iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: https://doi.org/10.22214/ijraset.2026.85324

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue X Oct 2026- Available at www.ijraset.com

AI-Based Intelligent HSRP Compliance Screening and Decision Support System Using Computer Vision

Madhur Madankar[1], Karan Shahu[2], Nikhil Netam[3], Divyansh Verma[4], Yash Karme[5], Vedant Mankar[6], Sumalatha Bandari[7]
Department of Computer Science and Engineering, G H Raisoni University, Amravati, India

Abstract: ***HSRPs are designed to make it uniform and to facilitate quick identification of vehicles but the surveillance of vast amounts of traffic manually is expensive and unfeasible. In this report, an AI-powered HSRP compliance screening and decision-support system is proposed that focuses on traffic-video screening and is based on computer vision. The infrastructure uses a YOLO-based traffic video analysis module for vehicle and HSRP detection, and Byte Track for vehicle tracking to associate vehicles across frames. Vehicle detections are later improved through best-frame selection, and multiple observations are filtered to accumulate evidence on each vehicle before OCR recognition. PaddleOCR is selected as the main OCR engine, and OCR results are fused across frames using temporal evidence while considering aspect ratio and registration-format constraints to reduce the impact of misclassification. The visible evidence of HSRPs, including IND-band evidence, plate quality, aspect ratio, and OCR confidence, is combined into a screening assessment. Tracking, OCR, HSRP screening, explainable prediction and reasoning, and human image-rejection and failure mode checks are combined within a streamlined vehicle-screening tool, distinguishing vehicles that are likely to be compliant from those that are potentially subject to further human review. All plate and video evidence, and screening information, reasoning, and failure modes are presented via an API-powered display available for download and human review. The entire system is validated through end-to-end traffic video analysis and automated testing, though it does not seek to achieve comprehensive legal authentication of HSRP security features.***

Keywords: HSRP, Automatic Number Plate Recognition, Computer Vision, YOLO, Byte Track, OCR, Intelligent Transportation System, Explainable AI, Decision Support System.

I. INTRODUCTION

Vehicle recognition has become an essential part of present-day traffic regulation, enforcement, toll collection, parking, access control and city security. Given the steady increase of vehicles on the streets, it is no surprise that human observation has become difficult, especially when multiple cars need to be supervised in a limited time frame in a traffic hotspot. Computer vision and AI could be a cost-effective approach to automate some of this traffic supervision.

Automatic Number Plate Recognition (ANPR) are widely used for license plate detection and registration recognition using optical character recognition (OCR). There is also indication that this technique may be negatively affected by the real-world conditions an ANPR system operates in, such as motion blur, low illumination, multiple camera angles, distance, partially blocked license plates and low resolution image. The registration plate that the system is attempting to detect may not be visible in one frame in the video but may become more visible as the vehicle approaches the camera, whereas if the system uses just the one frame recognition results could be unstable.

The primary purpose of the High-Security Registration Plates (HSRPs) is to introduce a standard, higher security registration system. HSRP has many attributes physically and visibly but they cannot be manually checked against each vehicle in traffic, and additionally, some of the security features cannot be reliably verified by passing traffic-camera footage. Consequently, it would be wiser to use an automated mechanism as a screening and decision support tool rather than a blanket substitute for official HSRP verification.

The proposed work presents an AI-Based Intelligent HSRP Compliance Screening and Decision Support System using Computer Vision. The system processes traffic video using a multi-stage computer vision pipeline. Vehicles are detected using a YOLO-based object detector and are tracked across consecutive frames using ByteTrack. Number plates are then detected and associated with their corresponding tracked vehicles.

©IJRASET: All Rights are Reserved

Multiple plate observations are collected over time, and plate-quality analysis with best-frame selection is used to retain clearer evidence for recognition. PaddleOCR is used as the primary OCR engine, while temporal OCR analysis combines recognition results from multiple frames to improve registration-number reliability.

After plate detection, the system proceeds to visible-feature-based HSRP screening. It considers plate quality, OCR confidence, time consistency, registration-format reasonableness, plate aspect ratio, and visible Indian Band (IND) evidence, to generate a screening outcome at one of two levels: a vehicle likely complies, or the vehicle needs manual verification. The outcome is then stored in an SQLite database and rendered in the FastAPI dashboard, where the human operator can examine Indian Band evidence, interpret screening factors, edit OCR output if needed, and perform manual verification. The dashboard provides interpretability by displaying "Why This Result?" explanation text listing evidence that supports the positive screening outcome, as well as any relevant warning signs.

The main research contribution of this work lies in bringing together vehicle detection, multiple object tracking, plate detection, best-frame selection, temporal OCR validation, visible HSRP screening, explainable decision support, database storage, and a human-in-the-loop validation under a single framework for traffic video analysis. Unlike traditional techniques which are focused solely on plate detection or OCR, this system is capable of aggregating evidence from consecutive video frames and providing a readable interpretation for human consideration. Also, it is aimed at assisting in screening and decision-making instead of automatic and judicial enforcement of HSRP compliance.

II. RELATED WORK

Several studies have explored automatic license plate recognition, vehicle tracking, optical character recognition, High-Security Registration Plate analysis, and explainable artificial intelligence for transportation applications.

Al-Batat et al. developed an end-to-end Automated License Plate Recognition system that combines vehicle detection, license plate detection, vehicle classification, and plate recognition within a single pipeline. Their work used YOLO-based models and emphasized that the reliability of an ANPR system depends on the performance of every processing stage rather than on OCR alone. The study also highlighted difficulties in maintaining recognition performance under complex real-world traffic conditions. This work is relevant to the proposed system because it supports the use of an integrated multi-stage ANPR architecture.

