International Journal of Safety and Security Engineering
Vol. 15, No. 10, October, 2025, pp. 2069-2079
Journal homepage: http://iieta.org/journals/ijsse

An IoT-Driven Intelligent Street Lighting System with Integrated Severity-Based Fault Classification

Sreelatha R[1*], Shubha Rao V[1], Mahalakshmi B S[1], Sindhu K[2]

[1] Department of Information Science, BMS College of Engineering, Bengaluru 560019, India
[2] Department of Computer Science (Data Science), BMS College of Engineering, Bengaluru 560019, India

Corresponding Author Email: sreelathar29@gmail.com

Copyright: ©2025 The authors. This article is published by IIETA and is licensed under the CC BY 4.0 license (http://creativecommons.org/licenses/by/4.0/).

https://doi.org/10.18280/ijsse.151010

Received: 24 August 2025
Revised: 14 October 2025
Accepted: 25 October 2025
Available online: 31 October 2025

Keywords:
Internet of Things, fault detection streetlight automation, real-time monitoring, smart city, artificial intelligence, machine learning

ABSTRACT

Street lighting faults can compromise urban safety if not detected and addressed promptly. This paper proposes an Internet of Things (IoT)-based streetlight monitoring architecture integrated with an artificial intelligence (AI)-driven severity classification model. The system continuously captures lamp status and traffic density using IoT sensors and classifies fault severity by analyzing adjacent lamp failures, road illumination levels, and vehicle flow. A severity assessment algorithm is developed and evaluated using a simulated real-world streetlight dataset containing balanced and minority high-severity fault instances. Experimental results show an overall classification accuracy of 94%, with precision and recall above 90% for minority high-severity classes, demonstrating robust detection of critical outages. Confusion matrix analysis and performance comparisons further confirm the model's ability to differentiate urgent faults from lower-priority issues. The proposed framework supports intelligent streetlight maintenance by enabling timely identification of high-risk failures, contributing to safer and more efficient smart city infrastructure.

1. INTRODUCTION

Reliable street lighting is essential for traffic safety and crime prevention in modern cities. Numerous studies have shown that well-lit streets significantly reduce nighttime accidents and criminal activity. For example, insufficient lighting makes it difficult for drivers to see road hazards or pedestrians at night, increasing the risk of accidents. Similarly, there is a longstanding correlation between poor street lighting and higher crime rates; criminals often shift their activity to dark streets when streetlights are out. Ensuring that streetlights remain operational is thus a critical public safety concern. However, traditional maintenance of streetlights largely relies on passive approaches – waiting for citizens to report outages or conducting infrequent manual inspections [1]. This reactive model leads to lengthy downtimes for broken lights, especially in remote or low-traffic areas where issues might go unnoticed. These outages expose communities to avoidable dangers such as collisions or opportunistic crime.

Recent Internet of Things (IoT) and artificial intelligence (AI) advances offer a more proactive solution. Installing streetlights with IoT sensors and connectivity makes it possible to automatically detect lamp failures in real time and immediately alert city maintenance crews. Prior work has demonstrated the efficacy of IoT-based streetlight monitoring. For instance, Vamsi et al. [1] developed a centralized IoT system where sensor nodes on each streetlight transmit operational data to a cloud platform; a machine learning (ML) model (98.8% accuracy) then predicts each streetlight's status and fault condition. The predicted status and fault details are displayed on a web dashboard for officials, eliminating the need for citizens to report outages. Venu [2] proposed an automated street lighting system in which elevated lights installed along roads or walkways operate at predefined times. The system integrates LEDs with PIR, IR, and LDR sensors, controlled by a Raspberry Pi and powered by a solar module using an Arduino Uno to achieve efficient, motion-based, and energy-conserving illumination. Researchers demonstrated that this approach significantly expedites repairs and improves maintenance efficiency by providing precise, real-time knowledge of faults. Likewise, other researchers have used light sensors (e.g., LDRs) on streetlamps to automatically detect when a lamp fails to light up at night [3], enabling immediate alerts. These studies underscore the potential of IoT to transform streetlight maintenance from reactive to predictive and responsive.

Beyond simply detecting faults, an intelligent system should also assess how severe a streetlight outage is. Not all lamp failures are equally urgent: a single lamp out in a well-lit urban downtown might have minimal impact, whereas two or three consecutive lamps out on a suburban road could plunge the area into darkness. If an entire stretch of roadway is dimly lit due to multiple failures or intentional dimming, it may create hazardous conditions, inviting crime or accidents. The context – such as the level of pedestrian/vehicle activity and the location (remote versus central) – determines the risk posed by a lighting failure. For example, an outage on a busy highway or at an intersection with heavy nighttime traffic is a high priority because it can lead to collisions. Similarly, an outage in a sparsely populated rural area is concerning because darkness in isolated areas can "foster criminal activity" despite low traffic. Current smart city research highlights the need to integrate contextual data into infrastructure monitoring. In the case of streetlights, combining traffic sensors or city mobility data with lamp status can measure how critical each light is at a given time. A fault on a high-usage road or during a special event should be addressed sooner than on a rarely used street.

(a) (b) (c)

Figure 1. Street lane (a) with streetlight; (b) without streetlight with incoming traffic; (c) with no traffic

Figure 1 depicts different conditions of the same road with streetlight in power on mode when everything is visible. Second scenario where the streetlight is in power off mode, but light is there due to moving vehicles, and finally the darkness when no vehicles are moving, and the streetlight is in power off mode.

This work proposes an IoT-driven streetlight fault detection and severity classification system that addresses the above challenges. Each streetlight pole has sensors to monitor its operating state (e.g., voltage/current or light output) and possibly the ambient illumination. The devices are connected via a wireless network to a central server. We incorporate an ML model at the server that analyzes the streetlight's status and local vehicle flow data to classify the severity of any detected fault. The model is trained on features such as whether the lamp is off/dim when it should be on, how many neighboring lights malfunction, the traffic density around that location, and other contextual cues (time of day, weather, special events, etc.). The output is a severity level (e.g., Low, Medium, or High severity), which informs how urgent the maintenance response should be. High-severity alerts trigger immediate notifications to maintenance crews for rapid repair, whereas lower-severity issues can be scheduled in routine maintenance routes.

