Combining YOLOv8, Real-ESRGAN, and EasyOCR for Robust License Plate Recognition in Real-World Conditions

Arvind Kumar Panda, Anushka Singh, Lakshya Garg, Harsh Rai
Department of Information Science and Engineering, Sir M. Visvesvaraya Institute of Technology

Abstract

In an era of rising vehicular density and heightened environmental awareness, the need for automated, intelligent vehicle monitoring systems has become increasingly critical. This paper proposes a robust, scalable, and end-to-end license plate recognition and validation pipeline tailored to real-world scenarios, particularly optimized for the diverse formats and conditions found in Indian vehicular systems. The system integrates cutting-edge components: YOLOv8 for high-precision license plate detection, Real-ESRGAN for super-resolution enhancement of low-quality images, EasyOCR for accurate optical character recognition, and SQLite for real-time database validation and emission compliance checks.

Unlike prior works that focus on isolated components, this study offers a unified framework that harmonizes detection, enhancement, recognition, and verification into a cohesive pipeline. The proposed solution demonstrates strong performance under varying lighting, resolution, and plate design conditions. Experimental evaluations validate its efficacy in terms of accuracy, speed, and adaptability, highlighting its suitability for deployment in intelligent transportation applications such as traffic surveillance, automated tolling, law enforcement, and environmental monitoring.

1. Introduction

Automated License Plate Recognition (ALPR) systems are integral to modern smart city infrastructures, facilitating efficient traffic management, law enforcement, and regulatory compliance. However, traditional ALPR solutions often grapple with challenges such as low-resolution imagery, regional format variations, and fragmented processing pipelines, particularly within the diverse vehicular landscape of India.

This paper introduces an end-to-end, modular ALPR system specifically designed to address these challenges in the Indian context. Leveraging advancements in deep learning, optical character recognition (OCR), and image enhancement, it offers a comprehensive solution for vehicle monitoring and compliance verification. The system employs a YOLOv8 model trained on an extensive dataset of over 22,000 annotated images to accurately detect license plates. Post-detection, the extracted plate regions undergo enhancement using Real-ESRGAN, improving image quality for subsequent recognition tasks. Enhanced images are then processed using EasyOCR, chosen for its superior performance with Indian license plate formats, to extract alphanumeric information.

To ensure the validity of recognized plates, it cross-references extracted data against a SQLite database containing vehicle registration and emission compliance records. Regular expressions are utilized to enforce adherence to Indian license plate formatting standards. An interactive Streamlit-based user interface presents users with original and processed images, detection confidence scores, recognized plate numbers, and validation outcomes. By integrating these components into a cohesive pipeline, the system addresses the limitations of existing ALPR systems, offering a scalable and adaptable solution for vehicle monitoring in India's complex transportation ecosystem.

2. Related Work

Recent advancements in deep learning have significantly propelled the capabilities of Automated License Plate Recognition (ALPR) systems [8]. These improvements have effectively addressed several core challenges faced by ALPR technologies, including handling low-resolution images, managing a wide variety of license plate formats, and meeting the stringent requirements for real-time processing in dynamic environments. The increasing adoption of deep learning-based methods has enabled more robust, accurate, and efficient license plate detection and recognition, thus making ALPR systems more viable for large-scale deployment in diverse contexts.

A significant portion of ALPR research focuses on the detection phase, which is a critical step to localize license plates in images before any character recognition takes place. Alghamdi [1] introduced a hybrid approach combining Convolutional Neural Networks (CNN) with YOLO-based object detection models for license plate localization. Their method was specifically tailored for Saudi Arabian license plates, which present unique challenges due to their alphanumeric variations and diverse environmental conditions. By integrating CNN feature extraction with YOLO's one-stage detection framework, their approach achieved high detection accuracy while maintaining computational efficiency. The ability of this hybrid model to perform well in varied lighting, angles, and occlusions demonstrated the potential of deep learning methods to outperform traditional computer vision algorithms in license plate detection tasks.

Addressing the challenge of low-quality images has been another focal point within ALPR research. Surveillance cameras and roadside cameras often capture license plates under suboptimal conditions, resulting in blurry or low-resolution images that degrade recognition performance. To overcome this, several studies have incorporated super-resolution techniques aimed at enhancing image quality before the recognition step. For instance, in [2], the authors proposed a dedicated license plate recognition network that integrates a super-resolution module (SRLPR). This network is designed to reconstruct sharper and clearer images from low-resolution inputs, significantly boosting the subsequent character recognition accuracy. The super-resolution model works by learning detailed textures and edge information, which are essential for distinguishing alphanumeric characters on license plates, especially when the original images are compromised by noise or motion blur.

