MS-Platenet: Design and Performance Evaluation of an AI-Based License Plate Recognition System for Automated Fuel Access Simulation

Muhammad Khairul Afiq Abdullah., *Norfadzlia Mohd Yusof

Fakulti Kecerdasan Buatan dan Keselamatan Siber, Universiti Teknikal Malaysia Melaka, Jalan Hang Tuah Jaya, 76100 Durian Tunggal, Melaka

DOI: https://doi.org/10.47772/IJRISS.2026.101100214

Received: 13 November 2026; Accepted: 18 November 2026; Published: 24 November 2026

ABSTRACT

Automated number plate recognition (ANPR) can support access-control decisions based on a vehicle's country of registration. This study presents MS-PlateNet, a web-based prototype that combines YOLOv8 licence plate detection, optical character recognition (OCR), and rule-based classification of Malaysian and Singaporean plates. Three YOLOv8 variants were evaluated across four dataset configurations derived from 1,351 original images. YOLOv8 Medium trained on the augmentation-only configuration achieved the highest mean average precision at an intersection-over-union threshold of 0.50 (mAP50), with a score of 0.8892. Among four OCR engines evaluated on 100 Malaysian and Singaporean licence plate crops, Fast Plate OCR achieved the highest exact-match accuracy (86.0%) and mean character similarity (96.1%). In a separate string-level benchmark, the rule-based classifier correctly classified 95 of 100 labelled plate strings (95.0%). The prototype used the predicted country labels to control a simulated fuel pump, which responded as expected in all five functional test conditions. The results demonstrate the feasibility of the simulated workflow. End-to-end country classification accuracy and processing latency on images and video remain to be evaluated.

Keywords: Automated number plate recognition; licence plate detection; optical character recognition; registration-country classification; YOLOv8; access-control simulation.

INTRODUCTION

ANPR is used in applications such as parking management, toll collection, and vehicle access control. A typical ANPR pipeline locates a licence plate in an image and converts the visible characters into machine-readable text. In settings where vehicles from different countries use the same facilities, the system may also need to infer the country of registration from the recognised plate text. This additional step introduces a separate source of error because a correct plate detection can still lead to an incorrect country prediction if the OCR engine misreads a character or if the recognised format is ambiguous.

This study focuses on Malaysian and Singaporean licence plates. Many standard Malaysian registration numbers use a letter prefix followed by digits, with a trailing letter in some series. A common Singaporean private-vehicle format uses a prefix beginning with S, followed by digits and a suffix letter. These structures provide cues for rule-based classification. Their usefulness depends on the range of plate series covered by the rules and the quality of the recognised text. Glare, blur, oblique viewing angles, and partial occlusion can affect detection and recognition, while a missing prefix or suffix character can change which rule matches.

MS-PlateNet integrates these stages into a prototype that accepts webcam, image, or video input. The application uses the predicted registration country as an input to a simulated RON95/RON97 fuel selection workflow. The simulation provides a concrete setting for examining how a recognition output affects an access decision, including cases where the system returns an Unknown result.

The study has three objectives:

1. Compare YOLOv8 Nano, Small, and Medium across four dataset preparation configurations for license plate detection.

2. Benchmark four OCR engines on Malaysian and Singaporean license plate crops.

3. Implement rule-based registration-country classification and verify the simulated pump responses to the resulting labels.

The study covers Malaysian and Singaporean plate recognition and software-based pump simulation. Its scope excludes driver identification, vehicle-type recognition, and physical integration with fuel-dispensing equipment.

Related Work

Deep Learning for License Plate Detection

Automatic licence plate recognition (ALPR) typically involves locating a licence plate in an image and recognising the characters within the detected region. The quality of the detected crop affects recognition because an imprecise bounding box can omit characters or include background features. Studies have therefore evaluated both detector performance and the conditions under which the resulting plate crops remain usable.

A comparison of five YOLOv8 variants on 2,432 images of Qatari licence plates identified YOLOv8s as offering a favourable balance between accuracy and inference time. The model achieved a mean average precision of 99.4% at an intersection-over-union threshold of 0.50 (mAP50) and 77.7% across thresholds from 0.50 to 0.95 (mAP50-95), with a detection time of 6.2 ms per image (Al-Hasan et al., 2024). A separate comparison of YOLOv8, Faster R-CNN, and SSD reported stronger detection performance for YOLOv8 on a small dataset and improved video detection after training with a larger dataset that covered more variation in lighting, background, and vehicle pose (Shyaa & Hashim, 2024). In a multi-dataset evaluation, YOLOv8n achieved mAP50 values of 99.3% on CCPD, 99.0% on UFPR-ALPR, and 91.0% on RodoSol-ALPR, with reported detection throughput of 30–36 frames per second (Satya et al., 2025). The variation across these datasets shows how strongly detector performance depends on the image source and capture conditions.