Adak et al. proposed an ANPR system using YOLOv3 and a Convolutional Neural Network for number plate detection and recognition. Their work considered practical challenges such as viewing angle, lighting variation, motion blur, and customized or multi-line license plates. The results showed that clear and directly visible plates were recognized more reliably, while recognition accuracy decreased under difficult environmental conditions. These findings are relevant to the present work because traffic-video frames frequently contain blurred or poorly angled plates, motivating the use of best-frame selection and multi-frame evidence.

Dhyani and Kumar presented a real-time license plate detection and recognition system using YOLOv7x and EasyOCR. The detector was trained on Indian vehicle license plates, while OCR was used to recognize the registration characters from the detected plate regions. Their work demonstrated that deep-learning-based detection combined with OCR can be effective for Indian license plate recognition. However, the authors also observed that the recognition performance depends strongly on image quality and plate visibility.

You et al. proposed a multi-object vehicle detection and tracking approach based on an improved YOLOv8 model and ByteTrack. Their work focused on maintaining vehicle identity across consecutive frames under challenging traffic conditions. The study addressed issues such as occlusion, missed detections, false detections, scale variations, and identity switching. This work directly supports the use of ByteTrack in the proposed system, where a persistent vehicle identity is required to collect multiple plate observations belonging to the same vehicle.

Sugiyono et al. studied vehicle registration plate information extraction using Tesseract OCR. Their work included plate localization, image processing, character segmentation, and OCR-based recognition. The study reported that OCR performance was affected by plate background, font variation, image quality, and plate format. This observation supports the use of image preprocessing, plate-quality analysis, and temporal OCR validation in the proposed system instead of relying on a single recognition result.

Gianey and Goyal investigated the localization and recognition of High-Security Registration Plates for Indian vehicles. Their study discussed HSRP-specific characteristics such as the chromium-based hologram, permanent identification number, and the blue IND marking. The work is important for the present study because it highlights both the visible HSRP characteristics that may be analyzed from images and the limitations of verifying certain physical security features using ordinary traffic-camera footage.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Badamikar et al. proposed an HSRP detection and classification approach using YOLOv8 and Federated Learning. Their work demonstrated that artificial intelligence techniques can be used specifically for HSRP-oriented detection and classification while also considering privacy-preserving distributed learning. However, such an approach requires an HSRP-specific dataset and a dedicated federated training infrastructure. In contrast, the proposed system focuses on visible-feature screening, temporal evidence, explainability, and human verification rather than complete HSRP authenticity classification.

Kuznietsov et al. presented a systematic review of Explainable Artificial Intelligence for safe and trustworthy autonomous driving. Their study emphasized that highly complex AI models can create trust and safety concerns when users cannot understand the reasons behind predictions. The authors highlighted the importance of explanation mechanisms in intelligent transportation systems. This directly motivates the explainability component of the proposed system, where the dashboard provides a "Why This Result?" explanation showing the evidence that supports or weakens each screening decision.

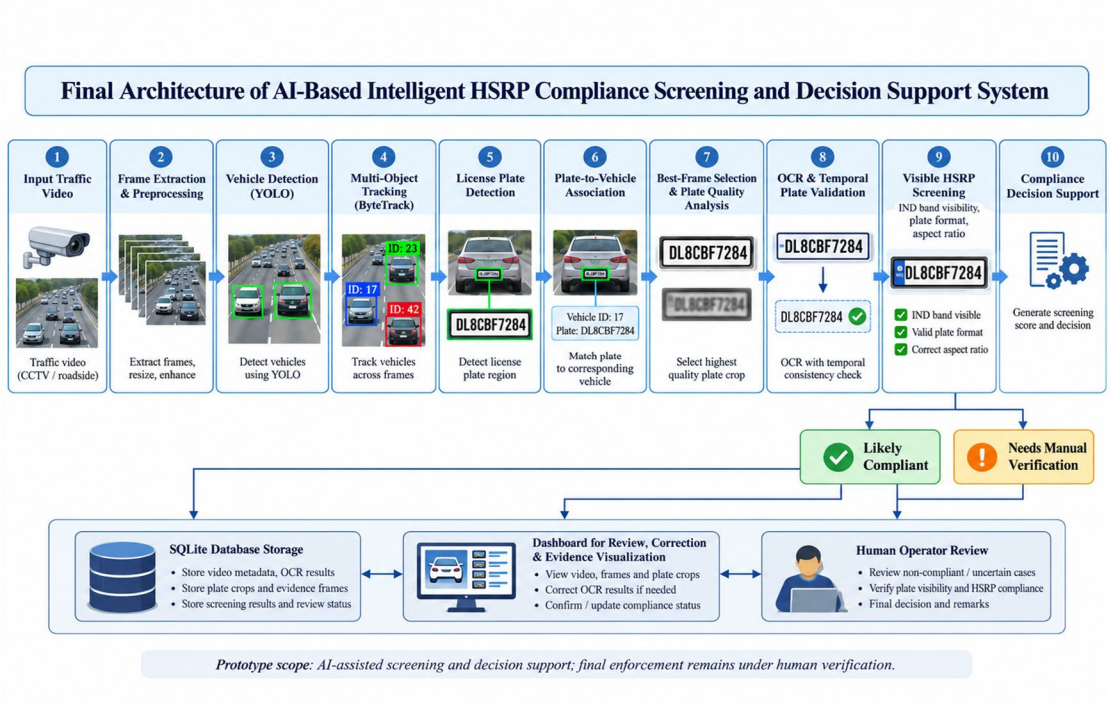
A. Research Gap

All the above problems have been successfully addressed and solved separately and independently in various literature related to vehicle detection, license plate recognition, optical character recognition (OCR), object tracking, HSRP detection, as well as explainable AI. However, most existing methods and solutions consider individual parts and single frame recognition; a single frame in traffic video may experience motion blurring, poor illumination, low plate resolution, occlusion, inappropriate camera angle, all of which will lead to reduced recognition accuracy; additionally, some physical security features on the HSRPs could not be reliably examined in normal traffic-video, and many automated systems do not have enough explanation or human feedback in cases of uncertain decisions.