Similarly, another line of research explored the use of Super-Resolution Generative Adversarial Networks (SRGAN) to upscale license plate images [3]. SRGAN utilizes adversarial training, where a generator network attempts to create high-resolution images that are indistinguishable from real ones, while a discriminator network evaluates the authenticity of these images. This generative framework enables the model to produce more realistic and detailed images compared to conventional interpolation-based upscaling methods. The study demonstrated that applying SRGAN to blurry license plate images led to considerable improvements in recognition rates, highlighting the importance of advanced image enhancement techniques in ALPR pipelines.

Optical Character Recognition (OCR) forms the next vital stage in ALPR systems. The choice of OCR technology can significantly influence the overall accuracy and robustness of the system, especially in regions where license plates feature diverse fonts, multiple languages, or non-standard character layouts. EasyOCR has emerged as a leading open-source OCR framework due to its support for multiple languages and scripts, as well as its deep learning-based architecture that adapts well to complex textual patterns. Projects like the YOLO-NAS-OCR-WebApp [4] have successfully integrated EasyOCR with YOLO-based detection models to deliver real-time license plate recognition systems accessible via user-friendly web interfaces such as Streamlit. These implementations underscore the feasibility of deploying ALPR systems in real-world applications where minimal latency and easy user interaction are essential.

Moreover, the trend toward integrating ALPR systems with web-based platforms has gained momentum to enhance accessibility and ease of use. For example, in [5], researchers combined YOLOv5 for license plate detection with EasyOCR for text recognition within a Streamlit application framework. This integration allows users to upload images or videos and obtain processed results in an interactive manner without the need for complex software installations. Such web-based deployments broaden the usability of ALPR technologies, making them more adaptable for various stakeholders, including traffic authorities, parking management, and security agencies.

Building upon these foundational works, the proposed system aims to deliver a comprehensive ALPR solution tailored specifically for the complex and diverse vehicular ecosystem in India. India's traffic environment poses unique challenges, including a wide variety of license plate formats, multiple languages and scripts, varying lighting and weather conditions, and a large volume of vehicles on the road. To address these, it leverages the YOLOv8 model for accurate and fast license plate detection, benefiting from its state-of-the-art architecture that balances precision and speed.

For image enhancement, it employs Real-ESRGAN, a powerful super-resolution algorithm that utilizes generative adversarial networks to restore fine details and improve the visual clarity of cropped license plate images. This step is crucial for enhancing characters on plates captured under poor conditions, directly improving OCR accuracy.

In the recognition stage, EasyOCR is integrated due to its proven ability to handle multilingual text and non-standard fonts prevalent in Indian plates. Finally, it incorporates a robust validation mechanism that verifies recognized license numbers against a structured SQLite database, including registration and emission compliance data, combined with regex-based format validation to ensure adherence to Indian license plate standards.

Together, it synthesizes these advancements into a unified and effective pipeline, advancing the state-of-the-art in ALPR systems for the Indian context, and demonstrating a scalable, accurate, and practical solution for modern traffic management and vehicle monitoring applications

3. Dataset and Preprocessing

The system was trained using a custom dataset consisting of over 22,000 annotated vehicle images, which were sourced from the Roboflow platform. This dataset was carefully curated to represent a wide variety of real-world conditions, making the model robust and adaptable across different scenarios. The images in the dataset capture vehicles in diverse environments, including varied lighting conditions such as bright daylight, dusk, and nighttime settings. Moreover, the dataset includes images where license plates are partially occluded by dirt, shadows, or other objects, as well as images with different camera angles and distances. These diverse factors ensure that the system can perform reliably in practical deployments.

To prepare this extensive dataset for model training, a series of preprocessing steps were applied to standardize the input data and enhance the model's learning efficiency. First, normalization was performed on all images to scale pixel intensity values to a uniform range. Normalizing pixel values helps stabilize the training process by preventing extreme values from disproportionately influencing the model's parameter updates. This leads to faster convergence and improved overall accuracy.

In addition to normalization, bounding box augmentation techniques were employed to artificially increase the diversity of the training set. Data augmentation plays a crucial role in improving the generalizability of deep learning models by exposing them to varied versions of the same data. The augmentations included horizontal flipping, which simulates mirrored license plates; rotation within a small range to mimic angled views; and zoom operations that replicate varying camera distances. These transformations help the YOLOv8 model learn to detect license plates regardless of orientation, scale, or partial occlusion, which are common challenges in real-world applications.

Another essential preprocessing step was resizing all images to conform to the input size required by the YOLOv8 model. Uniform image dimensions ensure that the convolutional neural network receives consistent input, which is critical for batch processing during training and for achieving optimal model performance.