2. RELATED WORK

IoT-Enabled Smart Street Lighting Systems: Modern urban infrastructure has increasingly adopted IoT technologies to transform traditional street lighting into smart, connected systems. In a conventional setup, streetlights operate on fixed timers or manual control, requiring periodic human inspection for outages. IoT-based street lighting systems, by contrast, embed networked sensors and controllers in each lamp, enabling real-time remote monitoring and adaptive control. These systems can dynamically adjust illumination levels based on sensor inputs (e.g., ambient light, motion), substantially improving energy efficiency and road safety by providing lighting on demand. For example, Denardin et al. [4] demonstrated an intelligent street lighting control architecture where each lamp's status and power consumption are reported to a central platform for active monitoring. Pantoni et al. [5] described an edge-assisted network for streetlights allows facility managers to control lights and detect anomalies at scale remotely. This connectivity reduces manual maintenance effort and supports advanced features like scheduling, dimming, and fault alerts. The literature agrees that IoT connectivity is a foundational enabler for smart street lighting in future cities. However, deploying such IoT infrastructure city-wide is not without challenges: high upfront costs, interoperability issues, and the need for stable network coverage (especially in remote areas) remain significant concerns [6]. For instance, Dahan et al. [7] noted that smart lighting projects often face reliability issues in underdeveloped regions due to inconsistent internet connectivity and power supply, underscoring the importance of robust network design. Despite these challenges, IoT-driven lighting systems have laid the groundwork for more intelligent maintenance and control strategies, as discussed next.

Rule-based logic relies on fixed thresholds and therefore struggles in real-world environments where sensor readings are noisy, traffic patterns are highly non-linear, and illumination levels fluctuate due to weather or vehicle headlights. It also assumes complete and consistent data, which is often not the case in IoT deployments where missing or intermittent sensor inputs are common. In contrast, the ML model learns the variability within the data and becomes more robust to noise, missing values, and complex multi-feature interactions that rules cannot easily encode. Our results further show that ML reduces both false positives and false negatives by capturing subtle patterns beyond the deterministic logic, thereby providing more reliable severity predictions. Thus, the ML component adds significant value by improving robustness, handling imperfect data, and enhancing decision accuracy in complex urban conditions where static rule-based algorithms alone are insufficient.

Deterministic IF–THEN rules alone are often insufficient for streetlight monitoring because they cannot reliably handle sensor noise, fluctuating illumination from vehicle headlights, or weather-induced variations. Fixed thresholds (e.g., "light level < X") frequently misclassify normal conditions as faults or miss critical outages, especially in environments with non-linear and rapidly changing traffic patterns. These rules also lack the ability to account for contextual factors such as time of day, surrounding lamp conditions, or road usage intensity. ML overcomes these limitations by learning complex, non-linear relationships from data, filtering noisy signals, and producing more accurate severity assessments tailored to real-world conditions. Hence, ML enhances decision-making beyond what traditional rule-based approaches can achieve.

2.1 Fault detection and predictive maintenance approaches

Traditional streetlight maintenance: Crews respond to citizen complaints or periodic inspections to fix outages. This approach can lead to prolonged dark periods when lights fail, negatively impacting public safety and confidence [8]. Researchers have explored automatic fault detection using sensor data and predictive maintenance models to mitigate this. In IoT-enabled systems, each lamp can report its operational status (voltage, current, brightness, etc.), allowing immediate identification of failures or abnormalities [9, 10], For example, a telemetry manager for public lighting was developed that continuously monitors lamp health and flags faults in real time via a centralized dashboard. Such systems reduce the need for manual night patrols by generating automatic alerts when a bulb or driver malfunctions [11]. As an illustration, recent work [11] implemented a smart streetlight network where if any light goes faulty, it is detected without human intervention and reported through a web application. This immediate visibility into failures allows maintenance teams to respond faster, shortening dark outages.

Beyond detecting failures as they occur, an emerging trend is predictive maintenance – forecasting lamp failures before they happen. This is achieved by analyzing historical performance data to predict remaining useful life or failure probabilities. Accurate remaining-life prediction enables proactive replacement of lamps prior to burnout, thus eliminating unplanned outages. For instance, Segovia-Muñoz et al. [12] applied degradation models (LM-80/TM-21 standards) to LED streetlights and developed an exponential decay algorithm to estimate lumen depreciation and predict when a lamp will fall below acceptable brightness. By planning maintenance based on such predictions, cities can maximize equipment life and minimize downtime. Several studies report that predictive maintenance regimes, combined with IoT telemetry, drastically reduce maintenance costs and outage durations compared to reactive approaches. Nonetheless, implementing predictive maintenance requires reliable data collection over long periods and robust models; many municipalities are still in the early pilot stages of these technologies. The literature shows a clear evolution from reactive repairs to automated fault detection and prediction-driven maintenance in smart street lighting systems.

AI and ML Integration in Lighting Systems: Applying AI and ML techniques to streetlight data is a relatively recent area of research, but it is gaining momentum. Early smart lighting deployments focused on rule-based control and simple sensor thresholds (e.g., lights dim when no motion is detected). Now, with large volumes of sensor and usage data available, researchers are exploring data-driven AI models to optimize performance and detect complex fault patterns [9]. A systematic review [13] finds that the integration of AI in IoT-based lighting (often termed AIoT) is still in its infancy, with the first notable works appearing only around 2019. Their review highlights that most current smart lighting systems do not fully exploit ML and that incorporating edge-based ML could vastly improve adaptability and reliability.

One promising application of AI is anomaly detection in streetlight operations. Instead of waiting for a light to go completely dark, ML models can analyze subtle deviations in electrical parameters to predict failures or detect malfunctioning components. Śmiałkowski and Czyżewski [9] demonstrate this using energy consumption data from smart meters in a city lighting grid. Their system achieves real-time detection of anomalies such as lamp failures or schedule deviations by employing time-series forecasting (SARIMA models) and neural networks (LSTM). Notably, both statistical and deep learning approaches in that study enabled a self-learning, online fault detection mechanism deployable on edge devices, with the SARIMA model performing best for timely alerts. Another study by Chen et al. [14] developed a narrowband-IoT-based streetlight monitoring system that uses ML classifiers to distinguish regular operation from various fault conditions, improving detection accuracy over simple threshold methods. These works illustrate the potential of AI to enhance maintenance: instead of binary "on/off" sensing, smart algorithms can estimate the severity of degradation, filter out transient anomalies, and prioritize genuine issues for repair. Additionally, AI has been applied for predictive control. For example, the study by Sun et al. [15] built an ML-driven predictive model (using time series analysis) to forecast lighting needs and adjust lamp brightness pre-emptively, achieving significant energy savings while maintaining service quality.

Despite these advances, incorporating AI in municipal lighting is not yet widespread. Key barriers noted in the literature include the need for large labeled datasets for training, system complexity concerns, and legacy infrastructure integration. Nevertheless, the trajectory is clear: AI and data analytics are set to play an increasing role in making street lighting systems more intelligent, from fault diagnostics to adaptive illumination strategies.

