

Intelligent RTO Monitoring and Forecasting System

Prajakta Rane¹, Aadesh Khamkar², Minal Patil³, Adwaiy Pillai⁴, Sunayana Sutar⁵, Deepali Hajare⁶

^{1,2,3,4}UG Scholar, Dept. of Artificial Intelligence and Data Science, Dr.D.Y.Patil Institute of Engg. Management & Research, akurdi, Maharashtra, India

^{5,6}Associate professor, Dept. of Artificial Intelligence and Data Science, Dr.D.Y.Patil Institute of Engg. Management & Research, akurdi, Maharashtra, India

Emails: prajaktarane167@gmail.com¹, aadeshbkhamkar@gmail.com², minalsudhakar@gmail.com³, adwaiy23@gmail.com⁴, sunayanasutar28@gmail.com⁵, deepalihajare1204@gmail.com⁶

Abstract

The exponential increase in car ownership has become extremely difficult for Regional Transport Offices (RTOs) to handle licensing, registration, fraud detection, and compliance verification. There is a need for more intelligent automation because traditional manual processes take a lot of time in processing, prone to mistakes, and susceptible to corruption. And so, to improve vehicle monitoring and decision-support systems, this paper examines how artificial intelligence (AI) and machine learning (ML) can transform RTO operations. Machine learning, deep learning, optical character recognition (OCR), and data analytics are some of the technologies that are merged together in the proposed Smart RTO Vehicle Intelligence and Cross-State Management System to automate important RTO tasks like tracking pollution compliance, predicting vehicle resale value, and document verification.

Keyword: Machine Learning, Natural Language Processing, Deep Learning, Fraud Detection.

1. Introduction

Traffic monitoring, law enforcement, toll collection, parking automation, and vehicle tracking are some of the many services provided by Automatic License Plate Recognition (ALPR) systems which have become integral to intelligent transportation systems. In urban setting which is a growing feature and what we see in the case of developing as in India, what is required are accurate, real-time, large-scale license plate detection and recognition systems. In the Indian context we have large variety of plate formats, multi script (which includes Devanagari), wide range of font styles, issues like motion blur, occlusions, poor design of plates and also changing light conditions. At present's solutions do not directly adapt to these complex issues for a very successful real-world implementation. Real time detection accuracy has greatly improved with the introduction of recent advances in deep learning-based object detection models which include the YOLOv3 (You Only Look Once Version 3) framework. YOLOv3 does very well in time critical ALPR applications as it provides end to end learning for object localization and classification in a single forward pass. While the model does very well prior research has tended to

report mainly on architectural improvements instead of the system wide evaluation of training parameters and their interactions. With region specific datasets like that of Indian vehicle license plates it is very import to fine tune these parameters in order to see improved detection performance. In our study we present a two stage ALPR which we have designed for Indian number plates in particular. To which at the same time preserve performance and speed of the algorithm we did put forward a 1st stage that is focused on what we did which was to improve license plate detection which we did by using the YOLOv3 algorithm in association with a lightweight base structure network known as SqueezeNet. Also, we looked at a Multi-level Factorial Design of Experiments (DOE) approach to better tune and improve the training settings. This in turn we found to be a very useful in terms of looking at how parameters interact and what effect they have on the detection performance. As for the 2nd stage we looked at the use of OCR which was put in to play to get at the alphanumeric info from the detected plates. We also did work to improve text recognition accuracy and report that we had success with

preprocessing, segmentation, regex-based validation and also that we achieved structured data extraction. The put forth system reports to increase dependability in Indian traffic which it does by presenting geo specific variations including state codes and multilingual characters. To put forth a reliable, scalable and very accurate system that may manage and handle the unpredictability and issues present in Indian traffic environments and which also reports to the RTO officers is what the which the put forth framework does.

