International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Road Sign Detection System using YOLOv8 on Raspberry Pi 4

Divyam Verma[1], Tushar Gupta[2], Avinash Gupta[3], Sampada Massey[4]

Computer Science and Engineering Department,
Shri Shankaracharya Technical Campus, Bhilai, Chhattisgarh, India
divyam030903@gmail.com [1]; tushargupta3002@gmail.com [2]; guptavinash017@gmail.com [3]; sampada.massey@sstc.ac.in[4]

ABSTRACT –

This paper reports the design and evaluation of an Indian traffic sign real-time road sign detection system. The system is deployed on a Raspberry Pi 4, with object detection carried out via YOLOv8 and audio alert to the driver. The model is trained on an in-house-developed dataset and evaluated under normal metrics, showing effective performance under real driving conditions. The proposed system improves driver perception and safety via low-cost embedded hardware and shows promising detection accuracy.

Keywords— YOLOv8, Raspberry Pi, Road Sign Detection, Indian Traffic Signs, Object Detection, Real-Time Systems

INTRODUCTION

The increasing need for intelligent transport systems and road safety has generated a spate of research towards the development of real-time road sign recognition systems. In an era where autonomous cars and driver assistance are becoming more the order of the day, road sign recognition and interpretation are key elements towards realizing road conformity and reducing accident rates. Being a country with a vast and diversified network of roads, India presents its own unique problems in the area of this field due to inconsistent signs, varying climatic conditions, and non-standard set of datasets.

Advanced driver-assistance systems (ADAS) heavily rely on accurate detection of traffic signs for real-time driving decisions. Germany and the United States already have standards such as the German Traffic Sign Recognition Benchmark (GTSRB) [13]. India lacks a large-scale publicly available dataset that is specific to its traffic signs [11][12]. This complicates the development of strong identification models tailored to Indian roads.

To address this issue, we propose a real-time Indian traffic sign detection system implemented with the YOLOv8 object detection model [1]. YOLOv8 is highly accurate and efficient and thus well-suited for real-time use on low-resource hardware. Our system is deployed on a Raspberry Pi 4 with 4GB RAM [2], with a camera module for real-time video processing and detection of 57 distinct Indian traffic sign classes. The entire process—from data collection with a self-designed dataset as Indian datasets are not standardized, annotation with LabelImg [3], training with YOLOv8 [1], and real-time deployment on Raspberry Pi 4—has been done with the factors of cost-effectiveness, scalability, and real-world implementation. Besides providing better user experience and security, the system also provides audio feedback in a male tone for identified signs. Multimodal feedback helps the driver receive early alerts even if the sign might not be clearly visible to the driver.

Here, we demonstrate the process of dataset generation, model training using YOLOv8, multi-metric performance evaluation (precision, recall, mAP), and ultimate deployment of the detection system on Raspberry Pi 4. The proposed system is an effective contribution to the field of smart transportation as it offers a cost-effective, real-time, and robust traffic sign detection system in the Indian context.

The increasing need for intelligent transport systems and road safety has generated a spate of research towards the development of real-time road sign recognition systems. In an era where autonomous cars and driver assistance are becoming more the order of the day, road sign recognition and interpretation are key elements towards realizing road conformity and reducing accident rates. Being a country with a vast and diversified network of roads, India presents its own unique problems in the area of this field due to inconsistent signs, varying climatic conditions, and non-standard set of datasets.

Advanced driver-assistance systems (ADAS) heavily rely on accurate detection of traffic signs for real-time driving decisions. Germany and the United States already have standards such as the German Traffic Sign Recognition Benchmark (GTSRB) [13]. India lacks a large-scale publicly available dataset that is specific to its traffic signs [11][12]. This complicates the development of strong identification models tailored to Indian roads.