Studies in Indonesia extend this evidence to regional vehicle and plate formats. A YOLOv11 system trained on images of both motorcycles and cars achieved 94.2% mAP50, 97.7% precision, and 86.7% recall on a 71-image test set (Martadinata et al., 2025) . In a university parking study, a YOLOv8-based system demonstrated plate reading on moving vehicles and black licence plates. Its confusion matrix also showed that some licence plates were classified as background (Al Matiin & Arifin, 2024). These findings make recall particularly relevant to live monitoring, where a missed plate prevents recognition and any subsequent decision.

Character Recognition and Image Enhancement

YOLO-based ALPR systems commonly use a separate optical character recognition (OCR) stage to convert detected plate crops into text. One YOLOv8–EasyOCR approach used k-means clustering, thresholding, and morphological opening to enhance the plate crop and subsequently extracted a country identifier from the recognised text (Moussaoui et al., 2024). The Indonesian YOLOv11–EasyOCR study reported a character recognition accuracy of 91.0% (Martadinata et al., 2025). Character recognition accuracy and complete plate recognition measure different outcomes, since an error in a single character can make the full registration number incorrect.

The distinction is evident in an Indonesian study that paired YOLOv8 with PaddleOCR. The detector achieved 99.4% mAP50 and successfully detected and cropped all 106 images used for the reported recognition evaluation. PaddleOCR correctly recognised 42 of those plates, giving a complete plate recognition rate of 39.62% (Syihabuddin et al., 2025). A Malaysian YOLOv8–PaddleOCR study reported an overall recognition accuracy of 98.64% across 74 images covering indoor, outdoor, low-light, and rainy conditions (Lim Zheng You et al., 2025). The studies used different image collections and evaluation settings, which limits direct comparison of

their results. Their findings show why an ALPR evaluation should report the proportion of complete plates read correctly alongside detector metrics.

Image enhancement can help when it addresses a specific source of recognition error. On CCPD2019 and CRPD-single, adding a spatial transformer network to LPRNet to correct tilted plate crops increased recognition accuracy from 83.07% to 91.45% and from 88.62% to 92.00%, respectively (Wang et al., 2025) . Another study reported that image enhancement improved OCR performance under low light, although recognition rates still varied across datasets. Its pipeline achieved 83.1% on UFPR-ALPR and 71.7% on RodoSol-ALPR, where Rosetta achieved 89.0% (Satya et al., 2025). In the Qatari study, several preprocessing methods yielded no improvement over the baseline mAP50 of 0.87 in the reported experiment, while training with images that simulated challenging conditions increased it to 0.92 (Al-Hasan et al., 2024). These results support evaluating enhancement at the stage where the error occurs, such as augmenting training images for variable capture conditions or correcting the geometry of a plate crop before OCR.

Generalisation, Deployment, and Research Gap

Cross-dataset testing illustrates the difficulty of transferring ALPR models between settings. When an improved YOLOv5s–LPRNet system was trained on CCPD2019 and evaluated on CRPD-single, recognition accuracy decreased from 91.45% to 89.27%. In the reverse direction, accuracy decreased from 92.00% to 85.31% (Lim Zheng You et al., 2025). The unequal decreases suggest that the range of plate appearances and capture conditions represented in the training data affects performance on unseen images. The regional studies likewise show the importance of evaluating the vehicle types, plate formats, and environmental conditions relevant to the intended deployment (Al Matiin & Arifin, 2024; Lim Zheng You et al., 2025; Martadinata et al., 2025; Syihabuddin et al., 2025).

Deployment studies also measured performance at different points in the pipeline. The Qatari system reported approximately 166 ms from image capture through server-side recognition, including 100 ms for transmission (Al-Hasan et al., 2024). The Indonesian YOLOv11–EasyOCR web system reported an average processing time below eight seconds per image on its virtual private server (Martadinata et al., 2025). These timings describe different deployments and processing paths. For a live monitoring application, the operationally relevant interval extends from image capture to the decision presented by the system.

