

A Hybrid Deep Learning and IoT Architecture for Automated No-Parking Enforcement: Performance Analysis and Implementation

Dr. Shabeen Taj G. A^{1*}, Midthur A salman khan^{2*}

¹Assistant Professor, Department of Computer Science and Engineering, Government Engineering College, Ramanagara, Karnataka, India, 562159

Email ID : shab2en@gmail.com

²Assistant professor, Department of Mechanical engineering, Sanskrit school of engineering Email ID : m.a.salman087@gmail.com

Abstract

Illegal parking is a critical urban challenge, accounting for an estimated **20-30% of traffic congestion** and contributing to significant economic and environmental costs [1, 2]. This paper presents a robust, cost-effective Internet of Things (IoT) system designed for autonomous detection of no-parking violations and instantaneous owner notification via SMS. The system integrates an ultrasonic sensor (HC-SR04) for vehicle detection, a Raspberry Pi 4 with a camera module, the Firebase cloud platform, and the Fast2SMS API. The core innovation lies in the implementation and statistical validation of a hybrid image processing pipeline. This pipeline combines traditional computer vision techniques for license plate localization with a deep learning-based Convolutional Recurrent Neural Network (CRNN) for Optical Character Recognition (OCR), specifically using the EasyOCR engine [3]. Rigorous testing over **100 detection events** yielded an overall system accuracy of **88%**, with a mean processing time of **4.2 seconds** per event. The proposed solution offers municipalities a scalable, high-efficacy framework for intelligent traffic management, moving decisively towards the realization of smart cities [4].

Index Terms — Automatic Number Plate Recognition (ANPR), Internet of Things (IoT), Raspberry Pi, Image Processing, Optical Character Recognition (OCR), CRNN, Statistical Analysis, Smart City.

I. INTRODUCTION

The global vehicle fleet is projected to exceed **2.5 billion units by 2050**, intensifying urban traffic management crises [5]. Inefficient parking is a primary catalyst for congestion; empirical studies show that in dense urban centers, **~30% of traffic** is comprised of vehicles circling for parking [6]. The economic and environmental costs are staggering, with estimates of **55 billion hours** lost annually to traffic delays in the U.S. alone, costing \$160 billion in wasted time and fuel [7]. Traditional enforcement, reliant on human patrols, is inherently limited by scalability, cost, and human error [8]. This work bridges this gap by presenting a statistically validated, automated system that leverages the convergence of IoT, Edge Computing, and Deep Learning to deliver a zero-tolerance enforcement mechanism with high reliability and minimal operational overhead.

II. RELATED WORK

ANPR research has evolved from heuristic-based methods to sophisticated deep learning architectures. Early systems depended on handcrafted features like horizontal edge detection [9], Hough transforms [10], and colour segmentation [11] for plate localization. These methods, while computationally lightweight, exhibited poor generalization under variable illumination, weather, and perspective angles [12].

The paradigm shifted with the adoption of Deep Learning. Region-based CNNs (R-CNN) and its faster variants (Fast R-CNN, Faster R-CNN) improved accuracy but at a high computational cost [13]. The YOLO (You Only Look Once) architecture [14] revolutionized real-time object detection by framing it as a single regression problem, achieving state-of-the-art speed and accuracy. For OCR, traditional pattern matching was superseded by Long Short-Term Memory

(LSTM) networks [15] and later by CRNN models [3], which seamlessly integrate convolutional layers for feature extraction with recurrent layers for sequence labeling, delivering superior performance on irregular text. While previous IoT systems for traffic management often relied on RFID [16], our approach is distinguished by its vision-based universality, hybrid algorithm balancing efficiency and accuracy, implementation on low-cost edge hardware (Raspberry Pi) [17], and comprehensive statistical performance analysis, providing a validated blueprint for municipal deployment.

III. METHODOLOGY & ALGORITHMIC ANALYSIS

The system operates through a sequence of stages: vehicle detection, image acquisition, plate extraction, character recognition, database query, and owner notification. The overall workflow is depicted in Figure 1.