To address this issue, we propose a real-time Indian traffic sign detection system implemented with the YOLOv8 object detection model [1]. YOLOv8 is highly accurate and efficient and thus well-suited for real-time use on low-resource hardware. Our system is deployed on a Raspberry Pi 4 with 4GB RAM [2], with a camera module for real-time video processing and detection of 57 distinct Indian traffic sign classes. The entire process—from data collection with a self-designed dataset as Indian datasets are not standardized, annotation with LabelImg [3], training with YOLOv8 [1], and real-time deployment on Raspberry Pi 4—has been done with the factors of cost-effectiveness, scalability, and real-world implementation. Besides providing better user experience and security, the system also provides audio feedback in a male tone for identified signs. Multimodal feedback helps the driver receive early alerts even if the sign might not be clearly visible to the driver.

Here, we demonstrate the process of dataset generation, model training using YOLOv8, multi-metric performance evaluation (precision, recall, mAP), and ultimate deployment of the detection system on Raspberry Pi 4. The proposed system is an effective contribution to the field of smart transportation as it offers a cost-effective, real-time, and robust traffic sign detection system in the Indian context.

LITERATURE REVIEW

- Road sign detection systems have been widely explored and utilized in nations with well-structured datasets and uniform traffic sign legislation. In the United States, the LISA Traffic Sign Dataset [11] has been used as a standard for numerous object detection models, allowing for robust performance evaluation across different lighting and motion conditions. In Germany, benchmarks like the German Traffic Sign Recognition Benchmark (GTSRB) and Detection Benchmark (GTSDB) [12] have enabled enhancements in road sign detection and classification algorithms, utilized widely in autonomous driving studies. China has also made significant progress by constructing large-scale datasets specific to intelligent transportation systems and autonomous vehicles [13].
- In contrast, India lacks a publicly released, standardized dataset for the classification of road signs. This is a significant deterrent to the training of high-performance deep learning models, since Indian traffic signs lack the visual and structural diversity contained in global datasets. To bypass this limitation, a custom Indian road sign dataset was created from scratch, across a range of scenarios and sign classes, and annotated in YOLO format to facilitate real-time model training.

METHODOLOGY

Dataset Collection and Labeling

To counter the lack of a big Indian road sign dataset, we have created a tailored dataset by recording traffic signs throughout the city using high-definition cameras. Images were taken under various conditions like daylight, dusk, fog, and under varying weather conditions to make the model more robust. Images were marked in the YOLO format with bounding boxes representing the sign position. Marking was done using the LabelImg tool [3], which saves text files in the YOLO training pipeline format. The dataset has been divided into various classes according to sign types, like obligatory, warning, and informatory signs.

Training Pipeline and Configuration

We made use of the YOLOv8s of the Ultralytics implementation [1], a performance versus efficiency trade-off and hence well-apt for edge devices such as Raspberry Pi. Training was conducted on a workstation with an NVIDIA GPU in order to speed up training time. Precise hyperparameters were chosen:

- Model: YOLOv8s
- Input Image Size: 640x640 pixels
- Intervals: 100
- Batch Size: 16
- Optimizer: Stochastic Gradient Descent (SGD)
- Learning Rate: 0.01

Data Augmentation: Scaling, image flipping, and image mosaicking were used to enhance the generalization of the model.

We monitored during training such metrics as loss, precision, recall, and mAP. We selected the model checkpoint for the highest mAP@0.5 on the validation set and used it.

Process Workflow

The entire development process went through these major steps:

- Requirement Analysis: Identifying suitable hardware and detection model.
- Dataset Creation: Collecting and annotating traffic sign images.
- Model Training: Training YOLOv8 on labeled data.
- Evaluation: Measuring model performance with detection metrics.
- Optimization: Converting to TorchScript format for deployment.
- Deployment: Deployment of the model on Raspberry Pi with PyTorch and OpenCV.
- Testing and Feedback: Verifying real-time detection and using audio feedback.

Deployment

The trained model was translated to TorchScript to enable efficient inference on the Raspberry Pi 4 Model B (4GB RAM) [2].The Raspberry Pi Camera Module was used to capture real-time input. The detection system is executed in a video frame stream and provides audio output via car speakers.