The proposed work addresses these gaps by integrating vehicle tracking, best-frame selection, temporal OCR aggregation, visible HSRP screening, explainable decision support, database storage, and human-in-the-loop verification within a single system. Instead of making an automatic enforcement decision, the system classifies vehicles as Likely Compliant or Needs Manual Verification, thereby supporting safer and more transparent decision making.

III. PROPOSED METHODOLOGY

The proposed system follows a multi-stage computer vision and decision-support pipeline for analysing traffic video. The complete workflow is:



Photograph : Overall Architecture of the Proposed HSRP Compliance Screening and Decision Support System

A. Vehicle Detection and Tracking

An object detection framework based on YOLO is employed on every frame of the recorded traffic video to conduct detection. Only categories associated with the road-vehicle group will be considered; namely cars, motorcycles, buses and trucks. A bounding box, together with a confidence and a class label, is produced for each vehicle. After detection, Byte Track is used to maintain vehicle identity between frames, preventing vehicles from being treated as different objects when they reappear in a subsequent frame. The details of maintaining identity (i.e., assigning a unique tracking ID) to each vehicle is not included, but involves the system inferring that multiple detections over multiple frames refer to a single vehicle. Maintaining vehicle identity is important for the best frame selection and multi-frame OCR application.

B. License Plate Detection and Plate-to-Vehicle Association

A dedicated license plate detector is used to find registration plates associated with tracked vehicles. When a plate is found, it is matched to the corresponding vehicle based on the position relationship of the vehicle's bounding box and the plate's bounding box, so that the plate evidence collected by the vehicle in multiple frames is associated with the vehicle being tracked. The corresponding plate region is then cropped and saved for further processing.

C. Plate Quality Analysis and Best-Frame Selection

The quality of a plate image may vary significantly as the vehicle moves through the camera view. Therefore, the system does not immediately perform OCR on every detected plate. Instead, each plate observation is evaluated using factors such as detection confidence, plate size, sharpness, aspect ratio, and overall visual quality.

A generalized plate-quality score can be expressed as:

$$Q = w_1 C + w_2 W + w_3 H + w_4 A + w_5 R + w_6 S_h$$

where Q represents the overall plate-quality score, C represents plate detection confidence, W and H represent normalized plate width and height, A represents plate area, R represents the aspect-ratio score, S represents image sharpness, and w_1 to w_6 represent the corresponding weights.

The plate observation with stronger quality is retained as better evidence for OCR. This reduces the effect of blurred, distant, or poorly visible frames.

D. OCR and Temporal Plate Recognition

Before applying OCR to the selected plate crop, it may be preprocessed by rescaling, graying, sharpening, contrast-enhancement, thresholding, etc. The main OCR engine used by our system is PaddleOCR, EasyOCR is an optional secondary support. The OCR strings are normalized by removing excess spaces and symbols, e.g. "MH 49-BK 0504" can become "MH49BK0504". We combine recognition evidence accumulated over multiple frames instead of considering only the single best OCR output. The final registration candidate depends not only on the OCR confidence and the number of evidence frames, but also on the quality of the plate, the validity of registration-format, and the repeated observation, etc. This temporal technique hopefully reduces errors from a number of blurred or inferior quality image frames.

E. Visible HSRP Feature Analysis

After obtaining a reliable registration candidate, the system performs visible HSRP-oriented screening. The current implementation considers evidence that can reasonably be observed from normal traffic video, including: plate quality, OCR reliability, temporal consistency, registration-format plausibility, plate aspect ratio, visible IND-band evidence.

F. Screening Decision

Based on the accumulated evidence, the system produces one of two outcomes:

Likely Compliant

This result is generated when sufficient supporting evidence is available.

Needs Manual Verification

This result is generated when the available evidence is weak, inconsistent, or insufficient for a confident screening decision.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue X Oct 2026- Available at www.ijraset.com

G. Database and Dashboard

The generated vehicle records, OCR results, plate evidence, screening scores, and review information are stored in an SQLite database. SQLite is used because it is lightweight, easy to integrate, and appropriate for the current academic prototype. A FastAPI-based web dashboard presents the stored information to the operator. The dashboard allows users to inspect processed runs, tracked vehicles, plate images, OCR outputs, screening status, annotated video, and review information. It also supports OCR correction and manual verification.

H. Explainability and Human-in-the-Loop Verification

Explainability is provided by a "Why This Result?" panel. Instead of only providing the final label, the dashboard would theoretically show the screening evidence, with the positive signs being a valid registration format, reliable temporal OCR agreement, good plate quality, or multiple IND evidence for example. Indicators for being a warning indicator would be poor OCR confidence, poor plate quality, too few supporting frames, and lack of consensus on visual evidence. The human operator would be responsible for reviewing any cases where there was one or more warning indicators. The operator would have access to the plate crop, OCR, score, confidence information, and explanation factors before noting a manual decision. From this, the system would appear to be a decision-support system and not an automated enforcement system.

IV. SYSTEM IMPLEMENTATION

The HSRP compliance screening and decision-support system of the proposed framework was deployed as a modular software pipeline using Python. The implementation combines modules of computer vision, optical character recognition, database, and web-based visualization, which are built and validated separately but work together as part of an entire end-to-end pipeline.

