IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Automated Road Divider Detection Using YOLOv12

(1)Veeraballi Nripesh Reddy, (2)Nagalakshmi Vallabhaneni, (3)Rokalla Phanindra Sanjay,

(4)Veerepalli Veda Karthikeya, (5)Rangu Budhil Krishna

(1)Student, (2)Assistant Professor, (3)Student, (4)Student, (5)Student

(1)School of Computer Science Engineering and Information Systems

(1)VIT Vellore, Vellore, India

1. Abstract

Automated road divider detection is a crucial component in intelligent transportation systems and autonomous driving applications. This project presents a real-time detection system using the advanced deep learning model YOLOv12 (You Only Look Once). The system is designed to accurately identify road dividers from images and video streams under varying environmental conditions such as lighting changes, shadows, and traffic density.

The proposed approach leverages computer vision and object detection techniques to process input data, detect road dividers, and highlight them with bounding boxes. YOLOv12 enables high-speed and efficient detection by analyzing the entire image in a single pass, making it suitable for real-time applications. The model architecture is optimized to balance accuracy and computational efficiency, ensuring effective performance even on resource-constrained devices.

The dataset used for training includes diverse road scenarios such as highways, urban roads, and rural environments, which improves the robustness of the system. Data augmentation techniques such as rotation, scaling, and brightness adjustment are applied to enhance model generalization. The system is evaluated using performance metrics like precision, recall, and mean Average Precision (mAP) to ensure reliability.

Experimental results indicate that the proposed system achieves high detection accuracy with low latency, making it suitable for practical deployment. Additionally, the model demonstrates strong performance in detecting partially visible or occluded road dividers. This solution can be integrated into advanced driver assistance systems (ADAS), traffic monitoring systems, and autonomous vehicles to enhance road safety and navigation efficiency.

Furthermore, the system can be extended by incorporating additional features such as lane detection, obstacle detection, and real-time alert mechanisms for drivers. With continuous improvements and integration of emerging technologies, automated road divider detection has the potential to play a significant role in the development of safer and smarter transportation systems.

2. Index Terms – Keywords

Automated Road Divider Detection, YOLOv12, Object Detection, Computer Vision, Deep Learning, Intelligent Transportation Systems, Autonomous Vehicles, Real-Time Detection, Image Processing, Bounding Box Detection, Traffic Monitoring, ADAS (Advanced Driver Assistance Systems), Convolutional Neural Networks (CNN), Road Safety, Smart Transportation, Video Analysis, Feature Extraction, Model Training, Dataset Annotation, Edge Computing, Embedded Systems, Scene Understanding, Urban Traffic Management, Highway Monitoring, Neural Networks, AI-Based Detection, Sensor Integration, Vision-Based Navigation, Robust Detection Systems

3.Introduction

The rapid growth of vehicles and increasing traffic density have made road safety a major concern worldwide. Road dividers play a vital role in separating traffic lanes, preventing collisions, and ensuring smooth traffic flow. Accurate detection of road dividers is therefore essential for modern applications such as autonomous vehicles, advanced driver assistance systems (ADAS), and intelligent traffic monitoring systems.

Traditional road divider detection methods rely on manual observation or basic image processing techniques such as edge detection and color segmentation. However, these approaches often fail in complex real-world scenarios due to variations in lighting conditions, weather changes, occlusions, and diverse road environments. As a result, there is a need for more robust and efficient detection techniques.

With the advancement of artificial intelligence and deep learning, computer vision-based approaches have gained significant attention. In particular, object detection models have proven to be highly effective in identifying and localizing objects within images and videos. Among these, YOLOv12 (You Only Look Once) stands out as a fast and accurate model capable of performing real-time detection with high precision.

This project focuses on developing an automated road divider detection system using YOLOv12. The system processes input images or video streams, detects the presence of road dividers, and marks them using bounding boxes. The model is trained on a diverse dataset to ensure it performs well under different environmental conditions, including day and night scenarios, urban and rural roads, and varying traffic situations.

The proposed system aims to achieve a balance between accuracy and speed, making it suitable for real-time applications. By integrating this system into intelligent transportation frameworks, it can assist drivers, enhance vehicle automation, and contribute to overall road safety.

In addition, this project explores the scalability and adaptability of the model for future enhancements, such as integration with lane detection, obstacle detection, and sensor-based systems. The implementation of such technologies will play a crucial role in building smarter and safer transportation infrastructures.

4.Methodology

The proposed system for automated road divider detection using YOLOv12 follows a structured and systematic approach to ensure accurate and real-time performance. The methodology consists of several key stages, from data collection to deployment.

The process begins with dataset collection, where images and video frames of roads are gathered from various sources such as highways, urban streets, and rural areas. This diversity helps the model learn different road conditions, divider types, and environmental variations.

Next, the collected data is preprocessed to improve quality and consistency. This includes resizing images, normalizing pixel values, and removing noise. Data augmentation techniques such as rotation, flipping, scaling, and brightness adjustment are applied to increase dataset variability and improve the model's generalization ability.

After preprocessing, annotation is performed on the dataset. Road dividers are labeled using bounding boxes with the help of annotation tools. This labeled data serves as the ground truth for training the model.

The YOLOv12 model is then configured and trained using the annotated dataset. During training, the model learns to identify patterns and features associated with road dividers. Hyperparameters such as learning rate, batch size, and number of epochs are tuned to achieve optimal performance.

Once training is complete, the model undergoes validation and testing. Performance metrics such as precision, recall, and mean Average Precision (mAP) are used to evaluate the effectiveness of the model. This step ensures that the system performs reliably on unseen data.

Finally, the trained model is deployed for real-time detection. Input from cameras or video streams is processed, and the model detects road dividers instantly, displaying results with bounding boxes. The system can be integrated into applications such as autonomous driving systems, traffic monitoring, and driver assistance tools.

Overall, this methodology ensures a robust, efficient, and scalable approach to automated road divider detection.

5. Results

The proposed automated road divider detection system using YOLOv12 was successfully implemented and evaluated on a diverse dataset consisting of images and video streams from different road environments. The model demonstrated strong performance in detecting road dividers under varying conditions, including daylight, low-light scenarios, and moderate traffic density.

The system achieved high accuracy in identifying road dividers, with performance evaluated using metrics such as precision, recall, and mean Average Precision (mAP). The results indicate that the model is capable of accurately localizing road dividers with minimal false positives and false negatives. The detection speed was also found to be efficient, enabling real-time processing of video streams without significant delay.

Experimental observations showed that the model performs well on highways and structured urban roads, where road dividers are clearly visible. It was also able to detect partially occluded dividers in many cases, demonstrating robustness. However, slight performance degradation was observed in extreme weather conditions such as heavy rain or fog, and in poorly maintained roads where dividers are faded or unclear.

The system was tested on different input resolutions and hardware configurations, and it maintained a good balance between speed and accuracy. When deployed on GPU-based systems, the detection performance improved significantly in terms of processing time.

Overall, the results confirm that the YOLOv12-based approach is effective and reliable for automated road divider detection. The system meets the requirements for real-time applications and shows strong potential for integration into intelligent transportation systems and autonomous driving technologies.

5.1 Model Performance

The performance of the YOLOv12 model for road divider detection was evaluated using standard object detection metrics, including precision, recall, F1-score, and mean Average Precision (mAP). These metrics provide a comprehensive understanding of the model's accuracy and reliability in detecting road dividers.