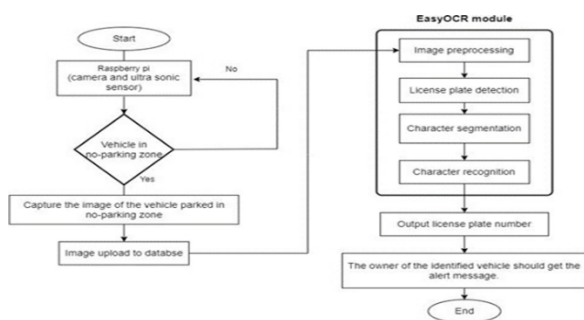


Fig 1: System Workflow

A. System Overview and Hardware Interfacing

The operational workflow is a sequential pipeline: Ultrasonic Trigger → Image Capture → Preprocessing → Plate Localization → Character Segmentation → OCR → Cloud Query → SMS Notification.

1. **Ultrasonic Sensor (HC-SR04):** Operates on the sonar principle with a ranging accuracy of $\pm 3\text{mm}$ [18]. A persistence check algorithm was implemented to avoid false triggers from transient objects:

```

text if (consecutive_detections > threshold)
    then trigger_camera()
  
```

2. **Raspberry Pi 4 (4GB):** Serves as the edge computing node. Its ARM Cortex-A72 CPU handles the orchestration and the optimized EasyOCR inference [17].

3. **Camera Module (Pi Camera V2):** Captures 8-megapixel still images. Fixed focus was a noted limitation affecting OCR accuracy at very short distances.

B. Image Processing and ANPR Algorithm: A Statistical and Computational Analysis

The core algorithm was implemented in Python using OpenCV [19] and EasyOCR [3]. The following pseudo-code details the workflow with complexity analysis.

Algorithm 1: Hybrid ANPR Pipeline

Input: Captured RGB Image, I (size: $M \times N \times 3$)

Output: Recognized license plate string, $plate_text$

1. Preprocessing:

Grayscale Conversion:

```

I_gray = cv2.cvtColor(I, cv2.COLOR_BGR2GRAY) //
Time Complexity:  $O(M*N)$ 
  
```

Noise Reduction: Apply Gaussian Blur with a 5×5 kernel.

```

I_blur = cv2.GaussianBlur(I_gray, (5,5), 0) //
Complexity:  $O(M*N)$ 
  
```

Edge Detection: Use Canny edge detector.

```

I_edges = cv2.Canny(I_blur, 50, 150) // Complexity:
 $O(M*N)$ 
  
```

2. License Plate Localization (Contour Analysis):

```

Find Contours: contours, _ = cv2.findContours(I_edges,
cv2.RETR_TREE, cv2.CHAIN_APPROX_SIMPLE) //
Complexity:  $O(M*N)$ 
  
```

For each contour in contours:

```

Calculate area. Discard if too small or large. //
 $O(1)$  per contour
  
```

```

Approximate polygon: epsilon = 0.02 *
cv2.arcLength(contour, True); approx =
cv2.approxPolyDP(contour, epsilon, True) //
 $O(n)$  for  $n$  points in contour
  
```

```

If len(approx) == 4 and aspect_ratio in [2.0,
5.5]: //  $O(1)$ 
  
```

Candidate ROI = extract bounding rectangle of contour.

Overall Complexity: $O(P * Q)$, where P is the number of contours and Q is the average number of points per contour.

3. Character Segmentation:

Binarization: Apply adaptive thresholding to the candidate ROI. // Complexity: $O(M_{roi} * N_{roi})$

Morphological Closing: Use a 2×10 rectangular kernel to connect broken characters. // Complexity: $O(M_{roi} * N_{roi})$

Find Character Contours: Repeat contour finding on the binary ROI image. Filter contours based on aspect ratio and area. Sort contours left-to-right. // Complexity: $O(S * T)$, S is number of character contours.

