International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

CivicEye: A Digital Framework for Urban Civic Issue Reporting

K B Pavan
SVR College of Commerce & Management Studies, Bengaluru, Karnataka, India
Email: kbpavan2005@gmail.com

Abstract - Urban civic issues such as unmanaged waste, potholes, drainage overflow, and nonfunctional streetlights can significantly affect the quality of urban living when their reporting and resolution are not systematically managed. Conventional reporting mechanisms, including telephone calls, in-person complaints, and social media, often lack structured data collection, status transparency, duplicate handling, and effective coordination between citizens and local authorities. This paper presents CivicEye, a digital framework designed to provide an integrated mechanism for reporting, managing, monitoring, and analyzing urban civic issues. The framework enables citizens to submit complaints containing textual descriptions, photographic evidence, and location information. Submitted complaints are subsequently verified and routed to the appropriate administrative department and field personnel for resolution. CivicEye incorporates duplicate-complaint identification to reduce redundant records and improve the prioritization of recurring civic issues. The framework also maintains resolution evidence and provides analytical summaries of complaint patterns, frequently affected locations, and resolution timelines to support data-driven resource planning. In addition, an optional image-based issue classification component can assist citizens in identifying the appropriate category while submitting a complaint. By integrating citizen participation, structured complaint management, duplicate identification, location-based information, and analytical capabilities within a unified framework, CivicEye aims to improve the transparency, traceability, and efficiency of urban civic issue management. The proposed framework provides a foundation for developing more responsive and data-driven mechanisms for citizen-authority collaboration in smart urban environments.

Keywords: Civic issue management, citizen participation, urban governance, complaint management, duplicate detection, smart cities, location-based services, data-driven decision making.

1. Introduction

Cities in developing economies are growing faster than the institutions that maintain them. Roads, drains, street lighting, and solid-waste services must serve larger and denser populations every year, yet the mechanisms through which residents tell municipal bodies that something is broken have changed little. A resident who notices an overflowing drain typically has three options: call a helpline, visit a ward office, or post about it on social media. Each of these channels produces information in a different shape, and none of them guarantees that the report reaches the right officer, that the resident will learn what happened next, or that the municipality will gain any lasting insight from the complaint.

The consequences are practical. A pothole reported by ten people through ten different channels may appear to the administration as ten unrelated records, or as none at all. A garbage point that is cleared and then refills within days may be closed in the records as resolved each time, hiding the fact that the underlying problem is recurring. Without photographic evidence at the time of resolution, supervisors cannot verify that work was completed. Without location data, field staff may spend time searching for the problem. Without analytics, budgets and workforce are allocated by habit rather than by evidence.

Digital platforms for civic reporting have been introduced in many cities worldwide and in India, and they have shown that smartphones can lower the effort of reporting. However, a reporting interface alone does not make a complaint-management system. Three problems recur in practice: the same issue is reported

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

many times and inflates the backlog; complaints are routed manually and inconsistently; and the data collected is seldom turned into planning information for the authorities.

This paper proposes CivicEye, a digital framework that treats civic issue reporting as an end-to-end process rather than a form-submission feature. The framework covers the complete life of a complaint: structured submission with text, photograph, and location; verification and routing to the responsible department; field resolution with evidence; duplicate identification across the whole process; and analytical reporting for decision-makers. An optional image-based classification component helps citizens select the correct category when they are unsure.

The main contributions of this paper are as follows.

1. A layered architecture for civic issue management that unifies citizen reporting, administrative workflow, duplicate handling, and analytics in one framework.
2. A defined complaint lifecycle with explicit states, roles, and evidence requirements, designed to make each complaint traceable from submission to closure.
3. A multi-signal duplicate identification approach that combines geographic proximity, time window, category, textual similarity, and optional image similarity into a single, tunable score.
4. An optional image-assisted categorization module, positioned as a suggestion tool for the citizen rather than an automated decision-maker.
5. A set of analytical indicators, their operational definitions, and an evaluation methodology that municipalities and researchers can use to assess the framework in a pilot deployment.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 states the problem and objectives. Section 4 presents the CivicEye framework. Section 5 describes duplicate identification, Section 6 the image-based classification component, and Section 7 the analytics layer. Section 8 discusses system design and implementation considerations. Section 9 presents the evaluation methodology, Section 10 discusses implications and limitations, and Section 11 concludes.

2. Related Work

2.1 Citizen Participation and Digital Governance

The idea that citizens should take part in decisions affecting their surroundings has a long scholarly history. Arnstein's ladder of citizen participation distinguishes levels ranging from token consultation to genuine citizen control, and it remains a common reference for judging whether a participation mechanism gives people real influence or merely the appearance of it [1]. Civic reporting tools sit in the middle of this ladder: they let residents inform the authorities, but their value depends on whether the authorities respond visibly.