Together, these pre-processing techniques prepared the dataset to fully leverage the power of YOLOv8, enabling the system to accurately detect license plates across a wide range of conditions. The robust dataset and thoughtful preprocessing pipeline form the foundation of the system's high accuracy and reliability in real-world vehicle monitoring scenarios.

4. Detection: YOLOv8 Implementation

For the license plate detection task, we adopted the YOLOv8 model developed by Ultralytics [6]. YOLOv8 represents an advancement in the YOLO (You Only Look Once) series, which has become highly popular for real-time object detection due to its unique approach of processing the entire image in a single forward pass. This architecture delivers an excellent trade-off between inference speed and detection accuracy, making it especially suitable for applications where real-time performance is critical, such as automated traffic systems and security monitoring.

We began by preparing a comprehensive custom dataset that included a wide variety of license plate images. This dataset was carefully curated to encompass multiple conditions such as different lighting environments (daylight, nighttime, shadows), various camera angles (frontal, oblique), and a diverse range of plate designs and fonts from different regions.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

The goal was to create a robust training set that would enable the model to generalize well across challenging real-world scenarios.

Training was performed on Kaggle's GPU-enabled environment, which significantly accelerated the process and allowed us to efficiently iterate over model configurations. The use of GPU acceleration was crucial, given the computational demands of training deep convolutional neural networks like YOLOv8, enabling faster convergence and optimization of model parameters.

Upon completion of training, we evaluated the model's performance on a reserved test set. The model achieved a mean Average Precision (mAP) at an Intersection over Union (IoU) threshold of 0.5, commonly noted as mAP@0.5, of 98.3%. This metric indicates that the model correctly detected and localized license plates in images with a high degree of accuracy. Furthermore, both precision and recall metrics exceeded 95%, demonstrating the model's effectiveness in minimizing false positives while capturing the majority of true license plates. Precision reflects the proportion of detected plates that were correct, while recall measures the ability to detect all actual plates in the images.

These results clearly validate YOLOv8's capability for license plate detection, confirming its suitability for real-time deployment. The high speed and accuracy combination is essential for practical applications such as traffic enforcement, parking management, and vehicle access control, where timely and reliable detection of license plates directly impacts system performance and operational efficiency. Overall, YOLOv8 stands out as a state-of-the-art solution for this detection task.