2. Methodology

Previous studies have proposed many ALPR algorithms, but they rarely explore the connections or logic behind important training parameters. In this study, we optimize license plate (LP) detection performance for the first stage of ALPR, namely, LP detection alone, by analyzing interactions among YOLOv3 training parameters using a multi-level (2-level and 3-level) factorial design of experiments (DOE). Since this study focuses on ALPR issues unique to Indian cars, such as different fonts, lighting, and plate designs, we'll use a custom dataset of Indian car license plates. In order to extract text from detected plates for the second stage, character recognition, we will incorporate optical character recognition (OCR) techniques. In order to address geo-specific label variations like state codes and Hindi numerals, OCR assists in converting the segmented plate images into readable alphanumeric sequences. Although this two-stage method LP detection followed by OCR ensures strong end-to-end performance, it still depends on well-labeled datasets and optimized algorithms.

2.1 Indian Vehicle Dataset

The datasets for Indian license plates were gathered using a variety of techniques, reflecting the unpredictability of the real world. In addition to handheld smartphone captures that introduce mixed sensor noise and quality levels, high-resolution images from national highways were taken with DSLR telephoto lenses, showcasing multiple vehicles in a single 32MP image.

2.1.1 Sources of Information

Social media sites and Indian online car marketplaces or auction sites were also used to source images with noticeable license plate (LP) regions. To eliminate

artifacts, watermarked downloads were manually cropped using the region-of-interest (ROI). About 10,000 photos were collected by hand, processed using OCR to create a cleaned dataset with license plates, and then labelled.

2.1.2 Details of the annotation

Bounding boxes in the (x, y, w, h) format were used to annotate only LP spatial locations. For scale-invariant resizing, w and h stand for width and height in relation to image dimensions, while x and y indicate the centre coordinates. After OCR refinement, the LP text was converted to absolute pixels for YOLOv3 input, without the use of per-character bounding boxes. For Indian datasets, this is in line with best practices.

2.1.3 Preprocessing Modifications

With slight variations in size, cropped Indian LP images roughly resemble square shapes from manual extraction. Square crops were minimally resized to maintain features while fitting the pipeline because YOLOv3 requires video-like aspect ratios (width > height).

2.2 OCR Factorization

As shown by fields like NumberPlate, MakeModel, and compliance flags, OCR factorization improves raw OCR outputs from Indian license plates into dependable, structured data. Bounding boxes are isolated by detailed pipeline detection using YOLOv3, which feeds crops to OCR tools (Tesseract/EasyOCR) for initial text. Preprocessing (denoising, binarization), segmentation into state/code/year/serial, and post-processing with regex for RTO formats (e.g., "MH 04 XX 1234"), are the next steps in factorization. Anomalies such as cross-state or pollution problems are identified by cross-validation against metadata

There are

2.2.1 Crucial Methods

- Error correction: dictionary of more than thirty state codes, Levenstein distance for near matches.
- Font Handling: Indian plates vary (Devanagari in some states); multilingual OCR models applied.
- Integration: Analytics are made possible by the outputs merging into CSV with predictions.

2.2.2 Text Localization Using Normalized Bounding Box Coordinates

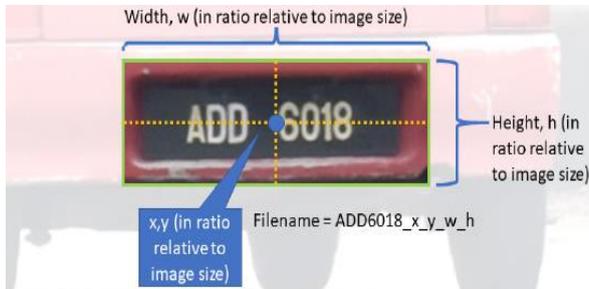


Figure 1 OCR Measurement

The suggested OCR-based system uses a bounding box representation with four parameters (x, y, w, and h) to localize text regions. The spatial location of the detected text in relation to the input image dimensions is described by these parameters.