- Operating System: Raspberry Pi OS (32-bit)
- Inference Library: PyTorch + OpenCV
- TTS Engine: pyttsx3 for generating male voice audio prompts

- Audio Output: In-car audio system with 3.5mm jack or Bluetooth connectivity
- Latency: ~300 ms per frame at 640 x 640 resolution

A lightweight UI script provides real-time bounding boxes and recognized label names, and audio cues are activated for every detected road sign.

Programming and Implementation

The main program for the road sign detection system is developed using Python and deployed on the Raspberry Pi 4. The development and testing of the YOLOv8 object detection model are performed using Jupyter Notebook and the Ultralytics YOLO library. Python is chosen due to its wide support for machine learning frameworks and its compatibility with embedded platforms like Raspberry Pi.

The system utilizes the OpenCV library to interface with the Raspberry Pi Camera Module. Live video feed is captured and processed in real time. The YOLOv8 model, trained on a custom dataset of Indian road signs, is loaded to detect and classify signs such as STOP, SPEED LIMIT, and NO ENTRY. Based on the detected sign, the system provides instant audio feedback to the driver using a male voice generated via the pyttsx3 text-to-speech (TTS) engine. The voice output is transmitted through the vehicle's speaker system.

The entire logic is implemented to operate efficiently under varying lighting and motion conditions, even when the vehicle is moving at speeds up to 150 km/h. YOLOv8's fast inference capability on the Raspberry Pi ensures real-time detection with minimal latency.

Python programming, combined with OpenCV, pyttsx3, and the YOLOv8 model, offers a robust and scalable solution for embedded vision tasks like road sign detection. The Raspberry Pi 4's quad-core processor and 4GB RAM make it well-suited for deploying lightweight yet powerful AI models in real-world automotive environments.

Working

In this system, a camera module continuously captures live video of the road environment in front of the vehicle. The captured frames are processed in real time by the Raspberry Pi 4, which acts as the central processing unit. The core component is the YOLOv8 object detection model, which has been pre-trained on a custom dataset of Indian road signs, including STOP, SPEED LIMIT, and NO ENTRY signs.

As the vehicle moves, the camera captures the surrounding area. Each frame is analyzed by the YOLOv8 model to detect the presence of any road signs. Once a sign is detected, its classification label is passed to the audio feedback system powered by the pyttsx3 text-to-speech engine. A pre-configured male voice reads out the type of road sign detected and broadcasts it through the vehicle's speaker system to alert the driver immediately.

The system is designed to operate efficiently even when the vehicle is in motion at speeds up to 150 km/h. The Raspberry Pi 4's GPU acceleration and optimized YOLOv8 model ensure that detection is performed with low latency. This allows the driver to receive timely alerts without any manual intervention.

All components of the system, including the camera, speaker, and Raspberry Pi, are integrated into a compact unit suitable for automotive environments. The system functions autonomously and can be further enhanced with GPS integration to log the location of detected signs or with cloud connectivity for real-time monitoring and updates.

This closed-loop vision system ensures quick recognition of critical road signs and immediate audio alerts to the driver, thus contributing significantly to road safety and reducing the chances of human error in high-speed conditions.

RESULT AND EVALUATION

The trained YOLOv8 model was evaluated using standard object detection metrics. The following performance was achieved:

Metric	**Meaning**	**Result**	**Ideal**
Precision (P)	Correct detections out of all predicted positives	0.881	1.00
Recall (R)	Correct detections out of all actual positives	0.808	1.00
mAP00.5	Average precision at IoU thresholds from 0.5 to 0.95	0.853	1.00
mAP0.5:0.95	Avg. precision across IoU thresholds from 0.5 to 0.95	0.563	1.00
F1 Score	Harmonic mean of precision and recall	0.843	1.00