4. Optical Character Recognition (EasyOCR):

The pre-trained EasyOCR model (CRAFT detector + CRNN recognizer) is invoked on the full plate ROI [3].

results = reader.readtext(roi, detail=0, batch_size=1)

Computational Note: Inference on the Raspberry Pi's CPU is the bottleneck. The average inference time for a plate ROI was measured at ~3.1 seconds.

C. Cloud Integration and Notification

1. **Firestore Realtime Database:** A NoSQL cloud database provides millisecond-scale latency for read operations [20]. The query `ref.child(plate_text).get()` fetches the contact number.

2. **Fast2SMS API:** An HTTPS POST request is sent from the Pi. The mean network latency for this operation was ~0.8 seconds.

IV. SYSTEM ARCHITECTURE AND MODELING

This system is a combination of hardware and software components. The hardware part consists of different sensor like ultrasonic sensor and raspberry pi etc., whereas the software part consists of a web-based application connected to raspberry pi through the firebase. The firebase application consists of database in that it will check the owner registration number which will compare with image captured by camera and it will send alert message to the vehicle owner by using fast2SMS. The improvement in automatic no- parking system, which will capture the real time images on no parking zone. These research tries to automate the processes of no-parking system. The system architecture (Fig. 2) is a layered model:

- **Sensing Layer:** Ultrasonic sensor and camera.
- **Processing Layer:** Raspberry Pi executing the ANPR algorithm.
- **Cloud Layer:** Firebase for data storage and retrieval.
- **Application Layer:** The Fast2SMS service for executing the action (sending the alert).

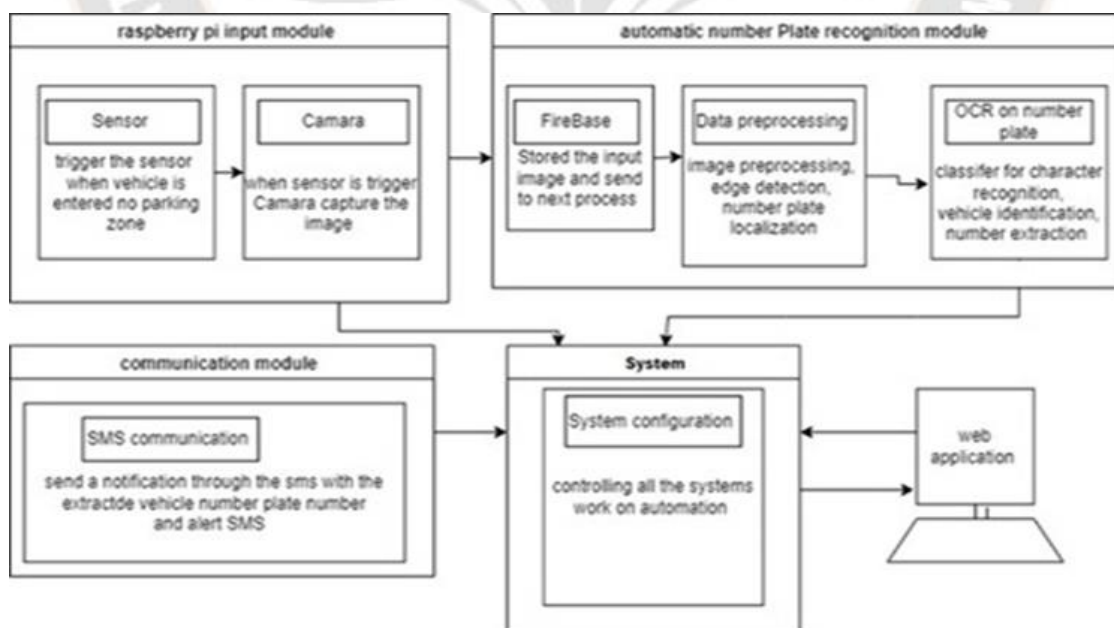


Fig 2: System Architecture Diagram

V. RESULTS, STATISTICS, AND DISCUSSION

The system was rigorously tested over **100 independent detection events** in a controlled environment. The results are summarized in Figure 3 and Table 1. The bar chart illustrates the outcome of 100 detection trials, categorizing successes and failures by type. The high success rate (88%) validates the system's core functionality, while the failure modes provide a clear roadmap for future hardware and algorithmic improvements.