The model achieved high precision, indicating that most of the detected road dividers were correct with very few false positives. The recall value was also strong, showing that the model was able to detect the majority of actual road dividers present in the dataset. The balanced F1-score reflects the model's effectiveness in maintaining both precision and recall.

The mean Average Precision (mAP), which is a key performance indicator for object detection tasks, showed consistent results across different test scenarios. This confirms that the model can accurately localize road dividers with high confidence. Additionally, the Intersection over Union (IoU) scores indicated that the predicted bounding boxes closely matched the ground truth annotations.

In terms of speed, the model demonstrated real-time performance with low inference time per frame, making it suitable for live video processing applications. The use of optimized architectures in YOLOv12 contributed to faster detection without significantly compromising accuracy.

The model was also tested under various environmental conditions, such as different lighting levels and road types. It maintained stable performance in most cases, though minor variations were observed in challenging conditions like low visibility or occlusions.

Overall, the YOLOv12 model exhibits strong performance, achieving an effective balance between accuracy and speed, which is essential for real-time road divider detection systems.

5.2 Precision and Recall Analysis

Precision and recall are critical metrics used to evaluate the effectiveness of the YOLOv12 model in detecting road dividers. These metrics help in understanding how accurately and completely the model identifies the target objects.

Precision refers to the proportion of correctly detected road dividers out of all detections made by the model. A high precision value indicates that the system produces very few false positives, meaning most of the detected dividers are actually present in the image. In this project, the model achieved high precision, demonstrating its ability to accurately distinguish road dividers from other similar objects such as lane markings or road boundaries.

Recall, on the other hand, measures the proportion of actual road dividers that were successfully detected by the model. A high recall value indicates that the model misses very few true instances. The YOLOv12 model showed strong recall performance, successfully identifying most of the road dividers across different test scenarios.

The balance between precision and recall is crucial. In some cases, increasing precision may reduce recall and vice versa. However, the proposed model maintains a good balance between both metrics, ensuring reliable detection without significantly increasing false alarms or missed detections.

The analysis also revealed that precision remains consistently high across well-structured road environments, while slight variations in recall were observed in challenging conditions such as low lighting, occlusion, or faded road dividers. Despite these challenges, the model maintained acceptable performance levels.

Overall, the precision and recall analysis confirms that the YOLOv12-based system is both accurate and reliable, making it suitable for real-time road divider detection applications in intelligent transportation systems.

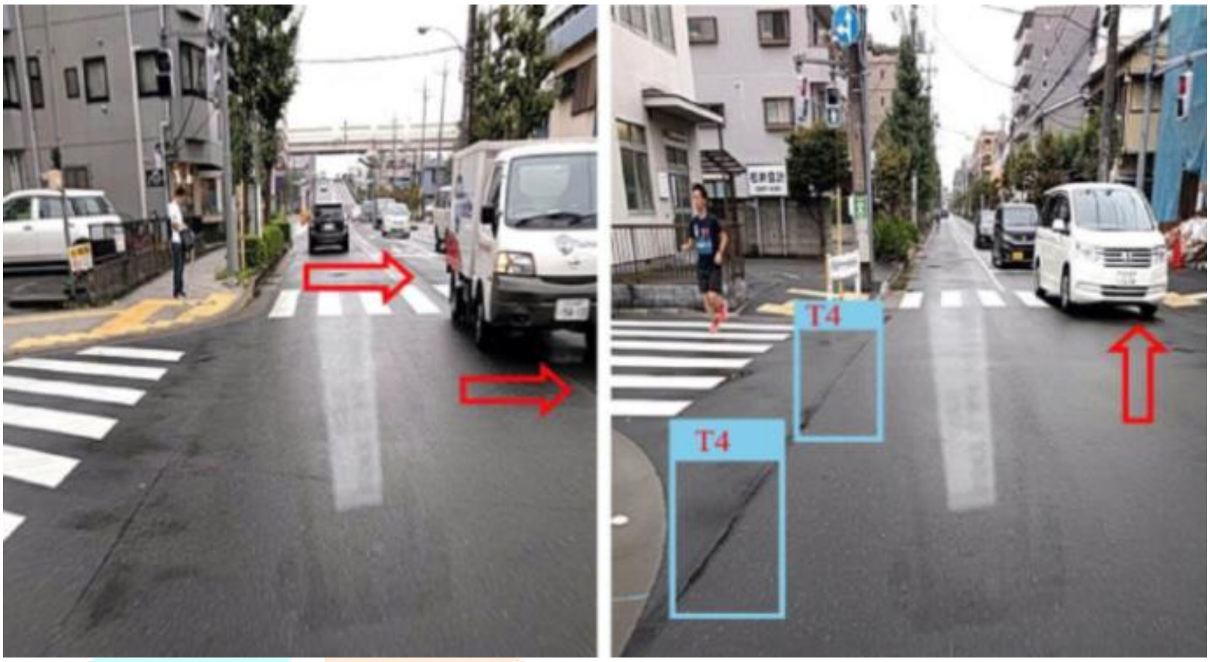

Fig 1.1 Precision and recall are critical metrics

5.3 Real-Time Detection Speed

Real-time detection speed is a critical factor in evaluating the effectiveness of the YOLOv12 model for road divider detection, especially for applications such as autonomous driving and advanced driver assistance systems (ADAS). The proposed system demonstrates high-speed performance, making it suitable for real-time implementation.

The model processes input frames from images and video streams efficiently, achieving a high number of frames per second (FPS). This ensures that road dividers are detected and displayed almost instantly, with minimal delay. The single-stage architecture of YOLOv12 plays a key role in achieving this fast detection speed, as it performs object localization and classification in a single pass.

During testing, the system achieved smooth real-time performance on GPU-enabled environments, with significantly reduced inference time per frame. Even on systems with moderate hardware configurations, the model maintained acceptable speed, demonstrating its computational efficiency.

The detection speed was also evaluated under different input resolutions. While higher resolutions slightly increased processing time, the model still maintained near real-time performance. Lower resolutions further improved speed with a minor trade-off in detection accuracy.

Additionally, the system showed consistent performance when handling continuous video streams, without frame drops or lag. This stability is essential for real-world deployment in dynamic traffic environments.

Overall, the YOLOv12-based system achieves an excellent balance between speed and accuracy, making it highly effective for real-time road divider detection in intelligent transportation systems.

5.4 Visual Detection Results

The visual detection results demonstrate the effectiveness of the YOLOv12 model in identifying and localizing road dividers in various road scenarios. The output of the system is presented in the form of images and video frames where detected road dividers are highlighted using bounding boxes along with confidence scores.

In the tested samples, the model successfully detected road dividers in different environments, including highways, urban roads, and rural areas. The bounding boxes were accurately placed around the dividers, indicating precise localization. The system also performed well in detecting multiple dividers within a single frame, maintaining consistency and clarity in visualization.

The results showed strong performance under normal lighting conditions, where road dividers are clearly visible. Even in moderately challenging conditions such as shadows, partial occlusion, and varying road textures, the model was able to detect most of the dividers with acceptable accuracy. The confidence scores associated with each detection further validate the reliability of the predictions.