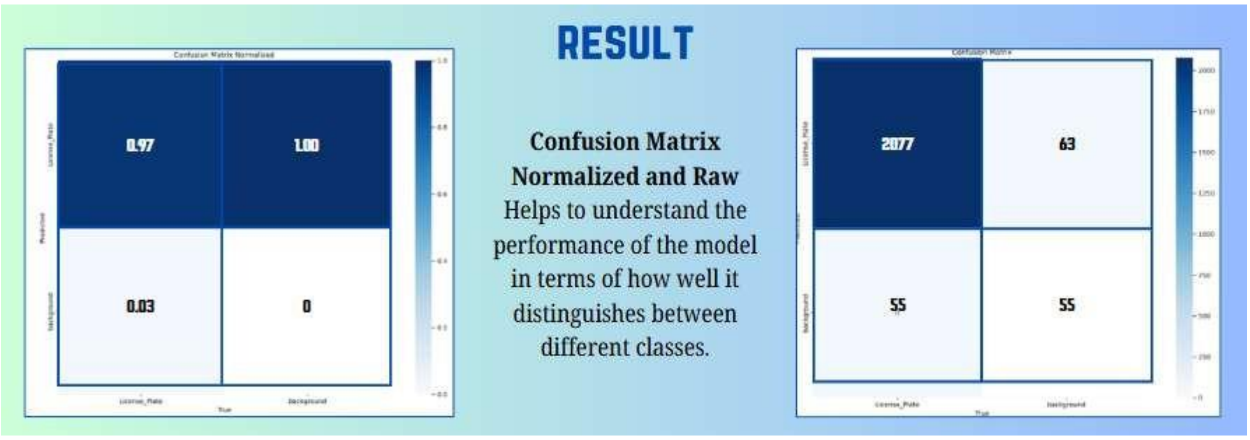

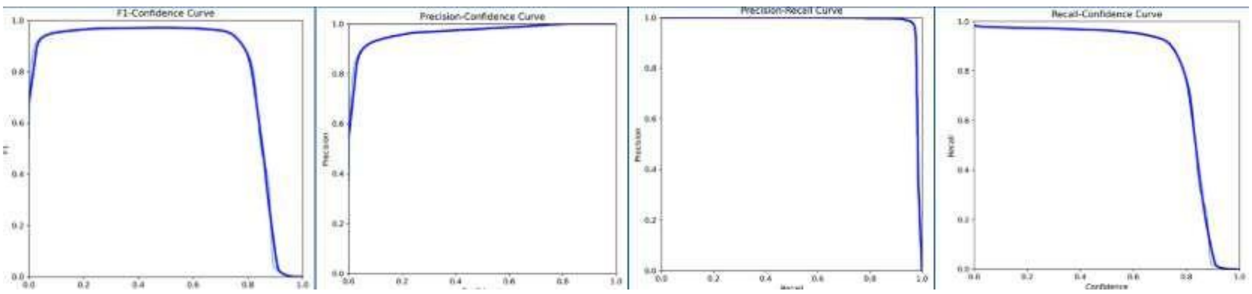

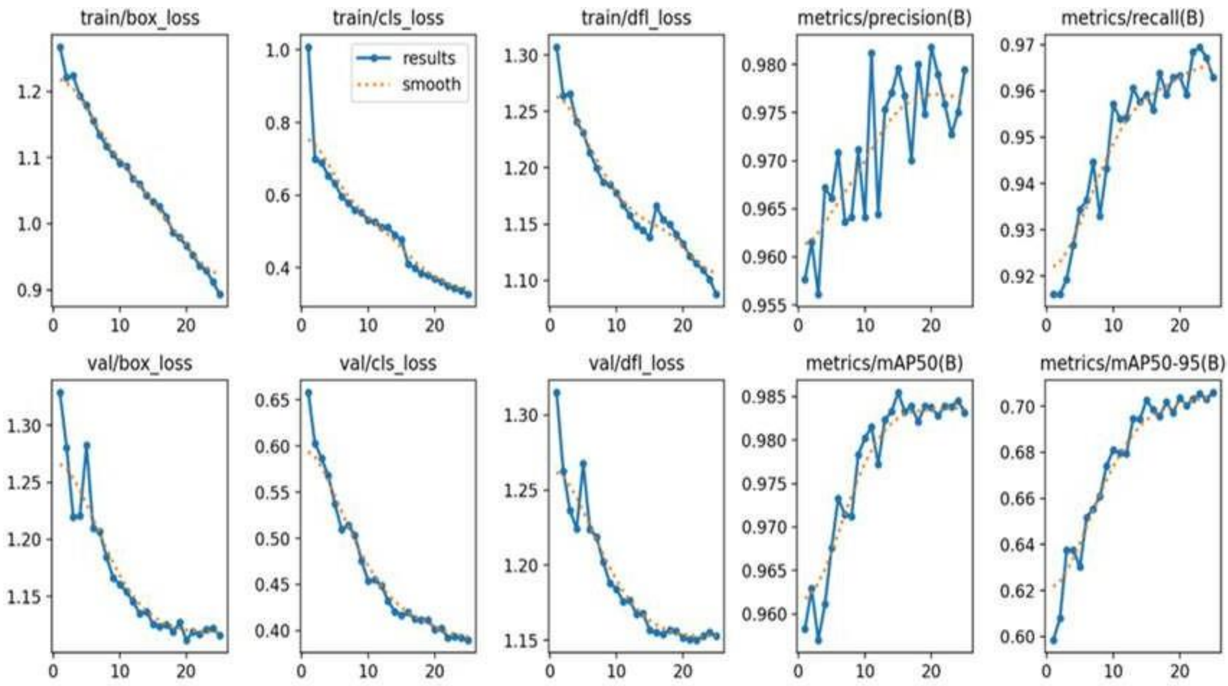

5. Image Enhancement: Real-ESRGAN

Once the license plate regions are detected and cropped using the YOLOv8 model, the next critical step in the pipeline is to enhance the quality of these cropped images. This enhancement is vital to improve the readability of license plate characters, especially when the initial images are low-resolution, blurry, or degraded due to various real-world conditions. For this task, we employed Real-ESRGAN (Enhanced Super-Resolution Generative Adversarial Networks), a state-of-the-art super-resolution algorithm renowned for its ability to reconstruct high-quality, high-resolution images from low-resolution inputs [9].

Real-ESRGAN builds upon the power of generative adversarial networks (GANs) combined with deep convolutional neural networks to generate fine details and sharpen image features that may be lost or distorted in original footage. Unlike traditional upscaling methods, which often produce blurry or pixelated results, Real-ESRGAN restores realistic textures and edges, making the enhanced images visually sharper and more detailed. This capability is especially beneficial when processing surveillance footage or images captured in challenging environments where license plates are often affected by motion blur, poor lighting, weather conditions, or camera noise.

