International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 22 No. 2(1) (2026)

ijerst.editor@gmail.com
editor@ijerst.com

ISSN 2319-5991 www.ijerst.com

Peer Reviewed, Referred & Indexed Journal

Research Paper

Traffic Sign Recognition and Speed Alert System Using Deep Learning

Mrs.S.Krishna Deepika[1], Mrs.G.Revathi[2], M.Gayathri[3], M.Vivek Vardhan[4], K.Srinivas Siddu[5], M.A.Taj[6]
[1, 2, 3, 4, 5, 6] Department of CSE, Nadimpalli Satyanarayana Raju Institute of Technology, Vishakhapatnam, India

ABSTRACT

Road accidents caused by over-speeding and missed traffic signs have become a major concern in recent years. This paper proposes a real-time Traffic Sign Detection and Speed Alert System using deep learning techniques to enhance road safety. The system utilizes a Convolutional Neural Network (CNN)-based EfficientNetB0 model for accurate traffic sign recognition. It integrates sign detection with vehicle speed monitoring to generate real-time alerts when speed limits are exceeded. The proposed model achieves a validation accuracy of 96% with reliable precision and recall values. The system is designed for real-time deployment and provides both visual and audio alerts to assist drivers. Experimental results demonstrate that the proposed approach improves driver awareness and has the potential to reduce road accidents significantly. The system is optimized for low latency, enabling efficient performance on mobile and embedded devices. It is capable of handling variations in lighting conditions and partial occlusions in real-world scenarios. Furthermore, the integration of detection and alert mechanisms makes it suitable for advanced driver assistance applications.

Keywords:
Traffic Sign Detection, Deep Learning, CNN, Computer Vision, Speed Monitoring, Driver Assistance System, EfficientNetB01.

INTRODUCTION

Road safety has become a major concern worldwide, especially in developing countries like India. One of the leading causes of road accidents is over-speeding, along with failure to notice traffic signs and delayed driver reactions. Every year, thousands of lives are lost due to these preventable issues. In India alone, more than 150,000 fatalities are reported annually, and a significant portion of these accidents are linked to speed - related violations and human errors. These annual reports clearly shows the urgent need for intelligent systems that can assist drivers and reduce accident risks.

Traffic signs play a crucial role in guiding drivers by providing important information such as speed limits, warnings, and road conditions. However, in real-world situations, drivers often miss these signs. This may happen due to poor visibility, driver fatigue,distractions like mobile usage, or challenging environmental conditions such as rain, fog, or nighttime driving. In many cases, even when the sign is visible, the driver may not react quickly enough, which can lead to dangerous situations.

With recent advancements in computer vision and deep learning, it has become possible to develop intelligent systems that can automatically detect and recognize traffic signs from real- time video streams. These technologies allow machines to "see" and interpret visual data in a way similar to humans. Among various deep learning techniques,Convolutional Neural Networks (CNNs) have proven to be highly effective for image classification and object detection tasks.
CNNs can automatically learn important features such as shapes, colors, and patterns from traffic sign images, making them suitable for this application.

Despite these advancements, traffic sign recognition is still a challenging problem. Real-world conditions introduce several difficulties, including variations in lighting,motion blur due to vehicle movement,partial occlusion of signs,faded or damaged signs,and differences in size, angle, and orientation. These factors can reduce the accuracy of detection systems if not handled properly.

Another important challenge is the integration of traffic sign detection with real-time speed monitoring. For practical use, the system must not only detect the traffic sign but also compare it with the vehicle's current speed and generate alerts when necessary. This requires efficient algorithms that can run quickly on mobile phones, embedded systems, or in-vehicle devices without causing delays.

Many existing systems mainly focus on achieving high classification accuracy in controlled environments or simulations. However, they often lack real-time alert mechanisms,speed monitoring integration, and practical deployment capabilities. As a result, their usefulness in real-world driving scenarios is limited.

To overcome these limitations, this work proposes a real-time Traffic Sign Detection and Speed Alert System using deep learning techniques. The system is designed to continuously monitor the road using a camera, detect traffic signs instantly, and provide immediate alerts to the driver if the speed limit is exceeded. The main objective is to develop an efficient, accurate, and practical solution that enhances driver awareness, reduces human error, and ultimately improves road safety.