The nine reviewed studies cover plate detection, character recognition, environmental robustness, and deployment. One study also extracted a country identifier from recognised text (Moussaoui et al., 2024), while the Malaysian study evaluated recognition across several local capture conditions (Lim Zheng You et al., 2025). The studies reviewed leave open how plate recognition and country classification affect the correctness of a downstream fuel-access decision. MS-PlateNet addresses this application gap by integrating plate detection and recognition with classification into Malaysia, Singapore, or Unknown and using the classification to control fuel options in a pump simulation. A correctly located plate can still lead to an incorrect simulated decision if its text is misread or its country is misclassified. The study evaluates detection on images, OCR on plate crops, and country classification on labelled plate strings. The functional tests verify how the simulated pump responds to the assigned labels.

METHODOLOGY

System Architecture and Processing Pipeline

Figure 1 shows the architecture and processing pipeline of MS-PlateNet. The React frontend, built with Vite, accepts uploaded images and videos and supports live camera input. The frontend communicates with the FastAPI backend through a REST API and transmits live camera frames via WebSockets. The backend uses YOLOv8 to locate licence plates and return their bounding boxes and detection confidence scores. It crops the plate region with the highest detection confidence and passes the crop to Fast Plate OCR for transcription. The recognised text is normalised and classified as Malaysia, Singapore, or Unknown using the rules described in Section 3.5. The assigned label drives the software-based fuel access simulation described in Section 3.6.

PostgreSQL stores detection records and recording metadata for display in the Dashboard and Recording Archive.

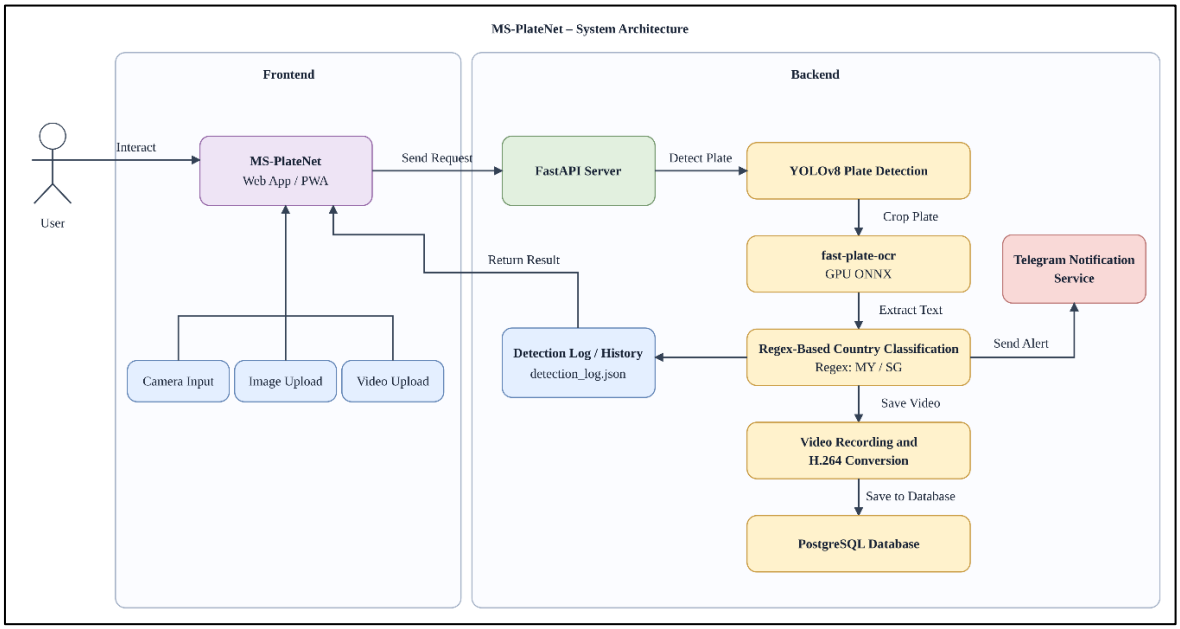


Figure 1. System architecture and processing pipeline of MS-PlateNet.

