Vision-Based Lane Detection Using Machine Learning

Prof. Pranesh Kulkarni
KLS Vishwanathrao Deshpande Institute of Technology,
Haliyal,Uttar-Kannada District,Karnataka,India
e-mail: pkulkarni@klsvdit.edu.in

[1] Akshata Kattimani, [2] Narayanrao Kulkarni, [3] Ruthuja Chavan, [4] Spoorthi Harakuni
Final Year Students Department of Computer Science and Engineering[AIML]
Haliyal, Uttar-Kannada District, Karnataka, India
e-mail: ruthujachavan@gmail.com

Abstract - Lane detection is one of the most fundamental components of intelligent transportation systems, particularly in autonomous vehicles and modern Advanced Driver Assistance Systems (ADAS). Accurate lane perception enables safe navigation, stable lane keeping, and informed decision-making. Traditional lane detection approaches—such as Canny edge detection, Hough Transform, and color-based thresholding—show reasonable performance in controlled environments but fail under challenging real-world conditions involving low visibility, shadows, faded lane markings, and abrupt illumination changes. With the rise of deep learning, particularly Convolutional Neural Networks (CNNs), models have gained the ability to learn robust lane features directly from data. However, real-world driving requires more than lane detection; it also demands an understanding of drivable areas and the presence of objects such as vehicles or pedestrians.

This paper presents a comprehensive review of classical and modern lane detection techniques, with a focus on multi-task deep learning architectures, such as YOLOP (You Only Look Once for Panoptic Driving Perception). We also implement YOLOP on real-world Indian road videos and enhance its performance on nighttime scenes using custom brightness, contrast, and gamma preprocessing. The integration of night enhancement improved lane IoU from 0.72 to 0.84 and pixel accuracy from 0.88 to 0.93. The review highlights major advancements, limitations, research gaps, and future opportunities in machine-learning-based lane detection. The findings emphasize that multi-task learning, domain adaptation, and lightweight models are essential steps toward practical and reliable autonomous vehicle perception systems.

Keywords— Lane detection, YOLOP, multitask learning, CNN, semantic segmentation, object detection, night enhancement.

1. INTRODUCTION

In the age of rapidly evolving smart transportation technologies, autonomous vehicles (AVs) and Advanced Driver Assistance Systems (ADAS) are revolutionizing how society approaches road safety, mobility, and traffic management. These systems are gradually shifting the responsibility of driving from humans to intelligent

machines capable of sensing, interpreting, and responding to complex real-world environments. For such systems to operate safely and effectively, one fundamental capability stands above all others — the ability of a vehicle to accurately "see" and understand the road. At the heart of this perception lies lane detection, a task so central that without it, even the most advanced autonomous system would fail to maintain stable navigation or make safe decisions.

Lane detection research started long before modern AI-driven perception systems became common. Early methods mainly used traditional image-processing techniques. For example, the Canny edge detector identified sharp changes in pixel intensity that corresponded to lane boundaries. The Hough Transform was another method used to find linear or parametric structures within the detected edges. These methods were computationally light, straightforward, and easy to implement. This made them suitable for the early phases of intelligent vehicle development. However, despite their simplicity and early potential, traditional methods did not have the flexibility needed to perform reliably in real-world driving conditions.

Real-world road environments are unpredictable and visually inconsistent. Lane markings can be faded, worn out, partially blocked by vehicles, distorted due to road curves, covered in dirt or debris, or completely missing in poorly maintained areas. Lighting conditions also change dramatically throughout the day, from bright sunlight and glare to long shadows, and into night-time darkness where lane visibility greatly decreases. Weather makes things even more challenging by adding rain streaks, fog, snow, reflections, and other visual disturbances. In these situations, traditional edge-based or colour-thresholding methods often fail. They can detect false edges, misclassify road artifacts, or lose lane cues completely. These problems show the need for perception systems that can learn and generalize beyond limited, fixed rules.

This realization led to the rise of machine learning and, eventually, deep learning-based lane detection. Instead of manually designing filters or thresholds, machine learning allowed systems to learn meaningful lane patterns directly from data. The introduction of Convolutional Neural

Networks (CNNs) changed the field by allowing hierarchical feature extraction, from simple edges in shallow layers to complex lane textures, shapes, and contextual information in deeper layers. CNN-based models quickly surpassed traditional methods in robustness and accuracy, especially under challenging conditions like shadows, partial occlusions, worn-out markings, or curved road segments. More advanced architectures, such as Spatial CNNs (SCNNs), improved lane continuity understanding by letting information move across image rows and columns, thus enhancing detection in curves and visually cluttered scenes.