A. Software Environment

Python was used as the main language for software development. The main libraries used for implementing the main pieces of the software are a YOLO-based object detector, Byte Track for tracking the vehicles and OpenCV for video processing and image processing. Additionally, for number plate recognition, PaddleOCR was used, local data storage was done using a database - SQLite and the web platform dashboard was built with FastAPI.

Each of these components such as the detector, OCR engine, decision logic or storage layer can be individually replaced or changed.

B. Vehicle Detection

Each frame of the traffic-video is used to detect typical road vehicles – car, bike, bus, and truck using a YOLO-based object detection model. The object detector outputs the class of the vehicle, the confidence of detection, and the bounding box.

Only those vehicle types relevant to the relevant vehicle categories are passed onto the subsequent plate- and plate-processing tracking process. This eliminates the continued processing of irrelevant objects, which are not to be processed to screen for vehicle compliance.

C. Multi-Object Vehicle Tracking

ByteTrack is used to maintain the identity of detected vehicles across consecutive frames. Each tracked vehicle is assigned a unique identifier, which remains associated with the vehicle while it is visible in the video. Tracking is important because the same vehicle may appear in many frames with different plate visibility. Maintaining a persistent vehicle ID enables the system to collect multiple number-plate observations for a single vehicle and supports later operations such as best-frame selection and temporal OCR analysis.

D. Number Plate Detection and Association

A dedicated license plate detection model is used to detect registration plates associated with tracked vehicles. Plate detections are linked to the corresponding vehicle using the spatial relationship between the detected vehicle region and plate bounding box. The detected plate region is cropped from the original frame and stored as evidence for that vehicle. Multiple plate observations may therefore be collected for the same tracked vehicle as it moves through the video

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

E. Plate Quality Assessment and Best-Frame Selection

Not all plate observations are suitable for OCR. Some observations may be blurred, tiny, improperly lit or misaligned. The implementation checks the observations of a plate before choosing one for OCR. Detection confidence, plate size, area, aspect ratio, sharpness, and quality are factors. More confident detections are kept to provide better evidence on which to apply OCR, instead of using the first plate detection. This part is especially important for traffic video since as the vehicle moves, the plates appear and disappear all the time.

F. OCR Processing

PaddleOCR recognizes the registration number from each plate crop, which has been extracted using the previously described plate detection algorithm. A variety of pre-processing operations can be applied to the images before a registration number is recognised, including resize, grey conversion, contrast, sharpening and thresholding. The OCR results have been normalize, removing unnecessary spaces and characters, e.g., MH 49-BK 0504 is normalized to MH49BK0504. If necessary, an optional second OCR engine can be configured to serve as an additional source of recognition.

G. Temporal OCR Validation

The system is not rely on a single OCR observation. OCR observations from more than one frame, in which the same tracked vehicle has been detected, are combined. Selection of the last plate candidate is based on the following metrics: OCR confidence, registration-format validity, number of contributing observations, quality of the plate-image, and consistency across frames. This makes the approach more robust since a bad or blurry OCR result from one frame can be ignored if it has better and more consistent results in subsequent frames.

H. HSRP Screening and Decision Logic

Following registration number recognition, visible-feature-based HSRP screening is conducted. The implementation accounts for information like plate quality, OCR confidence, temporal consistency, registration format plausibility, aspect ratio of the plate, and visible IND-band support. These factors contribute to a screening score and decision. The system produces two main outputs:
Likely Compliant - when the strength of the HSRP-like features in the evidence was arguably demonstrated.
Needs Manual Verification - evidence is weak, missing, inconsistent or cannot be verified.
The system has been designed to not declare that any vehicle is legally non-compliant, by default, as there are a number of HSRP security features that cannot be tested accurately in an everyday traffic-video image.

I. Database Storage

SQLite database. This database is responsible for recording information about processed runs, tracked vehicles, OCR results, plate evidence, HSRP screening status, and human-review information. SQLite database was chosen as it is relatively lightweight and serverless and hence easily ported over to the existing Academic prototype. A more scalable database system can be used during the final deployment, leaving the computer-vision pipeline unchanged.

J. Web Dashboard

A FastAPI-based web application is used to present the stored results through a browser interface. The dashboard allows the operator to view processed runs, tracked vehicles, recognized plate numbers, plate images, screening status, and associated evidence. The dashboard also supports manual OCR correction and human-review actions. This enables the operator to correct recognition errors and record a final review decision when required.

K. Explainability and Human Review

The "Why This Result?" feature in the dashboard provides an explanation of the basis for each screening decision.
For example, supporting evidence could be: valid registration format, strong OCR confidence, temporal OCR agreement, good plate quality, visible IND evidence.
Warning factors may include: weak OCR confidence, insufficient supporting frames, poor plate quality, inconsistent OCR results, weak visible HSRP evidence.
This is an explainability feature allowing the operator to view the reasoning of the system for reaching the screening outcome instead of looking just at the outcome label. The final workflow still includes human review. So, the system in use helps with legal decisions using AI, but it isn’t a fully independent legal enforcement agent.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue X Oct 2026- Available at www.ijraset.com

V. EXPERIMENTAL SETUP AND VALIDATION

A. Experimental Setup

The prototype developed was tested with a recorded traffic-video data input. The experiment was conducted to see the feasibility of the entire processing pipeline, right from detection of vehicle to final decision-support output. The workflow of the system is: detecting vehicles, tracking based on Byte Track algorithm, detecting number plates, associating number plates with vehicles, selecting the best frame for plate recognition, running optical character recognition, performing temporal checker on OCR results, screening the HSRP plates in the visible range, saving the data to a database, dashboard display and human-reviewer.

