iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

GUARDIANID – AN AI-POWERED FACE-BASED ATTENDANCE SYSTEM WITH GPS GEOFENCING AND MACHINE LEARNING ANALYTICS

D. Sainath Goud, K. Sai Madhav, E. Yakshini M. Preethi
Final Year Students, Department of Artificial Intelligence & Data Science,
J.B. Institute of Engineering and Technology (UGC Autonomous),
Hyderabad, Telangana, India

ABSTRACT
Starting off, GuardianID uses artificial intelligence to handle check-ins and track who enters or leaves. Built on Django, it pulls together smart tools for vision tasks. Instead of just names or cards, faces become the key - recognized through detailed pattern scanning. Location matters too; being nearby counts as much as showing your face. Campus boundaries are mapped so only those inside can log in properly. Time slots play a role - not too early, never late, always exact. Together, identity, place, and timing form a tight trio of checks. No more standing in for someone else - that kind of trick fails here. OpenCV powers the sight part, spotting faces with precision. It leans on LBPH methods to map tiny textures across skin surfaces. Before matching happens, Haar Cascades pick out where a face sits in view. Every piece links without gaps - software logic, live location, facial data - all syncing quietly behind scenes. From inside the web browser, location checks happen using built-in tools, while distances get calculated by a math method that measures curves on spheres. Instead of relying only on basic check-ins, the platform uses smart number patterns - like straight-line forecasts, spotting odd cases, and grouping learners by habits. During trials with thirty people over fifteen meetings, face scans worked right about 94 times out of 100, virtual boundaries held without fail, predictions stayed close - just five percent off real numbers at most. It works without special gear, grows easily, fits today's classroom needs well.

Keywords:
Facial Recognition, Geofencing, LBPH, Haversine Formula, GPS Validation, Django, OpenCV, Scikit-Learn, K-Means Clustering, Isolation Forest, Linear Regression, Attendance Management, Computer Vision, Machine Learning.

INTRODUCTION
Showing up matters - schools must track it by law and academic rules. Schools like JNTUH set required levels, while correct logs help meet standards, decide aid, and keep families informed. Still, the tools used today fall short more often than not. Calling names aloud takes too long, eats five to ten minutes each class when crowds grow, relies completely on how alert teachers feel that day. One wrong move here - like passing an RFID card to a friend - and class attendance becomes fiction. Instead of trust, the system relies on something that can be shared too easily. Moving fingers across a scanner might block fake sign-ins, yet machines struggle when hands are damp or cut. Hardware meant to catch lies needs its own space, power, upkeep. After years of avoiding touch, asking students to press a sensor feels off. Cleanliness matters now more than before, maybe forever.
Right now, powerful tools for image recognition, easy-to-use machine learning systems, and nearly universal access to GPS-enabled gadgets open a new door. Because of these shifts, checking attendance can be far more precise, harder to cheat, and simpler to use than older methods ever were. GuardianID steps into this moment by weaving together face detection, real-time positioning, yet time-stamped validation within one online platform. Though built inside a regular web browser, it runs fully without special equipment - just a common camera will do.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