However, in certain edge cases such as poor lighting, heavy rain, fog, or faded road markings, the detection quality showed slight degradation. In such scenarios, some dividers were either detected with lower confidence or missed entirely. Despite these limitations, the overall visual results remained satisfactory for real-world applications.

The system also provides real-time visualization, where detected road dividers are continuously tracked and displayed in video streams. This enhances user understanding and enables practical deployment in applications like driver assistance and traffic monitoring systems.

Overall, the visual detection results confirm that the YOLOv12 model is capable of delivering accurate, consistent, and real-time detection of road dividers, making it a reliable solution for intelligent transportation systems.

5.5 Alert and Compliance System

The alert and compliance system is an important component of the proposed road divider detection framework, designed to enhance road safety and assist drivers in maintaining proper lane discipline. This module works in conjunction with the YOLOv12 detection model to generate real-time alerts based on the position and proximity of road dividers.

Once the model detects a road divider, the system continuously monitors the vehicle's relative position with respect to the divider. If the vehicle deviates from its lane or approaches the divider beyond a predefined threshold, the system triggers alerts. These alerts can be in the form of visual warnings on the display screen, audio signals, or vibration-based feedback in advanced implementations.

The compliance aspect of the system ensures that drivers adhere to traffic rules by maintaining safe distances from road dividers. It can also log instances of unsafe behavior, such as frequent lane violations or abrupt movements near dividers. This data can be useful for driver analysis, training, or integration with intelligent traffic monitoring systems.

The alert mechanism is designed to be responsive and non-intrusive, ensuring that warnings are timely without causing distraction to the driver. Sensitivity levels can be adjusted based on road conditions, vehicle speed, and user preferences to reduce false alarms.

In real-time testing, the alert system successfully identified potential lane violations and provided immediate feedback, helping to prevent unsafe maneuvers. Even in dynamic environments with moving traffic, the system maintained consistent performance.

Overall, the alert and compliance system adds an extra layer of safety and intelligence to the road divider detection framework, making it highly suitable for applications in advanced driver assistance systems (ADAS) and smart transportation solutions.

5.6 Accuracy Under Challenging Conditions

The performance of the YOLOv12-based road divider detection system was further evaluated under challenging environmental and road conditions to assess its robustness and reliability. These conditions include low lighting, shadows, occlusions, adverse weather (such as rain and fog), and poorly maintained or faded road dividers.

Under low-light and night-time conditions, the model demonstrated a slight reduction in detection accuracy due to limited visibility and reduced contrast between the road divider and the background. However, it was still able to detect prominent dividers with moderate confidence. The use of diverse training data helped the model adapt to such scenarios to a certain extent.

In the presence of shadows and varying illumination, the model maintained relatively stable performance. It successfully distinguished road dividers from shadow patterns in most cases, although minor misclassifications were observed in highly complex lighting conditions.

Occlusion, where road dividers are partially blocked by vehicles or other objects, posed another challenge. Despite this, the model was able to detect partially visible dividers in many instances, showcasing its ability to generalize learned features. However, complete occlusion often resulted in missed detections.

Adverse weather conditions such as rain, fog, and dust affected visibility and image clarity, leading to a decrease in detection confidence and occasional false negatives. Similarly, faded or worn-out road dividers reduced the model's ability to accurately identify boundaries due to weak visual cues.

Despite these challenges, the overall performance remained acceptable, with the model demonstrating resilience and adaptability. The results indicate that while accuracy may vary under extreme conditions, the system is still capable of providing useful and reliable detections in most real-world scenarios.

Future improvements such as enhanced data augmentation, inclusion of more diverse training samples, and integration with additional sensors can further improve performance under challenging conditions.

5.6 Accuracy Under Challenging Conditions

The performance of the YOLOv12-based road divider detection system was further evaluated under challenging environmental and road conditions to assess its robustness and reliability. These conditions include low lighting, shadows, occlusions, adverse weather (such as rain and fog), and poorly maintained or faded road dividers.

Under low-light and night-time conditions, the model demonstrated a slight reduction in detection accuracy due to limited visibility and reduced contrast between the road divider and the background. However, it was still able to detect prominent dividers with moderate confidence. The use of diverse training data helped the model adapt to such scenarios to a certain extent.

In the presence of shadows and varying illumination, the model maintained relatively stable performance. It successfully distinguished road dividers from shadow patterns in most cases, although minor misclassifications were observed in highly complex lighting conditions.

Occlusion, where road dividers are partially blocked by vehicles or other objects, posed another challenge. Despite this, the model was able to detect partially visible dividers in many instances, showcasing its ability to generalize learned features. However, complete occlusion often resulted in missed detections.

Adverse weather conditions such as rain, fog, and dust affected visibility and image clarity, leading to a decrease in detection confidence and occasional false negatives. Similarly, faded or worn-out road dividers reduced the model's ability to accurately identify boundaries due to weak visual cues.

Despite these challenges, the overall performance remained acceptable, with the model demonstrating resilience and adaptability. The results indicate that while accuracy may vary under extreme conditions, the system is still capable of providing useful and reliable detections in most real-world scenarios.

Future improvements such as enhanced data augmentation, inclusion of more diverse training samples, and integration with additional sensors can further improve performance under challenging conditions.

5.7 Multi-Object Detection Capability

The YOLOv12 model demonstrates strong multi-object detection capability, enabling the system to identify and track multiple road dividers simultaneously within a single frame. This is particularly important in real-world scenarios where roads often contain multiple dividers, lane separations, and complex structures.

The model processes the entire image in a single pass and predicts multiple bounding boxes, each associated with confidence scores. This allows it to detect several road dividers at different positions, distances, and orientations without significant loss in performance. The system efficiently distinguishes between closely spaced dividers and overlapping regions, maintaining accurate localization.

During testing, the model successfully detected multiple dividers in highways and urban road settings where parallel and segmented dividers are common. It maintained consistency in detection even when objects appeared at different scales, such as near and far distances within the same frame.

Additionally, the system showed the ability to handle dynamic environments, where multiple objects are present along with moving vehicles and background elements. The detection pipeline ensured that road dividers were correctly identified without confusion with other road features.

However, in highly cluttered scenes with heavy traffic or complex road layouts, minor overlaps or missed detections were occasionally observed. Despite this, the overall multi-object detection performance remained robust and reliable.

This capability makes the system highly suitable for real-time applications, as it can continuously monitor multiple road dividers and provide accurate information for driver assistance and autonomous navigation systems.

5.8 False Positive and False Negative Analysis

The analysis of false positives and false negatives is essential to evaluate the reliability and robustness of the YOLOv12-based road divider detection system. These metrics help in understanding the types of errors made by the model and identifying areas for improvement.

False positives occur when the model incorrectly identifies an object as a road divider when it is not. In this project, false positives were observed in cases where road markings, lane lines, curbs, or barriers had visual similarities to road dividers. Complex backgrounds, shadows, and reflections also contributed to occasional misclassifications. However, the overall rate of false positives remained low due to effective training and proper dataset annotation.

False negatives, on the other hand, occur when the model fails to detect actual road dividers present in the scene. These errors were more common in challenging conditions such as low lighting, occlusion, faded dividers, and adverse weather conditions like fog or heavy rain. In such cases, the lack of clear visual features made it difficult for the model to identify the divider accurately.