The existing prototype is designed to feed in recorded traffic video rather than a live CCTV/RTSP feed. This was a deliberate decision to allow for controlled prototype validation, with live-camera integration as future work.

B. Automated Software Validation

In addition to traffic-video processing, the implementation was validated using automated software tests. A total of 52 automated tests were executed successfully. The automated tests covered important functional areas of the application, including: pipeline lifecycle handling, OCR candidate selection, vehicle scanning logic, screening and decision logic, persistent database storage, integration between major system modules. All 52 out of 52 tests passed, providing evidence that the principal software components were correctly integrated and functioning as expected.

C. End-to-End Video Validation

A verified end-to-end test run was conducted using 60 traffic-video frames. During this run, the complete system processed the video through the detection, tracking, plate-analysis, OCR, and decision-support stages.

The following implementation-level results were obtained:

TABLE: IMPLEMENTATION RESULTS

Validation Parameter	Results
Frames Processed	60
Tracked vehicles	21
Best-plate updates	14
OCR executions	13
Automated tests	52/52 passed

The experiment confirmed that the system was able to maintain multiple vehicle identities across frames, collect plate observations, update the selected plate evidence when a better crop became available, and execute OCR on suitable plate candidates.

D. Validation of Best-Frame Selection

An important observation during testing was that the first detected number plate was not always the clearest plate image available for a vehicle. As a vehicle moved through successive frames, later observations sometimes produced sharper and more readable plate crops. The best-frame mechanism allowed a weaker observation to be replaced by a better-quality plate image based on factors such as image sharpness, plate dimensions, confidence, and visual quality. This confirms the practical value of using multi-frame evidence instead of immediately performing recognition on the first available frame.

E. OCR and Temporal Validation

PaddleOCR was used as the primary OCR engine during the validation process. OCR text obtained from selected plate images was normalized and evaluated together with plate quality, registration-format validity, and repeated support across multiple frames. During testing, it was observed that OCR confidence alone was not always sufficient to determine whether a recognized registration number was reliable. In some cases, an OCR output could have relatively strong confidence while still containing an incorrect or incomplete character. For this reason, the system compares multiple OCR observations associated with the same tracked vehicle. Recognition candidates supported consistently across several frames are given greater importance than isolated OCR results.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

F. HSRP Screening and Dashboard Validation

After OCR processing, the system used the available visual and recognition evidence to generate an HSRP screening result. The two possible outputs were:

- Likely Compliant when sufficient supporting evidence was available.
- Needs Manual Verification when the available evidence was weak, inconsistent, or insufficient.

The resulting vehicle records, OCR information, plate evidence, and screening status were stored in the SQLite database and displayed through the FastAPI-based dashboard. The dashboard was also verified to support plate-evidence viewing, OCR information, screening status, explainability, and human-review controls.

G. Scope of the Validation

The form factor validation is mainly for functional validation and end-to-end behavior validation of the prototype. The form factor validation used to prove that detection, tracking, plate processing, temporal OCR, visible HSRP screening, database persisting and dashboard based review can be combined together to work as an integrated system. At this time, the validation was not able to specify an overall percentage accuracy for HSRP classification or OCR recognition over a significant labelled dataset. A much larger dataset of manually labelled samples could be used to calculate accuracy, precision, recall, F1-score and a confusion matrix etc. It should be listed as a limitation instead of claiming accuracy that hasn't been measured.

VI. RESULTS AND DISCUSSION

The final prototype was evaluated through automated testing and end-to-end traffic-video processing. The results show that the major components of the proposed system were successfully integrated and were able to operate as a complete workflow, including vehicle detection, tracking, number plate detection, best-frame selection, OCR, temporal validation, HSRP screening, database storage, dashboard visualization, and human review.

A. Overall Validation Results

A verified run consisting of 60 video frames produced 21 tracked vehicles, 14 best-plate updates, and 13 OCR executions. In addition, 52 out of 52 automated tests passed successfully. These tests covered important areas such as pipeline lifecycle management, OCR candidate selection, scanning logic, decision logic, persistent storage, and integration between major system modules.

B. Vehicle Detection and Tracking Results

The vehicle-detection stage successfully identified road vehicles in the input video. ByteTrack maintained unique vehicle identities across successive frames so that repeated appearances of the same vehicle were not treated as separate vehicles. This tracking capability is important because it creates a vehicle-wise history. Instead of processing each frame independently, the system can collect multiple observations belonging to the same vehicle. This provides the temporal information required for best-frame selection and OCR validation.

C. Number Plate Detection and Best-Frame Results

License plates associated with tracked vehicles were successfully detected and linked to their corresponding vehicle records. An important observation during testing was that the first available plate image was not always the best one. As a vehicle moved through later frames, clearer and sharper plate crops were sometimes obtained. The system therefore updated the stored plate evidence when a stronger observation became available. The 14 best-plate updates obtained during the verified run demonstrate that plate quality changed over time and that retaining improved observations can be more useful than relying only on the first detection. This supports the use of best-frame selection in traffic-video applications.

D. OCR and Temporal Recognition Results

We employed PaddleOCR to detect license plate number. The recognized strings were normalized and tested together with license plate quality, whether the plate is registered according to the format, and the support time. Testing indicated that on its own OCR confidence should not be considered conclusive evidence of a valid registration number. On occasion an OCR engine will accept a result with only a modicum of confidence even when the registration contains at least one incorrect character. The system can further increase the importance of register candidates that are supported over time, comparing several tracking observations of the same vehicle. This validation over time can reduce the impact of unstable OCR results.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

E. HSRP Screening Results

Using plate-quality assessment and OCR, the unit then screens using visible-feature-based HSRP. The screening phase will use any and all available evidence currently possible:

OCR reliability, temporal agreement, registration-format plausibility, plate quality, aspect ratio, visible IND-band evidence. The system then produces one of two outputs: Likely Compliant or Needs Manual Verification. This conservative approach is driven by the fact that ordinary traffic-video footage will not be able to prove the physical security components of an HSRP beyond a doubt. As a result, questionable instances are passed through to a human instead of being automatically classified as legal violations.

