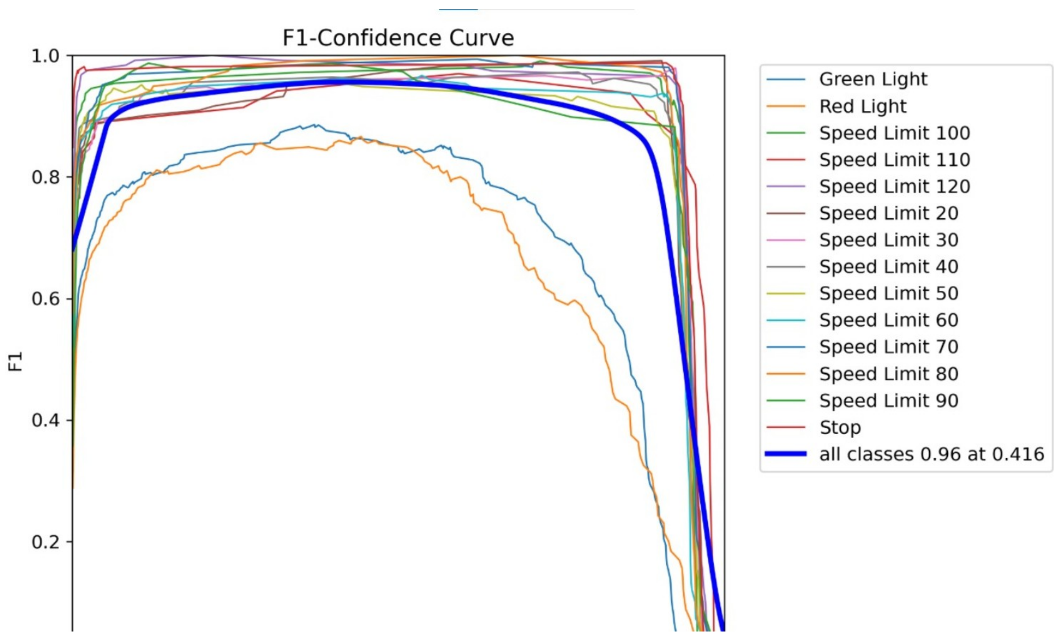

Figure 5. F1-confidence curve showing the balance between precision and recall.

The F1-confidence curve shows how the harmonic mean of precision and recall changes with the detection threshold. The overall curve reaches an F1 score of approximately 0.96 at a confidence value of about 0.416. This indicates that a moderate confidence threshold provides a strong balance between retaining detections and avoiding incorrect predictions.

11.3 Recall-Confidence Analysis

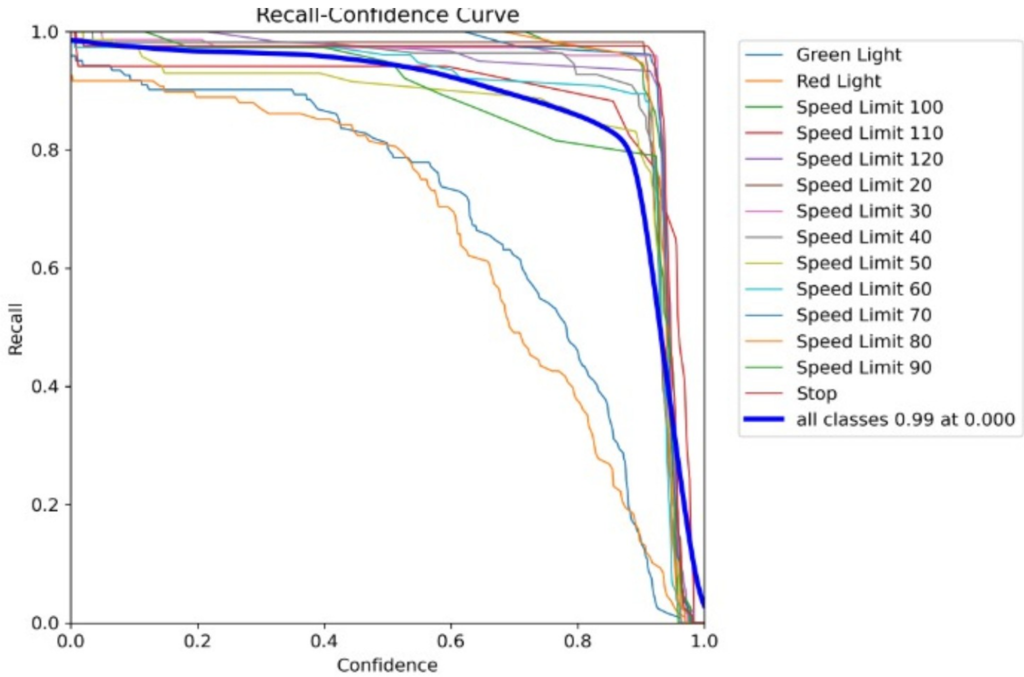

Figure 6. Recall-confidence curve for the traffic-sign classes.