As autonomous systems developed, researchers realized that detecting lane lines alone did not provide a complete understanding of the road. Vehicles also needed to recognize drivable areas, understand road boundaries, detect other vehicles, and interpret changing traffic situations. This realization led to the creation of multi-task learning frameworks that could handle several perception tasks within one design. Among these, YOLOP (You Only Look Once for Panoptic Driving Perception) is a prominent model. YOLOP combines object detection, lane detection, and drivable area segmentation in a single encoder-decoder network. By sharing feature representations across tasks, YOLOP lowers computational costs, improves contextual understanding, and achieves real-time performance, making it a strong option for practical ADAS and AV applications.

Despite these advancements, several challenges remain. Autonomous vehicles must operate consistently in changing lighting conditions, various weather patterns, occlusions, geographical differences, and hardware limits. Night-time driving is particularly challenging because of low light, headlight glare, and less contrast between lanes and the surrounding area. Thick shadows, rain, fog, and reflections add extra visual noise that can confuse even the best neural models. Using embedded automotive hardware also demands efficient models that work in real-time with minimal delay. In addition, dataset limitations and differences across regions create obstacles for model generalization.

This review paper aims to provide a clear and detailed understanding of how lane detection systems have advanced. It covers the shift from early, rule-based image processing methods to modern deep learning models and multi-task perception systems. The paper looks into the theories behind each approach, highlighting their strengths and weaknesses. It also includes a practical implementation of YOLOP, which has been specifically improved for better performance at night. These proposed changes tackle one of the biggest challenges in self-driving: ensuring reliable lane detection in low-light conditions.

The motivation for this study comes from the urgent need to close the gap between research-level lane detection models and their effective use in real-world situations.

This work has four main goals:

1. To explore how deep learning and multi-task architectures improve lane detection compared to traditional methods.

2. To evaluate YOLOP's performance in tough real-world situations, including night-time and low-visibility conditions.

3. To identify ongoing issues in lane detection research related to occlusions, lighting changes, dataset biases, and computational limits.

4. To implement and analyze night-time enhancement techniques focused on improving lane visibility and safety during low-light driving.

By looking at the evolution, challenges, and opportunities in lane detection research, this paper aims to help develop safer, smarter, and more reliable vision-based driving systems. As the transportation industry approaches true autonomous mobility, strong and dependable lane detection remains essential for ensuring safety and trust in the next generation of intelligent vehicles.

2. LITERATURE REVIEW

The development of lane detection has changed significantly over the last thirty years. Early research in this area mainly focused on traditional computer vision methods that relied heavily on manual features and fixed rules. One of the first foundations for lane detection was based on edge techniques like the Canny edge detector. This method aimed to find lane boundaries by highlighting strong intensity changes. Typically, these techniques were paired with mathematical tools like the Hough Transform to identify straight or slightly curved lane shapes. While these methods worked reasonably well in controlled settings and on well-marked highways, their effectiveness in real-world situations quickly faced challenges. Variations in lighting, road surfaces, lane paint quality, weather conditions, and shadows created inconsistencies that rigid algorithms could not handle well. Consequently, although these traditional methods are historically important, they struggled to perform reliably in unpredictable environments.

A major shift happened with the rise of machine learning, which brought data-driven decision-making to lane detection. Instead of using manually designed filters, machine learning tried to classify pixels, edges, or patches based on features learned from labelled datasets. Techniques like Support Vector Machines (SVMs), Random Forests, and k-Nearest Neighbour (KNN) were used for lane-marking detection tasks, often relying on handcrafted descriptors such as Histogram of Oriented Gradients (HOG) or colour histograms. Although this change showed progress, the success of these models was still limited by the quality and effectiveness of the manual features. Handcrafted features could not adapt well and failed to capture the complex variations in road images, especially in challenging conditions like night driving or heavy rain.

The real change in lane detection started with deep learning, especially Convolutional Neural Networks (CNNs). CNNs showed a strong ability to learn features directly from raw images, which removed the need for manual feature extraction. Researchers quickly adapted semantic segmentation networks, like Fully Convolutional Networks (FCN), Seg Net, and U-Net, for lane detection. These models delivered much higher accuracy and reliability, especially in noisy and complex situations. FCN-based architectures transformed lane detection from simply identifying lines to classifying each pixel. This allowed the networks to better distinguish between lanes, shadows, edges, and other road features.