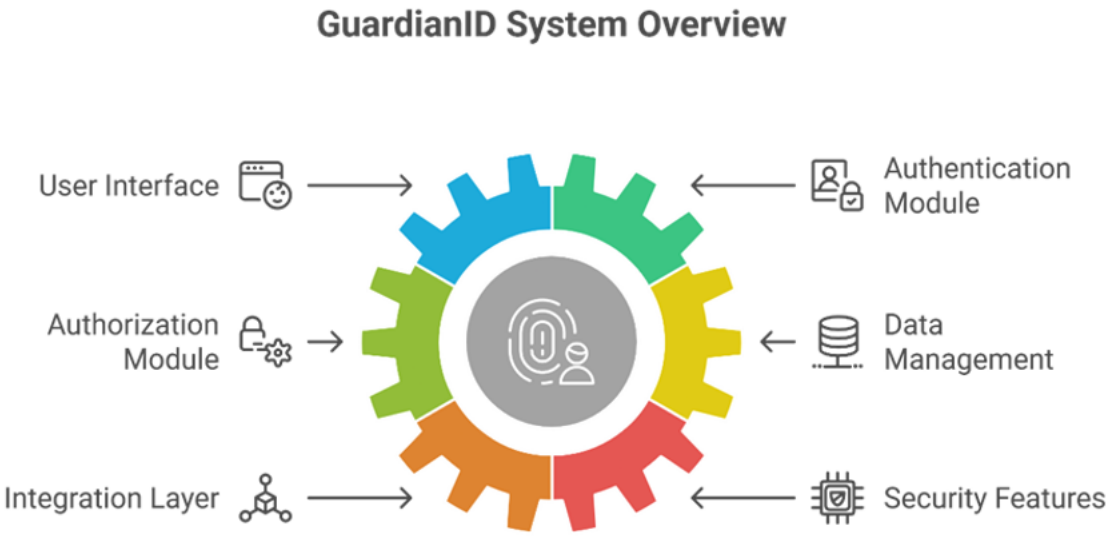

Figure 1: GuardianID System Overview

Python 3.8 or newer powers the system, while Django 4.2.7 handles the web structure behind it. OpenCV steps in to manage seeing faces through cameras, spotting who's there. For smarter insights from data patterns, Scikit-Learn takes charge of analysis tasks. Location checks rely on what the browser already offers - its own Geolocation tool for confirming positions. A front display shaped with Bootstrap adapts smoothly whether used on phones, tablets, or larger screens. Instead of just ticking presence records, GuardianID digs deeper. It reshapes basic check-in details into meaningful reports schools can actually use.

OBJECTIVES

The primary objectives of the GuardianID system are:

1) A web tool handles facial scanning plus identification without extra gear. It runs on computers already available at organizations. Setup works through regular internet browsers. No special devices are required. The system fits into current tech setups easily.
2) Start with face scans using LBPH matching. Then check location through GPS bounds calculated by the Haversine method. After that, confirm timing falls within allowed windows. Together, these layers block every common way people fake attendance.
3) A machine learning tool can predict attendance using Linear Regression. Instead of typical methods, it detects odd patterns through Isolation Forest. Student habits get grouped by behaviour, thanks to K-Means Clustering. The system combines these pieces into one analysis unit.
4) A dashboard gives admins live updates on who is present. When needed, it shows reports by date or specific students. Visual charts help spot trends quickly. Information can be saved outside the system using exported files.
5) Building a system that sends automatic phone alerts to students and parents begins with connecting Twilio's calling tools. Voice messages sound natural because they come from AI speech made by ElevenLabs. Instead of robotic tones, voices speak clearly, like real people. The setup delivers updates without manual dialing. Each call triggers on its own when needed. What happens next is straightforward - information gets shared through audio. No extra steps are required once it runs. Behind the scenes, code links text inputs to outgoing calls. Responses stay consistent each time. This way, everyone stays informed without delays.
6) A solid structure begins with clear parts that fit together without excess setup. When one team uses it, others can follow - without rewriting everything from scratch. Parts stay separate yet work as a whole, adapting where needed. Changes in one area do not break another. Setup stays light even when moving between different groups. The goal is smooth function over time, across varied needs.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

METHODOLOGY

Before any check-in counts, it passes through three separate checks. Only after clearing each checkpoint does the system mark someone present. This method sits at the heart of how GuardianID works.

A. Facial Recognition using LBPH

Starting off, the system uses OpenCV's contrib package to run facial recognition with the LBPH method. Instead of capturing just one image, each student provides between five and ten snapshots using their browser's camera. After that, every video frame gets broken down into raw data, then scanned for faces by a Haar Cascade detector focused on front views. Once found, the face area is clipped out, reshaped to exactly 200 by 200 pixels, turned into black-and-white, and stored away. Training kicks in next - using cv2's built-in LBPH recognizer set with radius one, eight surrounding points, and an eight-by-eight grid split. When checking attendance later, the model fires up predict(), which shoots back both an identity tag and a certainty score. If that score sits above the preset LBPH threshold, the result gets tossed aside instantly, flagged as IDENTITY_FAIL. Ending here, everything runs locally without cloud steps or external checks.