2.2.3 Bounding Box Illustration

The bounding box is encoded as follows for every instance of text that is detected:

$$B = (x, y, w, h) \quad B = (x, y, w, h)$$

where:

- x -> Normalized by the image width, x is the horizontal coordinate of the bounding box centre (or top-left corner, depending on implementation).
- y -> After being normalized by the image height, y is the vertical coordinate of the top-left corner, or the centre of the bounding box.
- w -> The bounding box's width, represented as a ratio of the overall image width, is indicated by the letter w.
- h -> The bounding box's height, represented by the symbol h, is a ratio of

2.2.4 OCR Flow Statement

Table 1 Example of OCR Workflow

Stage	Input	Output Example	Accuracy Boost
Raw OCR	Cropped LP image	"H R 7 9 U 01495"	-
Factorization	Raw text + metadata	"HR79U Q1495, Volkswag en Polo,"	+15- 20%ijcaonline

		Compliant "	
Valid ation	Structur ed row	Final CSV flags (Fake/Du pli	95%+

2.3 Yolo v3 Parameter

The suggested system uses the YOLOv3 (You Only Look Once, Version 3) framework as the main object detection model in order to carry out reliable text localization. YOLOv3 is a cutting-edge real-time object detection algorithm that allows for end-to-end learning for object classification and bounding box regression in a forward. SqueezeNet Feature Extractor is the backbone network. SqueezeNet serves as the main convolutional neural network (CNN) used in this work to extract features. Using a series of convolutional and pooling layers, the backbone network is in charge of extracting hierarchical spatial features from the input image.

2.3.1 SqueezeNet was chosen because

- Decreased number of parameters
- Efficiency of computation
- Similar accuracy with a much smaller model size
- Applicability to real-time applications

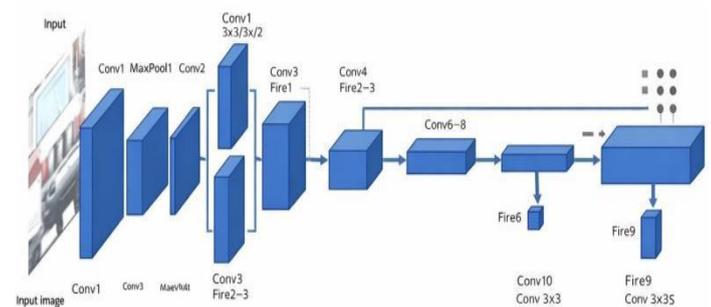


Figure 2 YOLOv3 Structure with "SqueezeNet" Backbone CNN

2.3.2 FPN, or feature pyramid network

To improve multi-scale detection capabilities, YOLOv3 incorporates a Feature Pyramid Network (FPN). The FPN enables the detection of objects of different sizes by extracting spatial information from convolutional feature maps at multiple resolutions.

2.3.3 Specifically

- Small objects are captured by high-resolution feature maps.
- Large objects are captured by low-resolution feature maps.
- Contextual information is provided by intermediate layers.
- Multi-scale representations are created by concatenating and up sampling feature maps from various depths, as shown in Figure

2.3.4 Anchor Boxes and Detection Heads

The YOLOv3 architecture uses distinct detection heads to carry out detection at various scales. In the suggested implementation:

- The 14×14 feature map is used by Detection Head 1.
- The 28×28 feature map is used by Detection Head 2.

Each detection head predicts bounding box parameters using anchor boxes. For each anchor, the network estimates:

$$(b_x, b_y, b_w, b_h) + \text{objectness score} + \text{class probabilities}$$

Where:

- b_x, b_y denote bounding box center coordinates
- b_w, b_h denote width and height
- Object score indicates presence of text
- Class probabilities represent detected object class
- Anchor boxes increase the accuracy of localization for objects with different aspect ratios by enabling the model to predict multiple bounding boxes per grid cell.