Dataset and preparation

The BaoVisionPlate dataset comprised 1,351 original vehicle images featuring Malaysian and Singaporean licence plates captured at varying distances, viewing angles, and lighting conditions. The plates were manually annotated in Roboflow using bounding boxes under a single class, licence plate. The original images were divided into training, validation, and test sets before augmentation. Only the training images were augmented using random rotation (±15°), brightness adjustment (±25%), Gaussian blur (up to 1.6 pixels), and pixel noise (up to 1.81% of pixels). The final dataset contained 2,297 images, comprising 1,892 training images, 270 unaugmented validation images, and 135 unaugmented test images. The dataset was exported in YOLOv8 PyTorch format.

Licence Plate Detection

The YOLOv8 Nano, Small, and Medium variants were initialised with COCO-pretrained weights and trained using the training split of the final BaoVisionPlate dataset. All three variants used an input resolution of 640 × 640 pixels, a batch size of 16, and a maximum of 100 epochs, with early stopping after 20 epochs without improvement. Training used automatic mixed precision on an NVIDIA RTX 4070 GPU. Detection performance was assessed using precision, recall, mAP50, and mAP50–95. The study used mAP50 as the primary criterion for selecting the detector to integrate with the OCR stage.

Optical Character Recognition

Four OCR engines, Fast Plate OCR, EasyOCR, PaddleOCR, and RapidOCR, were evaluated on the same set of 100 real-world Malaysian and Singaporean licence plate crops. The samples varied in character font and spacing and included examples with glare and slight motion blur. Each engine's output was compared with the reference plate transcription. Exact-match accuracy was calculated as the percentage of crops for which the complete predicted string matched the reference. Mean character similarity was calculated by averaging the similarity scores between predicted and reference strings across the 100 crops. Fast Plate OCR was selected for integration into MS-PlateNet based on its performance on these measures. Because the benchmark used plate crops, the reported scores measure OCR performance on cropped inputs.

Rule-Based Registration-Country Classification

MS-PlateNet infers a vehicle's registration country from the plate string returned by Fast Plate OCR. Before classification, the system converts the string to uppercase and removes non-alphanumeric characters. It then applies the regular-expression rules in Table 1 to assign a label of Singapore, Malaysia, or Unknown.

Table 1. Regular-expression rules and matching precedence for licence plate registration-country classification in MS-PlateNet.

Assigned label	Regular expression	Format matched
Singapore	^S[A-Z]{1,3}\d{1,4}[A-Z]$	An initial S, followed by one to three letters, one to four digits, and one final letter
Malaysia	^[A-Z]{1,3}\d{1,4}[A-Z]{0,2}$	One to three letters, followed by one to four digits and up to two trailing letters
Unknown	Neither expression matches	A string outside the supported formats

The Singapore rule takes precedence when a string matches both expressions. This overlap can produce incorrect country labels for plates with similar formats. The assigned label is passed to the fuel access simulator described in Section 3.6. The classifier was evaluated on 100 labelled plate strings comprising 50 Malaysian, 35 Singaporean, and 15 samples grouped as Unknown, foreign, or corrupted. Predicted labels were compared with the reference labels to calculate overall accuracy and the proportion of correctly classified samples in each category. The evaluation used plate strings as inputs, so the reported results describe the classification stage. Performance on full images and video also depends on plate detection and OCR.

Fuel Access Simulation and Functional Testing

MS-PlateNet uses the registration-country label produced by the licence plate recognition pipeline to trigger a software-only fuel access decision. The Pump Simulation module changes the availability of the simulated fuel options, while the Dashboard and Recording Archive display detection events and associated recordings. The prototype operates independently of physical dispensers, payment systems, and fuel transaction records. When the classifier assigns a Singapore label, the simulator enters the ARMED state, locks RON95, keeps RON97 available, and starts video recording. Selecting RON97 while the pump is armed leaves the RON95 restriction and recording active. When the classifier assigns a Malaysia or Unknown label, both fuel options remain available. After approximately five seconds without a plate detection, the simulator stops recording and restores both fuel options. Video recordings are encoded as H.264 MP4 files, and their metadata are stored in PostgreSQL for retrieval through the Recording Archive. Functional testing covered five scenarios: Malaysia, Singapore, and Unknown classification outputs, RON97 selection while the pump was armed, and reset after plate detection ceased. For each scenario, the observed fuel-option availability, pump state, and recording behaviour were compared with the expected response. These tests assessed whether the simulator correctly applied the rules to the assigned label. The appropriateness of the resulting access decision also depends on whether the pipeline correctly identifies the vehicle's registration country.