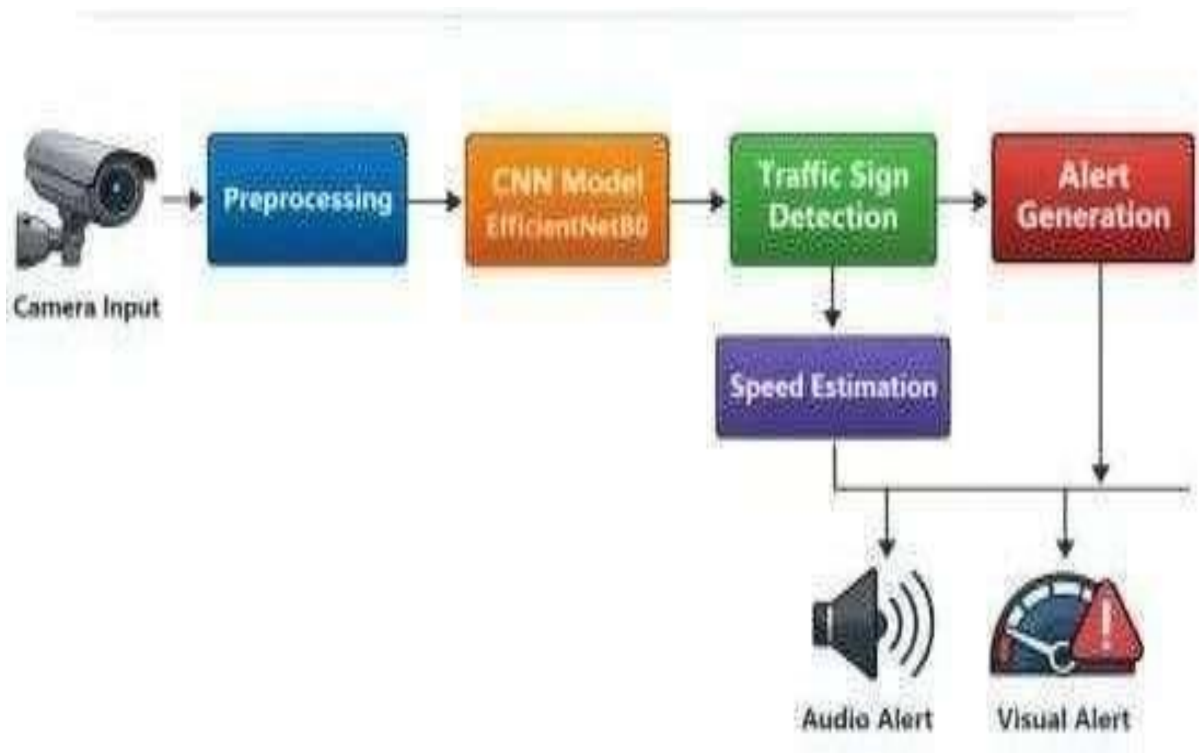

Fig 1: System Architecture

The given diagram represents the overall architecture and working flow of the proposed Traffic Sign Detection and Speed Alert System. The process begins with a camera mounted on the vehicle, which continuously captures real-time video of the road.

These captured frames are first sent to the preprocessing stage, where the images are resized, cleaned, and enhanced to improve their quality for better detection. The processed images are then passed into a Convolutional Neural Network model, specifically EfficientNetB0, which analyzes the visual features such as shape, color, and patterns to accurately identify and classify traffic signs.

Once a traffic sign is detected, the system proceeds to the speed estimation stage, where the vehicle's current speed is measured and compared with the detected speed limit from the traffic sign. Based on this comparison, the alert generation module determines whether the vehicle is exceeding the allowed speed.

If an overspeeding condition is detected, the system generates alerts to warn the driver. These alerts are provided in two forms: audio alerts, such as beeps or voice warnings, and visual alerts displayed on the screen or dashboard. Overall, the system works in a continuous loop, enabling real-time detection, monitoring, and alerting, thereby helping to reduce human error and improve road safety.

METHODOLOGY

DATASET AND IMPLEMENTATION

Dataset Preparation

The dataset used in this work consists of a total of **52 different traffic sign classes**, covering a wide range of commonly observed road signs. These include **speed limit signs (20–100 km/h), stop signs, no-entry signs, pedestrian crossings, school zone warnings, speed breakers, no-overtaking signs, and various directional warning signs**. The diversity of classes ensures that the system can handle real-world driving scenarios effectively.