Smart city literature frames technology as only one of several dimensions of urban improvement, alongside people and institutions [2]. A platform that digitizes complaints without changing how departments assign, track, and close work addresses only the technological dimension. Goodchild's notion of citizens as sensors describes how ordinary people, equipped with location-aware devices, can contribute geographic information that formal agencies cannot collect at the same density [3]. Civic reporting is a direct application of this idea, since each complaint is a geotagged observation of the city.

2.2 Existing Civic Reporting Platforms

Several platforms illustrate the range of approaches to issue reporting. FixMyStreet, developed by the non-profit mySociety, has operated across the United Kingdom since 2007: residents place a pin on a map, add a description and an optional photograph, and the platform matches the report's location and category to the responsible council, sends the report to that council, and publishes it publicly; publishing reports online is itself credited with reducing duplicate reports [8]. The FixMyStreet platform is open source and has been deployed by communities and governments in more than twenty countries [8]. SeeClickFix, deployed by

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

many municipalities, provides a comparable experience in which the resident picks a category, drops a pin, and adds a photograph; the request is routed to the responsible department and both resident and staff track its progress to closure, with duplicate handling limited to directing residents to follow existing requests [9]. In India, the Swachhata application, the official grievance-redressal application of the Ministry of Housing and Urban Affairs under the Swachh Bharat Mission (Urban), lets citizens photograph an issue, attaches the geo-location of the photograph automatically, and forwards the complaint to the responsible municipal body [10]. Bengaluru's civic body has also operated its own complaint channels for roads, waste, and street lighting.

These systems have demonstrated that mobile reporting is feasible at scale. They are, however, built around different priorities. Some emphasize public visibility on a map, others emphasize internal ticketing, and others focus on a single service such as sanitation. The properties that this paper considers essential together, namely structured evidence, duplicate handling, resolution proof, and analytics for planning, are not uniformly present, and where they are present they are usually implemented as separate features rather than as an integrated design. Table 1 summarizes this comparison.

Property	FixMyStreet [8]	SeeClickFix [9]	Swachhata [10]	CivicEye (proposed)
Photo and location capture	Yes (photo optional)	Yes	Yes, geo-tag from photo	Yes, mandatory
Public map of reports	Yes	Yes	Not emphasized	Optional, privacy-protected
Routing to responsible body	Automatic, by location and category	By category to departments	Sanitation category workflows	Configurable category-area-to-department table
Administrator verification before routing	No	No	Not documented	Yes
Duplicate handling	Public visibility discourages duplicates; no automatic merging	Resident directed to follow existing requests	Not documented	Multi-signal score with review band
Mandatory resolution evidence	No	No	Not documented	Yes, photograph after work
Explicit reopen state on recurrence	No	Via new report	Not documented	Yes, with history
Planning analytics for administrators	Not documented	Not documented	Not documented	Eight operationally defined indicators

Table 1. Comparison of existing civic reporting platforms with the proposed framework, based on publicly available documentation (accessed October 2026). Closed administrative features are harder to verify from the outside, so the comparison is indicative rather than exhaustive; its purpose is to show where the essential properties are jointly absent, not to rank the platforms.

2.3 Duplicate and Near-Duplicate Report Detection

Duplicate detection is well studied in bug tracking, customer support, and information retrieval, where text similarity measures are used to link reports describing the same fault. Civic complaints differ from these settings in two ways. First, the text is typically short, informal, and often multilingual, which weakens purely lexical matching. Second, civic complaints carry strong non-textual signals, particularly geographic coordinates and photographs, that can be used alongside text. Sentence embedding models such as Sentence-BERT produce vector representations in which semantically similar short texts lie close together, and they can be compared efficiently using cosine similarity [4]. For the spatial dimension, density-based clustering such as DBSCAN can group nearby reports without requiring the number of clusters in advance and can identify areas where complaints concentrate [5].

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

2.4 Image-Based Classification of Urban Issues

Convolutional and transformer-based image models have become standard tools for visual classification. Residual networks made it practical to train deep models that achieve strong accuracy on image recognition tasks [6], and lightweight architectures such as MobileNets were designed to run on devices with limited computing resources [7]. These properties suit civic reporting, where the user is on a smartphone, connectivity may be unreliable, and the required decision is a choice among a handful of categories such as garbage, pothole, drainage, and streetlight. A recurring concern in the literature on applied image classification is that models trained on curated datasets can perform poorly on real user photographs with varied lighting, angles, and backgrounds, so classification outputs should be treated as aids rather than final decisions.

2.5 Research Gap

Previous work and existing platforms establish that digital reporting is useful, that duplicate detection techniques exist, and that image classification is technically feasible. Recent practitioner and academic systems, such as CivicSync [11] and comparable crowdsourced reporting platforms described in the current literature [12], already combine several of these pieces, which confirms the direction in which the field is moving. What these systems generally lack, and what existing international platforms do not provide jointly, is a formally specified duplicate score with tunable thresholds and a human review band, an audit trail attached to every lifecycle transition, operationally defined planning indicators with explicit denominators, and an evaluation methodology for testing the whole design. CivicEye is proposed to fill this gap by specifying how citizen participation, workflow, duplicate identification, location intelligence, and analytics fit together, and by defining in Section 9 how the specification is to be tested.