RESULTS

License plate detection

Table 2 compares the licence plate detection performance of the YOLOv8 Nano, Small, and Medium variants. All three models were trained on the same training split of the final BaoVisionPlate dataset and achieved mAP50 scores ranging from 0.8830 to 0.8892, exceeding the study's target of 0.85. YOLOv8 Medium achieved the highest mAP50 (0.8892), followed by YOLOv8 Small (0.8866) and YOLOv8 Nano (0.8830). YOLOv8 Nano achieved the highest precision (0.9069) and the lowest recall (0.8257), indicating that a greater proportion of its

detections were correct, while it detected a smaller proportion of the ground-truth plates. YOLOv8 Small achieved the highest recall (0.8606) and mAP50–95 (0.5970), giving it the best performance across the range of intersection-over-union (IoU) thresholds used to calculate mAP50–95. YOLOv8 Medium was selected for the MS-PlateNet prototype because mAP50 was the study's primary detection metric. Its mAP50 exceeded that of YOLOv8 Small by 0.26 percentage points, indicating a narrow lead on the metric used for model selection.

Table 2. License plate detection performance of YOLOv8 Nano, Small, and Medium trained on the final BaoVisionPlate dataset.

Model	Precision	Recall	mAP50	mAP50–95
YOLOv8 Nano	**0.9069**	0.8257	0.8830	0.5946
YOLOv8 Small	0.8772	**0.8606**	0.8866	**0.5970**
YOLOv8 Medium	0.8816	0.8552	**0.8892**	0.5952

Optical Character Recognition

Table 3 compares four optical character recognition (OCR) engines evaluated on 100 cropped images of Malaysian and Singaporean licence plates. Exact-match accuracy measures the proportion of plates for which the entire character sequence was transcribed correctly. Mean character similarity measures how closely the transcribed text matched the ground truth, including cases with partial matches.

Fast Plate OCR achieved the highest exact-match accuracy (86.0%) and mean character similarity (96.1%). EasyOCR ranked second on both measures, with an exact-match accuracy of 61.0% and a mean character similarity of 87.8%. Fast Plate OCR therefore led the next-best engine by 25 percentage points in exact-match accuracy and 8.3 percentage points in mean character similarity. These results support the use of Fast Plate OCR in the MS-PlateNet prototype.

The effect of an OCR error on downstream country classification depends on the character involved and its position in the plate sequence. Substituting one digit for another may preserve the format required by the regular expression, allowing the plate to be classified correctly despite the transcription error. Omitting a required prefix or suffix character may cause a different rule to match or lead the system to return an Unknown result. A separate evaluation of country classification is therefore needed to determine how often OCR errors change the assigned country.

Table 3. OCR performance on 100 Malaysian and Singaporean license plate crops.

OCR engine	Exact matches (out of 100)	Mean character similarity
Fast Plate OCR	**86**	**96.1%**
EasyOCR	61	87.8%
RapidOCR	59	80.7%
PaddleOCR	57	74.4%

Rule-based Registration-Country Classification

As shown in Table 4, the classifier correctly labelled 95 of 100 benchmark strings, yielding an overall accuracy of 95.0%. It correctly classified 48 of 50 Malaysian strings (96.0%), 34 of 35 Singaporean strings (97.1%), and 13 of 15 Unknown / Foreign strings (86.7%). These results measure classification performance on labelled plate strings. End-to-end performance on images or video also depends on plate detection and OCR.

The five errors illustrate two limitations of the rules. First, formats can overlap across categories. The Sabah plate SAA1234K matched the Singaporean pattern, while the Singaporean commercial plate GB1234A matched the Malaysian pattern. Two foreign strings, XY1234Z and A1234B, also matched accepted patterns. Second, valid registrations can fall outside the rules. The Malaysian commemorative plate PUTRAJAYA1234 exceeded the three-letter prefix limit. These cases show how a format-based classification error can carry through to the simulated pump decision.

Table 4. Quantitative Benchmark Results of the Rule-Based Country Classifier (100 Labelled Samples).

Category	Tested	Correct	Misclassified	Accuracy (%)	Primary Error Analysis
Malaysia (JPJ)	50	48	2	96.0	The Sabah plate SAA1234K matched the Singaporean pattern. The commemorative plate PUTRA-JAYA1234 exceeded the prefix character limit.
Singapore (LTA)	35	34	1	97.1	The commercial plate GB1234A, which has a non-S prefix, matched the Malaysian pattern.
Unknown / Foreign	15	13	2	86.7	The foreign alphanumeric strings XY1234Z and A1234B matched standard prefix–digit patterns.
Overall	**100**	**95**	**5**	**95.0**	

Simulated Fuel Access Behaviour

Table 5 summarises the pump simulation responses observed during functional testing. When a plate was classified as Malaysia or Unknown, both RON95 and RON97 remained available. A Singapore classification locked RON95, kept RON97 available, and initiated video recording.