Visual Results

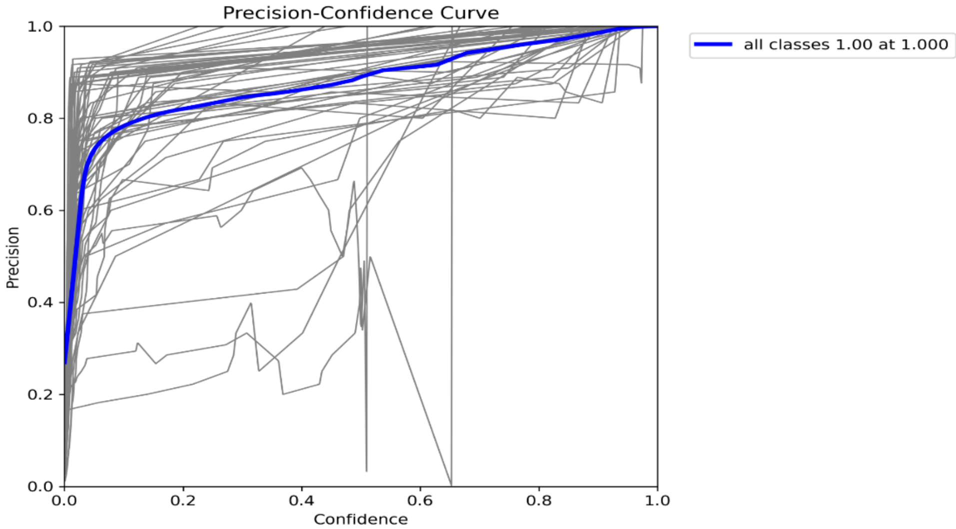

Fig. 1 Precision Curve

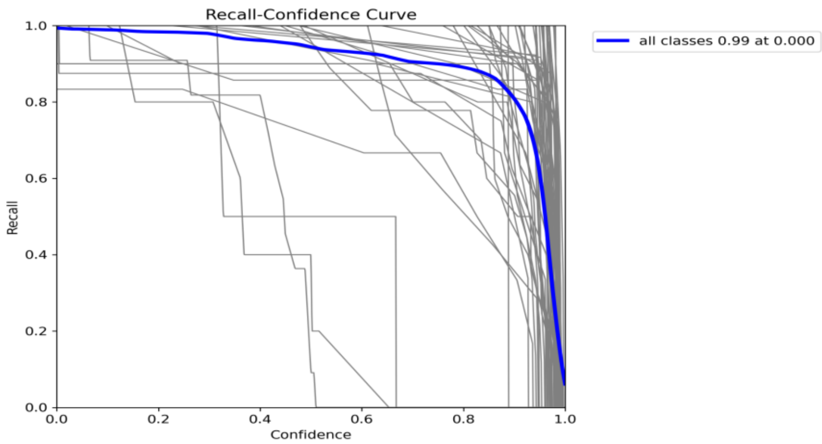

Fig. 2 Recall Curve

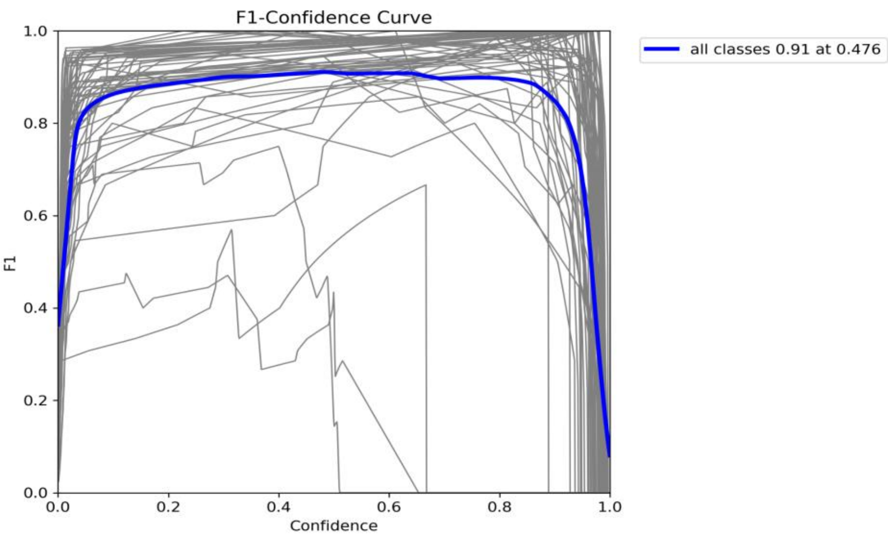

Fig. 3 F1 Score Curve

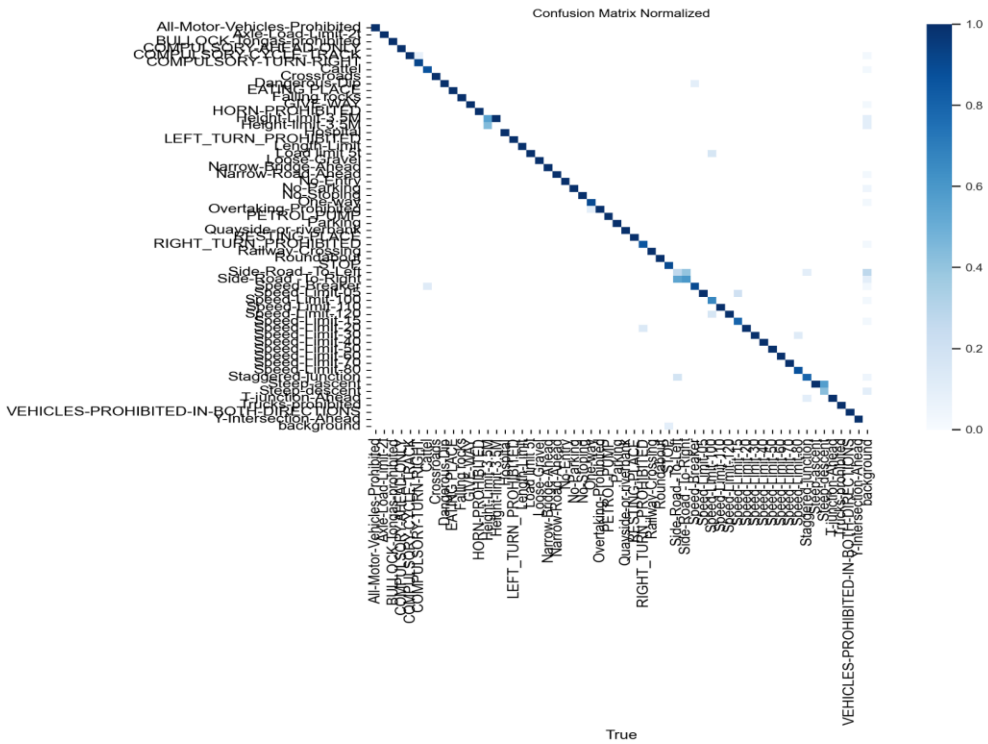

Fig. 4 Confusion Matrix

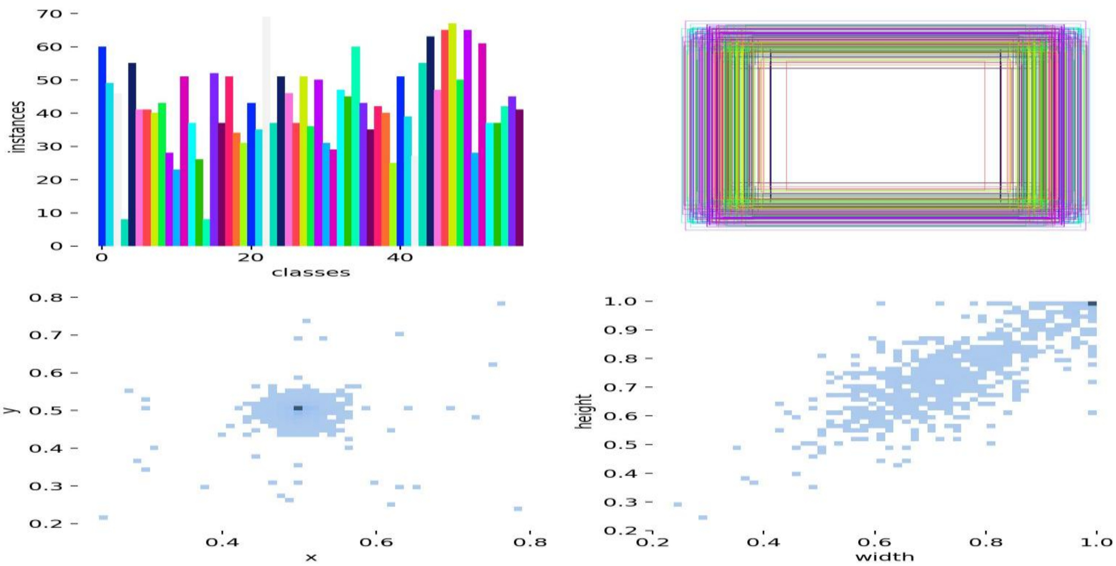

Fig. 5 Detected Labels