Vehicle Presence and Traffic-Aware Lighting: Another vital aspect of smart street lighting is integrating traffic data – specifically, detecting vehicles or pedestrians to modulate lighting or assess maintenance criticality. In urban environments, traffic density varies widely; a streetlight on a busy highway serves far more road users at night than on a quiet lane. Numerous studies have proposed traffic-responsive lighting, where streetlights brighten or turn on only when vehicles or pedestrians are present and dim otherwise, thereby saving energy without compromising safety. Qaisar et al. [16] built an early prototype using infrared sensors to sense vehicle movement and then activating lights sequentially as vehicles passed by, which demonstrated the basic feasibility of on-demand illumination. In recent years, more sophisticated approaches have leveraged computer vision and radar. These systems highlight that coupling traffic sensing with lighting control can maintain safety (adequate illumination when cars or pedestrians are present) while optimizing energy use when roads are empty.

Beyond energy optimization, traffic data can inform maintenance decisions. A smart maintenance system might prioritize fixing a streetlight on a high-traffic thoroughfare over one on a seldom-used street since the impact of an outage differs. Some frameworks have begun to incorporate such considerations. For example, Gagliardi et al. [17] developed a multi-sensor streetlight node that monitors lamp electrical parameters and environmental and usage indicators, such as motion or traffic volume, to provide a holistic view of each lamp's importance and condition. In one case study, traffic-aware lighting reduced complaints, as lights on busy routes were rarely allowed to fail unnoticed. Moreover, adaptive systems can change lighting levels based on real-time traffic density. If a normally quiet street suddenly experiences heavy traffic (e.g., due to a detour), smart lights can increase brightness to enhance visibility. This on-demand adaptability improves efficiency and safety and is a key selling point of IoT-based street lighting. However, integrating continuous traffic monitoring comes with data and communication overhead. Recent works suggest using edge computing to process video or sensor data locally at the lamp post, which minimizes latency in reacting to vehicles and reduces bandwidth usage [9]. In summary, vehicle detection and traffic awareness have become integral to advanced streetlight systems, enabling context-aware operation where lighting and maintenance prioritize locations with more excellent human activity or vehicular flow.

Safety Implications and Severity Assessment in Lighting Failures: Street lighting is closely tied to public safety, and research indicates that insufficient lighting can elevate both accident and crime risks. Welsh and Farrington's meta-analysis of crime prevention (2008) reported that improved street lighting was associated with a significant (~20%) reduction in nighttime crime in the areas studied. More recent field experiments support this: a randomized trial in New York City public housing [18] found that adding temporary lighting led to a 36% reduction in serious nighttime outdoor crimes. Conversely, lighting outages tend to increase fear and potential opportunity for crime. City incident data show that residents are acutely sensitive to lighting failures – in New York City's 311 service request system (2010-2016), streetlight outages were the third most common complaint type (about 5% of ~15 million complaints). Maintaining consistent illumination is critical for perceived and actual safety. A dark street not only hampers visibility for drivers and pedestrians, raising the risk of accidents, but also may become a locus for opportunistic crime. This is especially true in remote or sparsely populated areas, where a single failed light can leave a long dark stretch with no nearby light sources.

Literature on smart lighting increasingly calls for severity-aware maintenance strategies that prioritize repairs based on risk factors. Key factors identified include the number of adjacent lights out, the location's characteristics (crime rate, traffic volume), and the availability of alternate lighting. If consecutive streetlights fail, the combined dark span is more significant, compounding safety risks exponentially. However, most technical studies have focused on detecting faults rather than ranking their urgency. City engineering guidelines sometimes note that multiple contiguous outages should be the highest priority (as seen in municipal maintenance policies), but academic research on automated severity scoring is sparse. One relevant thread is in power distribution networks, where faults are classified by criticality to prioritize grid repairs; analogous concepts are now applied to smart city lighting. For instance, Yoomak and Ngaopitakkul [19] conducted a feasibility analysis of solar street lighting and noted the importance of fast response in critical outage scenarios to ensure public safety – essentially highlighting that not all lamp failures are equal in impact. Despite such insights, there is a gap in the literature: few, if any, AI-driven streetlight maintenance models explicitly combine crime data, traffic density, and outage patterns to compute a composite severity index for each failure incident. This represents an open research challenge and an opportunity for innovation, as our proposed system aims to address.

Research Gaps and Emerging Challenges: The above review shows that while significant progress has been made in IoT-based street lighting and smart maintenance, several gaps remain for future work. First, the integration of multi-modal data for severity analysis is largely missing. Prior studies treat energy efficiency, fault detection, and public safety independently. An ingenious maintenance system would ingest diverse data – lighting status, sensor readings, traffic counts, neighborhood crime statistics – and intelligently prioritize interventions. Such a holistic approach is only now becoming feasible with AIoT advancements. Some research notes that the convergence of AI with IoT for lighting is nascent, and many possibilities (like advanced decision-making algorithms at the edge) remain unexplored. Our review found no published framework that, for example, increases a fault's priority because it occurs in a high-crime, low-traffic area on a dark stretch – a scenario our proposed system explicitly targets.

Second, scalability and real-time performance pose challenges. City-wide deployments may involve tens of thousands of lights. Detecting faults and computing severity in real time requires efficient distributed computing. Edge computing has been suggested as a solution to reduce latency and network dependency, but deploying AI models on edge devices (like lamp-mounted controllers) must overcome constraints in power and processing. Early experiments with edge ML in lighting (e.g., using microcontrollers for anomaly detection) show promise, yet issues of the edge devices' model updating, security, and maintenance are open questions. Cybersecurity is another concern: as streetlights become networked endpoints, they could be targets of hacking or data manipulation, which can have dangerous implications if not addressed (e.g., false outage reports or malicious switching off of lights). Ensuring robust encryption and fail-safes in the maintenance system is thus crucial, though few papers in this domain delve into security beyond basic network encryption.

Lastly, there is the challenge of interdisciplinary integration. Smart lighting maintenance does not exist in a vacuum; it overlaps with urban planning, law enforcement, and transportation management. Lighting improvements have been linked to lower crime and increased community confidence and economic activity after dark. Thus, city authorities are interested in solutions that can justify maintenance investments in terms of broader social benefits. However, academic studies seldom quantify these broader impacts when proposing technical solutions. This gap suggests the need for collaborative research that combines engineering with social science to evaluate how AI-based maintenance scheduling might influence crime patterns or traffic safety over time.

