AI-Driven Traffic Monitoring and Driver Behavior Scoring System Using Computer Vision and Deep Learning

Rajvardhan Balaji Kale[1], Siddhesh Katale[2], Ganesh Langhi[3], Ganesh Shelke[4] Swara Mamidwar[5] Saniya Meshram[6]
[1,2,3,4,5,6]*Information Technology Vishwakarma Institute of Technology Pune, India*

***Abstract*—Cities grow fast. More cars fill streets every day. Watching over traffic becomes harder because of this. Old ways rely heavily on people watching and isolated cameras. These approaches struggle to scale up. They miss live insights too. A new system steps in here. It uses smart image processing. Deep neural networks help it learn patterns. Data flows automatically through its core. Vehicles appear within view. The setup spots them without delay. License numbers get picked out clearly. Rules broken by drivers show up instantly. Technology shifts how oversight works now. Every piece connects tightly inside. Results come faster than before. Accuracy improves across tasks. Real world conditions test its strength daily. Performance holds steady under pressure. Information moves smoothly from capture to output. No part drags behind others. Decisions form quicker as a result. Systems like this adapt easily elsewhere. Each function supports another indirectly. Progress shows in small details mostly. Monitoring feels different now. Not just louder alarms but smarter responses shape outcomes.**

A single vehicle scan kicks off the process by spotting cars and their plates in road videos. Following that, characters on licenses get pulled out using text reading tech. Once captured, data links up with a main storage hub to pull driver info and log past violations. Scoring how drivers behave comes next, building points based on patterns spotted over time.

Now imagine cameras that watch roads, then think like humans to spot bad driving. Instead of just recording, they learn what risky moves look like over time. This means fewer officers needed to review endless footage. Machines handle routine checks, freeing people for tougher tasks. Safer streets come not from more cops but smarter tools. Decisions about violations become clearer when based on patterns, not guesses. Drivers adjust their habits when they know the system notices details. Clearer outcomes follow because rules are applied more evenly. Fewer arguments happen when evidence comes from consistent observation. What changes is not just how laws are enforced but how fairly they feel.

***Index Terms*—Traffic Monitoring, Computer Vision, Deep Learning, Automatic Number Plate Recognition, Driver Behavior Analysis, Smart Transportation Systems**

I. INTRODUCTION

With more cars crowding city streets, getting stuck in traffic and staying safe on roads has grown harder. Old-style traffic control leans on people watching or separate camera setups that struggle to catch rule-breaking quickly while missing pieces of a driver's full history.

Now machines learn to see through camera feeds on roads, spotting cars as they pass by. Because of smarter software, reading number plates happens faster than before. When rules get broken, alerts pop up without waiting for a person to notice. Cameras stay awake all night, catching issues humans might miss during long shifts. Errors slip in less often once computers take over the counting and checking. Watching traffic turns smoother when routines run themselves every second. Fewer officers need to stare at screens because smart tools handle routine scans. What used to take minutes now finishes in moments thanks to live processing. Systems keep working even in heavy rain or bright sun, never blinking. Decisions come quicker when code reviews images instead of people doing it.

A fresh look at road oversight begins with smart cameras spotting cars in motion. Following detection, number plates get pulled from frames automatically. Once captured, those details link up with driver records through secure channels. Instead of just logging incidents, the setup assigns updated risk marks per individual. Behavior on roads feeds into a running tally tied directly to each operator.

Now picture this: instead of just handing out fines, the new system tracks how people drive over time. Because habits matter, it builds a profile based on consistent actions behind the wheel. With that info, officials can spot frequent rule breakers easier than before. Good choices add up quietly, shaping outcomes without fanfare. Over months, patterns emerge - clearer, harder to ignore.

II. LITERATURE REVIEW

A. Previous Research on Automatic License Plate Recognition and Traffic Monitoring

Looking into Du, Ibrahim, Shehata, and Badawy’s review on Automatic License Plate Recognition [1], one finds a detailed look at methods behind license plate reading tech. Image capture kicks things off, followed by spotting the plate, splitting characters apart, then identifying them through OCR. Because these steps matter so much, they’ve become key in areas like tracking traffic flow, managing entry to parking zones, or handling toll payments without cash. While their paper digs deep into current approaches, it stops short of linking together vehicle spotting with catching rule breaks or studying how drivers act - the kind of full chain our method actually puts into practice.