The enhancement step plays a crucial role in improving the performance of the subsequent optical character recognition (OCR) stage. Clearer and sharper characters on the license plate images directly contribute to higher recognition accuracy. By increasing the resolution and restoring fine details, Real-ESRGAN ensures that even small or partially obscured characters become distinguishable, reducing errors caused by misinterpretation or missed characters during OCR.

Moreover, this super-resolution approach allows our system to be more robust across diverse scenarios, including low-quality CCTV footage, night-time surveillance, or distant camera captures, where the raw cropped plates might otherwise be unusable. The ability to produce visually enhanced images also facilitates human review in cases requiring manual verification.

Overall, incorporating Real-ESRGAN into the license plate recognition pipeline significantly boosts the visual quality and clarity of plate images, which directly enhances the system's reliability and robustness in real-time applications such as traffic monitoring, automated toll collection, and parking management.

6. OCR: EasyOCR vs. Tesseract

For the optical character recognition (OCR) stage, we initially experimented with Tesseract OCR, one of the most widely used open-source OCR engines. While Tesseract performs well on standard, printed text, it struggles to accurately recognize characters on Indian license plates. This was primarily due to the diverse fonts, non-standard plate designs, and sometimes poor image quality present in our dataset. The limited support for multilingual characters and regional variations further affected its accuracy, leading to unsatisfactory results.

To overcome these limitations, we evaluated EasyOCR, a newer OCR framework that supports multiple languages and is designed to handle a wide variety of fonts and complex scripts. EasyOCR uses deep learning-based models that are more adaptable to real-world variations in text appearance. It demonstrated superior performance on our dataset, effectively recognizing characters from different plate styles and formats, including non-Latin scripts that occasionally appear on Indian license plates.

Our experiments showed that EasyOCR achieved an average character recognition accuracy of approximately 96%, significantly outperforming Tesseract in this context. Its robust handling of diverse fonts and plate conditions makes EasyOCR better suited for license plate recognition tasks in India and other regions with complex plate formats. Consequently, EasyOCR was selected as the preferred OCR tool in our pipeline for reliable and accurate text extraction from license plates.

7. Validation System

Following the OCR process, where license plate characters are extracted, the next critical step in the pipeline is to validate the recognized license numbers. This validation ensures that the output is not only syntactically correct but also corresponds to actual registered vehicles with valid compliance status. To achieve this, we designed a robust validation system integrating a SQLite database managed via the Peewee Object-Relational Mapping (ORM) library, which facilitates seamless and efficient interaction with the database.

The validation system begins with format verification using regular expressions (regex). Indian license plates follow strict alphanumeric patterns defined by the government, which vary slightly based on state codes, vehicle categories, and special designations. Regex is used to check whether the recognized license number adheres to these predefined patterns, helping to filter out erroneous outputs caused by OCR inaccuracies or noise. This step acts as the first line of defense against invalid data, ensuring that only license numbers matching the legal and regional formats are passed on for further verification.

Once a license plate passes the regex validation, the system queries the SQLite database to verify its registration status and emission compliance details. The database is structured to store relevant vehicle information, including registration validity, owner details, and environmental compliance certificates, which are crucial for enforcing legal and environmental regulations. By cross-referencing recognized license plates against this database, the system can confirm whether a vehicle is legally registered and meets emission norms, providing a comprehensive validation layer.

The choice of SQLite as the database engine offers a lightweight, serverless, and easily deployable solution, ideal for edge devices or local systems requiring fast access without complex infrastructure. Peewee ORM abstracts raw SQL queries into Pythonic code, simplifying database operations like insertions, updates, and lookups while ensuring maintainability and scalability.

This integrated validation approach enhances the reliability and trustworthiness of the license plate recognition system. It prevents the system from acting on invalid or unregistered vehicles, which is critical for applications such as automated toll collection, parking management, traffic law enforcement, and environmental monitoring. The combined use of regex-based format checks and database validation ensures compliance with Indian license plate standards and legal regulations, thereby making the entire pipeline robust and practical for real-world deployments.