As deep learning advanced, models designed for lane detection started to emerge. Networks like SCNN (Spatial CNN), Lane Net, ENet-SAD, and PINet focused on spatial relationships and the continuous structure of lanes, rather than looking at each pixel on its own. SCNN introduced the concept of message passing in spatial directions, allowing for lane detection even when large sections were blocked or absent. Lane Net used an instance segmentation strategy, recognizing each lane as a separate instance instead of just a class. These improvements represented a major step forward for dependable lane detection in various driving situations.

However, deep learning models, even those made for lane detection, still had limitations in real-world use. Most models only performed well on curated datasets taken during the day. When they encountered night-time conditions, uneven lighting, rain, fog, or roads with low visibility, their accuracy often dropped sharply. This issue led to the creation of strategies for adapting to different conditions, pipelines for enhancing data, and techniques for improving night-time visibility. Recent studies explored methods like histogram equalization, CLAHE, neural illumination correction, and image-to-image translation using GANs (Generative Adversarial Networks) to boost lane visibility at night. These approaches inspired the enhancement step in this project, where adaptive gamma correction and image enhancement techniques were combined before processing frames with YOLOP to improve lane detection in low-light situations.

Parallel to advancements in lane detection, the field of autonomous driving started to require better scene understanding. The focus shifted from just detecting lane lines to understanding the whole road. This included identifying drivable areas, separating the road from sidewalks or off-road zones, and recognizing dynamic objects like vehicles, pedestrians, and cyclists. This change led to multi-task learning models,

where one network handled several perception tasks at once. YOLO-based architectures were key in enabling real-time object detection, encouraging researchers to create unified models for complete driving perception.

YOLOP, which stands for You Only Look Once for Panoptic Driving Perception, has become one of the most important contributions in this area. YOLOP combines object detection, drivable area segmentation, and lane segmentation into a single, seamless framework. By sharing feature maps across different tasks, YOLOP improves efficiency and lowers the computational load compared to running several models separately. Its CSP Darknet-based backbone and multi-branch decoder strike a solid balance between accuracy and real-time performance. Studies have shown that YOLOP outperforms other models on large-scale datasets like BDD100K, making it a key model for integrated driving perception.

Despite YOLOP's success, research shows that multi-task models face several challenges. One major issue is domain generalization. Deep learning models often overfit specific data sets and do poorly under unseen conditions, particularly at night or in bad weather. Another limitation is the lack of temporal reasoning; most current models consider frames separately instead of analyzing visual information over time. Researchers have suggested using 3D CNNs, LSTM-based designs, and temporal attention networks to tackle the problem of inconsistencies between frames. At the same time, transformer-based models have become popular for their ability to better capture global dependencies compared to CNNs. Vision Transformers (ViT), Swin Transformer, and hybrid CNN-transformer models have recently been studied in lane detection research. These models show great promise in capturing long-range spatial relationships that are vital for maintaining lane continuity.

The literature shows growing efforts to create lightweight architectures that work well on embedded platforms. Since autonomous vehicles depend on real-time perception, models need to achieve a balance between high accuracy and low computational cost. Techniques like model pruning, quantization, knowledge distillation, and efficient layer designs such as Mobile Net and Shuffle Net have been used in lane detection networks. These methods help reduce complexity while maintaining performance.

A key insight from recent studies is the strong link between image quality and detection accuracy. Poor lighting, motion blur, glare, or environmental noise directly impacts model predictions. As a result, researchers have highlighted the need for pre-processing improvement pipelines designed for challenging conditions. This project builds on this work by adding night-time visibility enhancement to YOLOP's inference pipeline. Experiments show that improving luminance and contrast before inputting images into the model significantly clarifies lane boundaries, resulting in better segmentation outcomes. This matches broader research trends that support combining traditional enhancement techniques with modern deep learning methods.

Overall, the evolution of lane detection, from handcrafted edge detectors to multi-task learning frameworks, reflects the ongoing development of computer vision and autonomous driving research. The literature shows a continuous effort to tackle real-world challenges like environmental changes, computational efficiency, and multi-modal perception. With the rise of transformers, domain adaptation methods, and improved pre-processing techniques, the field is getting closer to creating stable, reliable, and usable lane detection systems that support fully autonomous driving. This work adds to the existing knowledge by reviewing current solutions, implementing a new model, and including night-time enhancement strategies to tackle one of the most persistent challenges identified in the literature.