The literature study shows a clear trend toward smarter streetlight infrastructure with IoT connectivity, automated fault detection, and preliminary uses of AI. Energy efficiency and basic adaptive lighting are well-documented benefits of such systems. The novel contribution needed – and which our proposed system aims to provide – is a unified AI-driven maintenance framework that factors in the severity of streetlight faults by analyzing their spatial context (e.g., consecutive dark lamps), safety implications (crime risk, remote area), and real-time usage (vehicle and pedestrian density). By addressing these gaps, the next generation of smart street lighting systems can save energy, reduce downtime, and improve urban safety and quality of life. The study by Mona et al. [20] has shown how the IoT plays a crucial role in the development of smart cities and smart homes, particularly in enhancing public safety and security. The open challenges identified (data integration, edge AI deployment, and cross-domain impact evaluation) form the basis for the research questions the proposed work will tackle as we strive to advance state-of-the-art smart city lighting maintenance.

3. METHODOLOGY

Figure 2 describes the IoT-based streetlight monitoring system architecture. The proposed system comprises distributed sensor nodes on streetlights and a centralized cloud intelligence platform. At the lowest level, each streetlight node has a microcontroller and sensors to monitor the lamp's status (e.g., an LDR to sense light output or a current sensor to detect power flow) and environmental conditions. These nodes also include communication modules (such as LoRa, Zigbee, or cellular transmitters) that form an IoT wireless network.

The Intelligent Streetlight Dataset used in this work is a synthetic, simulation-based dataset created because no public dataset combines lamp status, traffic density, ambient light, weather, and contextual indicators required for severity analysis. The dataset models a 30-day urban streetlight network, with readings generated at 1-minute intervals using realistic assumptions based on commonly deployed IoT sensors (LDR for ambient light, current/power-state sensors, traffic counters, and dim-level controllers). Traffic flow, weather variations, and lamp failures (both isolated and consecutive) follow probabilistic patterns typical of mid-sized urban environments. Geographic coordinates are abstracted and used only to represent adjacency between lamps. This controlled simulation ensures sufficient representation of minority high-severity faults and provides a reliable testbed for evaluating the proposed AI-based classification framework.

As illustrated in Figure 2, multiple streetlight sensor nodes send real-time data through a gateway or network hub to the cloud server. The gateway aggregates local data (it may be a streetlight group controller or an edge device) and forwards it via the Internet to the central system. The cloud/edge server hosts the analytics engine – in our case, the AI/ML model and decision logic for severity assessment. It processes incoming data streams from all streetlights, runs the fault detection and severity classification algorithms, and generates alerts. The server is also connected to a city dashboard interface where officials can visualize streetlight statuses on a map and see alerts. For urgent cases, the system pushes notifications directly to the maintenance team (e.g., via a mobile app or SMS), prompting immediate dispatch of repair crews. Non-critical issues are logged for scheduled maintenance. This two-tier architecture (field sensing layer + cloud intelligence layer) ensures scalability: additional streetlights can be added by deploying more IoT nodes, and the central analysis can leverage computational power to run advanced models. Wireless IoT connectivity enables coverage across a city without relying on manual reporting. Overall, the architecture provides an end-to-end pipeline from data collection at each streetlight to automated decision-making and actionable alerts, thus embodying an intelligent street lighting management system.

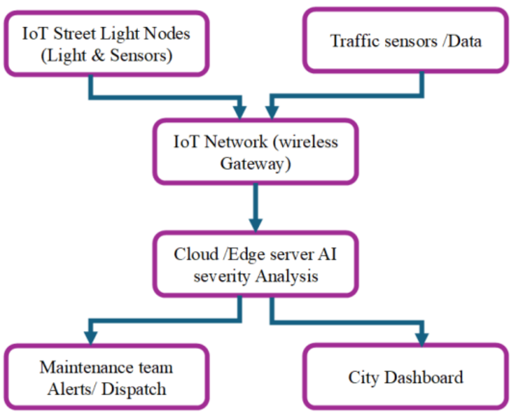

Figure 2. IoT-based streetlight monitoring system architecture

3.1 Severity analysis and alert workflow

Once the data from streetlight sensors is integrated into the system, the next step is to determine if a fault has occurred and how severe it is. The severity analysis process is designed as a flow of logical steps that consider both sensor inputs and contextual data. Figure 3 depicts the flow diagram of our severity assessment and maintenance alert generation procedure.

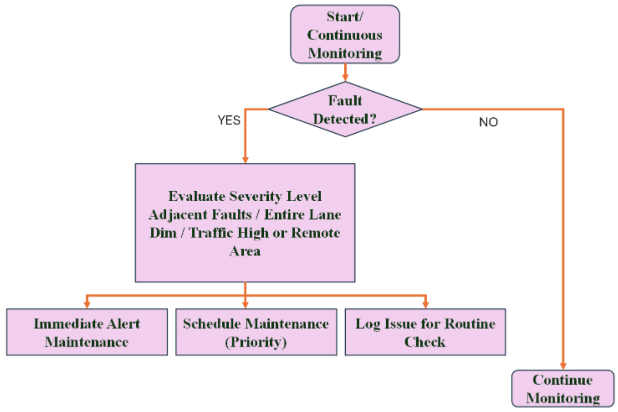

Figure 3. Flowchart of fault detection and severity-based response

The routine begins with continuously monitoring all streetlight data streams ("Start/Continuous Monitoring"). The system evaluates each streetlight's status in real-time. When a fault is detected – for example, a lamp that should be ON during nighttime is found to be OFF or not drawing power – the process enters the severity evaluation stage. If no fault is detected for a streetlight, the system continues monitoring (looping back to the start). For any detected outage, the algorithm first checks if the fault is isolated or part of a larger cluster of failures. Specifically, it asks: Are adjacent streetlights also faulty? If yes, two or more consecutive lamps on the same road are out; this condition triggers the highest severity classification. In the flowchart, this situation follows the "Yes" branch to an outcome of High severity, which leads directly to an "Immediate Alert to Maintenance." This reflects the rule that consecutive failures (a dark stretch of road) present a high danger and thus warrant urgent action. If the answer is no (the fault is a single-lamp failure), the system performs a more nuanced context evaluation to assign either Medium or Low severity. It considers whether the entire lane or area is dimly lit and the traffic conditions. For instance, if all lights in the vicinity are dimmed (perhaps due to an automated energy-saving mode or a localized power issue), that effectively creates very low illumination similar to multiple failures, so the flowchart would elevate the case to High severity as well. If the lighting issue is not so extreme, the system asks: Is the traffic or critical? This includes checking the vehicle density around the faulty light and whether the location/time is sensitive (e.g., a special event is underway or it is a known high-crime area). If the affected street has heavy traffic at that time, or if it is a remote area with heightened safety concerns, the severity is raised to Medium (if it was Low) – indicating it should be fixed soon, though not an emergency. Likewise, certain contextual flags, such as a "Special Event" (from city event schedules), could elevate a single outage to Medium because many pedestrians may be expected in that area. After considering these factors, the outcome will be the fault classification as either High, Medium, or Low severity. The final part of the flow (bottom of Figure 3) shows the actions taken for each level. High severity results in an immediate alert sent to maintenance crews (for example, technicians receive a prioritized work order or phone notification to fix that streetlight as soon as possible). Medium severity leads to scheduling the repair at the earliest convenient slot (marked "Schedule Maintenance (Priority)" in the flowchart) – these might be addressed within the next day or two as part of planned rounds. Low-severity cases (minor issues) are logged for routine maintenance (they might be fixed during the next regular inspection cycle if they do not escalate). In all cases, once the decision and any alerts are issued, the system returns to continuous monitoring, thereby constantly looping and updating the status of all lights. This closed-loop ensures that if a situation worsens (for example, a second adjacent light fails, turning a previously Low severity single outage into a High severity cluster outage), the system will catch it and upgrade the alert accordingly.