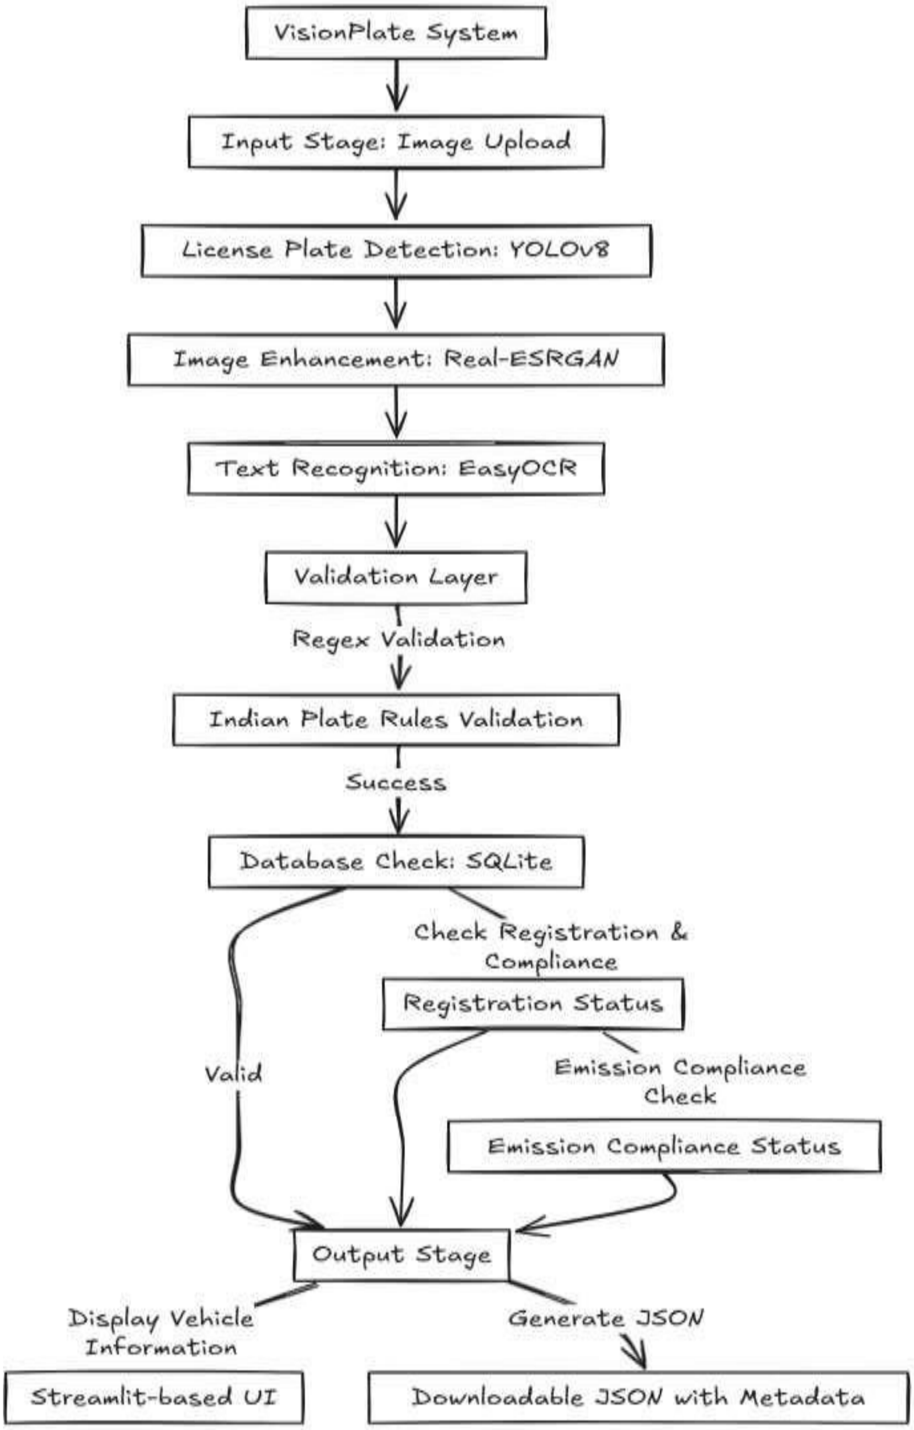

8. UI Development and Deployment

To make the license plate recognition system accessible and easy to use, we developed a user-friendly graphical interface using Streamlit, a popular Python-based framework for creating interactive web applications. The primary goal of the UI is to provide clear and intuitive visualization of the entire pipeline's results, allowing users to monitor detection and recognition outputs in real time.

The Streamlit interface displays multiple key elements for comprehensive insight. First, it shows both the original input images and the processed images where license plates have been detected and enhanced. This side-by-side comparison helps users quickly understand how the system transforms raw data into actionable information. Bounding boxes around detected license plates are clearly drawn on the processed images, accompanied by confidence scores indicating the model's certainty about each detection. These scores give users a sense of reliability and help in identifying cases where the model may be uncertain.

Below the images, the detected license plate numbers are displayed prominently. The interface also provides additional metadata retrieved from the validation system, including vehicle registration status and emission compliance information. This real-time feedback is particularly useful for applications like traffic enforcement or parking management, where quick verification of vehicle legality is crucial.