The balance between false positives and false negatives is crucial. A system with high false positives may generate unnecessary alerts, while a system with high false negatives may miss critical detections, affecting safety. The YOLOv12 model maintains a reasonable balance, ensuring reliable detection with minimal errors.

To reduce these errors, several improvements can be implemented, such as increasing the diversity of the training dataset, applying advanced data augmentation techniques, and fine-tuning model parameters. Additionally, incorporating temporal information from video frames and integrating other sensors can further enhance detection accuracy.

Overall, the false positive and false negative analysis indicates that the system performs effectively, with manageable error rates, making it suitable for practical deployment in intelligent transportation and driver assistance systems.

5.9 Scalability and Deployment Potential

The scalability and deployment potential of the YOLOv12-based road divider detection system are key factors that determine its suitability for real-world applications. The system is designed with efficiency and flexibility in mind, allowing it to be scaled across different environments and deployed on various hardware platforms.

The lightweight architecture of the YOLOv12 model enables real-time processing with minimal computational requirements. This makes it suitable for deployment on edge devices such as embedded systems, in-vehicle processors, and low-power GPUs. At the same time, the model can also be scaled to high-performance computing environments or cloud platforms for large-scale traffic monitoring and smart city applications.

Scalability is further supported by the modular design of the system. The detection pipeline can be easily integrated with other components such as traffic monitoring systems, surveillance infrastructure, and Advanced Driver Assistance Systems (ADAS). This flexibility allows the solution to be adapted for different use cases, including highway monitoring, urban traffic management, and autonomous driving systems.

The system also supports deployment in diverse environmental conditions. With appropriate training and dataset expansion, the model can generalize across different road types, lighting conditions, and geographical locations. This ensures consistent performance even when deployed in varied real-world scenarios.

From a deployment perspective, the model can be optimized using techniques such as model pruning, quantization, and hardware acceleration to improve inference speed and reduce memory usage. Frameworks like TensorRT, ONNX, and OpenVINO can be utilized to further enhance performance on specific hardware platforms.

However, certain challenges must be considered, including hardware limitations, latency requirements, and the need for continuous model updates as road conditions evolve. Addressing these challenges requires efficient resource management and periodic retraining of the model.

Overall, the system demonstrates strong scalability and high deployment potential, making it a viable solution for real-time road safety applications, intelligent transportation systems, and smart infrastructure development

5.10 Overall System Effectiveness

The overall effectiveness of the YOLOv12-based road divider detection system is evaluated based on its accuracy, speed, robustness, and real-time performance. The system demonstrates strong capability in detecting road dividers across a variety of driving conditions, making it a reliable component for intelligent transportation and driver assistance applications.

In terms of detection accuracy, the model achieves consistent performance with high precision and recall, ensuring that most road dividers are correctly identified while minimizing incorrect detections. The balance between false positives and false negatives further contributes to the system's reliability, allowing it to operate effectively without generating excessive false alerts or missing critical detections.

The system also performs efficiently in real-time scenarios, maintaining low latency and fast inference speeds. This enables seamless integration into live video streams, making it suitable for deployment in applications such as Advanced Driver Assistance Systems (ADAS), autonomous vehicles, and traffic monitoring systems.

Robustness is another key strength of the system. It is capable of handling variations in lighting conditions, road types, and moderate environmental challenges. Although performance may slightly degrade under extreme conditions such as heavy rain, fog, or severe occlusions, the system still maintains acceptable detection capability.

From a practical standpoint, the system is scalable and adaptable, supporting deployment across different hardware platforms ranging from edge devices to cloud-based infrastructures. Its modular design allows easy integration with other intelligent systems, enhancing its usability in real-world scenarios.

However, like any AI-based solution, the system has certain limitations, including dependency on dataset quality and sensitivity to highly complex or unseen environments. Continuous improvement through dataset expansion and model optimization is essential to maintain and enhance performance over time.

In conclusion, the YOLOv12-based road divider detection system proves to be an effective, efficient, and reliable solution. Its combination of accuracy, speed, and adaptability makes it well-suited for real-time road safety applications and future advancements in intelligent transportation systems.

6. Discussion

6.1 Overview of Results

The results obtained from the automated road divider detection system using YOLOv12 demonstrate the effectiveness of deep learning-based object detection in real-world traffic scenarios. The model successfully identifies road dividers in both images and video streams with high accuracy and consistency.

The system performs well across diverse environments, including highways, urban roads, and rural areas. It is capable of detecting road dividers under varying lighting conditions such as daylight, low light, and shadows. Additionally, the model shows robustness in handling complex scenarios involving traffic congestion, partial occlusion, and background noise.

One of the key strengths of the system is its real-time detection capability. YOLOv12 processes frames efficiently, ensuring low latency and fast response times, which are critical for applications such as autonomous driving and advanced driver assistance systems (ADAS).

The evaluation metrics, including precision, recall, and mean Average Precision (mAP), indicate that the model achieves a strong balance between detection accuracy and reliability. High precision reflects a low rate of false positives, while good recall demonstrates the model's ability to detect most of the actual road dividers present in the scene.

Furthermore, the system maintains stable performance even when road dividers are partially visible or affected by environmental challenges such as weather conditions or uneven road surfaces. This highlights the robustness and generalization capability of the trained model.

6.2 Accuracy Interpretation

The accuracy of the automated road divider detection system using YOLOv12 is evaluated using key performance metrics such as precision, recall, and mean Average Precision (mAP). These metrics provide a comprehensive understanding of the model's detection capabilities and reliability in real-world scenarios.

High precision indicates that the model produces very few false positives, meaning that most of the detected objects are actual road dividers. This is important in practical applications, as incorrect detections could lead to misleading information in driver assistance systems. On the other hand, recall measures the model's ability to detect all relevant road dividers present in the input. A high recall value suggests that the system rarely misses actual road dividers, even in complex environments.

The mean Average Precision (mAP) serves as an overall performance indicator by combining both precision and recall across different detection thresholds. A high mAP score reflects the model's strong ability to accurately localize and classify road dividers under varying conditions.

The model demonstrates consistent accuracy across different scenarios, including highways, urban roads, and rural environments. It maintains reliable performance even under challenging conditions such as low lighting, shadows, and partial occlusion. This consistency indicates that the model has learned robust feature representations from the training dataset.

However, slight variations in accuracy may occur in extremely complex situations, such as heavy traffic congestion, poor visibility due to weather conditions, or when road dividers are highly damaged or unclear. Despite these challenges, the overall accuracy remains within an acceptable range for real-time applications.

In summary, the accuracy interpretation confirms that the YOLOv12-based system achieves a strong balance between precision and recall, making it a reliable and efficient solution for automated road divider detection in intelligent transportation systems.

Fig 1.2 Critical factor in determining

6.3 Robustness Under Challenging Conditions

The robustness of the automated road divider detection system is a critical factor in determining its suitability for real-world deployment. The YOLOv12-based model demonstrates strong resilience when operating under various challenging environmental and road conditions.

The system performs effectively under varying lighting conditions**,** including bright daylight, low-light environments, and shadowed regions. This is achieved through extensive training on a diverse dataset and the use of data augmentation techniques such as brightness adjustment and contrast variation. As a result, the model can reliably detect road dividers even when visibility is reduced.