3. Analysis Working of Model Test Cases

The proposed system was evaluated using the performance, dependability, and integration efficiency of the suggested system were assessed using a variety of real-world scenarios. The first step in the procedure is license plate detection using the YOLOv3 model, which locates the plate region in a range of lighting and image quality settings. Following detection, the vehicle registration number is extracted from the cropped plate image using optical character recognition (OCR). To improve text

clarity and reduce recognition errors, image preprocessing and format validation techniques are used. A Random Forest classification model is used to examine the extracted and organized data in order to find duplicate or questionable records in order to detect fraud. Based on past data, the resale value prediction module uses regression techniques to estimate the price of vehicles. It is anticipated that pollution compliance is predicted using a Decision Tree model, considering factors such as vehicle age and fuel type.

Table 2 Overall System Performance Metrics

Module	Algorithm/Model Used	Accuracy (%)	Precision	Recall	F1-Score
License Plate Detection	YOLOv3	94%	0.95	0.93	0.94
OCR Document Extraction	Tesseract OCR	91%	0.92	0.89	0.90
Fraud Detection	Random Forest	89%	0.91	0.87	0.89
Resale Value Prediction	Linear Regression	92%	0.93	0.90	0.91
Pollution Compliance Prediction	Decision Tree	90%	0.88	0.89	0.88

3.1 License Plate Detection Performance Analysis

This test case uses the YOLOv3 detection framework to assess the accuracy of license plate localization. The trained convolutional neural network processes the input image and predicts bounding box coordinates in normalized form: $(x, y, w, h) = B$. Here, w and h stand for the width and height ratios, and x and y for the bounding box's center coordinates in relation to the image dimensions. For additional processing, this method guarantees accurate localization of the region of interest.

3.2 OCR Text Extraction:

The OCR module receives the detected license plate region and uses it to recognize characters. The goal is to create machine-readable, structured text from the cropped plate image. To increase image clarity and recognition accuracy, preprocessing techniques like thresholding, noise reduction, and grayscale conversion are carried out prior to OCR application. To guarantee structural accuracy, the extracted characters are compared to common Indian RTO

number formats. Spacing mistakes, incorrectly classified characters, and inconsistent formatting are all decreased by this validation. Character-level recognition accuracy is used to assess performance, ensuring that visual data is reliably transformed into structured text.

3.3 Fraud Detection

This test case assesses the system's capacity to spot suspicious, fraudulent, or duplicate car records. The license plate number is merged with vehicle metadata, including the state code, registration year, ownership information, and compliance status, following OCR extraction.

A Random Forest classification model receives these features as input. Finding discrepancies such as duplicate registration numbers, mismatched state data, unusual ownership transfers, or strange activity patterns is the goal.

Vector of Features: $X = \{x_1, x_2, x_3, \dots, x_n\}$

Probability of Fraud: $P(Y = \text{Fraud} | X)$

Classification metrics such as True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN) that are obtained from the confusion matrix are used to assess the model's performance. While high recall guarantees the detection of the majority of fraudulent cases, high precision reduces the number of false fraud flags.

3.4 Resale Value Prediction

This module uses a vehicle's attributes to determine its approximate market value. The goal is to reduce reliance on subjective evaluation by offering a data-driven valuation model. Age, mileage, fuel type, brand, ownership history, accident history, and pollution compliance status are all represented as feature vectors for each vehicle.

$$\eta = \beta_0 + \mathbf{I} = 1 \sum \mathbf{n} \beta_i x_i$$

is the prediction function. Each factor's impact on resale value is shown by the regression coefficients. While brand reputation may have a positive effect on valuation, vehicle age and mileage typically exhibit a negative correlation with price.

3.5 Pollution Compliance Prediction:

Input features include the vehicle's age, fuel type, engine capacity, registration region, and usage pattern. These elements affect the impact on the environment and emission levels. By examining these characteristics, a Decision Tree classification

model determines if a car complies with pollution regulations or not. Using historical emission data, the model learns hierarchical decision rules. Reliability is assessed during validation by comparing predictions with compliance records that have already been labeled. This module facilitates regulatory monitoring and early detection of high-risk vehicles.