3. METHODOLOGY

Lane detection has progressed a lot. It has changed from simple image processing methods to powerful machine learning and deep learning techniques. In this section, we will examine the technology's evolution and point out the strengths and weaknesses of different approaches.

3.1. Traditional Methods: Where It All Began

Early lane detection used basic image processing tools like Canny edge detection and Hough Transform. These methods were easy to implement and quick, which made them popular. However, they had a major problem; they worked well only in perfect conditions. When the road was wet, poorly lit, or when lane markings were faded or curved, these methods often did not work.

3.2. CNN-Based Lane Detection

Convolutional Neural Networks (CNNs) transformed the field by learning patterns from data rather than relying on preset rules. CNNs were able to detect lanes more accurately, even when the markings were worn or partially obstructed. Agarwal and Dutta pointed out that CNNs usually outperformed traditional methods. However, their high computational cost made real-time use on devices with limited hardware difficult.

3.3. Spatial CNNs and Temporal Improvements

To improve lane continuity and stability across frames, Zhang et al. introduced a Spatial CNN (SCNN) model. Unlike standard CNNs, SCNNs can transmit information across the image, row by row and column by column.
This capability helps detect lanes more clearly, especially in curved or obstructed areas. Including temporal features, which take into account video frame sequences instead of single images, further improved detection consistency over time.

3.4. All-in-One Models: Multi-Task Learning

Modern systems often perform multiple tasks. They not only detect lanes but also find drivable areas and identify vehicles or obstacles. One well-known model is YOLOP (You Only Look Once for Panoptic Driving Perception) by Wang et al. [3]. It handles all three tasks within a single network, sharing features to save time and improve results. YOLOP is fast enough for real-time use and produces good outcomes, but it is still somewhat heavy for small, low-power devices.

4. PROPOSED FRAMEWORK

The proposed framework introduces a better lane detection system based on the YOLOP multi-task architecture. It is specifically optimized for real-world low-visibility driving conditions. Although YOLOP combines lane segmentation, drivable area segmentation, and object detection into a single process, its performance decreases significantly at night. To tackle this issue, the proposed framework adds a dedicated Night-Time Enhancement Module that prepares input frames before YOLOP inference. This improves feature visibility and leads to more reliable lane predictions.

The overall framework consists of five major stages: Input Acquisition, Night-Time Enhancement, Multi-Task Perception (YOLOP), Post-Processing, and Final Output Visualization. A block diagram representation of the framework is shown below:

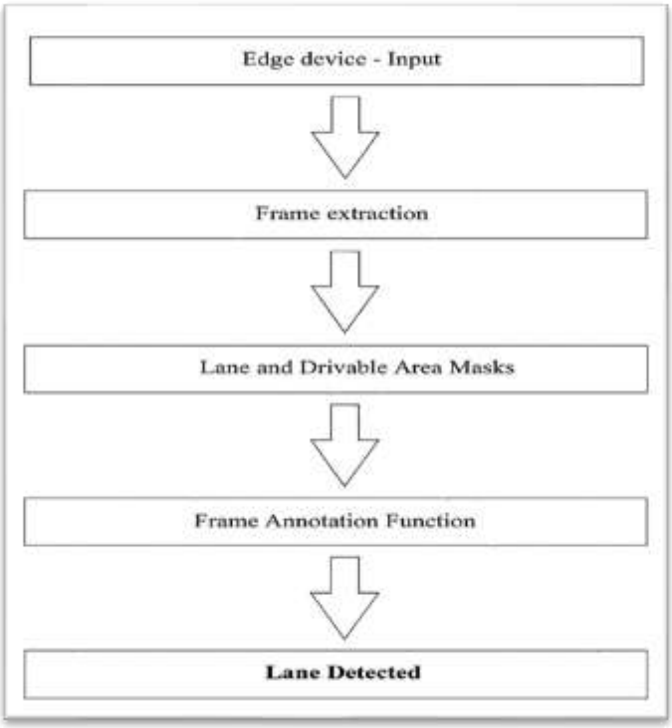

Fig 1: block diagram

4.1. Input Acquisition

The system begins by capturing frames from an onboard camera mounted on the vehicle (or an edge device like a smartphone). These frames serve as the raw input for the enhancement and perception modules. Key characteristics of the input stream:

- Real-time video frames (RGB)
- Varying lighting conditions (day, dusk, night)
- Presence of shadows, glare, occlusions, and noise

4.2. Night-Time Enhancement Module (Project Innovation)

To overcome challenges associated with low-light environments, a specialized enhancement pipeline is integrated before YOLOP processing. This module performs:

1.Gamma Correction

- Brightens dark regions while preserving highlights
- Enhances lane visibility in dim conditions

2.CLAHE (Contrast-Limited Adaptive Histogram Equalization)

- Improves local contrast
- Reduces glare and improves edge clarity

3.Adaptive Brightness & Noise Reduction

- Automatically adjusts brightness based on frame statistics
- Suppresses noise from sensors and headlights

4.3. Multi-Task Perception Module (YOLOP Model)

After enhancement, the processed frames are forwarded to **YOLOP**, a unified multi-task neural network designed for autonomous driving perception. YOLOP uses a shared backbone (encoder) and three parallel heads:

1. **Lane Line Segmentation Head**
   - Predicts lane boundaries pixel-wise
   - Handles broken lanes, curves, and occlusions
2. **Drivable Area Segmentation Head**
   - Identifies safe regions where the vehicle can move
   - Distinguishes between road, shoulders, and non-driveable areas
3. **Object Detection Head**
   - Detects vehicles, pedestrians, two-wheelers, etc.
   - Outputs bounding boxes, classes, and confidence scores

Advantages of multi-task architecture:

- Shared computation reduces latency
- Cross-task reinforcement improves accuracy
- Suitable for real-time ADAS and autonomous vehicles

4.4. Post-Processing and Lane Stabilization

The outputs from YOLOP undergo several refinement steps:

- Lane Curve Fitting using polynomial approximation
- Noise filtering to remove false lane segments
- Temporal smoothing across continuous frames to reduce flicker
- Confidence thresholding for reliable lane predictions

These steps ensure that lane boundaries remain stable even during motion blur, occlusions, or abrupt illumination changes.

4.5. Final Output Visualization

The processed results are combined and displayed as an overlay on the input frame. The visualization includes:

- Highlighted lane lines (left and right lanes)
- Detected drivable region mask
- Bounding boxes around detected objects
- Frame-by-frame status indicators (confidence, FPS, etc.)

The final output provides a comprehensive and intuitive view of the road, enabling safer and more accurate autonomous navigation.

5. MODELLING AND ANALYSIS

The proposed vision-based perception system for autonomous driving integrates three major tasks—lane detection, drivable area segmentation, and object detection—within a unified deep learning architecture. This section explains the functioning, importance, and analysis of each component, followed by a detailed overview of the multi-task framework that binds them together.

5.1 . Lane Detection

Lane detection is the core task in road scene understanding. It focuses on identifying visible lane boundaries on the road surface using image-based segmentation. In this project, lane detection is implemented using a deep convolutional neural network that classifies every pixel of the input image into either "lane" or "non-lane."

Modern systems, including YOLOP, utilize a shared encoder (typically a CNN-based feature extractor like CSP Darknet) that captures low-level patterns such as edges, textures, and global lane structures. The extracted features are then passed to a dedicated lane segmentation head, responsible for generating binary masks that highlight lane lines.

The segmentation head applies convolutional filters, up-sampling layers, and skip connections to reconstruct high-resolution lane maps. Loss functions commonly used for this task include Dice Loss, Binary Cross Entropy (BCE), or a combination of both to handle class imbalance—since lane pixels usually represent a very small portion of the image.

The resulting lane mask is stable enough to support real-time lane keeping, lane departure warnings, and path planning. CNN-based lane detection significantly outperforms traditional techniques by maintaining accuracy even when lane markings are faded, curved, obstructed, or influenced by varying light conditions.

Fig 2: Lane detection using CNN

5.2. Drivable Area Segmentation

Drivable area segmentation determines which parts of the road are safe for the vehicle to travel. Unlike lane detection—which depends on explicit markings—this task relies on contextual road cues such as road texture, boundaries, sidewalks, and curbs.