F. Dashboard and Explainability Results

The processed results were successfully stored in the SQLite database and displayed through the FastAPI-based dashboard. The dashboard provides access to: processed runs, tracked vehicle records, recognized plate numbers, OCR confidence, plate images, screening status, annotated video, human-review information.

The “Why This Result?” feature provides additional interpretability by showing the factors that increased or decreased confidence in the screening decision. For example, a result may be supported by valid registration format, good plate quality, or temporal OCR agreement, while warnings may include weak IND evidence, poor plate quality, or inconsistent OCR observations.

G. Discussion

The experiment shows that using more than one frame works better than just using a single frame. Vehicle tracking enables us to build up several observations of the same vehicle over time. Best-frame selection captures clearer plate evidence, and temporal OCR validation offers an additional means of cross-referencing multiple recognition results before choosing a final candidate.

Results also show the value in having a human-in-the-loop approach. No one factor is solely adequate for identifying HSRP: the evidence in normal traffic video can be compromised due to image quality, angle of view, camera motion blur, occlusion, and plate resolution.

As a result, the system is unable to make an automatic legal enforcement call. These are listed under Needs Manual Verification and the operator can then review the evidence and justification.

H. Observed Limitations During Testing

The testing also revealed several practical limitations. The Quality of OCR and HSRP Screening The ability to reliably process HSRPs and number plates is dependent on how visible the number plate is. A small, blurred, overexposed, blocked or an angled up to sharp will make it harder to pass reliably. Another problem that you may encounter with OCR is a misconception of similar characters 0/O, 1/I, 5/S. Furthermore, several security features of HSRP as well as hologram validity, laser-etched identifier numbers, manufacturer-specific data and tamper-proof security cannot be identified using standard-traffic video images. Hence, what has been implemented can be viewed as an AI-assisted compliance screening and decision support prototype, but not an advanced legal HSRP authentication system.

VII. LIMITATIONS

Although the proposed system successfully integrates vehicle detection, tracking, number plate recognition, visible HSRP screening, database storage, explainability, and human verification, the current prototype has several practical limitations.

A. Dependence on Video and Plate Quality

System performance depends heavily on the quality of input traffic video. Number plates that appear very small, blurry, overexposed, underexposed, partially occluded, and viewed at a high viewing angle all present challenges when it comes to detection and recognition. These can be problematically used for optical character recognition (OCR), but while best-frame selection and temporal evidence both enhance reliability, they cannot compensate for information missing from the original video.

B. OCR Errors

Optical Character Recognition recognises characters that are very similar, for example: 0 and O 1 and I 5 and S. Recognition errors may also occur if the font of the plate is hard to read, the plate is dirty or there are strong reflections. This approach alleviates the problem to some extent by registration-format validation and multi-frame temporal support, but cannot completely resolve the error problem of OCR.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

C. Limited Verification of HSRP Security Features

One of the major limitations of the system is that it performs visible-feature-based HSRP screening, not complete legal HSRP authentication. Ordinary traffic-camera footage cannot reliably verify security-sensitive or microscopic characteristics such as: chromium-based hologram authenticity, laser-etched identification number, manufacturer-specific security information, tamper-resistant physical mechanisms. Therefore, the system should not be used as the sole source for legal enforcement.

D. Dependence on Visible IND Evidence

The presence or absence of the IND marking may contribute to the screening decision. However, the IND region may not always be clearly visible due to blur, angle, distance, reflection, or partial obstruction. For this reason, weak IND evidence is treated as a reason for manual verification rather than direct rejection.

E. Current Prototype Uses Recorded Traffic Video

The present prototype has primarily been validated using recorded traffic-video input. Direct integration with live CCTV or RTSP camera streams has not yet been fully implemented and validated. Therefore, live continuous deployment is considered future work.

F. Local Deployment

The existing code uses a local FastAPI server and SQLite database. This is fine for a class prototype but this isn't how you would do it for a city-wide system. A production system would need a more scalable infrastructure such as centralized servers, robust database, authentication, access control, secure communication, logging and distributed processing.

G. Limited Large-Scale Dataset Evaluation

This work at hand is mainly a proof of function for the integrated prototype. No large manually labelled HSRP dataset has been used to determine, for example, the final accuracy, the precision, the recall, the F1-score or the confusion matrix. Hence, the results are good indication of system integration and functional verification and do not suggest explicit large scale classification accuracy of HSRP.

H. Human Verification Is Still Required

The system is intentionally built as a decision support system. Vehicles with weak or inconclusive evidence are classified as Needs Manual Verification. This implies that for doubtful situations, especially when there is not enough visual information, there will still need a human operator to make the final decision.

VIII. CONCLUSION AND FUTURE WORK

This work presented an AI-Based Intelligent HSRP Compliance Screening and Decision Support System using Computer Vision for traffic-video analysis. The developed prototype integrates vehicle detection, multi-object tracking, license plate detection, plate-to-vehicle association, best-frame selection, OCR, temporal validation, visible HSRP screening, database storage, explainability, and human verification within a single workflow.

The use of ByteTrack allows the same vehicle to be followed across multiple frames, enabling the system to collect repeated plate observations instead of relying on a single image. Best-frame selection improves the quality of the evidence used for OCR, while temporal OCR analysis helps identify more stable registration candidates by comparing observations over time.