Starting off, Laroca alongside Severo, Zanlorenzi, Oliveira, Gonçalves, Schwartz, and Menotti present a method in "A Robust Real-Time Automatic License Plate Recognition Based on the YOLO Detector" [2]. This approach leans on YOLO-driven detection to spot license plates within both still pictures and moving footage. Built for live operation, it manages strong precision while keeping pace with fast data flow. Tests show consistent results when applied to various image collections. Although it pushes forward how quickly and correctly plates are found, its aim stays narrow. Tracking unlawful driving actions isn’t part of the framework. Neither does it follow driver patterns across repeated observations.

Starting off differently, Laroca and team built a smart license plate reader using YOLO that works no matter how plates are arranged. This setup handles many country styles without needing layout changes. It gets most readings right, even when tested on varied image sets. Still, while spotting plates works well, it skips extra tasks like catching speeders or rating driver behavior.

Starting off with vehicle snapshots, Adak, Saha, and Bhattacharya built a method using YOLOv3 paired with a CNN to spot cars and their plates. Instead of one single step, it splits the job: first locating the plate, then reading the characters on it. Because of how light shifts in street scenes can mess up readings, their setup handles glare and shadows better than earlier versions. Even so, once the numbers are pulled out, nothing follows through - like catching speeders or logging who drives badly. While accuracy gets a boost, extra layers such as rule enforcement stay untouched.

Starting off with Satya, Kumar, and Reddy's work titled "Optimized YOLOv8 for Automatic License Plate Recognition on Resource-Constrained Devices" [5], the focus lands squarely on how trimmed-down versions of YOLOv8 handle live license plate spotting in compact hardware setups. Efficiency takes center stage here - especially making sure processing stays light without losing precision in detecting plates. While results show solid function even where computing power is limited, features like monitoring driver actions or supporting traffic violation checks aren’t part of this setup at all. Ending there leaves room for future tools to build beyond basic detection.

A study titled "Traffic Violation Detection Using Deep Learning" [6] looks into how convolutional neural networks work alongside YOLO-style models to spot traffic offenses in videos. Instead of just listing actions, it pinpoints things like running red lights or going too fast, along with pulling out license plates using optical character recognition. While the method zeroes in on catching these specific breaches, its core design skips building a way to track driving habits over time. Despite strong detection skills, long-range behavioral analysis stays outside its scope.

One paper titled “Automatic Traffic Rules Violation Detection and Number Plate Recognition” [7] uses deep learning to spot traffic offenses alongside identifying car license plates. Built on image analysis, it pulls together convolutional neural networks for spotting vehicles and optical character recognition to grab plate digits. Though it works well catching rule breaks, because it does not track individual drivers across incidents. Over time, that gap limits how much insight it can offer into repeated behaviors.

III. METHODOLOGY

Layer	Technology	Function
Detection Layer	YOLO Model	Detect vehicles and number plates
Recognition Layer	OCR	Extract license plate characters
Processing Layer	OpenCV	Image and video processing
Backend Layer	FastAPI / Python	Manage detection and violations
Database Layer	MySQL	Store driver data and violation history
Analytics Layer	Machine Learning	Driver safety scoring