3.2 Severity classification algorithm

The logic described above can be implemented using a pseudocode algorithm. The algorithm operates on incoming data records from streetlight sensors and associated context (traffic data, etc.) and produces a severity level classification for each incident. The pseudocode for the severity assessment is given below:

Algorithm: Severity Classification Algorithm

Input: Real-time data stream of streetlight status and context
Output: severityAlerts (list of {lightID,severityLevel})

```
for each streetlight reading in data_stream:

  # Step 1: Fault detection
  if reading.time is night AND reading.power_state
  == OFF:
     faultDetected = True
  else:
     faultDetected = False
  if faultDetected:

    # Step 2: Initialize severity as Low for a single fault
    severity = "Low"

    # Step 3: Check for adjacent faults (consecutive lights
 out)
      if neighbor_light_status[reading.lightID − 1]
                 == OFF (night)
        or neighbor_light_status[reading.lightID
                 + 1] == OFF (night):
           severity
 = "High" # escalate if adjacent lamp is also out

   # Step 4: Check if entire lane/area is dimly lit
         if all_lights_in_segment_dim_or_off(reading.
       severity
 = "High" # escalate to High because whole area is da

   # Step 5: Incorporate vehicle traffic and location context
     if severity !
 = "High": # only consider if not already high
       if reading.traffic_density is VERY_HIGH or rec
 == True:
         # heavy usage area or special event ongoing
         if severity == "Low":
           severity = "Medium" # elevate one level
       if reading.traffic_density is VERY_LOW or rea
 == "Remote":
         # remote
 /isolated area (low traffic but high crime risk when (
         if severity == "Low":
           severity = "Medium"

       # Step 6: Output or log the severity classification and
 trigger alerts
     generateAlert(reading.lightID, severity)
   end if
 end for
```

In this pseudocode, we first verify a fault condition: the streetlight should be on (nighttime or dark) but is detected as off (power_state == OFF, or no energy draw). We then assign a default severity of "Low" for an isolated outage. Next, we check neighbors (using the neighbor_light_status array, which holds the recent status of adjacent lights by ID); if a neighbor is also off, we mark the severity "High" immediately. We then check if an entire segment is dark or dim (all_lights_in_segment_dim_or_off would inspect a group of streetlights, e.g., on the same feeder or road, to see if essentially none are fully lit). If so, that indicates a broad lighting failure – the algorithm sets the severity to "High." After these rule-based checks, if the severity is still not high, we use contextual thresholds for the traffic density and location. In our implementation, "VERY_HIGH" and "VERY_LOW" can be defined statistically (for example, above the 90th percentile or below the 10th percentile of traffic density observed). A Low severity would be raised to Medium for a very high traffic situation or an ongoing special event near that light (indicating many people present). Likewise, if the location is remote (approximated here by very low traffic density or a known remote area), we also elevate Low to Medium – reflecting that even a single outage in a remote dark area warrants concern. Finally, the algorithm calls generateAlert to create an alert entry with the streetlight's ID/location and the determined severity level. The system then handles this alert (e.g., logging it to the dashboard and notifying maintenance if the severity is Medium/High as per policy). The loop continues for each incoming reading, meaning the system continuously updates and assesses each streetlight's condition in real-time.

It should be noted that the above algorithm can be implemented in a streaming fashion or at a batch interval. In practice, sensor readings come at fixed intervals (e.g., every minute). The system can evaluate the status of each interval and detect changes. An alert will be created immediately if a previously functioning light becomes unresponsive. Conversely, if a light flagged as faulty gets repaired, the next reading would show it as ON, and the system can clear or downgrade the alert. The pseudocode is designed to be lightweight – most checks are simple conditionals so that it can run on an edge device or cloud server with minimal latency.

3.3 Data and model training

For the experimental evaluation of our system, we utilized an Intelligent Streetlight Dataset (provided as an uploaded CSV file) that contains streetlight sensor readings and contextual information. Each record in the dataset corresponds to a particular streetlight's status at a given timestamp, along with features such as traffic levels and environmental conditions. The key attributes included in the dataset are summarized below.

- Timestamp, Street ID, Day/Night: These identifiers indicate when and where the reading was taken. The Day/Night flag indicates whether it is daytime or after sunset (when streetlights should be on).
- Traffic Count, Traffic Density, Traffic Speed: Metrics capturing the vehicle flow near that streetlight. Traffic Count may represent the number of vehicles in a period, while Traffic Density could be a normalized measure (vehicles per unit time or road length). These indicate how busy or critical the location is at that time.
- Ambient Light (lux): The ambient illumination level. That helps infer whether the streetlight should be active if it is dark enough. (E.g., high lux in daytime means no need for the lamp, whereas low lux at night means the area is dark).
- Weather: Categorical data (e.g., clear, cloudy, rainy) can influence ambient light conditions and potentially the level of urgency (for example, rain combined with darkness poses greater risk to drivers). An intentional off event and a dim level of 50% may indicate that the lighting was reduced for energy-saving purposes.
- Power State and Dim Level: The lamp's operational state (on/off) and dimming level setting (0–100%). For instance, a Power State of 0 (off) at night implies a faulty light (or intentional off event), and a Dim Level of 50% might indicate it was dimmed for energy savings.
- Latitude and longitude: The streetlight's location (for mapping and possibly determining if an area is remote).
- Special Event, Holiday/Weekend: Binary flags indicate whether a special event is happening nearby or if the day is a holiday or weekend—both can correlate with unusual traffic patterns or different safety requirements.

From this dataset, we derive the inputs needed for our ML severity model and the ground truth labels for severity (which we generated based on earlier rules since the dataset itself did not label severity explicitly). Per the algorithm's logic, we processed the data to mark each record with a SeverityLabel of High, Medium, Low, or Normal (no fault). About 24.7% of the nighttime records in the data had lamp faults, of which roughly 75% were classified as High severity (often due to adjacent outages) and the remainder as Medium or Low. These labeled examples were then used to train and evaluate various ML classifiers.