The system also adds to traditional ANPR by estimating visibly viewable HSRP related evidence like plate quality, OCR accuracy, registration format plausibility, plate aspect ratio, time consistency and IND-band evidence. From these, the system outputs Likely Compliant or Needs Manual Verification decisions. The system uses this conservative approach to avoid turning any unknowns into legal violations, instead of handing off these to a human operator for review.

The implemented prototype was validated using automated tests and end-to-end traffic-video processing. A verified run processed 60 video frames, maintained 21 vehicle tracks, generated 14 best-plate updates, and performed 13 OCR executions. In addition, 52 out of 52 automated tests passed, confirming the functional integration of the major system modules.

The FastAPI-based dashboard and SQLite storage layer provide a practical interface for reviewing processed vehicles, OCR results, screening decisions, and supporting evidence. The “Why This Result?” functionality improves transparency by showing the evidence and warning factors that influenced each screening result, while manual correction and human review preserve human control over uncertain decisions.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The system is not designed to supplant HSRP authentication or law-enforcement processing. Some physical security attributes, including hologram authentication, laser-etched ID numbers, manufacturer information, and anti-tampering features, are difficult to confirm from plain-vanilla surveillance footage. Consequently, the system should be employed as an AI-assisted screening and decision support tool.

A. Future Work

Future developments can build on this prototype by adding: direct-to-live CCTV or RTSP camera stream integration for real-time traffic monitoring. Multi-camera vehicle trajectory tracking for following vehicles at different sites and for general traffic analytics. Recognition performance can be further enhanced by having larger and more diverse Indian number-plate and HSRP dataset, better preprocessing, motion deblurring, super-resolution and OCR models. A bigger evaluation dataset with manual annotations would make it easier to check the system's performance using numbers like accuracy, precision, recall, F1-score, and confusion matrix. As a first step towards deploying to a production environment, the application can be moved onto a cloud or distributed infrastructure. The application database can be upgraded to a scalable server database solution such as PostgreSQL, replacing the current SQLite database. Security enhancements can also be made, including authentication, encryption, audit logs and secure storage of evidence. Other modules for traffic-volume analysis, congestion estimation, route-movement analysis and city-wide multi-camera monitoring can further extend the system into a comprehensive smart transportation platform.

REFERENCES

[1] R. Al-Batat, A. Angelopoulou, S. Premkumar, J. Hemanth, and E. Kapetanios, "An End-to-End Automated License Plate Recognition System Using YOLO Based Vehicle and License Plate Detection with Vehicle Classification," *Sensors*, vol. 22, no. 23, Art. no. 9477, 2022, doi: 10.3390/s22239477.

[2] R. Adak, A. Kumbhar, R. Pathare, and S. Gowda, "Automatic Number Plate Recognition (ANPR) with YOLOv3-CNN," *arXiv preprint arXiv:2211.05229*, 2022.

[3] S. Dhyani and V. Kumar, "Real-Time License Plate Detection and Recognition System using YOLOv7x and EasyOCR," in *Proc. 2023 Global Conference on Information Technologies and Communications (GCITC)*, Bangalore, India, Dec. 2023, doi: 10.1109/GCITC60406.2023.10425814.

[4] L. You, Y. Chen, C. Xiao, C. Sun, and R. Li, "Multi-Object Vehicle Detection and Tracking Algorithm Based on Improved YOLOv8 and ByteTrack," *Electronics*, vol. 13, no. 15, Art. no. 3033, 2024, doi: 10.3390/electronics13153033.