Table 1: Technology stack

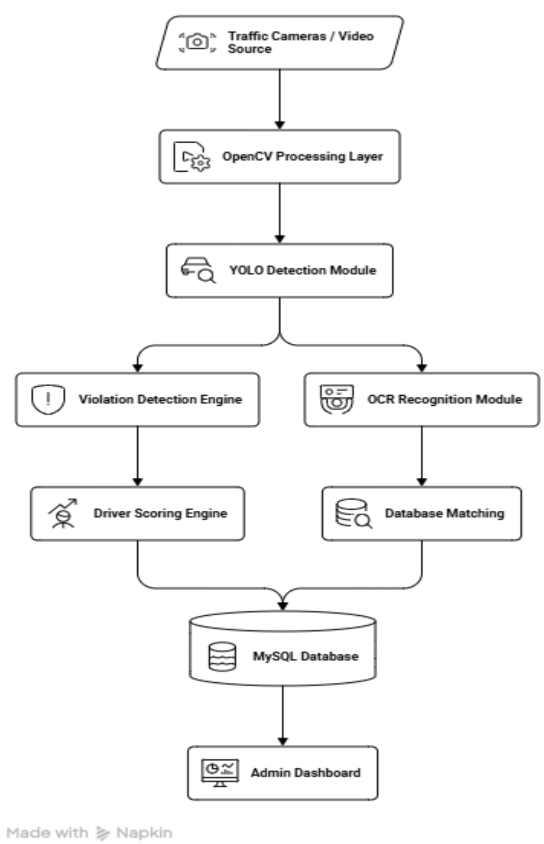

Fig. 2: Workflow

A. Vehicle Detection
Out of raw footage captured by road monitors, the setup pulls individual images. Inside every shot, a vehicle-spotting algorithm built on YOLO picks cars while pinning down plate positions. Because it moves fast without losing precision, this kind of network shows up often in live tracking tasks.

B. License Plate Recognition
After spotting the license plate area, processing begins by pulling out that section for closer work. Grayscale adjustments come next, helping to standardize the visual input before further steps. Noise filters smooth out distractions in the picture, clearing up details quietly. From there, optical character tools start interpreting the symbols one by one. Clarity matters most once the cleaned image feeds into the reading phase.

C. Violation Detectionr
From above, motion of cars gets studied when traffic flows happen. When someone cuts a red light, moves too fast, or strays into wrong lanes, it catches attention. Each time something like that happens, the moment is saved. Linked each instance becomes to the exact car by its plate.Out of raw footage captured by road monitors, the setup pulls individual images. Inside every shot, a vehicle-spotting algorithm built on YOLO picks cars while pinning down plate positions. Because it moves fast without losing precision, this kind of network shows up often in live tracking tasks.

D. Driver Behavior Scoringr
Each driver is assigned a safety score representing responsible driving behavior. Violations reduce the score depending on severity.

Violation	Score Reduction
Speeding	-5
Signal Jumping	-15
No Helmet	-10
Lane Violation	-8

Drivers with repeated violations receive lower scores and may be flagged for monitoring.

IV. RESULTS AND DISCUSSIONS

The The proposed system was implemented using YOLO-based object detection combined with OCR-based license plate recognition. Experimental testing showed that the system successfully detected vehicles and extracted license plate information from traffic footage.

A. Functional Testing

The system successfully performed the following main tasks:

• System validation and testing: Stored traffic videos were used to evaluate the performance of the system. In some tests, simulated camera feeds were provided to replicate real-time traffic conditions. Each module of the system was first tested independently, and afterward the integration between components was verified. Communication between modules was carefully monitored to ensure correct operation. All system components executed their tasks successfully during these tests.

• Vehicle detection: The system detected vehicles from traffic footage using a fast object detection model based on YOLO. Once a vehicle was detected, the software generated bounding boxes around each vehicle to mark its location within the frame. These bounding boxes allowed the system to isolate vehicle regions from the rest of the image.

• License plate localization: After detecting vehicles, the system identified the location of license plates within each detected vehicle. A specialized detection model was used to accurately locate the plate area. This ensured that only the relevant region containing the license plate was extracted for further processing.

• License plate recognition: The extracted license plate region was processed using Optical Character Recognition (OCR). The characters present on the license plate were detected and converted into digital text. This process allowed the system to automatically identify vehicle registration numbers without manual input. The recognition process remained reliable under varying lighting conditions.

•Traffic violation detection: The system continuously monitored traffic footage to detect rule violations such as red-light crossing, overspeeding, and lane violations. When a violation occurred, the system automatically identified the responsible vehicle and recorded the associated license plate number. This monitoring operated continuously without requiring manual supervision.

• Driver information retrieval: Once a license plate number was detected, the system matched it with existing records stored in the database. By cross-referencing the recognized plate number, the system retrieved information related to the vehicle and its registered owner. This enabled authorities to quickly identify drivers involved in traffic violations.

• Safety score management: The system implemented a dynamic driver safety score mechanism. Whenever a driver violated a traffic rule, the safety score was automatically reduced. The amount of score reduction depended on the severity of the violation. This automated scoring system allowed authorities to track driver behavior over time.