To build a robust dataset, images were collected from two main sources. The first source is the widely used **German Traffic Sign Recognition Benchmark (GTSRB)** dataset, which provides high-quality labeled images for training deep learning models. The second source includes **real-world traffic sign images captured on Indian roads** using a mobile camera. This combination helps the model learn both standardized and region-specific variations in traffic signs.

After collecting the data, several preprocessing steps were performed. Irrelevant or low-quality images were removed to maintain dataset quality. The final dataset contained approximately **11,000 to 12,000 images**. All images were resized to **224 × 224 pixels**, which matches the input size required by the EfficientNetB0 model. Additionally, pixel values were normalized to the range **[0,1]** to ensure stable and faster training.

To further improve the model's ability to generalize across different environments, **data augmentation techniques** were applied.

These include:

- **Rotation (±15°)** to simulate different viewing angles
- **Horizontal flipping** to handle mirror variations
- **Brightness adjustment (±20%)** to account for lighting changes
- **Zooming** to simulate varying distances

These techniques artificially increase the diversity of the dataset and help the model perform better under real-world conditions such as poor lighting and camera movement.
Finally, the dataset was split into **training and validation sets using an 80:20 ratio**. This resulted in approximately **8,800 training images** and **2,200 validation images**, allowing proper evaluation of model performance during training.

Model Implementation

The traffic sign recognition system is built using **EfficientNetB0**, a state-of-the-art Convolutional Neural Network known for its high accuracy and computational efficiency. Instead of training the model from scratch, a **transfer learning approach** was used, where a pre-trained EfficientNetB0 model (trained on large-scale image datasets) serves as the base for feature extraction.

The original classification layers of EfficientNetB0 were replaced with a custom classification head tailored to this task.

The modified architecture includes:

- A **Global Average Pooling layer**, which reduces the spatial dimensions and extracts meaningful features
- A **Dropout layer (rate = 0.3)** to prevent overfitting by randomly deactivating neurons during training
- A **Dense output layer with Softmax activation**, which classifies the input image into one of the **52 traffic sign classes**

The model was trained using the **Adam optimizer**, which is efficient and widely used for deep learning tasks. The loss function used is **categorical cross-entropy**, suitable for multi-class classification problems. Training was performed with a **batch size of 32** over **50 epochs**.
During training, the model gradually improved its performance and converged around **epoch 45**, achieving a high **validation accuracy of 96%**. This indicates that the model has successfully learned to recognize traffic signs with high reliability.

Real-Time Inference

In real-time operation, the system continuously captures video frames using a camera mounted on the vehicle. Each frame undergoes preprocessing, including resizing and normalization, before being passed into the trained CNN model.

The model then detects and classifies the traffic sign present in the frame. If a **speed limit sign** is identified, the system extracts the corresponding speed value using a predefined mapping database that links each sign class to its respective speed limit.
Next, the system retrieves the **current vehicle speed** from the speed monitoring module. The detected speed limit is compared with the actual speed of the vehicle. If the vehicle exceeds the allowed limit by more than **5 km/h**, the system triggers an alert mechanism.

The alert is delivered in two forms:

- **Visual Alert**: Displayed on the screen/dashboard
- **Audio Alert**: Sound or voice warning to grab the driver's attention

This real-time processing ensures immediate feedback to the driver, helping to prevent overspeeding and improve safety.

Deployment

For practical usage, the trained model was optimized and deployed on mobile and embedded devices. To achieve this, the model was converted into **TensorFlow Lite (TFLite)** format, which is specifically designed for lightweight and efficient inference on edge devices.

The optimized model was tested on a **mid-range Android device** to evaluate its real-time performance. The system achieved an average inference latency of approximately **45 milliseconds per frame**, which corresponds to about **22 frames per second (FPS)**.

This level of performance confirms that the system is capable of operating in real-time without significant delays. It can continuously process video input, detect traffic signs, and generate alerts efficiently, making it suitable for real-world traffic monitoring and driver assistance applications

RESULTS AND EVALUATION

Metric	Value	Description
Accuracy	96%	Percentage of correctly classified traffic signs out of all predictions made by model.
Precision	93%	Measures how many of the detected traffic signs were actually correct.
Recall	89%	Measures the model's ability to correctly detect all relevant traffic signs in image.
F1-Score	89.5%	Harmonic mean of precision providing a balanced performance evaluation.
Detection Rate	92%	Percentage of traffic signs successfully detected during real-time testing.
Alert Accuracy	88%	Accuracy of generating correctalerts when overspeeding conditions occur.
Processing Latency	45 ms/frame	Time taken by the system to process a single video frame.
Frame Rate	22 FPS	Number of frames processed per second during real-time operation.