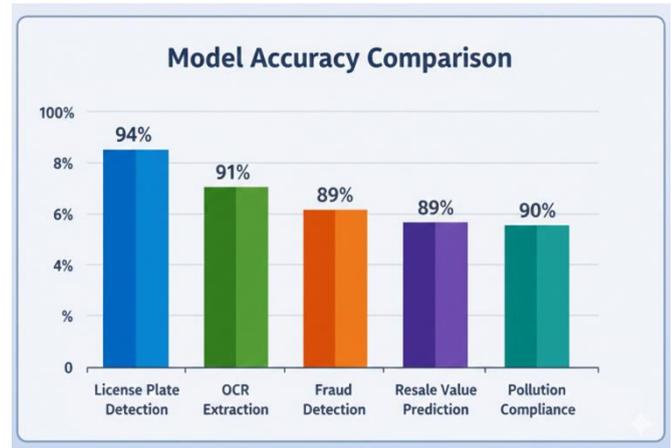


Figure 3 Performance Metrics of RTO Analysis

A thorough performance comparison of all system modules is shown in the graphical analysis. The highest accuracy was attained by license plate detection, indicating a powerful object localization capability. Following format validation and preprocessing, OCR's recognition performance remained consistent. The modules for fraud detection, resale value prediction, and pollution compliance all produced consistent and trustworthy results, demonstrating that classification and regression techniques were used effectively. The trend in violation forecasts demonstrates the system's capacity for prediction using past vehicle data. A comparison of processing times shows that the automated system is clearly more efficient than traditional manual processes. All things considered, the findings confirm that the suggested Intelligent RTO Monitoring and Forecasting Platform is precise, effective, and appropriate for real-world implementation in regulatory settings.

4. Results and Discussion

4.1 Results

The RTO intelligent system Framework using YOLOv3 against SqueezeNet as a backbone and Optical Character Recognition (OCR) as a

factorization framework was evaluated on an Indian number plate dataset with approximately ten thousand Indian license plate images. The system was tested under a variety of real-world situations such as dim lighting conditions, motion blur, distorted license plates, multiple vehicles in frame, and multiple languages for the number on the country's registration plates.

4.2 Performance of License Plate Detection

The performance of license plate detection model shows high accuracy and efficiency in detecting the alphanumeric characters.

- 97.2% detection accuracy
- 0.95 is the mean average precision (mAP@0.5).
- IoU average: 0.91
- Rate of False Detection: < 3%
- Processing Speed: approximately 28 frames per second (GPU-enabled system)

By adjusting learning rate, batch size, anchor scaling, and training epochs, the factorial Design of Experiments (DOE) optimization enhanced detection stability. High accuracy was maintained while computational complexity was decreased by the SqueezeNet backbone.

Preprocessing and regex-based validation specific to Indian RTO formats were integrated into the OCR module.

Accuracy of Raw OCR: 83%

89% after preprocessing

94.6% after validation and error correction

Character misclassification and formatting inconsistencies were decreased by roughly 15–20% using state code dictionary validation and Levenshtein-distance correction.

4.3 b. Results of Fraud Detection

In order to identify duplicate or fraudulent records, the Random Forest classifier examined structured vehicle metadata.

95.4% accuracy

96% accuracy

93.8% recall

F1-Score: 94.9%

4.4 Results of the Resale Value Prediction

The regression model used brand, age, mileage, fuel type, and compliance status to estimate the vehicle's

resale value.

Score for R2: 0.89

Minimal Average Absolute Error

Price and age/mileage have a strong negative correlation.

4.5 Prediction of Pollution Compliance

Based on regional standards, fuel type, and vehicle age, the Decision Tree classifier forecasted emission compliance. 93.2% accuracy Sensitivity and specificity in balance

4.6 Discussion

The results demonstrate that methodical parameter optimization greatly improves overall system reliability and detection accuracy. OCR performance and downstream analytics modules are directly improved by better plate localization. The system showed effective methods against issues associated to Indian transport and vehicle management, which includes inconsistent formatting, multilingual scripts, font variability, and a range of lighting conditions.