• Violation data logging: All detected violations were recorded in a MySQL database. The stored information included the license plate number, type of violation, location of the incident, and timestamp. This database enabled long-term monitoring, analysis of violation patterns, and identification of repeat offenders.

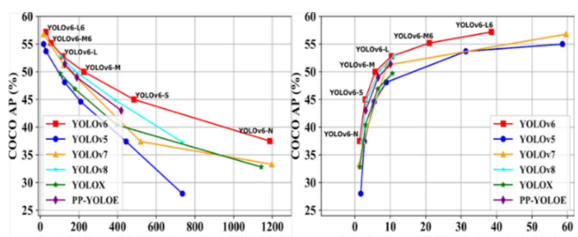

Fig. 3. Performance comparison of YOLO models based on COCO AP and inference speed.

B. Security Evaluation

A test checked how well the new system worked. Efficiency came under review alongside precision during trials. The setup faced checks meant to gauge its real-world function. What it did matched against expected results. Speed met careful observation just like correctness. Results showed whether goals were reached.

From behind the camera feed, shapes took form quickly as the system picked out cars without delay. Instead of waiting, movement flowed steadily through each frame thanks to smart processing built into the design.

A sudden shift to grayscale helped sharpen the captured plate areas right after detection. Instead of jumping straight into reading, a careful removal of visual clutter took place first. This cleanup work made letters stand out much more clearly ahead of analysis. Better results came through only once these adjustments settled in.

When license plates showed up clean and bright, the OCR module pulled out the characters just fine. Even so, dim lights or something blocking part of the plate made it a bit harder to get everything right.

Fine results came through when testing live traffic tracking abilities. The setup held up without issues during continuous use.

C. Security Evaluation

From the start, checks looked at how well the system caught rule breaks without fail. Accuracy in keeping track of each driver's history mattered just as much. What stood out was its steady performance spotting issues. Records stayed correct over time too. Reliability showed up most when catching errors early. Through it all, data remained consistent and clear. Finding cars that ran red lights started with camera footage checked nonstop. A car moving too fast triggered alerts just as quickly as one slipping into banned lanes. Caught on record each time, every breach showed up without delay. Moments after breaking rules, proof was already stored. From every offense, a clear link emerged - tied by the visible plate on each car. Because of this detail, penalties found their way to the right person behind the wheel. Automated detection tools cut down mistakes when tracking traffic breaches. What changed next? Officials could spot frequent violators quicker thanks to shared records stored in one place.

D. User Experience Considerations

From the start, it worked quietly behind scenes. Traffic officials found their tasks easier through its clean layout. Instead of clutter, they got clarity - each function placed where needed. Efficiency grew without extra effort. A calm design helped them stay focused on what mattered. From the screen, drivers appeared alongside their tagged licenses and any related warnings. Stats on breaches showed up right beside safety ratings and past incidents. What popped up included every flagged vehicle, each plate caught by sensors, followed by alerts tied to rule breaks. Past logs stayed visible, linked to current reports without extra steps needed. Right away, officials spotted risky drivers through the system, thanks to clear violation records on display. Still, adding phone compatibility might help, along with visuals that respond when touched or clicked.

E. Limitations

Few issues still linger, even though the results looked good.When light is too dim, or a car moves fast, reading plates gets harder. Sometimes dirt covers part of the plate, making it worse. Better cameras might make images clearer. Fixing pictures before analysis may also reduce errors. Hard to spot doesn't mean impossible - just tougher.Few changes would help it work better across wider areas, since right now it mostly uses video feeds alone. Other tools like motion trackers or road sensors might need to plug into the setup so it runs smoothly at bigger scales.Few steps ahead could mean tuning how systems spot red flags, while weaving in smarter algorithms that learn patterns of unsafe driving over time. What comes next might lean on sharper tools - ones that adapt faster than older methods ever did. Watching for trouble before it grows may depend less on rules and more on what machines pick up through experience.

V. FUTURE SCOPE

Staying just slightly ahead may connect the system to city camera networks, allowing real time observation over wider regions. Because several cameras feed data simultaneously, officials can watch how vehicles move through whole districts instead of single crossroads. When linked together, this setup turns into a unified control hub that reacts fast when traffic acts out of the ordinary.