Recall generally decreases as the confidence threshold becomes stricter because fewer detections are accepted. The overall curve begins near 0.99 at a confidence threshold of 0.000 and gradually decreases before falling rapidly near the highest confidence values. This behaviour is expected in confidence-threshold analysis.

11.4 Precision-Confidence Analysis

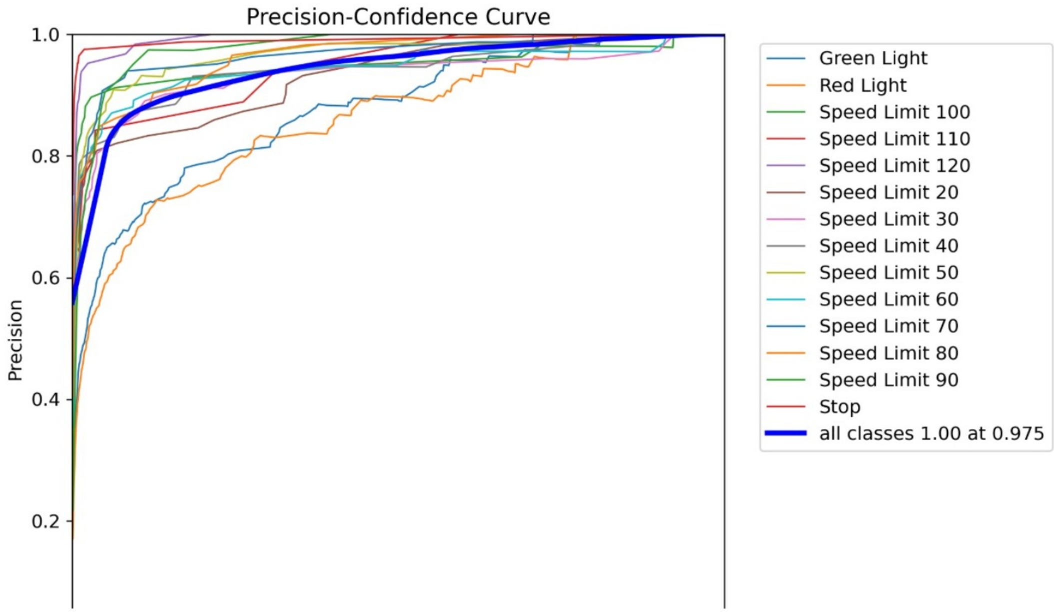

Figure 7. Precision-confidence curve for the traffic-sign classes.

The precision-confidence curve shows that precision increases as the confidence threshold becomes more selective. The overall curve reaches approximately 1.00 at a confidence value of about 0.975. This indicates that detections accepted at a high confidence threshold are highly precise, although using an excessively high threshold can reduce recall.

11.5 Class-Wise Performance

The precision-recall plot provides class-wise AP values. The strongest reported classes include Speed Limit 120 (0.995), Speed Limit 70 (0.995), Speed Limit 80 (0.995), Speed Limit 90 (0.994), Speed Limit 100 (0.993), Speed Limit 30 (0.991), Speed Limit 40 (0.991), Speed Limit 50 (0.988), and Speed Limit 20 (0.985). The comparatively lower AP values are observed for Green Light (0.914) and Red Light (0.874).

11.6 Interpretation of the Confusion Matrix

The normalized confusion matrix indicates that many classes have strong diagonal values, showing correct predictions. For example, Speed Limit 120 has a normalized diagonal value of approximately 0.93, Speed Limit 70 approximately 0.92, Speed Limit 40 approximately 0.82, Speed Limit 60 approximately 0.80, and Stop approximately 0.98. Some speed-limit classes show confusion with neighbouring speed-limit categories, which is understandable because these signs share similar shape, colour, and layout while differing mainly in the printed number.

Performance Summary

Observed result	**Value**
Project-reported accuracy	97.5%
Overall mAP@0.5	0.976 (97.6%)
Maximum overall F1 score	0.96 at confidence ≈ 0.416
Precision at selected confidence	1.00 at confidence ≈ 0.975
Recall at confidence 0.000	0.99
Prediction images	638
Target classes	15

12. Applications

• Advanced Driver Assistance Systems (ADAS): detected signs can be used to provide warnings and driver information.

• Autonomous Vehicles: traffic-sign information can contribute to road-scene understanding and decision making.

• Traffic Monitoring: cameras can automatically identify traffic signs in monitored road environments.