Following a Singapore classification, selecting RON97 did not release the RON95 lock or stop the recording. Both the restriction and recording remained active until the departure condition was met. After approximately five seconds without a plate detection, the simulation stopped recording and restored access to both fuel options.

The observed responses matched the expected behaviour in all five test conditions. The functional tests verified that the simulation applied the prescribed fuel access rules and reset correctly after the departure condition. In operation, the simulated pump responds to the classifier's assigned label, so the appropriateness of the fuel access decision depends on the accuracy of that label. The string-level country classification results are reported in Table 4.

Table 5. Simulated pump responses to license plate classification and subsequent events.

Test condition	Simulated response
Plate classified as Malaysia	RON95 and RON97 remained available.
Plate classified as Singapore	RON95 was locked, RON97 remained available, and recording began.
RON97 selected following a Singapore prediction	RON95 remained locked, and recording continued until the departure condition was met.
Plate classified as Unknown	RON95 and RON97 remained available.

Plate no longer detected following a Singapore prediction	After approximately five seconds without a plate detection, recording stopped and both fuel options became available again.

Pump Simulation and Recording Archive

Figure 3 shows the simulated pump response to a Singapore classification, one of the functional test conditions reported in Table 5. The interface highlights the detected licence plate in the live camera feed and displays the recognised plate number alongside the pump status and fuel options. The pump is in the ARMED state, with RON95 locked and RON97 available for selection.

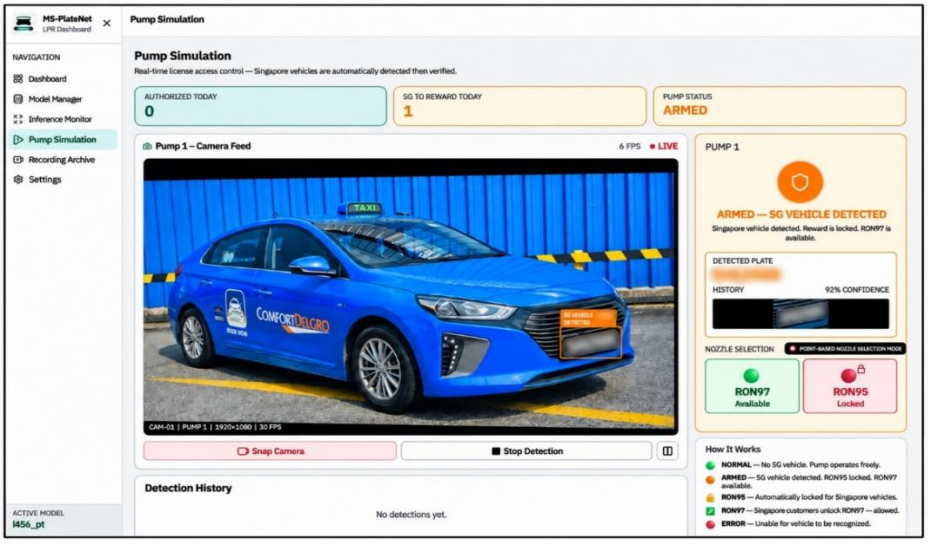


Figure 2. Pump Simulation Interface After Singaporean Vehicle Detected.

The Recording Archive complements the live pump view by allowing users to retrieve and review saved event recordings, as shown in Figure 4. Each recording is presented with a video thumbnail, the recognised licence plate number, the pump ID, the recording date and time, and the clip duration. Users can search by plate number, filter recordings by pump or date, and play or download individual clips.