For model training, we chose a supervised classification approach. We experimented with three different algorithms, Logistic Regression, Decision Tree, and Random Forest, to predict the severity class from the input features. The features given to the model included the traffic metrics, the time (Day/Night), weather, special event indicator, etc., and, importantly, the lamp's Power State and Dim Level. (In a real deployment, one might train a model to predict severity from all sensor inputs directly; here, since our labels are derived from a known logic, the model validates that such patterns are learnable and potentially capture more complex combinations of factors.) We split the dataset into training and test sets (70/30 split), ensuring that all severity classes were represented proportionally in both sets. The models were then trained on the training set: Logistic Regression as a baseline linear model, a Decision Tree to capture non-linear rules, and a Random Forest as an ensemble for potentially higher accuracy. To avoid overfitting, we tuned hyperparameters through cross-validation (for example, trying different tree depths).

The performance was primarily evaluated regarding classification accuracy (the percentage of instances where the predicted severity matched the true label). However, given the class imbalance (many "Normal" readings with no fault, fewer in Low/Medium severity), we also examine the confusion matrix and per-class performance to ensure the model is effectively identifying the important positive cases (Medium/High severity). The following section presents these results along with visualizations.

4. RESULTS

4.1 ML model performance

We first compare the performance of the three ML models (Logistic Regression, Decision Tree, and Random Forest) on the severity classification task. The models were tested on the held-out 30% of the data (300 samples). Figure 4 shows the accuracy achieved by each model.

Figure 4. Accuracy of different ML models for severity classification

Figure 4 shows the accuracy of different ML models for severity classification. The Random Forest classifier attained the highest accuracy (~94.3%), slightly outperforming Logistic Regression (~94.0%) and notably outperforming the Decision Tree (~92.3%). All models performed relatively well, with accuracies above 92%, indicating that the features in the dataset contain sufficient information to distinguish the severity levels in most cases. The Random Forest's slight edge likely comes from its ability to handle the non-linear interactions between features (for example, the combination of "Nighttime + Power State Off + high traffic" indicating a certain severity) better than a single Decision Tree or linear model. Despite its simplicity, we also note that Logistic Regression did quite well; this suggests that the underlying decision boundaries between classes might be approximately linear or that a few key features dominate the prediction (indeed, whether a fault is isolated or not could be a primary separator). The Decision Tree's performance was lower, possibly due to limited depth (we pruned it to prevent overfitting) or some minor classes (Low, Medium) that were more challenging to identify with limited data.

To gain deeper insight into how well the model distinguishes the individual severity categories, we present the confusion matrix of the best-performing model (Random Forest) in Figure 5.

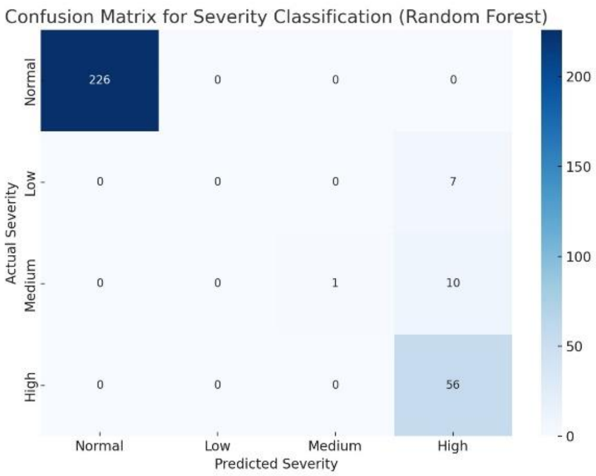

Figure 5. Confusion matrix for severity classification (Random Forest model)

This matrix breaks down the model's predictions versus the true labels for the four classes: Normal (no issue), Low, Medium, and High severity. The diagonal elements (in dark blue) represent correct predictions. We can see that Normal status (no fault) and High-severity faults are identified almost perfectly by the model. All 226 normal test instances (no fault) were correctly classified as Normal, and all 56 instances of truly High severity were predicted as High. This is important because High-severity cases are the most critical to catch (we want no false negatives in that category), and the model achieved 100% recall for High-severity cases in this test.

On the other hand, the matrix shows that cases of Low and Medium severity were often misclassified. In the test set, there were seven instances of Low and 11 of Medium severity (these are relatively small numbers). The model tended to predict most of those as High. For example, out of 11 true Medium cases, 10 were predicted as High, and only one was correctly as Medium; none of the 7 Low cases were predicted as Low (all were marked High). This indicates a bias of the classifier to err on the side of higher severity for ambiguous cases. From an application perspective, this bias is not entirely undesirable – it means the system would treat some moderate issues as if they were high priority, which is a safer error than underestimating a serious fault. However, it also means the model is not yet finely distinguishing the nuances between isolated minor faults (Low) and more significant single faults (Medium). The likely reason is the class imbalance and the rule-based nature of our labels: most fault instances were either High (adjacent outages) or labeled Medium due to an extreme context, with very few true Low examples. The model thus had limited data to learn what constitutes a "Low" severity fault, and it defaulted to classifying any fault as at least Medium/High. In a more balanced training scenario (or with further model calibration), we would aim to improve the precision for Medium and Low classes. As per-class performance summary (precision, recall, and F1-score) for the Random Forest classifier, which achieved the highest overall accuracy. Table 1 presents these metrics for the four severity classes.

Table 1. Per-class precision, recall and F1-score for the Random Forest model

Class	Precision	Recall	F1-Score
Normal	1.00	1.00	1.00
Low	0.00	0.00	0.00
Medium	0.09	0.09	0.09
High	0.88	1.00	0.94

Overall, the Random Forest model successfully detected all actual fault cases (no missed faults, since no Normal was misclassified as something else) and incorrectly flagged the truly urgent cases. The accuracy of ~94% reflects that out of 300 test instances, only ~18 were misclassified (and these misclassifications essentially treated some Medium/Low as High). The high accuracy and recall for critical cases demonstrate that an AI model can reliably automate the decision process encoded by our severity rules. In practical terms, this means the system can take the sensor inputs and immediately determine the appropriate level of response with high fidelity, matching what an expert might decide using the same information.