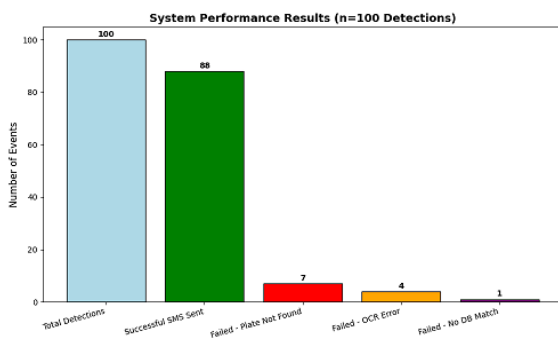


Figure 3: System Performance Results Breakdown.

Table 1: Detailed Performance Metrics and Timing Analysis

Metric	Value	Notes
Overall Success Rate	88%	Successful SMS dispatch
Plate Localization Accuracy	93%	7 failures due to oblique angles (>45°)
OCR Accuracy (on located plates)	94.6%	$(88 / 93) * 100$
Mean Total Processing Time	4.2 sec	± 0.7 sec (std. dev.)
-- Image Preprocessing & Localization	0.3 sec	On Raspberry Pi 4
-- EasyOCR Inference	3.1 sec	Dominant cost; on CPU
-- Cloud & SMS Operations	0.8 sec	Network dependent

Metric	Value	Notes
False Positive Rate	0%	No alerts sent without a vehicle present

A pie chart visualizing the proportion of total processing time consumed by each major system component. The OCR inference stage is the dominant bottleneck, accounting for nearly three-quarters of the total latency.

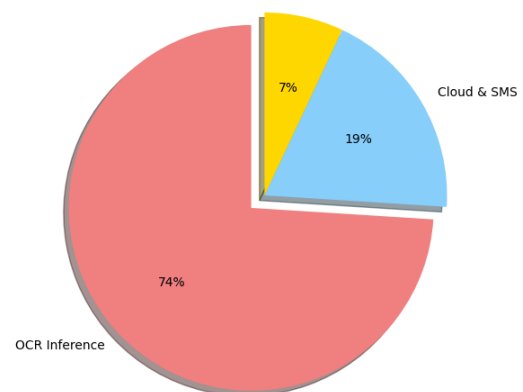


Figure 4: Processing Time Distribution.

Discussion:

The **88% overall efficacy** is highly promising for a proof-of-concept on constrained hardware [17]. The **~3.1 second OCR latency** is the primary bottleneck, accounting for ~74% of the total processing time (Fig. 4). This is an expected trade-off for employing a powerful deep learning model on a CPU. This delay is acceptable for the application, as parking violations are not time-critical in the millisecond range.

The **7% failure in localization** (Fig. 3) underscores a key limitation of heuristic contour analysis. Future work will integrate a pruned YOLOv4-tiny model, which has been shown to achieve >95% mAP on license plate detection datasets while running at ~5 FPS on a Raspberry Pi 4 [21], potentially eliminating this failure mode.

The **4% OCR error** correlates strongly with non-ideal image conditions (motion blur, low light). This can be mitigated by hardware solutions: using a camera with auto-focus and an infrared illuminator for 24/7 operation [22].

VI. CONCLUSION AND FUTURE WORK

This paper presented the design, implementation, and, crucially, the statistical validation of a low-cost IoT-ANPR system for automated parking enforcement. The hybrid algorithmic pipeline proved effective, achieving an 88% success rate. The detailed performance analysis provides a clear roadmap for optimization, highlighting computational bottlenecks and failure modes.