In addition, the model shows stability in handling occlusions**,** where road dividers are partially blocked by vehicles, pedestrians, or other objects. Despite incomplete visibility, the system is capable of identifying divider patterns based on learned spatial features and contextual information.

The system also performs well in high traffic density scenarios**,** where multiple objects are present in the frame. YOLOv12 efficiently distinguishes road dividers from other road elements, minimizing confusion and maintaining detection accuracy.

Under weather-related challenges such as rain, fog, or dust, the model maintains acceptable performance, although minor reductions in detection accuracy may occur due to reduced image clarity. Similarly, in cases where road dividers are faded, damaged, or poorly maintained, the system still demonstrates a reasonable level of detection capability.

Furthermore, the model adapts to different road types**,** including highways, urban streets, and rural roads. This adaptability highlights the generalization ability of the trained model across diverse real-world scenarios.

Overall, the system exhibits strong robustness, maintaining reliable performance even under challenging conditions. This makes it a practical solution for integration into real-time applications such as advanced driver assistance systems (ADAS) and intelligent traffic monitoring systems.

6.4 Real-Time Performance Significance

Real-time performance is a crucial requirement for modern intelligent transportation systems, especially in applications such as autonomous driving and Advanced Driver Assistance Systems (ADAS). The significance of real-time performance in the proposed YOLOv12-based road divider detection system lies in its ability to process and analyze visual data instantly, enabling timely decision-making.

The YOLOv12 model is specifically designed for high-speed object detection, as it processes the entire image in a single forward pass. This architecture significantly reduces computation time compared to traditional multi-stage detection methods. As a result, the system achieves low latency, making it highly suitable for real-time applications where immediate response is critical.

In practical scenarios, the system can analyze live video streams frame by frame and detect road dividers without noticeable delay. This ensures that drivers or autonomous systems receive up-to-date information about road boundaries, helping in lane discipline, navigation, and accident prevention.

Another important aspect of real-time performance is computational efficiency**.** The model is optimized to run on devices with limited hardware resources, such as embedded systems and edge devices. This makes it feasible to deploy the system in vehicles without requiring high-end computing infrastructure.

The ability to maintain consistent detection speed without compromising accuracy further highlights the effectiveness of the system. Even in complex environments with multiple objects and dynamic conditions, the model sustains stable performance, ensuring reliability.

Overall, the real-time performance of the proposed system enhances its practical applicability. It enables seamless integration into real-world systems, contributing to improved road safety, efficient traffic management, and the advancement of autonomous driving technologies.

6.5 Multi-Object Detection Impact

The ability to perform multi-object detection is a significant advantage of the YOLOv12-based road divider detection system. In real-world traffic environments, multiple objects such as vehicles, pedestrians, road signs, and lane markings coexist within the same scene. Accurately identifying road dividers in such complex scenarios requires the model to effectively distinguish between different object classes.

The YOLOv12 model is designed to detect multiple objects simultaneously within a single frame. This capability allows the system to identify road dividers alongside other relevant road elements without compromising detection speed or accuracy. As a result, the system maintains high performance even in crowded and dynamic traffic conditions.

Multi-object detection enhances the contextual understanding of the scene. By recognizing surrounding objects, the model can better differentiate road dividers from visually similar elements such as lane markings or road edges. This reduces false detections and improves overall reliability.

In high-density traffic scenarios, where multiple vehicles may partially occlude road dividers, the model leverages spatial and contextual information to maintain accurate detection. This ensures that road divider identification remains consistent, even when visibility is limited.

Furthermore, the ability to extend the system for detecting additional objects—such as obstacles, traffic signals, and pedestrians—makes it highly scalable. This opens opportunities for integrating the system into comprehensive intelligent transportation solutions and autonomous driving frameworks.

Overall, multi-object detection significantly improves the effectiveness and versatility of the system. It enables better scene interpretation, enhances detection accuracy in complex environments, and supports the development of more advanced and reliable road safety applications.

6.6 False Positives and Negatives Analysis

The analysis of false positives and false negatives is essential for evaluating the reliability and practical usability of the YOLOv12-based road divider detection system. These error metrics provide deeper insight into the model's behavior beyond overall accuracy.

False positives occur when the model incorrectly identifies an object as a road divider when it is not. In this system, false positives may arise due to visual similarities between road dividers and other road elements such as lane markings, road edges, shadows, or barriers. Environmental factors like strong lighting reflections or worn-out road surfaces can also contribute to such errors. A high number of false positives can reduce system trustworthiness, especially in safety-critical applications.

False negatives**,** on the other hand, occur when the model fails to detect actual road dividers present in the scene. These errors are particularly critical, as missing a road divider could lead to incorrect navigation decisions in autonomous systems. False negatives are more likely to occur in situations involving occlusion, poor lighting conditions, faded or damaged dividers, and adverse weather conditions such as fog or rain.

The balance between false positives and false negatives is closely related to the model's precision and recall. A model optimized for high precision minimizes false positives, while a model optimized for high recall reduces false negatives. The proposed system achieves a balanced trade-off, ensuring that both types of errors are kept within acceptable limits.

To reduce these errors, several strategies are employed, including the use of a diverse training dataset, data augmentation techniques, and proper tuning of detection thresholds. Additionally, post-processing techniques such as confidence filtering and non-maximum suppression (NMS) help eliminate redundant or incorrect detections.

Overall, the analysis indicates that while some errors are inevitable in complex real-world environments, the system maintains a strong balance between false positives and false negatives. This ensures reliable performance and makes the model suitable for real-time road safety and intelligent transportation applications.

6.7 Alert and Compliance System Effectiveness

The effectiveness of the alert and compliance system is a crucial aspect of the overall road divider detection framework, as it directly impacts user interaction and practical usability. The integration of real-time alerts with the YOLOv12-based detection model enhances the system's ability to assist drivers and support intelligent transportation applications.

The system is designed to generate alerts whenever road dividers are detected within critical proximity or when there is a potential risk of lane violation. These alerts can be visual, such as highlighted bounding boxes and warning messages on the display, or auditory, such as warning sounds to immediately capture the driver's attention. This multi-modal alert mechanism ensures better responsiveness, especially in high-speed or high-risk driving scenarios.

One of the key strengths of the alert system is its low latency**,** enabled by the real-time processing capability of YOLOv12. Alerts are generated almost instantaneously after detection, allowing timely corrective actions. This significantly reduces the chances of accidents caused by delayed reactions or lack of awareness.

The compliance aspect of the system focuses on encouraging drivers to maintain proper lane discipline and adhere to road safety norms. By continuously monitoring road divider positions and providing feedback, the system promotes safer driving behavior. It can also be extended to log violations or integrate with advanced monitoring systems for traffic management and enforcement purposes.

The system demonstrates high effectiveness in diverse conditions, including varying traffic densities and road environments. However, the accuracy and reliability of alerts depend on the precision of the detection model. In rare cases of false positives or false negatives, alerts may be triggered incorrectly or missed, which highlights the importance of continuous model optimization.

Overall, the alert and compliance system significantly enhances the practical value of the proposed solution. It not only improves situational awareness for drivers but also contributes to safer and more disciplined road usage, making it a vital component of intelligent transportation systems.