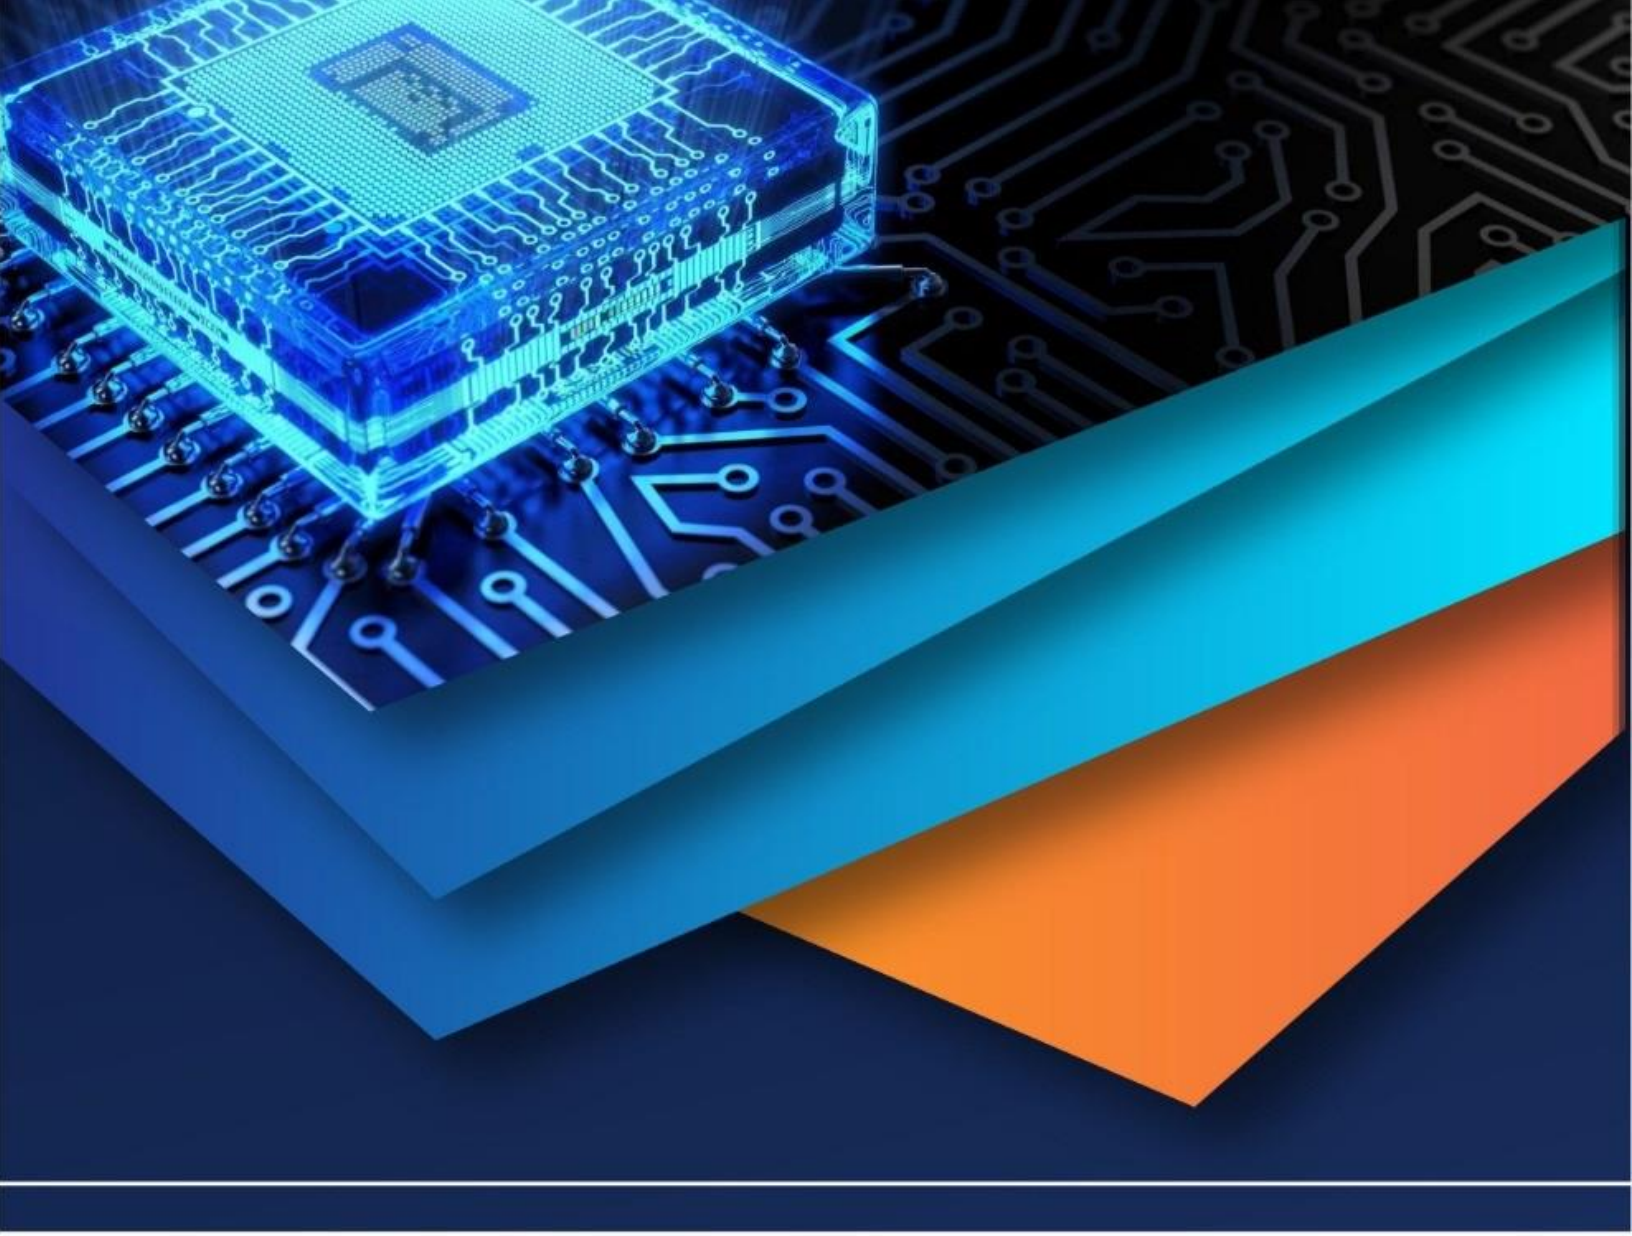
[5] A. Y. Sugiyono, K. Adrio, K. Tanuwijaya, and K. M. Suryaningrum, "Extracting Information from Vehicle Registration Plate using OCR Tesseract," *Procedia Computer Science*, vol. 227, pp. 932–938, 2023, doi: 10.1016/j.procs.2023.10.600.

[6] H. K. Gianey and A. Goyal, "Localization and Recognition of High Security Registration Plate for Indian Vehicle," *International Journal of Computer & Organization Trends*, vol. 4, no. 2, pp. 16–19, 2014, doi: 10.14445/22492593/IJCOT-V6P304.

[7] G. Badamikar, G. Parashetti, V. Handiganur, P. Singh, and M. M. Raikar, "High Security Registration Plate Detection and Classification Using Federated Learning," *Procedia Computer Science*, vol. 260, pp. 1129–1136, 2025, doi: 10.1016/j.procs.2025.03.298.

[8] A. Kuznietsov, B. Gyevnar, C. Wang, S. Peters, and S. V. Albrecht, "Explainable AI for Safe and Trustworthy Autonomous Driving: A Systematic Review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 12, pp. 19342–19364, Dec. 2024, doi: 10.1109/TITS.2024.3474469.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)