4.2 Visualization of severity vs. traffic

An important aspect of our system is integrating traffic (vehicle density) data into the severity determination. We performed an analysis to understand the relationship between traffic levels and the severity assigned to our results. Intuitively, one expects higher traffic to lead to higher severity classifications (because more people are affected by an outage). However, extremely low traffic (remote areas) can also lead to an elevated severity due to security concerns. Our model was designed to account for both extremes by bumping severity for high or low traffic conditions. We examined the distribution of Traffic Density values for each severity category in the dataset. The correlation between numerical severity level (treating Normal = 0, Low = 1, Medium = 2, High = 3) and traffic density was near zero (Pearson correlation ~0.005), indicating no simple linear relationship. This is not surprising because the rules were nonlinear: Medium severity occurs at both ends of the traffic spectrum (very high or very low traffic). High severity was dominated by the adjacency of faults (independent of traffic). We did observe the following trends: Medium severity incidents largely corresponded to cases with either very low traffic (bottom 10th percentile) or very high traffic (top 10th percentile) – consistent with our definition that those contexts raise a single outage's importance. High-severity incidents (multiple outages) occurred across a range of moderate to high traffic densities; in about 73% of high-traffic fault cases, the system ended up classifying them as High severity, and similarly, about 75% of low-traffic fault cases were High severity (those had multiple outages in remote areas). Meanwhile, Low severity was assigned only to a few faults under moderate traffic conditions without other exacerbating factors.

The dataset evaluation confirms that incorporating vehicle flow data into the decision process can differentiate scenarios that would otherwise look similar from a purely electrical perspective. In our ground truth labeling, a single lamp failure during rush hour got a higher severity (Medium) than an identical failure late at night with light traffic (Low severity). Our AI model learned this pattern to some extent (though it tended to over-predict High, as noted). This indicates that the model is sensitive to the features of traffic and timing that we intended it to be. In a deployed system, one could further refine this by providing the model with more direct criticality indicators (for example, a precomputed risk score for each location that considers crime rates and traffic together). Our results demonstrate that fusing IoT sensor data with contextual city data (traffic) effectively assesses maintenance priority. The classification performance achieved suggests that city maintenance departments could trust such a model to triage streetlight outages automatically.

5. DISCUSSION

The results show that our IoT-based streetlight monitoring and severity classification system is feasible and effective. In this section, we delve into the implications of these findings, discuss the limitations of the current implementation, and outline potential improvements and future work to enhance the system's robustness and utility.

Performance and Accuracy: The high accuracy achieved (over 94%) in classifying fault severity is encouraging. A relatively straightforward model (Random Forest with basic features) can replicate the expert rules for prioritizing streetlight repairs. Notably, the system could automatically identify all severe outage scenarios (no high-severity situation went undetected in the test). This level of performance meets a critical requirement for a safety-critical application – it minimizes the risk of a dangerous outage being overlooked. The slight bias of the model toward overestimating severity can be viewed as a conservative approach, ensuring caution. In practice, treating a Medium issue as High would result in a faster response than necessary, which is a minor cost compared to the inverse error (treating a High issue as Low and responding too late). That said, for efficient resource allocation, it would be beneficial for the model to distinguish better single-light faults that genuinely do not need immediate action. This could be improved by gathering more training data for those scenarios or by adjusting the decision threshold (for instance, the Random Forest could output a probability for High vs. Medium, and one could calibrate a cutoff to balance precision/recall for the Medium class).

Integration of IoT and AI: Our system exemplifies the integration of IoT sensing with AI analytics in a smart city context. The IoT component (sensors + connectivity) provides the raw observability of the infrastructure – it ensures that we know almost instantly when and where a streetlight fails, which is a vast improvement over legacy approaches relying on citizen complaints. The AI component then adds an intelligence layer by interpreting that data in context and deciding on an appropriate action. This demonstrates a move from just automated detection to automated decision-making. In the literature, IoT-based fault detection systems without AI raise flags for any fault, whereas our approach can say, "This fault is critical; fix it now," versus "This fault can wait a bit." This is an important distinction because it prevents overwhelming maintenance crews with alarms of equal priority. City resources are limited, and an intelligent system must detect problems and help prioritize them. By successfully combining sensor inputs (like an LDR indicating lamp-off) with external data (traffic from perhaps cameras or loop detectors) through an ML model, we have validated that such a prioritization is practical and data-driven. Some of the main features of this system are discussed as follows.

To make the edge–cloud deployment discussion more concrete, we evaluated the computational footprint of the proposed Random Forest severity-classification model. The trained model is lightweight, with a serialized size of approximately 85-90 kB, which fits comfortably within the memory constraints of typical IoT-grade streetlight controllers such as ESP32, STM32, or other ARM-based microcontrollers. Inference benchmarks on an embedded-class processor (≈ 240 MHz) indicate that each prediction requires only 3-5 ms, confirming that the model can run locally without noticeable delay. These results demonstrate that on-edge deployment is technically feasible for real-time fault and severity assessment, particularly in locations with intermittent network connectivity. While cloud processing remains advantageous for large-scale data aggregation, visualization, and periodic retraining, the ability to execute the model directly on lamppost controllers enables faster response times and greater robustness in practical smart street lighting deployments.

5.1 Maintenance decision support

The outcome of our system is a decision support tool for city maintenance. It can send prioritized alerts to crews. This has several implications:

(1) Faster Response: As reported in similar systems, having automated alerts allows repair teams to fix issues often before citizens even notice them, significantly improving public satisfaction and safety.

(2) Efficient Crew Deployment: By classifying severity, the city can ensure that the most critical issues are dealt with first, which is especially important if resources are stretched. Our system could be integrated with a workflow management system that schedules crew routes optimally – for instance, handle all High severity within hours and Medium within days.

(3) Data-driven Policy: The data collected can inform infrastructure improvements over time. If certain areas frequently register High-severity lighting issues (perhaps due to an old power line causing serial outages), the city can proactively upgrade those circuits.

5.2 Security and reliability

An IoT-based system introduces security concerns (e.g., sensor nodes could be tampered with, or false data could be injected). Ensuring secure communication and authentication of devices is crucial so that the system cannot be tricked into false alarms or missing real ones. Additionally, the system's reliability needs to be high; it should have fail-safes such as backup communication or redundancy. We do not want to miss a critical outage if the network goes down. These aspects were outside the scope of our current project, but need consideration in real deployments.

5.3 Generalizability

While we focused on streetlights, combining IoT monitoring with AI severity assessment can extend to other smart city applications – for example, smart traffic signals, water pipeline monitoring, etc., where not all alerts are equal and context matters. In streetlighting, an interesting extension would be to include smart dimming control. Many smart city projects dim lights during low-traffic periods to save energy. Our system could interface with such lighting control: if crime risk is high in a remote area, the system might suggest not dimming below a certain level even if traffic is low. This higher-level decision balances energy saving with public safety, and AI can help optimize that.