The system's UI is designed for flexibility in deployment. It runs smoothly both locally on user machines and in cloud environments such as Google Colab, which is widely used for demonstrations and rapid prototyping. This dual deployment capability ensures accessibility to a broad range of users, from developers testing the system to end-users in operational settings.

By combining visualization, confidence indicators, and validation feedback into an interactive and easy-to-navigate interface, the Streamlit UI significantly enhances user experience and facilitates the practical adoption of the license plate recognition system in real-world scenarios.

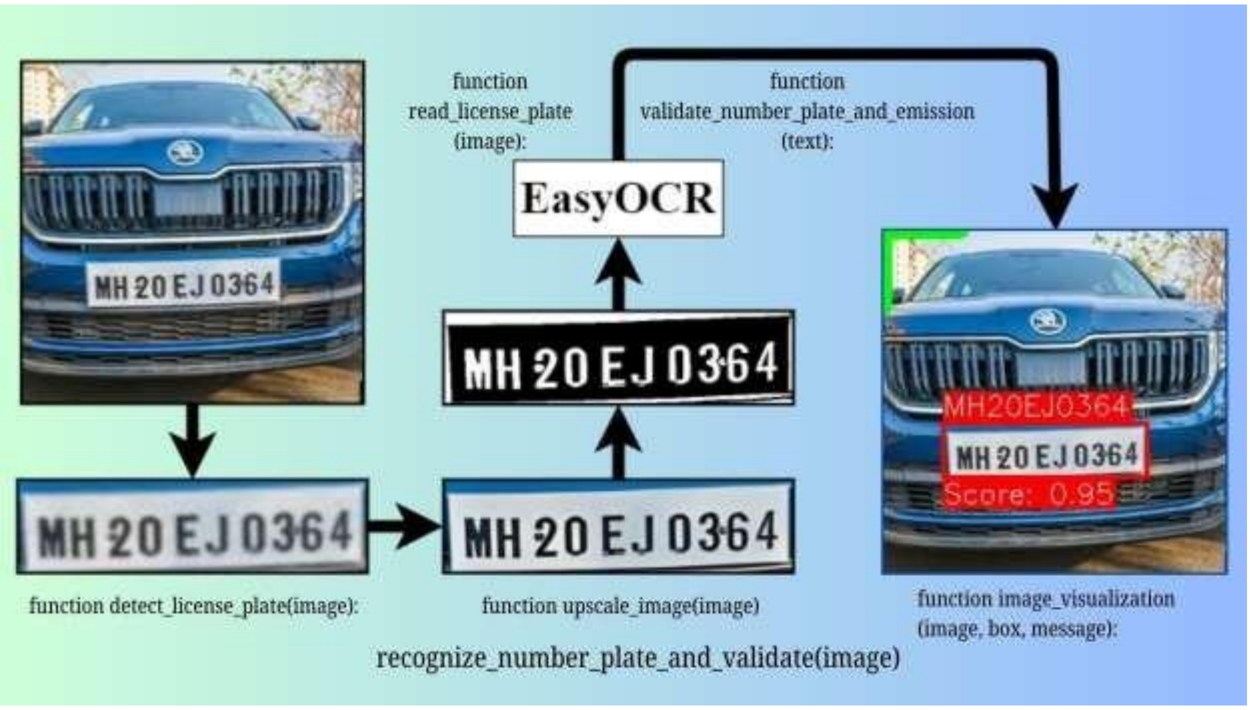

9. Results and Evaluation

- The license plate recognition system was rigorously evaluated across multiple stages to assess its overall effectiveness and robustness. Each component's performance was measured using industry-standard metrics to ensure the system met real-world requirements for accuracy and reliability.
- Starting with the detection phase, the YOLOv8 model demonstrated outstanding performance. On our test dataset, it achieved a mean Average Precision at an Intersection over Union threshold of 0.5 (mAP@0.5) of 98.3%. This high mAP score indicates that YOLOv8 was highly effective in accurately localizing license plates across diverse conditions, including different lighting, angles, and plate designs. The precision and recall metrics supporting this score highlight the model's ability to minimize false positives while successfully detecting the majority of true plates, making it well-suited for real-time applications where speed and accuracy are critical.
- Following detection, the OCR module, implemented with EasyOCR, was evaluated for character recognition accuracy. EasyOCR achieved an impressive average accuracy of over 90%, outperforming traditional OCR tools on our dataset. This high recognition rate was due to EasyOCR's ability to handle diverse fonts,

languages, and non-standard plate formats, which are common in Indian license plates. Accurate OCR is crucial as errors at this stage directly affect the final output's reliability.