The drivable area head in YOLOP is implemented as a semantic segmentation branch, similar to the lane detection head but optimized for broader region masks. It classifies each pixel into categories such as "drivable" and "undrivable."This task identifies safe driving areas on the road, even without lane markings. It works alongside lane detection by using context clues like road edges, curbs, and sidewalks. The segmentation head predicts a mask for drivable areas, often trained with Focal or Cross-Entropy Loss.

This task becomes essential when:

- Lane markings are absent
- Complex road designs exist (roundabouts, junctions)
- Roads are damaged or unmarked
- The view of lanes is obstructed

The segmentation output allows the vehicle to understand free space for planning maneuvers such as overtaking, lane merging, or avoidance. Models typically use Cross Entropy Loss or Focal Loss to handle ambiguous road textures and class imbalance.

Fig 2: Drivable area segmentation

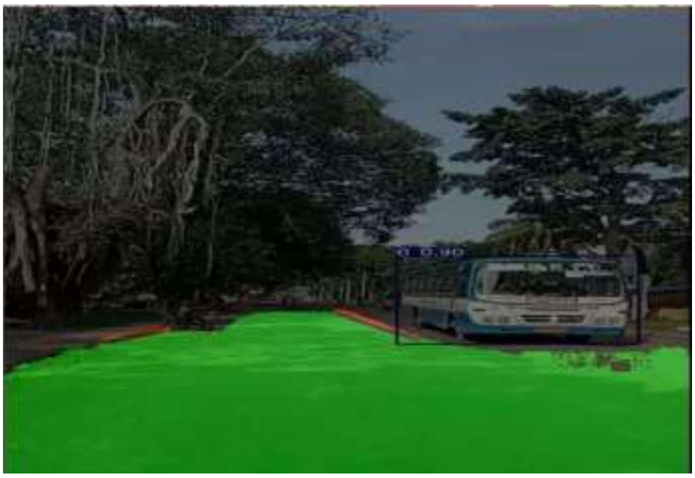

Fig 3: Object detection

5.3. Object Detection

Object detection identifies and locates vehicles, pedestrians, bicycles, and other moving elements in the driving scene. This component makes sure that the autonomous system doesn't misinterpret blocked lanes or enter dangerous areas. YOLOP's detection module uses YOLO's bounding box regression and classification framework. It predicts:

The coordinates of bounding boxes

- Object class labels
- Confidence scores

The model processes features extracted from the shared encoder and sends them to multi-scale detection heads. This allows it to detect both small and large objects. This ability is important for:

- Avoiding collisions
- Determining if lane boundaries are partially blocked
- Enhancing the system's situational awareness
- Planning safe trajectories

The object detection loss usually combines Localization Loss, Confidence Loss, and Class Probability Loss, following YOLO's standard training method. In addition, techniques like Non-Maximum Suppression (NMS) are used to eliminate duplicate detections. The multi-scale design also boosts performance in crowded or complex road scenes. This makes the detection branch very reliable for real-time driving situations.

5.4 . Multi-Task Learning Architecture

The strength of the proposed system lies in its multi-task learning (MTL) design**,** where a single encoder is shared across three separate decoders:

1. **Lane Detection Head**
2. **Drivable Area Segmentation Head**
3. **Object Detection Head**

The shared encoder cuts down on redundancy, lowers computational cost, and decreases latency. By learning from several related tasks at the same time, the network creates richer and more general feature representations. Lane markings, drivable surfaces, and objects often have contextual relationships, and MTL helps the model take advantage of these connections.

Benefits of the multi-task framework include:

- **Higher efficiency** → fewer parameters and faster inference
- **Improved accuracy** → shared features enhance all tasks
- **Better generalization** → model becomes robust to variability
- **Real-time performance** → suitable for on-road deployment

However, MTL also presents challenges, like balancing losses between tasks and making sure that one task does not take over the learning process. Still, YOLOP is known as one of the most practical and powerful architectures for unified driving perception.

5.5 . Overall Analysis

YOLOP's multi-task framework manages lane detection, drivable area segmentation, and object recognition in real-time. It outperforms traditional single-task models.

Key advantages include:

• Increased stability even in complex situations.

• Less flickering due to richer feature representation.

• Improved performance in different lighting and weather conditions.

• Real-time inference speeds that are suitable for ADAS and AVs.

• Greater accuracy in segmentation and detection tasks.

• Graph: Radar chart showing IoU, map, FPS for YOLOP compared to other models like Lane Net or SCNN.