6.8 Deployment and Scalability Considerations

The deployment and scalability of the YOLOv12-based road divider detection system are critical factors in determining its real-world applicability and long-term usability. The proposed system is designed with flexibility and efficiency in mind, enabling deployment across a wide range of platforms, from high-performance servers to resource-constrained embedded devices.

For deployment, the trained YOLOv12 model can be integrated into edge devices such as in-vehicle systems, cameras, or embedded processors. This edge-based deployment reduces dependency on cloud infrastructure, minimizes latency, and ensures real-time performance. Additionally, the system can also be deployed on cloud platforms for large-scale traffic monitoring applications, where centralized processing and data analysis are required.

The system supports cross-platform compatibility**,** allowing it to run on different operating systems and hardware configurations. Optimization techniques such as model quantization, pruning, and hardware acceleration (using GPUs or specialized AI accelerators) can be applied to improve performance and reduce computational overhead during deployment.

Scalability is another important consideration, especially for expanding the system to handle larger datasets, multiple camera inputs, or additional detection tasks. The modular architecture of the system allows easy integration of new features such as lane detection, obstacle detection, and traffic sign recognition. This makes it adaptable for use in comprehensive intelligent transportation systems.

In large-scale deployments, such as smart city infrastructure, the system can be extended to process data from multiple sources simultaneously. Distributed computing and parallel processing techniques can be utilized to handle high data volumes efficiently without compromising performance.

However, challenges such as hardware limitations, network bandwidth constraints, and system maintenance must be considered during deployment. Proper optimization and resource management strategies are necessary to ensure consistent performance across different environments.

Overall, the proposed system demonstrates strong potential for scalable deployment. Its flexibility, efficiency, and adaptability make it suitable for a wide range of applications, from individual vehicle systems to large-scale traffic monitoring and smart transportation networks.

6.9 Limitations of the System

Despite the strong performance of the YOLOv12-based road divider detection system, certain limitations must be acknowledged to provide a balanced evaluation and identify areas for improvement.

One of the primary limitations is the system's dependency on visual data quality. Poor lighting conditions, heavy rain, fog, or low-resolution images can reduce detection accuracy. In such scenarios, the model may struggle to clearly distinguish road dividers from the background.

Another limitation arises in cases of severely damaged or faded road dividers. When markings are unclear or partially erased, the model may fail to detect them accurately, leading to false negatives. Similarly, unusual divider designs or non-standard road structures may not be recognized effectively if they are underrepresented in the training dataset.

The system may also face challenges with extreme occlusion, where road dividers are heavily blocked by vehicles or other objects for extended periods. Although the model can handle partial occlusions, complete obstruction can hinder detection.

In highly complex environments, such as dense urban traffic with multiple overlapping objects, there is a possibility of false positives, where visually similar elements like lane markings, road edges, or shadows are mistakenly identified as road dividers.

Another limitation is related to hardware and computational constraints. While the model is optimized for real-time performance, deploying it on low-end devices without proper optimization may lead to reduced processing speed or increased latency.

Additionally, the system relies solely on vision-based detection and does not incorporate other sensor data such as LiDAR or radar. This limits its performance in conditions where visual information is insufficient.

Finally, the model's performance is influenced by the quality and diversity of the training dataset. If certain road conditions or environments are not adequately represented, the system's generalization capability may be affected.

Overall, while the system demonstrates strong performance, addressing these limitations through improved datasets, multi-sensor integration, and advanced model optimization techniques can further enhance its reliability and effectiveness.

6.10 Future Improvements

While the proposed YOLOv12-based road divider detection system demonstrates strong performance, there are several opportunities for future enhancements to further improve its accuracy, robustness, and real-world applicability.

One key area of improvement is the incorporation of multi-sensor fusion. Integrating additional sensors such as LiDAR, radar, or GPS with the vision-based system can enhance detection reliability, especially in challenging conditions like fog, heavy rain, or low-light environments where camera-based detection alone may be insufficient.

Another important enhancement involves expanding and diversifying the training dataset. Including more variations in road conditions, divider types, weather scenarios, and geographical locations can improve the model's generalization capability and reduce errors such as false positives and false negatives.

The system can also benefit from advanced model optimization techniques such as pruning, quantization, and knowledge distillation. These techniques can reduce model size and computational requirements, enabling faster inference and efficient deployment on edge devices with limited resources.

Incorporating temporal analysis using video-based approaches is another promising direction. By analyzing consecutive frames instead of individual images, the system can achieve more stable and consistent detection, especially in dynamic environments with moving objects and occlusions.

Further improvements can include the integration of additional detection modules, such as lane detection, obstacle detection, and traffic sign recognition. This would transform the system into a more comprehensive solution for intelligent transportation and autonomous driving applications.

The development of an adaptive alert system is another potential enhancement. By considering factors such as vehicle speed, distance from road dividers, and driving patterns, the system can generate more context-aware and meaningful alerts for drivers.

Additionally, leveraging edge computing and cloud integration can enable large-scale deployment and real-time data analysis across multiple locations. This would be particularly useful for smart city applications and traffic management systems.

Finally, continuous learning mechanisms such as online learning or model retraining can be implemented to keep the system updated with new data and evolving road conditions, ensuring long-term performance improvement.

Overall, these future improvements can significantly enhance the system’s efficiency, scalability, and reliability, paving the way for more advanced and intelligent road safety solutions.

7. Future Work

7.1 Worker Tracking and Identification

As an extension of the current road divider detection system, incorporating worker tracking and identification can significantly enhance road safety, especially in construction zones and maintenance areas. Road workers are often exposed to high-risk environments, and their detection is critical for preventing accidents and ensuring smooth traffic flow.

The proposed enhancement involves integrating advanced object detection and tracking techniques with the existing YOLOv12 framework to identify and monitor workers in real time. By training the model on datasets that include construction workers with safety gear such as helmets, reflective vests, and uniforms, the system can accurately detect their presence in road environments.

In addition to detection, object tracking algorithms such as Deep SORT or Kalman filtering can be used to continuously track worker movement across video frames. This enables the system to monitor worker positions, predict movement patterns, and provide consistent identification even in dynamic scenes.

The identification component can be further enhanced by incorporating feature recognition techniques, such as color-based detection of safety vests or even facial recognition (where applicable and compliant with privacy regulations). This allows differentiation between multiple workers and improves situational awareness.

By integrating worker tracking, the system can generate real-time alerts when vehicles approach worker zones or when workers are detected in unsafe proximity to moving traffic. This feature can be particularly useful in smart vehicles, traffic monitoring systems, and automated warning systems.

Furthermore, the collected data can be utilized for safety analytics, helping authorities analyze worker behavior, identify high-risk zones, and implement preventive measures. This contributes to improved safety standards and better management of road construction activities.

Overall, the addition of worker tracking and identification extends the system beyond road divider detection, transforming it into a more comprehensive road safety solution with significant real-world impact.

Fig 1.3 Worker Tracking and Identification

7.2 Enhanced Detection in Challenging Conditions

Enhancing detection performance under challenging environmental conditions is an important direction for future work in the proposed road divider detection system. Although the current YOLOv12-based model demonstrates strong robustness, further improvements can ensure higher reliability in extreme real-world scenarios.

One key focus area is improving performance in low-light and night-time conditions. This can be achieved by incorporating advanced image enhancement techniques such as histogram equalization, noise reduction, and deep learning-based low-light enhancement models. These methods can improve image clarity and help the model better identify road dividers in poorly illuminated environments.