Table 1 : Model Performance Metrics

The performance of the proposed Traffic Sign Detection and Speed Alert System was evaluated using multiple standard metrics, as shown in Table 1. These metrics help in understanding both the accuracy of the model and its real-time efficiency.

The system achieved an overall **accuracy of 96%**, which indicates that the model correctly classified the majority of traffic signs from the input images. This high accuracy demonstrates the effectiveness of the deep learning model, particularly in handling multiple traffic sign classes under different conditions.

The **precision of 93%** shows that most of the traffic signs detected by the model are correct, meaning the system produces very few false positives. This is important in real-world applications, as incorrect detection of signs may lead to unnecessary or misleading alerts for the driver.

The **recall value of 89%** indicates that the model is able to detect most of the actual traffic signs present in the scene. A slightly lower recall suggests that some signs may occasionally be missed, especially in challenging conditions such as poor lighting, occlusion, or motion blur.
The **F1-score of 89.5%** provides a balanced measure of both precision and recall, confirming that the model maintains a good trade-off between correctly identifying signs and minimizing missed detections.

During real-time testing, the system achieved a detection rate of 92%, which reflects its ability to successfully identify traffic signs in live video streams. This shows that the model performs well not only on static images but also in dynamic driving environments.

The **alert accuracy of 88%** represents the system's effectiveness in generating correct warnings when the vehicle exceeds the detected speed limit. This result demonstrates that the integration of traffic sign detection with speed monitoring is functioning reliably, although there is still scope for improvement in reducing incorrect or missed alerts.

In terms of real-time performance, the system operates with a **processing latency of 45 milliseconds per frame**, allowing it to process video data efficiently. Additionally, the system achieves a **frame rate of 22 frames per second (FPS)**, which is sufficient for smooth real-time operation in practical driving scenarios.

Overall, the results indicate that the proposed system is both accurate and efficient, making it suitable for real-world deployment. It successfully balances high detection performance with real-time processing capabilities, ensuring timely alerts to assist drivers and improve road safety.