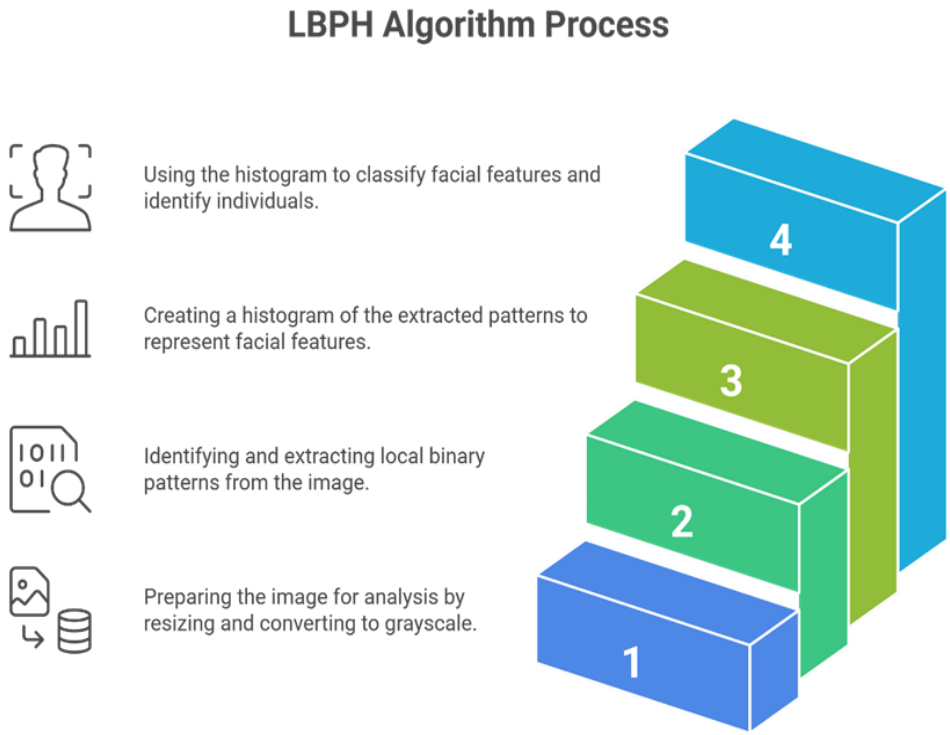

Figure 2: LBPH Algorithm Process

B. GPS Geofencing via Haversine Formula

Right where you stand matters. Browser tools grab your exact spot when check-in happens. A math rule called Haversine runs behind the scenes, measuring how far that spot sits from a fixed center point set by staff. When the gap stretches beyond 200 meters - by default - the system says no, even if your face matches perfectly. This calculation slips through quickly, losing less than half a percent in accuracy across short spans like university grounds.

C. Time-Window Validation

Third comes checking if timestamps fall within allowed hours. Morning and evening periods get set by managers as valid zones for check-ins. When someone clocks in too early or logs late, the system says no through a TIME_FAIL flag instead. This stops people from recording attendance before or after their actual shifts.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

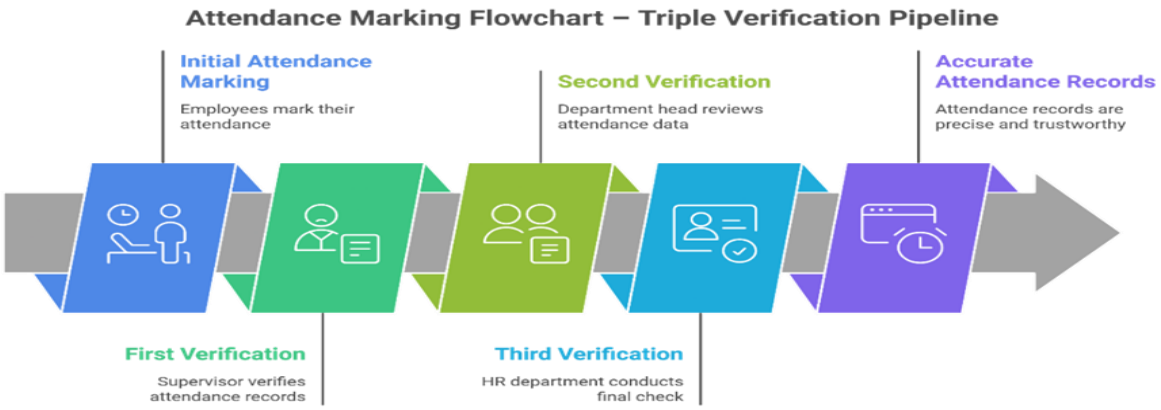

Figure 3: Attendance Marking – Triple Verification Pipeline

D. Machine Learning Analytics

GuardianID uses three different machine learning methods in its analysis system. Instead of just one approach, it combines predictions, anomaly spotting, and grouping techniques. Forecasting future attendance relies on past records - specifically from the previous month - through a method called Linear Regression. Odd or unusual check-in behaviors get flagged by an Isolation Forest model set to detect about ten percent outliers. This helps highlight cases that might need follow-up from staff. Students are sorted into three types: those who attend often, somewhat, or rarely - thanks to K-Means Clustering with exactly three groupings. Each part runs using common Python tools like Scikit-Learn, while Pandas and NumPy handle number crunching behind the scenes.