• Road Safety Systems: automatic recognition can support warning systems for speed limits and traffic controls.

• Intelligent Transportation Systems: detected traffic information can be combined with other computer-vision modules.

• Smart-City Infrastructure: roadside cameras can be used for automated road asset and traffic-sign monitoring.

• Educational and Research Applications: the trained model can serve as a practical example of deep learning based object detection.

13. Limitations

Although the model achieved strong results, traffic-sign detection remains sensitive to real-world conditions that may not be fully represented in a training dataset. Very small signs, severe blur, occlusion, poor illumination, unusual viewing angles, weather effects, and visually similar classes can reduce detection reliability.

The current project is also evaluated primarily through image-based prediction outputs. Real-time deployment would require additional analysis of inference speed, hardware requirements, video stability, and performance across long driving sequences.

14. Future Scope

• Integrate the trained model with a real-time camera or dash-camera video stream.

• Optimize the model for edge devices and embedded systems.

• Increase the diversity of training images, especially for night, rain, fog, glare, and low-light conditions.

• Add more traffic-sign categories and regional sign types.

• Improve recognition of small and distant signs through multi-scale feature techniques.

• Develop an alert system that converts detected signs into driver warnings.

• Combine traffic-sign detection with lane detection, vehicle detection, and traffic-light analysis for a complete road-perception system.

• Evaluate the model using additional datasets and real-world driving videos.

15. Conclusion

This chapter presented a YOLOv8-based traffic sign detection system designed to automatically identify and localize traffic-related objects in road images. The system recognizes fifteen classes covering traffic lights, speed limits from 10 to 120, and stop signs. The complete workflow includes data preparation, annotation, model training, inference, and quantitative evaluation.

The generated project results demonstrate strong detection performance. The project reported an overall accuracy of 97.5%, while the detailed YOLOv8 evaluation plot recorded an overall mAP@0.5 of 0.976. The F1-confidence curve reached approximately 0.96 at a confidence threshold of 0.416, and the precision-confidence curve reached 1.00 at approximately 0.975 confidence. The sample predictions also show high-confidence detections such as Speed Limit 80 at 0.95 and Green Light at 0.96.

The confusion matrices indicate that most classes are correctly identified, while the main sources of confusion occur among visually similar speed-limit categories and between traffic-light classes.

Overall, the project demonstrates that YOLOv8 is a practical and effective approach for automated traffic sign detection. With further improvements in data diversity, real-time processing, and edge deployment, the system can be extended toward intelligent transportation and advanced driver-assistance applications.

16. References

[1] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," Ultralytics, version 8.0.0, 2023. https://github.com/ultralytics/ultralytics

[2] J. Wang, Y. Chen, M. Gao, and Z. Dong, "Improved YOLOv5 network for real-time multi-scale traffic sign detection," arXiv:2112.08782, 2021.

[3] J. Zhang et al., "CCTSDB 2021: A More Comprehensive Traffic Sign Detection Benchmark," Human-centric Computing and Information Sciences, vol. 12, 2022, Art. 23. doi:10.22967/HCIS.2022.12.023.

[4] Y. Luo, Y. Ci, H. Zhang, and L. Wu, "A YOLOv8-CE-based real-time traffic sign detection and identification method for autonomous vehicles," Digital Transportation and Safety, vol. 3, no. 3, pp. 82–91, 2024. doi:10.48130/dts-0024-0009.

[5] "Traffic Sign Detection Algorithm Based on Improved YOLOv8," 2024 6th International Conference on Communications, Information System and Computer Engineering (CISCE), 2024. doi:10.1109/CISCE62493.2024.10653400.

[6] "Improved YOLOv8 for small traffic sign detection under complex environmental conditions," Fundamental Research, vol. 8, 2024, Art. 100167. doi:10.1016/j.fraope.2024.100167.

[7] "YOLOv8s-DDA: An Improved Small Traffic Sign Detection Algorithm Based on YOLOv8s," Electronics, vol. 13, no. 18, 2024, Art. 3764. doi:10.3390/electronics13183764.

[8] Y. Zhang, H. Liu, D. Dong, X. Duan, F. Lin, and Z. Liu, "DPF-YOLOv8: Dual Path Feature Fusion Network for Traffic Sign Detection in Hazy Weather," Electronics, vol. 13, no. 20, 2024, Art. 4016. doi:10.3390/electronics13204016.

[9] S. Naik, N. Paliwal, and T. Lele, "Real-Time Traffic Sign Detection Using YOLOv8 for Autonomous Driving Systems," Journal of Artificial Intelligence and Imaging, 2024. doi:10.48001/joaii.2024.1219-25.

[10] J. Luo et al., "ETSR-YOLO: An improved multi-scale traffic sign detection algorithm based on YOLOv5," 2023.