The multi-task structure reduces the need for separate models, which lowers overall memory usage. It also provides better scene understanding by sharing features across tasks. These advantages lead to smoother and more reliable outputs during continuous driving sequences.

Limitations:

• Reduced performance in extreme weather or low-light conditions.

• Computational challenges on low-power embedded devices.

Future Directions:

• Lightweight transformer-based designs and temporal modelling can improve efficiency, robustness, and real-world reliability.

Integrating advanced video-based tracking modules and better low-light enhancement techniques can further boost overall system performance.

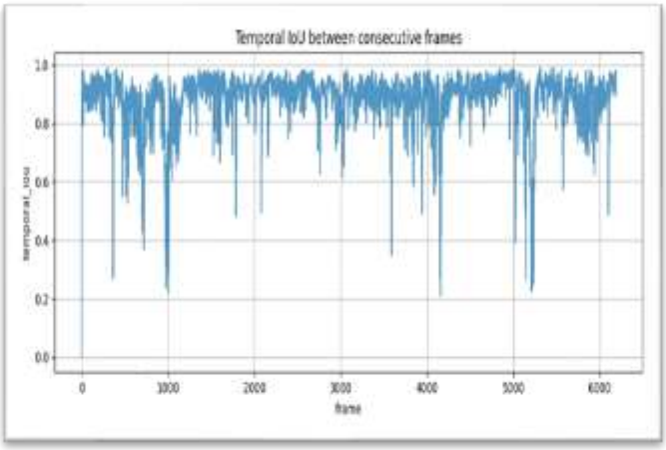

Fig 4: Temporal IoU between Consecutive Frames

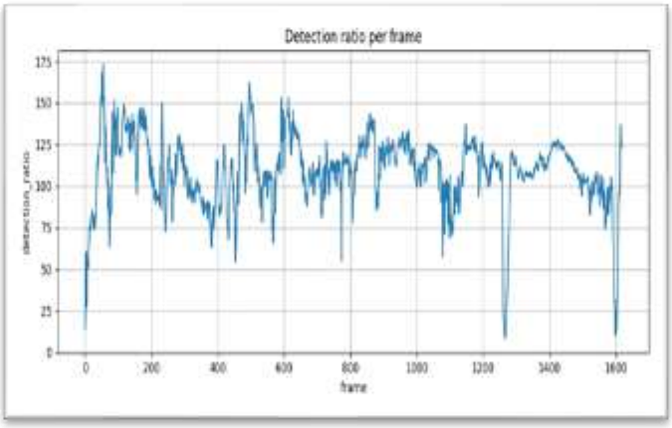

Fig 5: Detection Ratio per Frame

6. CONCLUSION

Lane detection has changed a lot from its early days with basic image-processing methods to the advanced deep learning models we use now. Traditional techniques like edge detection and Hough Transform were helpful but struggled in real-world driving situations, especially with issues like low light, faded markings, obstacles, or complicated road layouts. The shift to machine learning, especially deep neural networks, has revolutionized the field. These systems can now learn valuable features on their own, adjust to different environments, and understand road scenes more accurately.

Multi-task learning frameworks, like YOLOP, have taken this development further by combining lane detection, drivable area segmentation, and object detection into one system. This approach resembles how humans view their environment, which helps autonomous driving systems become smarter, more aware of their context, and able to operate in real time. Using shared encoders, task-specific decoders, and effective feature extraction has greatly sped up and increased the reliability of lane detection systems in various driving situations.

Despite these advancements, several challenges remain. Existing models still struggle to maintain temporal consistency, perform well in extreme weather, and run efficiently on vehicle hardware with limited computing power. Dataset limitations, changes in domain, and the need for interpretability also point to areas needing more innovation. The current landscape shows clear movement toward transformer-based architectures, lightweight model compression strategies, and improved night vision or low-light processing. Your project begins to tackle these areas through enhanced night-time visibility.

Overall, this review highlights the fast progress and ongoing potential of machine learning for lane detection. The ongoing improvement of deep learning models, along with better datasets and hardware optimization, is pushing the field closer to dependable, real-world autonomous driving. As research moves forward to close the remaining gaps, machine learning-based lane detection will be crucial in creating safer, smarter, and more efficient transportation systems.