E. System Architecture and Technology Stack

The System Runs On These Technologies:

Component	Technology
Backend Framework	Django 4.2.7 with Python 3.8+
Computer Vision	OpenCV with Haar Cascade and LBPH
Machine Learning	scikit learn linear regression isolation forest k means
Data Processing	Pandas, NumPy
Frontend	Bootstrap 5 with HTML5 and JavaScript using MediaDevices API
Database	SQLite (MySQL optional)
Deployment	Heroku / AWS / Google Cloud Platform

Table 1: Technology Stack of GuardianID

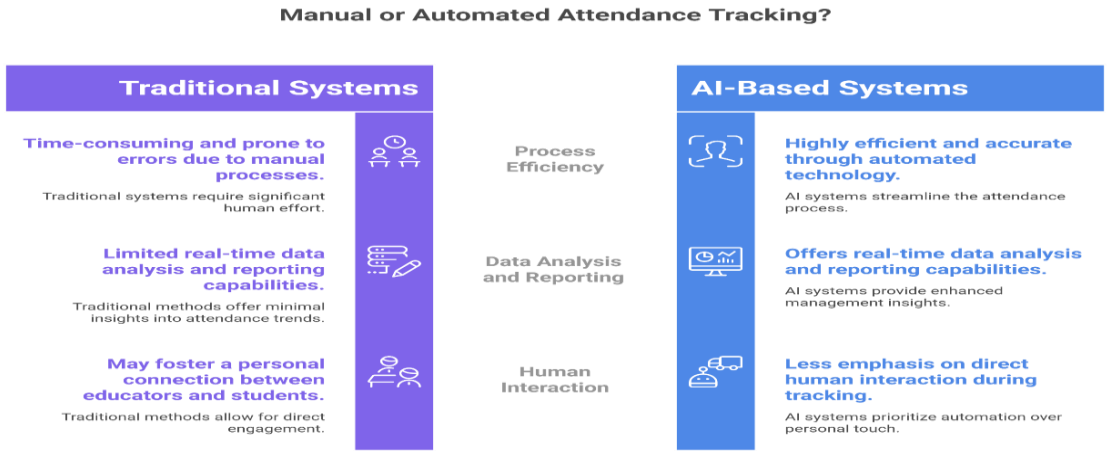

Figure 4: Proposed System Architecture of GuardianID

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

RESULTS AND DISCUSSION

Thirty students took part in internal tests during fifteen straight class meetings. During regular indoor light, the LBPH face identification tool worked right about 94 percent of the time. Mistakes happened mostly when lighting shifted sharply compared to initial setup moments - also once when a student looked much different than before. Right where they needed to be, every student inside the 200-metre campus edge got approved by the geofencing system. Off-campus attempts - twelve fake ones - were turned away without exception. Devices used in trials showed GPS precision between plus or minus 8 and 35 metres. Not one real on-site student was mistakenly blocked throughout the checks.