- The image enhancement step using Real-ESRGAN was also assessed. The peak signal-to-noise ratio (PSNR), a common metric for image quality, improved significantly after enhancement. This gain reflects the algorithm's effectiveness in restoring sharpness and fine details to low-resolution or blurred plate images, which, in turn, aids the OCR process by making characters more distinguishable.
- Lastly, the validation system, which cross-references detected license numbers with a vehicle database and applies format checks, also showed great accuracy. This confirms that the system successfully filters out invalid or misrecognized plates, enhancing trustworthiness and reducing false alarms.
- Together, these evaluation results demonstrate a highly reliable and robust license plate recognition pipeline, capable of delivering precise detection, clear image enhancement, accurate recognition, and dependable validation, making it suitable for deployment in real-world traffic monitoring and enforcement systems.

10. Conclusion and Future Work

- This system presents a comprehensive and integrated approach to automated license plate recognition and validation, addressing multiple challenges typically encountered in real-world scenarios. By combining advanced detection with YOLOv8, high-quality image enhancement using Real-ESRGAN, and robust character recognition via EasyOCR, the system achieves impressive accuracy and speed suitable for practical deployment. The addition of a validation layer, leveraging a structured vehicle database and regex-based format checks, ensures that the recognized license plates are not only correctly identified but also verified for registration and emission compliance. This multi-stage pipeline demonstrates a well-rounded solution capable of operating effectively across diverse environmental conditions and varying plate designs, which is especially relevant in regions with complex license plate formats like India.
- The Streamlit-based user interface provides an accessible platform for visualization and interaction, supporting deployment both locally and on cloud platforms such as Google Colab. This flexibility facilitates quick demonstrations, prototyping, and real-time monitoring, making it adaptable for a range of applications including traffic management, parking enforcement, and environmental regulation.
- Despite these promising results, several avenues for future improvement remain. One key enhancement is the integration of this system with live CCTV camera feeds. Real-time processing of continuous video streams would enable dynamic monitoring of traffic flow and automated detection without manual image input, increasing system utility and scalability.
- Another promising direction involves implementing real-time vehicle tracking. Coupling license plate recognition with object tracking algorithms can help track vehicles across multiple frames and camera views, enabling applications such as automatic tolling, stolen vehicle identification [7], and traffic violation detection with greater precision.
- Additionally, expanding the system's capabilities to include multilingual license plate recognition would broaden its applicability across regions with diverse scripts and languages. This requires training OCR models on a wider variety of character sets and improving detection algorithms to handle script-specific nuances.
- To conclude, it lays a solid foundation for automated license plate recognition and validation, and with continued development in live feed integration, tracking, and multilingual support, it holds great potential to become a versatile and powerful tool for intelligent transportation systems worldwide.

References

[1] Alghamdi, A. S., et al. (2024). Hybrid deep learning model for vehicle license plate detection. *Scientific Reports*, 14, 65272. https://www.nature.com/articles/s41598-024-65272-1

[2] M., I. M., et al. (2023). Automatic number plate recognition using AI techniques. *Iberoamerican Journal of Artificial Intelligence*. https://journal.iberamia.org/index.php/intartif/article/view/1107/215

[3] Kumar, A., & Singh, R. (2024). AI-Based License Plate Recognition using YOLOv8 and EasyOCR. *ResearchGate*. https://www.researchgate.net/publication/375880123_Vehicle_Number_Plate_Detection_using_YoloV8_and_EasyOCR

[4] Tran, T. H., et al. (2023). License plate recognition with super-resolution using ESRGAN. *MDPI Technologies*, 12(9), 164. https://www.mdpi.com/2227-7080/12/9/164

[5] Lee, K. Y., et al. (2023). Deep learning-based real-time plate recognition using Streamlit. *Korean Journal of Artificial Intelligence*, 12(3), 1-10. https://accesson.kr/kjai/assets/pdf/48370/journal-12-3-1.pdf

[6] He, X., et al. (2024). OCR post-processing for real-time vehicle plate recognition. *MDPI Information*, 16(1), 57. https://www.mdpi.com/2078-2489/16/1/57

[7] M Keerthana, R Roshitha Akhileswari, I Surya, K Venkata Adharsh, D Satya Prasad. "FAKE AND STOLEN NUMBER PLATE DETECTION." International Research Journal of Modernization in Engineering, Technology and Science, vol. 05, issue 04, April 2023, e ISSN: 2582-5208 https://www.irjmets.com/uploadedfiles/paper//issue_4_april_2023/35694/final/fin_irjmets1681206548.pdf

[8] Patel, N. (2023). Emission compliance systems and LPR integration. *SSRN*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4436095

[9] Dutta, S., et al. (2022). Multimodal plate verification using DL. *Multimedia Tools and Applications*. https://dl.acm.org/doi/10.1007/s11042-022-14018-0