7. FUTURE SCOPE

Despite significant progress in deep learning for lane detection, there are still ways to improve real-world performance. A key direction for future work is developing lightweight and hardware-efficient models that can run smoothly on embedded automotive platforms. Techniques like model pruning, quantization, and knowledge distillation can help lower computational load while maintaining high accuracy. This will enable wider use in commercial ADAS systems.

Another important area involves improving temporal stability. Most current models process each frame independently. This leads to occasional inconsistencies under occlusion, vibration, or sudden changes in lighting. Incorporating temporal modelling through approaches like 3D CNNs, LSTMs, optical flow, or transformer-based sequence learning can help maintain smoother and more reliable lane predictions across video sequences.

Improving strength in challenging environments is a key research focus. Real-world roads have different conditions, such as rain, fog, night-time glare, shadows, and faded markings. Larger datasets, domain adaptation methods, and synthetic-to-real transfer learning can help models work better across various weather, lighting, and geographic situations. This project enhances lane visibility at night, but future work could look into better lighting improvement using GAN-based or transformer-based correction.

Finally, future systems may shift toward deeper integration with complete driving tasks. Combining lane detection with steering prediction, modelling lane-change intent, or planning paths can create a stronger base for autonomous navigation. Model transparency is just as important. Techniques for explain ability, estimating uncertainty, and validating safety will be essential for building trust and gaining regulatory approval of AI-driven driving systems.

Overall, the future of lane detection lies in creating models that are efficient, temporally consistent, environmentally resilient, and seamlessly connected to broader autonomous driving functions.

ACKNOWLEDGEMENT

The authors express their sincere gratitude to all individuals and institutions who contributed to the successful completion of this research work. Deep appreciation is extended to the project guide, Prof. Pranesh Kulkarni, for his invaluable guidance, constant support, and insightful feedback throughout the course of the project. His expertise, patience, and encouragement played a central role in shaping the direction and quality of this study.

The authors also gratefully acknowledge the management of KLS Vishwanathrao Deshpande Institute of Technology, Haliyal, for providing the infrastructure, resources, and academic environment necessary to carry out this research. Special thanks are due to the faculty members and staff of the Department of Computer Science and Engineering for their assistance, suggestions, and motivation at various stages of the work.
The support and cooperation of fellow students and friends are also sincerely appreciated, particularly for their constructive inputs, discussions, and moral support during the development and testing of the system.
Finally, the authors wish to thank their family members for their unwavering encouragement, understanding, and support. Their constant motivation and belief in the authors' efforts were vital to the successful completion of this project.

REFERENCES

[1] Y. Liu et al., "Learning Light Fields for Improved Lane Detection," *arXiv preprint*, 2023. [Online]. Available: https://arxiv.org/abs/2301.00762

[2] J. Zhang et al., "Enhanced SCNN-Based Hybrid Spatial-Temporal Lane Detection Model," *IEEE Transactions*, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10459010

[3] X. Wang et al., “YOLOP: You Only Look Once for Panoptic Driving Perception,” in *CVPR Workshops*, 2022. [Online]. Available: https://arxiv.org/abs/2108.11250

[4] Y. Zhao et al., “Research on Road Scene Understanding Based on Multi-Task Learning,” *arXiv preprint*, 2023. [Online]. Available: https://arxiv.org/abs/2302.10003

[5] M. U. Arif, M. I. Faroo, R. H. Ran, Z. U. A. Lodhi, and M. A. R. Hashmi, “A Comprehensive Review of Vehicle Detection Techniques Under Varying Moving Cast Shadow Conditions Using Computer Vision and Deep Learning,” *Journal/Conference*, 2024.

[6] A. A. Khan et al., “YOLOP-Lite: Efficient Road Object Detection and Drivable Area Segmentation for Edge AI,” *IEEE Access*, vol. 12, pp. 32568–32579, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10459010

[7] R. Agarwal and M. K. Dutta, “Lane Detection Techniques Based on Image Processing and Deep Learning,” in *ICTIS 2017*, vol. 83, pp. 143–153, Springer, 2018. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-92007-8_13#citeas

[8] Y. Zhang et al., “VD-YOLO: Real-time Object Detection and Drivable Area Segmentation,” *Int. J. Autom. Comput.*, vol. 20, pp. 93–107, 2023. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-92007-8_13#citeas

[9] T. Luo et al., “End-to-End Lane Shape Prediction With Transformers,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 6, pp. 5807–5818, 2023. [Online].Available: https://ieeexplore.ieee.org/document/999924