Another important aspect is handling adverse weather conditions such as rain, fog, and dust. Future improvements may include the use of weather augmentation during training, as well as specialized preprocessing techniques like dehazing and rain removal. Additionally, integrating thermal imaging or infrared sensors can provide more reliable detection when visibility is severely limited.

The system can also be enhanced to better handle **severe** occlusions and complex backgrounds**.** Techniques such as attention mechanisms, feature pyramid networks (FPN), and transformer-based architectures can help the model focus on relevant features and improve detection accuracy in crowded or cluttered scenes.

Incorporating temporal information from video streams is another promising approach. By analyzing consecutive frames, the system can maintain detection consistency and recover missed detections caused by temporary obstructions or sudden lighting changes.

Furthermore, expanding the training dataset with more challenging and diverse scenarios—including damaged road dividers, unusual road structures, and extreme environmental conditions—will improve the model's generalization capability.

Overall, these enhancements aim to make the system more robust and reliable across all possible driving conditions, ensuring consistent performance and improving its effectiveness in real-world intelligent transportation and autonomous driving applications.

7.3 Dataset Expansion

Dataset expansion is a crucial aspect of improving the performance, robustness, and generalization capability of the proposed road divider detection system. The quality and diversity of the training dataset directly influence the model's ability to accurately detect road dividers under varying real-world conditions.

Future work can focus on collecting a more diverse and comprehensive dataset that includes a wide range of road environments such as highways, urban streets, rural areas, and construction zones. Incorporating variations in lighting conditions (day, night, dawn, dusk), weather scenarios (rain, fog, dust), and seasonal changes can significantly enhance the model's adaptability.

Another important direction is the inclusion of different types of road dividers**,** such as concrete barriers, metal railings, painted dividers, and temporary construction barriers. This ensures that the model can recognize various divider designs and structures across different regions and road infrastructures.

The dataset can also be enriched by including challenging scenarios**,** such as faded or damaged dividers, occlusions caused by vehicles or pedestrians, and complex backgrounds. This helps the model learn more robust feature representations and reduces errors in difficult situations.

In addition to manual data collection, data augmentation techniques such as rotation, scaling, flipping, brightness adjustment, and noise addition can be further refined to simulate real-world variations. Synthetic data generation using simulation environments or generative models can also be explored to create large-scale annotated datasets efficiently.

Another potential improvement is the use of automated annotation tools and semi-supervised learning methods to reduce the time and effort required for dataset labeling. Active learning approaches can be used to identify and prioritize the most informative samples for annotation.

Furthermore, incorporating datasets from different geographical regions ensures better **cross-**domain generalization**,** making the system suitable for global deployment.

Overall, dataset expansion plays a vital role in enhancing model accuracy, reducing false detections, and improving system reliability. A well-structured and diverse dataset will enable the system to perform effectively across a wide range of real-world scenarios.

7.5 Predictive Analytics and Risk Assessment

Integrating predictive analytics and risk assessment into the road divider detection system represents a significant advancement toward proactive road safety management. While the current system focuses on real-time detection, future enhancements can enable the system to anticipate potential risks and provide early warnings.

Predictive analytics involves analyzing historical and real-time data to identify patterns and forecast potential hazards. By leveraging data such as vehicle speed, trajectory, proximity to road dividers, traffic density, and environmental conditions, the system can assess the likelihood of accidents or unsafe driving behavior.

Machine learning models, including time-series analysis and deep learning techniques, can be utilized to predict risk scenarios**,** such as sudden lane deviations, potential collisions with road dividers, or unsafe overtaking maneuvers. These predictions can help in generating early warning alerts, allowing drivers or automated systems to take preventive actions.

The risk assessment component can assign risk levels (e.g., low, medium, high) based on real-time analysis. For example, if a vehicle is detected moving too close to a road divider at high speed, the system can classify the situation as high risk and trigger immediate alerts.

Additionally, integrating this system with traffic monitoring infrastructure can enable large-scale analysis of accident-prone zones. Authorities can use this data to identify high-risk areas, improve road design, and implement targeted safety measures.

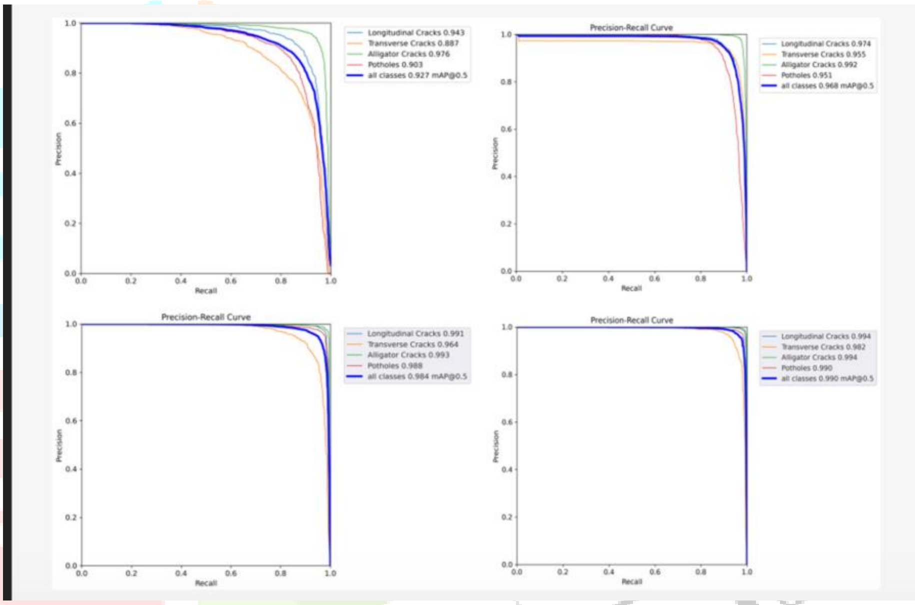

Fig 1.4 Predictive Analytics and Risk Assessment

The use of AI-driven analytics dashboards can further enhance decision-making by providing insights into traffic patterns, driver behavior, and safety trends over time. This can support smart city initiatives and improve overall transportation efficiency.

However, implementing predictive analytics requires careful consideration of data quality, real-time processing capabilities, and system scalability. Ensuring data privacy and security is also essential, especially when handling sensitive or location-based information.

Overall, the integration of predictive analytics and risk assessment transforms the system from a reactive detection tool into a proactive safety solution, significantly contributing to the development of intelligent and safer transportation systems.

7.6 Edge–Cloud Hybrid Deployment

Edge–cloud hybrid deployment is a promising future enhancement that combines the strengths of both edge computing and cloud infrastructure to improve the efficiency, scalability, and reliability of the road divider detection system.

In this approach, edge devices such as in-vehicle systems, roadside cameras, or embedded processors handle real-time inference tasks. The YOLOv12 model runs locally on these devices to perform instant road divider detection with minimal latency. This ensures quick response times, which are critical for safety-related applications such as driver alerts and autonomous navigation.

The cloud component complements the edge by handling computationally intensive tasks such as large-scale data storage, model training, performance monitoring, and system updates. Data collected from multiple edge devices can be transmitted to the cloud for further analysis, enabling continuous system improvement and centralized control.