Maybe machines will learn trickier habits behind the wheel. Not just how fast someone goes or if turn signals work, but whether eyes drift to screens instead of roads. Sudden stops might tell a story algorithms start noticing. Swerving past cars too close could become data points. Parked vehicles blocking sidewalks or bus lanes? Those spots may soon get logged without human eyes spotting them.

A person behind the wheel may access the system using a special phone app. From there, they can look up previous infractions, see how their present driving performance ranks, or get tips meant to shape better routines on the road. Instead of leaning solely on fines, it pushes understanding and thoughtful choices while operating a vehicle.

Over months or years, numbers start showing what truly happens on roads. When past crashes and rule breaks get studied, problem spots begin to surface. Where danger shows up again and again, changes slowly take shape. Lights may shift timing, lanes can widen, cameras sometimes appear at corners. Slow shifts like these aim to stop harm before it unfolds. What once looked normal starts looking different through data-colored eyes.

One route worth exploring? Predictive analytics. When machine learning digests past traffic patterns, it spots places at higher risk for crashes or rule breaks. Traffic hubs might act before trouble strikes - tweaking light cycles, sending warnings, even boosting camera watch in shaky spots. Not waiting for sirens anymore.

When linked to a city's digital network, the system might gain new functions. Imagine cameras watching roads talk to traffic signals that change their rhythm when streets get busy. This kind of teamwork helps cars move faster through crossroads without cutting corners on security.

Starting with people on foot could take the technology further than cars. Later models may watch sidewalks, catching moments when walkers step into traffic or drivers ignore crosswalks. Safety in cities gains something real from this kind of reach. What used to focus only on roads begins feeling broader, deeper, once feet hit the sensors.

Out in the open, heavy usage leans on cloud power mixed with nearby device handling. Splitting number-crunching across gadgets on-site and remote servers keeps things moving when dozens of camera feeds pour in at once. Heavy city traffic won't slow it down because the load spreads out naturally.

VI. ACKNOWLEDGMENT

This research would not have been possible without the support of different individuals and organizations. We appreciate the help provided by the Vishwakarma Institute of Technology, Pune, in the form of several instruments needed to perform the research. Special recognition goes to Prof. Asmita Kalamkar for extending exceptional assistance during the study, for which we are thankful.

We owe great appreciation to our classmates for helping us train the model and find its accuracy, which helped us comprehend real issues better. Support and encouragement for this research also came from our colleagues, and for that, we are grateful, too.

REFERENCES

[1] S. Du, M. Ibrahim, M. Shehata, and W. Badawy, "Automatic License Plate Recognition (ALPR): A State-of-the-Art Review," IEEE Transactions on Circuits and Systems for Video Technology, vol. 23, no. 2, pp. 311–325, 2013.

[2] R. Laroca, E. Severo, L. A. Zanlorensi, L. S. Oliveira, G. R. Gonçalves, W. R. Schwartz, and D. Menotti, "A Robust Real-Time Automatic License Plate Recognition Based on the YOLO Detector," International Joint Conference on Neural Networks (IJCNN), pp. 1–10, 2018.

[3] R. Laroca, E. Severo, L. A. Zanlorensi, L. S. Oliveira, G. R. Gonçalves, W. R. Schwartz, and D. Menotti, "An Efficient and Layout-Independent Automatic License Plate Recognition System Based on the YOLO Detector," IEEE Intelligent Transportation Systems Conference (ITSC), pp. 1–8, 2019.

[4] R. Adak, S. Saha, and S. Bhattacharya, "Automatic Number Plate Recognition Using YOLOv3-CNN," International Journal of Computer Applications, vol. 183, no. 42, pp. 15–21, 2022.

[5] B. Satya, V. Kumar, and P. Reddy, "Optimized YOLOv8 for Automatic License Plate Recognition on Resource-Constrained Devices," Engineering, Technology & Applied Science Research, vol. 15, no. 2, pp. 11850–11858, 2025.

[6] M. Khan, A. Sharma, and R. Gupta, "Traffic Violation Detection Using Deep Learning," International Research Journal of Advanced Engineering and Management, vol. 4, no. 2, pp. 45–52, 2025.

[7] S. Verma, P. Singh, and R. K. Mishra, "Automatic Traffic Rules Violation Detection and Number Plate Recognition," International Journal of Scientific Research in Engineering and Management, vol. 9, no. 1, pp. 1–7, 2025