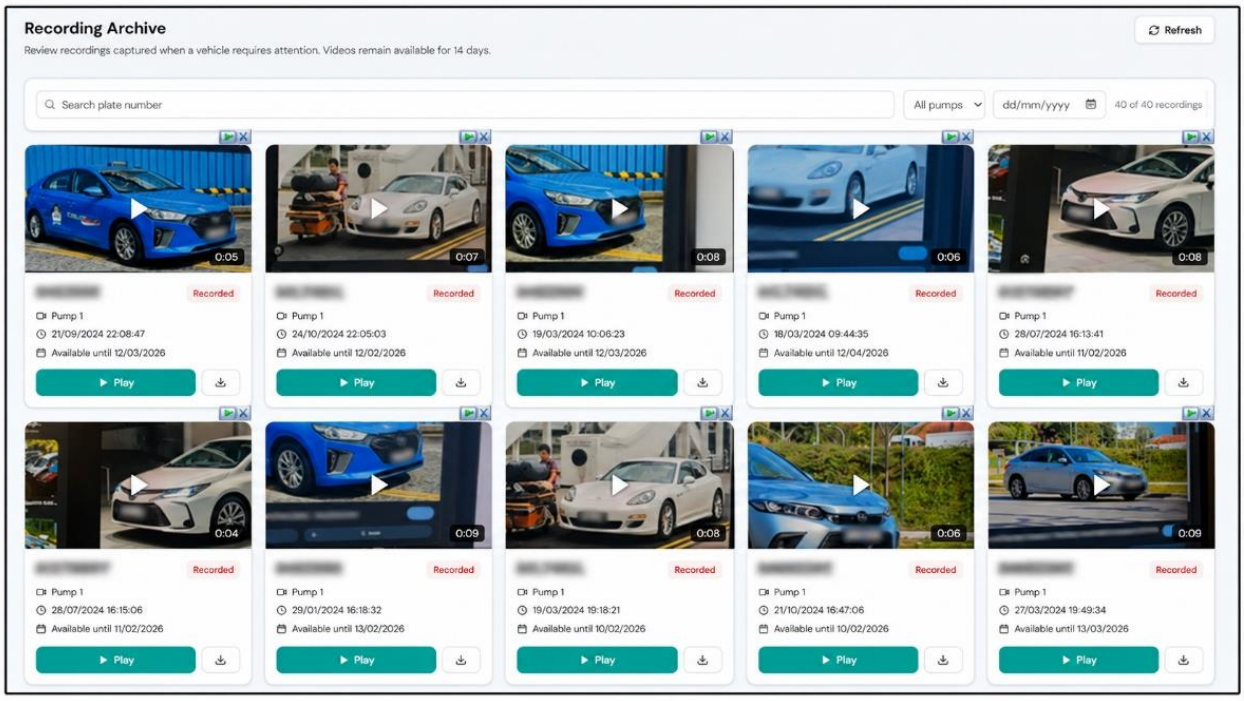


Figure 3. Recording Archive Interface.

DISCUSSION

Under the reported experimental setup, the augmentation-only configuration achieved the highest mAP50. The training images were augmented with rotation, brightness adjustment, blur, and noise to reflect variations in camera angle, lighting, image sharpness, and sensor noise. Because these transformations were evaluated together, their individual effects on detection performance could not be determined. Whether the higher mAP50 indicates better generalisation depends on two conditions: all configurations must be evaluated on the same held-out, unaugmented images, and images of the same plate or vehicle must be kept within a single split.

The OCR benchmark favoured the plate-specific engine on the 100 evaluated crops. The results also show why exact-match OCR and country classification should be reported separately. Some transcription errors preserve the plate's country-format cues, while others change the predicted label or produce an Unknown result. Evaluating the regex rules on ground-truth transcriptions and then on OCR output would help identify how much error arises from the rule coverage and how much comes from recognition. A further evaluation on full images would capture failures introduced by missed detections and inaccurate crops.

The simulated access decision makes the consequences of an Unknown prediction visible. Leaving both fuel options available preserves access when a Malaysian plate is unreadable. The same response also leaves RON95 available when a Singaporean plate is unreadable. Similarly, the five-second reset uses the absence of a detection as a proxy for departure. A temporary occlusion or a series of missed detections can trigger the same condition. These behaviours should be evaluated with labelled video sequences that include vehicle entry, temporary occlusion, and departure.

The OCR benchmark used plate crops, and the application tests focused on pump behaviour. End-to-end classification accuracy on original images, performance under continuous traffic, and end-to-end processing latency remain open evaluation questions. The deployment setting would also require explicit procedures for handling registration numbers, recorded video, and access to logs.