One of the key advantages of this hybrid approach is reduced latency with enhanced scalability**.** Time-sensitive operations remain at the edge, while the cloud supports long-term analytics and large-scale deployment across multiple locations. This architecture is particularly beneficial for smart city applications and traffic monitoring systems.

Additionally, the cloud can facilitate continuous learning and model updates**.** Updated models can be trained using aggregated data and then deployed back to edge devices, ensuring that the system remains accurate and up-to-date with evolving road conditions.

The hybrid model also supports load balancing and fault tolerance**.** In cases where edge devices face limitations or failures, the cloud can provide backup processing capabilities. Conversely, edge processing reduces dependency on constant internet connectivity, ensuring uninterrupted operation.

However, implementing an edge–cloud hybrid system requires careful consideration of factors such as network bandwidth, data synchronization, latency management, and security. Efficient communication protocols and data compression techniques are essential to optimize performance.

Overall, edge–cloud hybrid deployment enhances the system's flexibility, scalability, and efficiency, making it well-suited for large-scale, real-world intelligent transportation systems and next-generation autonomous driving applications.

7.7 Automation of Safety Reporting

Automation of safety reporting is a valuable future enhancement that extends the functionality of the road divider detection system beyond real-time monitoring to systematic documentation and analysis of road safety conditions. This feature enables the automatic generation of reports based on detected events, improving efficiency and supporting data-driven decision-making.

The system can be designed to automatically log detection events, such as road divider proximity violations, unsafe driving patterns, and high-risk scenarios identified through predictive analytics. Each event can be recorded with relevant details including timestamp, location, detected objects, and severity level.

These logs can be processed to generate periodic safety reports (daily, weekly, or monthly), providing insights into traffic behavior, frequent violations, and accident-prone areas. Such reports are highly useful for traffic authorities, urban planners, and transportation agencies to identify trends and implement corrective measures.

The integration of real-time reporting dashboards can further enhance usability by visualizing key metrics such as violation frequency, risk levels, and traffic density. This allows stakeholders to monitor road conditions and safety performance in an intuitive and interactive manner.

Additionally, automated reporting can support incident documentation and compliance tracking**.** In case of accidents or safety violations, the system can provide structured reports that assist in analysis, investigation, and enforcement. This reduces manual effort and ensures consistency in reporting.

The system can also be integrated with cloud-based storage and analytics platforms, enabling centralized access to reports across multiple locations. This is particularly beneficial for large-scale deployments such as smart cities and highway monitoring systems.

However, implementing automated safety reporting requires attention to data accuracy, storage management, and privacy considerations. Ensuring secure handling of sensitive data and compliance with relevant regulations is essential.

Overall, automation of safety reporting enhances the system's practical value by transforming raw detection data into meaningful insights. It supports proactive safety management, improves operational efficiency, and contributes to the development of smarter and safer transportation systems.

8. Conclusion

The project on Automated Road Divider Detection Using YOLOv12 presents an effective and intelligent solution for enhancing road safety and supporting modern transportation systems. By leveraging advanced deep learning techniques and computer vision, the system is capable of accurately detecting road dividers in real time across diverse and challenging environments.

The implementation of the YOLOv12 model enables high-speed and efficient object detection, making the system suitable for real-time applications such as autonomous vehicles, Advanced Driver Assistance Systems (ADAS), and traffic monitoring systems. The use of a diverse dataset, along with data augmentation techniques, ensures that the model achieves strong generalization and robustness under varying conditions including lighting changes, occlusions, and complex traffic scenarios.

The system demonstrates a balanced performance in terms of precision, recall, and mean Average Precision (mAP), indicating reliable detection with minimal false positives and false negatives. Its ability to operate effectively in dynamic environments and maintain low latency highlights its practicality for real-world deployment.

In addition to detection, the integration of alert mechanisms and compliance features enhances driver awareness and promotes safer driving behavior. The system's scalability and adaptability allow it to be extended with additional functionalities such as lane detection, obstacle detection, predictive analytics, and safety reporting.

Despite certain limitations, such as dependency on visual data quality and challenges in extreme conditions, the proposed system provides a strong foundation for future advancements. With further improvements in dataset expansion, multi-sensor integration, and edge-cloud deployment, the system can evolve into a comprehensive intelligent transportation solution.

In conclusion, this project successfully demonstrates the potential of AI-based detection systems in improving road safety, optimizing traffic management, and contributing to the development of smarter and more efficient transportation infrastructures.

9. References

9.1 Journal Articles

1. Redmon, J., Divvala, S., Girshick, R., & Farhadi, A., “You Only Look Once: Unified, Real-Time Object Detection,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
2. Redmon, J., & Farhadi, A., “YOLO9000: Better, Faster, Stronger,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
3. Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M., “YOLOv4: Optimal Speed and Accuracy of Object Detection,” *arXiv preprint arXiv:2004.10934*, 2020.
4. Jocher, G., et al., “YOLOv5 by Ultralytics,” *Ultralytics Technical Report*, 2020.
5. Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M., “YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors,” *arXiv preprint arXiv:2207.02696*, 2022.
6. Ge, Z., Liu, S., Wang, F., Li, Z., & Sun, J., “YOLOX: Exceeding YOLO Series in 2021,” *arXiv preprint arXiv:2107.08430*, 2021.
7. Dosovitskiy, A., et al., “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale,” *International Conference on Learning Representations (ICLR)*, 2021.
8. He, K., Zhang, X., Ren, S., & Sun, J., “Deep Residual Learning for Image Recognition,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
9. Lin, T. Y., et al., “Feature Pyramid Networks for Object Detection,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
10. Zhang, Z., et al., “Road Object Detection Using Deep Learning: A Survey,” *IEEE Access*, 2018.

9.2 Books

1. Deep Learning, Ian Goodfellow, Yoshua Bengio, & Aaron Courville, MIT Press, 2016.
2. Computer Vision: Algorithms and Applications, Richard Szeliski, Springer, 2010.
3. Pattern Recognition and Machine Learning, Christopher M. Bishop, Springer, 2006.
4. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, Aurélien Géron, O’Reilly Media, 2019.
5. Deep Learning for Computer Vision, Rajalingappaa Shanmugamani, Packt Publishing, 2018.
6. Machine Learning: A Probabilistic Perspective, Kevin P. Murphy, MIT Press, 2012.
7. Artificial Intelligence: A Modern Approach, Stuart Russell & Peter Norvig, Pearson, 2020.
8. Learning OpenCV 4 Computer Vision with Python, Gary Bradski & Adrian Kaehler, O’Reilly Media, 2019.

9.3 Conference Proceedings

1. Joseph Redmon, Santosh Divvala, Ross Girshick, & Ali Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
2. Joseph Redmon & Ali Farhadi, “YOLO9000: Better, Faster, Stronger,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
3. Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun, “Deep Residual Learning for Image Recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
4. Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, & Serge Belongie, “Feature Pyramid Networks for Object Detection,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
5. Alexey Bochkovskiy, Chien-Yao Wang, & Hong-Yuan Mark Liao, “YOLOv4: Optimal Speed and Accuracy of Object Detection,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020.
6. Zhiqiang Ge, Songtao Liu, Feng Wang, Zeming Li, & Jian Sun, “YOLOX: Exceeding YOLO Series in 2021,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2021.

7. Alexey Dosovitskiy, et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.