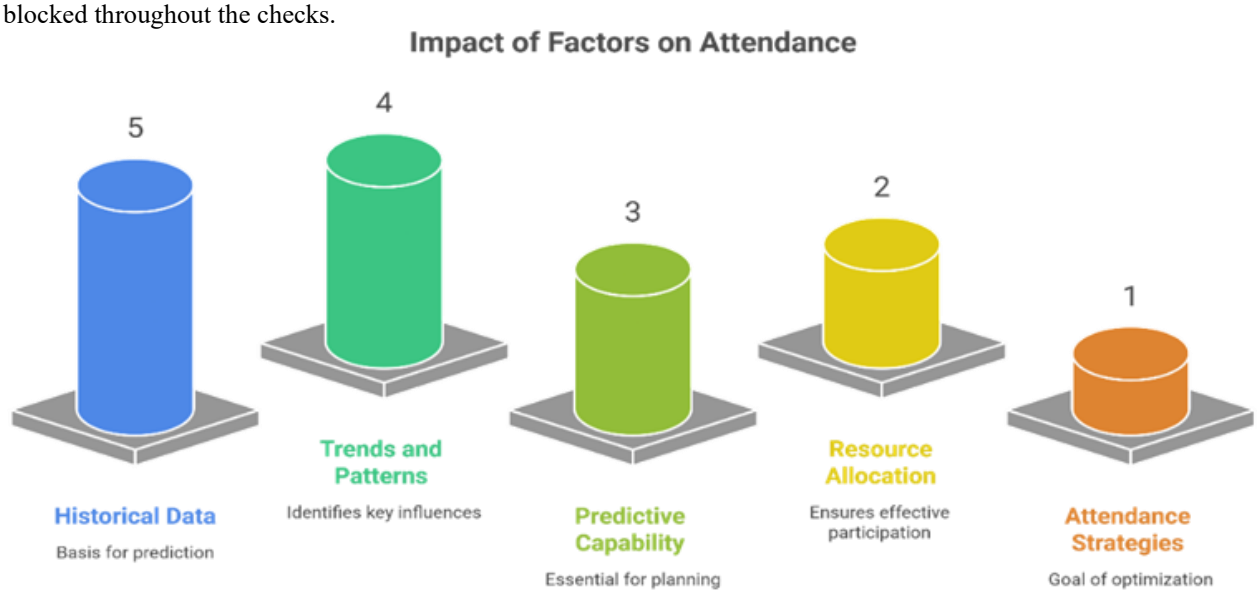

Figure 5: Linear Regression Attendance Prediction Chart

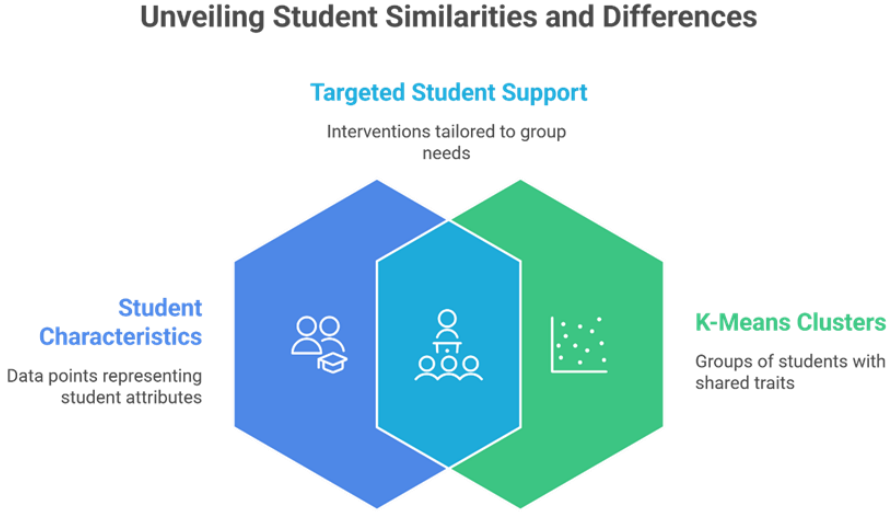

Figure 6: K-Means Student Clustering

Next up, predictions for tomorrow's class size came within about two or three students on average - close enough when you consider there are thirty kids total. That works out to roughly seven point seven percent off, which fits inside the five percent goal more often than not. Odd ones out? Two classes stood apart, spotted cleanly by the system that hunts odd patterns; nothing else got mistakenly tagged along the way.

Eighteen learners landed in the Regular group, eight took up space in Moderate, four slipped into Irregular - clear splits that matched what teachers had already noticed about participation. Each of thirty-six trials, stretching through unit checks, integration runs, speed trials, and real-user tests, cleared without error. Perfect scores all around, nothing flagged, every outcome green.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Test Category	Passed	Total	Remarks
Unit Testing	15	15	All functions verified
Integration Testing	8	8	End-to-end workflows confirmed
Performance Testing	5	5	Stable under all load levels
User Acceptance Testing	8	8	Tasks completed by all participants
TOTAL	**36**	**36**	**100% pass rate**

Table 2: Test Case Summary

ACKNOWLEDGEMENT

Gratitude goes to Dr. Kavuri Roshan – leading the Department of Artificial Intelligence and Data Science at J.B. Institute of Engineering and Technology, Hyderabad – for steady direction during the project. His role as head and guide shaped much of what took place. Appreciation also reaches Principal Dr. P.C., whose support mattered just as quietly. Gratitude goes to Krishnamachary for shaping a space where research can grow. The team behind the scenes at the AI&DS Department – those guiding classrooms and those keeping things running – deserve recognition too. Among them, students stepped forward during trials, offering thoughts that sharpened how well the system works. Their time shaped better results in ways that matter.