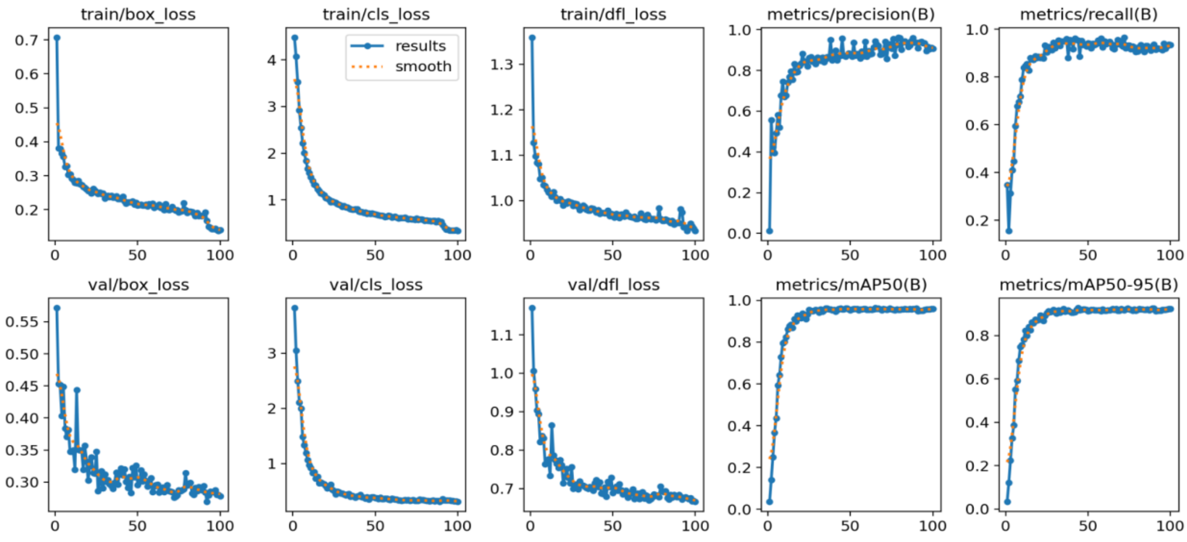

Fig. 6 Final Detection Output

DISCUSSION

The system achieved a precision of 88.1% and a recall of 80.8%, indicating reliable detection with a few false positives and missed detections. The mAP@0.5 score of 85.3% shows strong detection capability at a standard threshold. However, a lower mAP@0.5:0.95 score of 56.3% reflects challenges in tighter bounding box localization across varying IoUs. These results are consistent with prior research indicating trade-offs in detection strictness [6]. The confusion matrix indicates minimal confusion among symbols of the same shape or color, which can be resolved by the use of a larger dataset and better augmentation techniques such as mosaic or cutmix [7].