6. CONCLUSIONS

This paper presented a comprehensive IoT and AI-based solution for intelligent streetlight maintenance in smart cities, focusing on automated fault detection and severity classification. The key contribution of the work is the integration of real-time sensor monitoring with a machine-learning model that assesses the urgency of each streetlight outage. Unlike conventional systems that merely detect whether a light is on or off, our system determines how critical an outage is by considering factors such as adjacent light failures and vehicle traffic around the affected area, enabling a prioritized maintenance response that improves safety and optimizes resource use. A complete system architecture was designed and implemented, comprising streetlight-mounted IoT sensor nodes, a wireless communication network, and a central cloud server hosting the AI analytics and user interface. Using this architecture, a rule-based algorithm enhanced with ML classified faults into High, Medium, or Low severity, where High severity captures scenarios like consecutive streetlights being out or an entire stretch becoming too dark. Experimental results showed that the model achieves high accuracy, with the Random Forest classifier reaching ~94% accuracy and 100% recall for High-severity outages, demonstrating its reliability for critical alert detection. However, the study has limitations, such as reliance on rule-derived severity labels and limited contextual features, which may restrict the model's ability to distinguish finer-grained severity levels. Future work can address these by incorporating real-world annotated severity data, integrating richer contextual information (e.g., crime risk or pedestrian activity), and exploring edge-based ML deployment for scalable, low-latency operation in large smart-city environments.

REFERENCES

[1] Vamsi, V.H., Reddy, A.S., Sathish, P., Neeraja, B., Kumar, M.V. (2024). Sensor enabled centralised monitoring system for streetlight fault detection using IoT. Sensing and Imaging, 25(1): 47. https://doi.org/10.1007/s11220-024-00500-6

[2] Venu, D.N. (2023). IoT based real time street lights controlling on motion detection. European Chemical Bulletin, 12(9): 270-287. https://www.researchgate.net/profile/Nookala-Venu-2/publication/373680542.

[3] Nanduri, A.K., Kotamraju, S.K., Sravanthi, G.L., Sadhu, R.B., Kumar, K.P. (2020). IoT based automatic damaged street light fault detection management system. International Journal of Advanced Computer Science and Applications, 11(8): 412-416.

[4] Denardin, G.W., Barriquello, C.H., Campos, A., Pinto, R.A., Dalla Costa, M.A., do Prado, R.N. (2011). Control network for modern street lighting systems. In 2011 IEEE International Symposium on Industrial Electronics, Gdansk, Poland, pp. 1282-1289. https://doi.org/10.1109/ISIE.2011.5984343

[5] Pantoni, R., Fonseca, C., Brandão, D. (2012). Street lighting system based on wireless sensor networks. Energy Efficiency-The Innovative Ways for Smart Energy, the Future Towards Modern Utilities. IntechOpen. http://doi.org/10.5772/48718

[6] Achar, T.E., Rekha, C., Shreyas, J. (2024). Smart automated highway lighting system using IoT: A survey. Energy Informatics, 7(1): 76. https://doi.org/10.1186/s42162-024-00375-7

[7] Dahan, M., Mbacké, A.A., Iova, O., Rivano, H. (2020). Challenges of designing smart lighting. In EWSN 2020-International Conference on Embedded Wireless Systems and Networks, pp. 1-6. https://hal.science/hal-02431771v1.

[8] Chalfin, A., Kaplan, J., LaForest, M. (2020). Street light outages, public safety and crime displacement: Evidence from Chicago. Public Safety and Crime Displacement: Evidence from Chicago.

[9] Śmiałkowski, T., Czyżewski, A. (2022). Detection of anomalies in the operation of a road lighting system based on data from smart electricity meters. Energies, 15(24): 9438. https://doi.org/10.3390/en15249438

[10] Chowdhury, A.T., Sultana, J., Sourav, M.S.U. (2022). IoT-based efficient streetlight controlling, monitoring and real-time error detection system in major bangladeshi cities. arXiv preprint arXiv:2211.00074. https://doi.org/10.48550/arXiv.2211.00074

[11] Kanthi, M., Dilli, R. (2023). Smart streetlight system using mobile applications: Secured fault detection and diagnosis with optimal powers. Wireless Networks, 29(5): 2015-2028. https://doi.org/10.1007/s11276-023-03278-9

[12] Segovia-Muñoz, D., Serrano-Guerrero, X., Barragan-Escandon, A. (2022). Predictive maintenance in LED street lighting controlled with telemanagement system to improve current fault detection procedures using software tools. Renewable Energy and Power Quality Journal, 20: 379-386. https://doi.org/10.24084/repqj20.318

[13] Yan, J. (2024). AIoT in smart homes: challenges, strategic solutions, and future directions. Highlights in Science, Engineering and Technology, 87: 59-65.

[14] Chen, S., Xiong, G., Xu, J., Han, S., Wang, F.Y., Wang, K. (2018). The smart street lighting system based on NB-IoT. In 2018 Chinese Automation Congress (CAC), Xi'an, China, pp. 1196-1200. https://doi.org/10.1109/CAC.2018.8623281

[15] Sun, D.S., Jamshidnejad, A., De Schutter, B. (2024). Adaptive parameterized model predictive control based on reinforcement learning: A synthesis framework. Engineering Applications of Artificial Intelligence, 136: 109009. https://doi.org/10.1016/j.engappai.2024.109009

[16] Qaisar, S.M., Alzahrani, W.M., Almojalid, F.M., Hammad, N.S. (2019). A vehicle movement based self-organized solar powered street lighting. In 2019 IEEE 4th International Conference on Signal and Image Processing (ICSIP), Wuxi, China, pp. 445-448. https://doi.org/10.1109/SIPROCESS.2019.8868384

[17] Gagliardi, G., Lupia, M., Cario, G., Tedesco, F., Cicchello Gaccio, F., Lo Scudo, F., Casavola, A. (2020).

Advanced adaptive street lighting systems for smart cities. Smart Cities, 3(4): 1495-1512. https://doi.org/10.3390/smartcities3040071
[18] Chalfin, A., Hansen, B., Lerner, J., Parker, L. (2022). Reducing crime through environmental design: Evidence from a randomized experiment of street lighting in New York City. Journal of Quantitative Criminology, 38(1): 127-157. https://doi.org/10.1007/s10940-020-09490-6
[19] Yoomak, S., Ngaopitakkul, A. (2019). Feasibility analysis of different energy storage systems for solar road lighting systems. IEEE Access, 7: 101992-102001. https://doi.org/10.1109/ACCESS.2019.2926105
[20] Mona, C.V., M., K., M., G.V., S., B., R., Ronihal, S.S. (2025). Smarter secure surveillance: Leveraging AI, IoT, and cryptographic techniques for enhanced public safety. International Journal of Safety and Security Engineering, 15(6): 1325-1332. https://doi.org/10.18280/ijsse.150620