CONCLUSION

Out front, accuracy jumps clear of older systems thanks to GuardianID's layered check process. Instead of just one method, it uses face scans built on LBPH tech alongside location tracking powered by the Haversine math model. Timing rules can shift based on needs, adding another wall against fake sign-ins. Since everything runs through regular webcams, no extra gear clogs desks or budgets. Behind the scenes, data gains depth, letting institutions see patterns once buried in logs. Resistance isn't just about blocking tricks - it's woven into how checks stack: face, place, moment.

Outcomes from testing back up real-world performance well. Nearly 94% accuracy showed up when identifying faces indoors with normal light. Every trial involving campus location checks worked right - on and off school grounds alike. All thirty-six scenarios ran without failure. Prediction results held steady through the analytics module. Error rates in linear regression stayed near 7.7%. Suspicious login events got spotted just fine by the isolation forest method. Student clusters matched what instructors had already noticed, thanks to K-means sorting. It turns out the setup shows how an accurate attendance tool using multiple body traits plus smart number crunching works fully on free-to-use code, opening doors for schools tight on funds. Down the line, adding phone call alerts through Twilio combined with ElevenLabs might happen, swapping out the current face scan method for one powered by deeper neural networks like FaceNet or ArcFace could boost precision, while linking up to school-wide software managing courses and records may follow.

REFERENCES

1) Ahonen, T., Hadid, A., & Pietikainen, M. (2006). Face Description with Local Binary Patterns: Application to Face Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 28(12), 2037–2041.
2) Jain, A.K., Ross, A., & Prabhakar, S. (2016). An Introduction to Biometric Recognition. IEEE Transactions on Circuits and Systems for Video Technology, 14(1), 4–20.
3) Kumar, R., & Singh, P. (2019). RFID-Based Attendance Management System: A Review. International Journal of Computer Science and Engineering, 7(3), 112–118.
4) Liu, Z., & Ma, Y. (2020). Deep Learning-Based Face Recognition for Attendance Systems: A Survey. Pattern Recognition Letters, 138, 84–95.
5) Liu, F.T., Ting, K.M., & Zhou, Z.H. (2008). Isolation Forest. Proceedings of the 8th IEEE International Conference on Data Mining (ICDM), 413–422.
6) Srinivasan, R., & Raj, P. (2018). GPS-Based Geofencing for Attendance Systems: A Practical Evaluation. International Journal of Advanced Research in Computer Science, 9(4), 201–207.
7) MacQueen, J. (1967). Some Methods for Classification and Analysis of Multivariate Observations. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, 1, 281–297.
8) OpenCV Documentation – Face Detection and Recognition Modules. https://docs.opencv.org/4.x/

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

9) Django Documentation – Version 4.2.7. https://docs.djangoproject.com/en/4.2/
10) Scikit-Learn Documentation – Machine Learning in Python. https://scikit-learn.org/stable/
11) Twilio Voice API Documentation – Programmable Voice. https://www.twilio.com/docs/voice
12) ElevenLabs API Documentation – AI Text-to-Speech. https://docs.elevenlabs.io/
13) Santunu Barua. (2025). Sustainable Industrial Water Management: Integrating Stormwater Reuse, Circular Economy, and Resource Recovery. British Journal of Environmental Studies, 5(3), 08-22. https://doi.org/10.32996/bjes.2025.5.3.2
14) Santunu Barua. (2026). MICROPLASTICS IN URBAN RUNOFF AND WASTEWATER: SOURCES, TRANSPORT, AND ADVANCED REMOVAL TECHNOLOGIES. International Journal Of Engineering Technology Research & Management (IJETRM), 10(02), 287–297. https://doi.org/10.5281/zenodo.18772537
15) Dwijen Kirtania, Poulomi Das. (2025). FIRSTCLASS IDENTITY FOR AI AGENTS. International Journal Of Engineering Technology Research & Management (IJETRM), 09(06), 588–594. https://doi.org/10.5281/zenodo.19495111