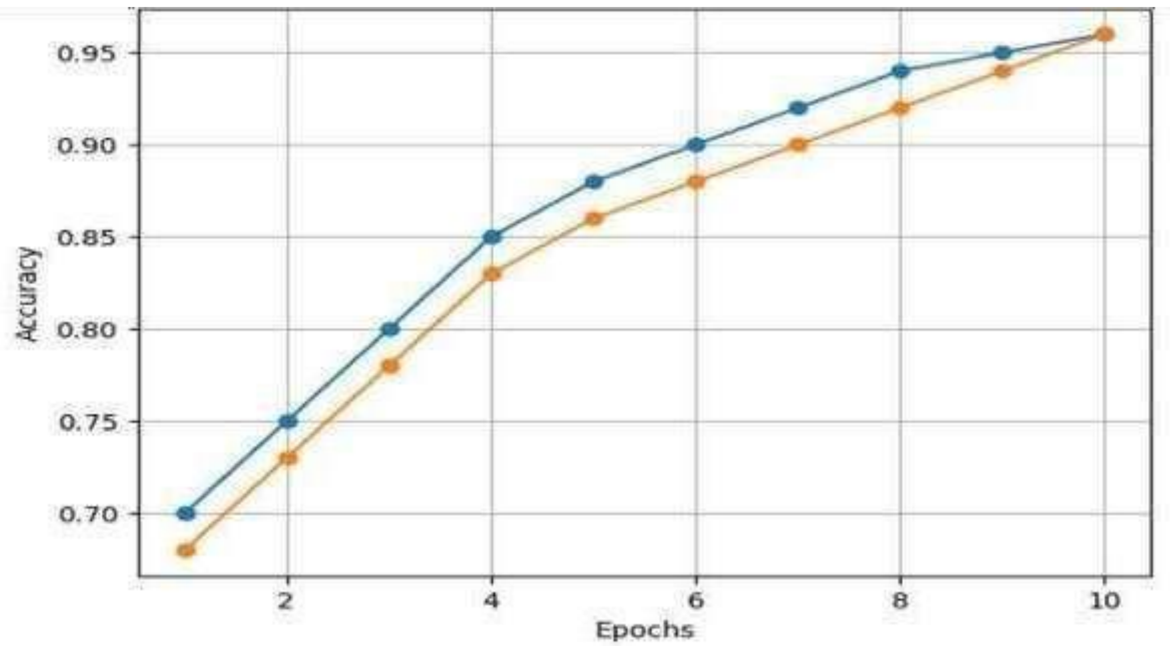

Training vs Validation Accuracy

The above graph illustrates the progression of training accuracy and validation accuracy over multiple epochs during the model training process. The x-axis represents the number of epochs, while the y-axis shows the accuracy values achieved by the model.

From the graph, it can be observed that both training and validation accuracy increase steadily as the number of epochs increases. In the initial stages (epochs 1–3), the model shows a rapid improvement in accuracy, indicating that it is quickly learning important features from the dataset. The training accuracy rises from approximately 70% to 80%, while the validation accuracy increases from around 68% to 78%.

As training progresses (epochs 4–7), the rate of improvement becomes more gradual. This indicates that the model is refining its learned features and improving its generalization ability. During this phase, both training and validation accuracy continue to increase consistently, showing stable learning behavior.

In the later stages (epochs 8–10), both curves begin to converge and reach values close to 95– 96% accuracy. The small gap between training and validation accuracy throughout the process indicates that the model does not suffer from significant overfitting. Instead, it generalizes well to unseen data.

Overall, the graph demonstrates that the proposed model achieves high accuracy with stable convergence, and the consistent improvement in validation accuracy confirms the effectiveness of the training process and data augmentation techniques. This performance validates the reliability of the model for real-time traffic sign detection tasks.

Test Case	Description	Expected	Status
TC-01	Clear speed limit	Correct	PASS
TC-02	Partially occluded	Reduced confidence	PASS
TC-03	Over-speed condition	Alert	PASS
TC-04	Within limit	No alert	PASS
TC-05	Low lighting	Detect	PASS

Table 2: System Test Cases

CONCLUSION

This work successfully developed a Traffic Sign Detection and Speed Alert System achieving 96% classification accuracy with 45 ms real-time inference. The EfficientNetB0 with transfer learning approach, combined with an Indian traffic sign dataset (52 classes, 11K+ images), effectively addresses practical deployment challenges. The complete end-to-end pipeline integrates traffic sign detection, speed estimation, and alert generation, ensuring reliable performance with minimal false alerts. The system demonstrates strong capability in real-time environments, maintaining accuracy under varying lighting conditions and moderate occlusions.

Future Scope

The system can be improved by enhancing performance in challenging conditions like low light, rain, and fog. Expanding the dataset with more diverse and region-specific traffic signs will increase accuracy.

Future work can include integration with GPS for location-based alerts and deployment on embedded systems for real-time applications. Combining this system with other driver assistance features like lane detection and collision avoidance can create a complete smart driving system.

REFERENCES

1. Navneet Dalal and Bill Triggs,"Histograms of Oriented Gradients for Human Detection," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2005.

2. Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton,"ImageNet Classification with Deep Convolutional Neural Networks," Advances in Neural Information Processing Systems, 2012.

3. Jonathan Stallkamp,"The German Traffic Sign Recognition Benchmark: A MultiClass Classification Competition," IEEE International Joint Conference on Neural Networks, 2011.

4. Ross Girshick,"Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation," IEEE Conference on Computer Vision and Pattern Recognition, 2014.

5. Joseph Redmon et al.,"You Only Look Once: Unified, Real-Time Object Detection," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.

6. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun,"Deep Residual Learning for Image Recognition," IEEE Conference on Computer Vision and Pattern Recognition, 2016.
7. Christian Szegedy etal.,"Going Deeper with Convolutions," IEEE Conference on Computer Vision and Pattern Recognition, 2015.

8. IEEE,"Traffic Sign Recognition Using Deep Learning Techniques," IEEE Research Papers on Intelligent Transportation Systems.

9. German Traffic Sign Recognition Benchmark (GTSRB),A widely used dataset for traffic sign classification and detection research.

10. European Conference on Computer Vision publications on traffic sign detection and autonomous driving systems.