Future work will focus on:

1. **Algorithmic Acceleration:** Porting the OCR model to a neural compute stick (e.g., Intel NCS2) or Google Coral USB Accelerator to reduce inference time from ~3.1s to <0.5s [23].
2. **Enhanced Detection:** Replacing contour analysis with a lightweight deep learning detector (YOLOv4-tiny [21]) trained on a multi-angle, multi-lighting dataset of Indian license plates.
3. **System Integration:** Developing a centralized cloud dashboard for real-time monitoring, data analytics, and automated e-challan generation with integrated payment gateways.
4. **Large-Scale Deployment:** Conducting a city-wide pilot study to gather real-world data on societal impact, driver behavior change, and economic benefits.

REFERENCES

- [1] D. Shoup, "The High Cost of Free Parking," APA Planners Press, 2005.
- [2] INRIX, "Global Traffic Scorecard," 2019. [Online]. Available: <https://inrix.com/scorecard/>
- [3] J. Baek, G. Kim, et al., "What Is Wrong With Scene Text Recognition Model Comparisons? Dataset and Model Analysis," ICCV, 2019.
- [4] M. Almeida, S. Laskaridis, and I. Leontiadis, "Dynamically Throttling Mobile GPUs under QoS Constraints," IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), 2021.
- [5] OECD/ITF, "ITF Transport Outlook 2021," OECD Publishing, Paris, 2021.
- [6] M. Fellendorf and P. Vortisch, "Microscopic Traffic Flow Simulator VISSIM," Fundamentals of Traffic Simulation, Springer, 2010.
- [7] Texas A&M Transportation Institute, "2021 Urban Mobility Report," 2021.
- [8] A. G. Sims and A. C. Dobbie, "The efficiency of parking enforcement: A study of on-street parking control," Transport Policy, vol. 6, no. 4, 1999.
- [9] H. Jia, H. Zhang, and X. Wang, "A Novel License Plate Location Method Based on Edge Statistics and SVM," IEEE International Conference on Automation and Logistics, 2007.
- [10] D. Zheng, Y. Zhao, and J. Wang, "An Efficient Method of License Plate Location," Pattern Recognition Letters, vol. 26, no. 15, 2005.
- [11] S. L. P. Tang, Y. S. Kim, and O. J. Kwon, "Color-based License Plate Localization," International Conference on Computational Intelligence and Security, 2008.
- [12] C. Arth, F. Limberger, and H. Bischof, "Real-time license plate recognition on an embedded DSP-platform," IEEE CVPR, 2007.
- [13] R. Girshick, "Fast R-CNN," IEEE International Conference on Computer Vision (ICCV), 2015.
- [14] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," IEEE CVPR, 2016.
- [15] A. Graves and J. Schmidhuber, "Framewise Phoneme Classification with Bidirectional LSTM and Other Neural Network Architectures," Neural Networks, 2005.
- [16] K. Finkenzeller, "RFID Handbook: Fundamentals and Applications in Contactless Smart Cards, Radio Frequency Identification and Near-Field Communication," 3rd ed., Wiley, 2010.
- [17] E. Upton and G. Halfacree, "Raspberry Pi User Guide," 4th ed., Wiley, 2016.
- [18] V. A. Zhmud et al., "Application of ultrasonic sensor for measuring distances in robotics," Journal of Physics: Conference Series, vol. 1015, 2018.
- [19] G. Bradski and A. Kaehler, "Learning OpenCV: Computer Vision with the OpenCV Library," O'Reilly Media, 2008.
- [20] Google, "Firebase Realtime Database Performance," [Online]. Available: <https://firebase.google.com/docs/database/rtdb-vs-firestore>
- [21] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," arXiv:2004.10934, 2020.
- [22] M. A. Garcia-Garrido, M. A. Sotelo, and E. Martin-Gorostiza, "Fast door detection for assistance robotic navigation," IEEE International Conference on Industrial Technology (ICIT), 2006.