CONCLUSION AND FUTURE WORK

This study developed MS-PlateNet, a web-based prototype that integrates licence plate detection, OCR, and rule-based registration-country classification with a simulated fuel-access workflow. YOLOv8 Medium trained on the augmentation-only configuration achieved the highest mAP50 of 0.8892. Fast Plate OCR achieved 86.0% exact-match accuracy and 96.1% mean character similarity on 100 plate crops. The rule-based classifier correctly labelled 95 of 100 plate strings (95.0%) in the string-level benchmark. The simulated pump responded as expected in all five functional test conditions, covering fuel-option availability, recording behaviour, and the reset after plate detection ceased.

Future work should evaluate end-to-end country classification and simulated pump decisions using held-out images and video with known registration-country labels. The evaluation should report missed detections, per-country recall, misclassifications, and Unknown outcomes. Tests involving glare, low light, motion blur, and temporary occlusion would help determine when recognition errors change the simulated pump response. Measuring latency from frame capture to the simulated response would establish how quickly the integrated system makes a decision under these conditions.

Ethical Considerations

The study evaluated licence plate recognition using vehicle imagery and a software-based fuel access simulation. It involved no recruitment of human participants or intervention with drivers. Because licence plate numbers can identify vehicles, they were blurred in the figures presented in this article, including on the vehicles, in the recognised-text fields, and in the recording archive. The prototype inferred registration country from plate format. It did not verify vehicle ownership, driver identity, or fuel eligibility. Its outputs were used to control simulated fuel options only.

Conflict of Interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

ACKNOWLEDGEMENT

The authors would like to thank Centre for Research and Innovation Management (CRIM), Universiti Teknikal Malaysia Melaka (UTeM) for supporting this work.

REFERENCES

1. Al Matiin, M. A. H., & Arifin, F. (2024). Evaluation of YOLOv8 Algorithm for Vehicle License Plate Detection System in UNY Integrated Parking Lot. Elinvo (Electronics, Informatics, and Vocational Education), 9(2), 220–229. https://doi.org/10.21831/elinvo.v9i2.70032
2. Al-Hasan, T. M., Bonnefille, V., & Bensaali, F. (2024). Enhanced YOLOv8-Based System for Automatic Number Plate Recognition. Technologies, 12(9). https://doi.org/10.3390/technologies12090164
3. Lim Zheng You, Bong Zi Jun, Pang Ying Han, Ooi Shih Yin, Khoh Wee How, & Hiew Fu San. (2025). INTEGRATION OF YOLOV8 AND PADDLEOCR FOR CAR PLATE RECOGNITION SYSTEM IN DIVERSE ENVIRONMENTS: MALAYSIA. Journal of Engineering Science and Technology, 20(1), 250–270.
4. Martadinata, F., Firdaus, A., & Amal, M. R. T. (2025). Detection and Identification of Vehicle License Plates in Indonesia Transportation System Based on Deep Learning Using YOLOv11 and Easyocr. Jurnal Sisfokom (Sistem Informasi Dan Komputer), 14(4), 573–578. https://doi.org/10.32736/sisfokom.v14i4.2524
5. Moussaoui, H., Akkad, N. El, Benslimane, M., El-Shafai, W., Baihan, A., Hewage, C., & Rathore, R. S. (2024). Enhancing automated vehicle identification by integrating YOLO v8 and OCR techniques for high-precision license plate detection and recognition. Scientific Reports, 14(1). https://doi.org/10.1038/s41598-024-65272-1
6. Satya, B., Manongga, D., Hendry, & Aminuddin, A. (2025). Optimized YOLOv8 for Automatic License Plate Recognition on Resource Constrained Devices. Engineering, Technology and Applied Science Research, 15(2), 21976–21981. https://doi.org/10.48084/etasr.9983
7. Shyaa, T. A., & Hashim, A. A. (2024). Superior Use of YOLOv8 to Enhance Car License Plates Detection Speed and Accuracy. Revue d'Intelligence Artificielle, 38(1), 139–145. https://doi.org/10.18280/ria.380114
8. Syihabuddin, E. R., Angela, W., Naufal, M., & Pramunendar, R. A. (2025). Pengenalan Plat Nomor Kenderaan dengan YOLOv8 dan PaddleOCR. Journal of Computer Science and Technology JCS-TECH, 5(2), 66–72.
9. Wang, X., Wang, M., Guo, H., Li, J., Wang, X., & Zhang, Y. (2025). License plate recognition system for complex scenarios based on improved YOLOv5s and LPRNet. Scientific Reports, 15(1). https://doi.org/10.1038/s41598-025-18311-4