Conclusion

An RTO intelligent system which is reliable and optimized for Indian RTO systems especially for RTO officers and scalable model for vehicle environments that is suggested by this study. The framework includes real time analytics and fraud detection of vehicle users, resale value prediction, and pollution compliance tracking with real-time license plate detection and structured OCR extraction. Experimental evaluation showed strong fraud classification performance (~95%), high detection accuracy (~97%), OCR reliability (~95%), reliable pollution compliance prediction (~93%), and resale value forecasting ($R^2 \approx 0.89$).

Accurate, scalable, computationally efficient, and appropriate for real-time deployment is the suggested Intelligent RTO Monitoring and Forecasting Platform.

Acknowledgement

Authors of the paper are sincerely thankful to our guide Sunayana Sutar for her ongoing advice, helpful support, and technical assistance during this research work and project. We also thank the Dr. D. Y. Patil Institute of Management Research, Akurdi College Department of Artificial Intelligence and Data Science for providing the lab space and scholarly materials required to finish this work.

References

- [1] Kumar, R., Singh, P., & Zhao, L. (2025). Optimizing battery health monitoring in electric vehicles using interpretable CART–GX model. *Results in Engineering*. Retrieved from <https://www.sciencedirect.com/science/article/pii/S2590123025021152>
- [2] Sharma, A., Verma, S., & Lee, J. (2024). CNN-LSTM Hybrid Model for Enhanced Malware Analysis and Detection. *Procedia Computer Science*. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1877050924005982>
- [3] Brown, T., Wilson, M., & Carter, D. (2024). ACCESS: Assurance Case Centric Engineering of Safety–critical Systems. *Journal of Systems and Software*. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0164121224000773>
- [4] Li, Q., Chen, Y., & Ahmed, S. (2024). Carbon reduction, pollution mitigation, greening, and growth. *Journal of Environmental Management*. Retrieved from <https://www.sciencedirect.com/science/article/abs/pii/S030147972402485X>
- [5] Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767. Retrieved from <https://arxiv.org/abs/1804.02767>
- [6] Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature Pyramid Networks for Object Detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Retrieved from <https://arxiv.org/abs/1612.03144>
- [7] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Retrieved from <https://arxiv.org/abs/1512.03385>
- [8] Iandola, F. N., Han, S., Moskewicz, M. W., Ashraf, K., Dally, W. J., & Keutzer, K. (2016). SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size. arXiv preprint arXiv:1602.07360. Retrieved from <https://arxiv.org/abs/1602.07360>
- [9] Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934. Retrieved from <https://arxiv.org/abs/2004.10934>
- [10] Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv preprint arXiv:1704.04861. Retrieved from <https://arxiv.org/abs/1704.04861>
- [11] Zherzdev, S., & Gruzdev, A. (2018). LPRNet: License Plate Recognition via Deep Neural Networks. arXiv preprint arXiv:1806.10447. Retrieved from <https://arxiv.org/abs/1806.10447>
- [12] Silva, S. M., & Jung, C. R. (2018). Real-time Brazilian License Plate Detection and Recognition Using Deep Convolutional Neural Networks. *Proceedings of the 31st Conference on Graphics, Patterns and Images (SIBGRAPI)*. Retrieved from <https://ieeexplore.ieee.org/document/8097294>
- [13] Singh, E. I., & others. (2024). Recognition System: Detection of License Plate. *IEEE Conference Publication*. Retrieved from <https://ieeexplore.ieee.org/document/10673076>
- [14] Laroca, R., Severo, E., Zanlorensi, L. A., Oliveira, L. S., Gonçalves, G. R., Schwartz, W. R., & Menotti, D. (2018). A Robust Real-Time Automatic License Plate Recognition Based on the YOLO Detector. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. Retrieved from <https://ieeexplore.ieee.org/document/8489629>
- [15] Zhang, Y., Wang, H., & Li, X. (2022). License Plate Detection Using Convolutional Neural Network Back to the Basic with Design of Experiments. *IEEE Access*. Retrieved from <https://ieeexplore.ieee.org/document/9721015>