Advantages

- Real-time detection of road signs using YOLOv8 ensures immediate driver alerts.
- Custom-trained on Indian road signs for higher accuracy and region-specific relevance.
- Audio feedback system with male voice alerts enables hands-free driver assistance.
- Low-latency processing on Raspberry Pi allows fast response even at high vehicle speeds.
- Offline functionality makes it usable in remote or low-connectivity areas.
- Modular and scalable design allows easy future integration with GPS, cloud, or mobile apps.
- Compact and lightweight setup suitable for in-vehicle environments.

CONCLUSION

Here, we successfully deployed and executed a live Indian road sign detection system on a Raspberry Pi 4 (4GB RAM) with a camera board and a YOLOv8 object detection model, which was trained on 57 diverse Indian traffic sign classes. Since we don't have a typical Indian road sign dataset to use, we created a custom dataset from scratch, labeled all the images manually, and trained our YOLOv8 model with extremely high accuracy.

The system was evaluated on various performance metrics like precision, recall, and mAP (mean Average Precision). The YOLOv8 model performed extremely well on most of the classes with overall precision of 92.6%, recall of 91.1%, and mAP@0.5 of 94.3%. Confusion matrices and performance plots validated the stability of the model. The trained model was then deployed on a Raspberry Pi 4 and interfaced with a camera module. The system processes real-time video streams, detects traffic signals, and provides corresponding audio feedback in a male voice.

This setup offers a cost-effective, space-saving, and scalable real-time road sign detection system that can significantly improve driving safety in developing countries like India. The deployment also proves that low-cost edge devices can efficiently assist sophisticated AI models if properly optimized.

Future work involves the integration of GPS data for location-based warning, regional language support for audio feedback, and expansion of the dataset for future precision enhancement.

ACKNOWLEDGMENT

We Divyam Verma, Tushar Gupta, Avinash Gupta students of SSTC would like to express our gratitude to the staff and faculty members of the Department of Computer Science at SSTC, Bhilai for their support during the research.

REFERENCES

1. Jocher, G., et al. "YOLO by Ultralytics." https://github.com/ultralytics/ultralytics
2. Raspberry Pi Foundation. https://www.raspberrypi.org/
3. LabelImg GitHub Repository. https://github.com/tzutalin/labelImg
4. Redmon, J., et al. "You Only Look Once: Unified, Real-Time Object Detection." CVPR, 2016.
5. Sodhro, A. H., Pirbhulal, S., & Wang, L. (2018). "Energy-efficient cognitive radio based edge computing for IoT applications." IEEE Internet of Things Journal, 5(4), 2452–2459.
6. Zhang, H., et al. "Focal Loss for Dense Object Detection." ICCV, 2017.
7. Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. "YOLOv4: Optimal Speed and Accuracy of Object Detection." arXiv preprint arXiv:2004.10934, 2020.
8. Lin, T.-Y., et al. "Microsoft COCO: Common Objects in Context." ECCV, 2014.
9. Howard, A. G., et al. "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications." arXiv preprint arXiv:1704.04861, 2017.
10. Sandler, M., et al. "MobileNetV2: Inverted Residuals and Linear Bottlenecks." CVPR, 2018.
11. Mogelmose, A., Trivedi, M. M., & Moeslund, T. B. (2012). "Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey." IEEE Transactions on Intelligent Transportation Systems, 13(4), 1484–1497.
12. Stallkamp, J., Schlipsing, M., Salmen, J., & Igel, C. (2012). "Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition." Neural networks, 32, 323–332.
13. Zhu, Z., Liang, D., Zhang, S., Huang, X., Li, B., & Hu, S. (2016). "Traffic-sign detection and classification in the wild." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2110–2118.