3. Problem Statement and Objectives

3.1 Problem Statement

Local authorities receive civic complaints through multiple unstructured channels, which leads to incomplete information, repeated records for the same issue, delayed or incorrect routing, little transparency for the complainant, and almost no systematic use of complaint data for planning. The problem is therefore how to design a unified digital framework that captures complaints in a structured, evidence-bearing form, manages them through a traceable workflow, identifies duplicates and recurring problems, and converts the resulting data into insight for resource planning.

3.2 Objectives

The framework is designed to meet the following objectives.

1.0 **O1, structured capture:** collect each complaint with a category, description, photograph, and geographic location so that it is actionable without further clarification.

1.0 **O2, accountable routing:** direct each verified complaint to the correct department and field officer, with a recorded owner at every stage.

1.0 **O3, transparency:** let the citizen follow the status of a complaint and see evidence of resolution.

1.0 **O4, redundancy reduction:** identify duplicate complaints and merge them into a single issue while retaining the count of reporters as a signal of severity.

1.0 **O5, evidence-based closure:** require proof of work before a complaint is marked resolved.

1.0 **O6, data-driven planning:** provide summaries of complaint volume, hotspots, and resolution times to administrators.

1.0 **O7, assisted reporting:** optionally suggest an issue category from a photograph to reduce miscategorization.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

3.3 Scope and Assumptions

The framework targets municipal-level issues that are physically observable and photographable. It assumes that users have access to a smartphone or web browser, that the municipality can designate departments and field personnel in the system, and that complaints are submitted voluntarily. Emergency situations that threaten life or safety are outside the scope and should continue to use dedicated emergency services.

4. The CivicEye Framework

4.1 Overview

CivicEye is organized as five cooperating layers: a citizen interface, a complaint management core, an intelligence layer, an administrative and field interface, and a data and analytics layer. The citizen interface accepts and displays complaints. The management core stores them and enforces the lifecycle. The intelligence layer performs duplicate identification and optional image classification. The administrative and field interface supports verification, assignment, and resolution. The analytics layer aggregates the stored data into indicators for decision-makers. Table 2 summarizes the layers.

Layer	Main responsibility	Typical users
Citizen interface	Submit complaint, track status, view resolution evidence	Residents
Complaint management core	Store records, enforce lifecycle, notify stakeholders	System
Intelligence layer	Duplicate identification, image-assisted categorization	System
Administrative and field interface	Verify, assign, update, resolve with proof	Administrators, department heads, field staff
Data and analytics layer	Hotspots, trends, resolution-time reports	Administrators, planners

Table 2. The five layers of the CivicEye framework.

4.2 Stakeholders and Roles

Four roles interact with the framework. Citizens report issues and follow their progress. Administrators verify submitted complaints, correct categories, and route them. Department heads oversee the workload of their department and monitor delays. Field personnel carry out the repair or cleanup and upload evidence of completion. Role-based access ensures that each user sees only the functions and data relevant to them.

4.3 Complaint Submission

A complaint is submitted through a mobile or web interface and contains four elements. The first is a category chosen from a defined list, such as waste, road damage, drainage, or street lighting. The second is a short text description. The third is one or more photographs taken at the location. The fourth is location information, captured automatically from the device where permitted and adjustable by the user on a map. Requiring a photograph and location at the time of submission is the central design decision of the capture stage, since these two elements allow the receiving officer to understand and find the issue without contacting the citizen.

Before final submission, the framework performs a nearby-issue check. If it finds open complaints of a similar type close to the chosen location, it shows them to the citizen, who may choose to support the existing issue instead of creating a new one. This reduces duplicates at the source and gives the citizen immediate feedback that the problem is already known.

4.4 Verification and Routing

After submission, the complaint enters a verification stage in which an administrator or an automated rule set checks that the evidence is legible, the location is plausible, and the category is correct. Verified

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

complaints are mapped to the responsible department through a configurable table that links category and administrative area, such as a ward or zone, to a department. Within the department, the complaint is assigned to a field officer, either manually or by a workload-aware rule. Invalid, abusive, or out-of-scope submissions are rejected with a recorded reason that the citizen can see.

4.5 Complaint Lifecycle

Each complaint moves through a defined set of states, listed in Table 3. The explicit state model is what makes the system traceable, because every transition is time-stamped and attributed to a user.

State	Meaning	Who triggers the transition
Submitted	Complaint received with evidence and location	Citizen
Under verification	Awaiting administrative check	System
Verified and routed	Accepted and sent to a department	Administrator
Assigned	Allocated to a named field officer	Department head
In progress	Work has begun	Field officer
Resolved	Work completed and proof uploaded	Field officer
Closed	Resolution confirmed by supervisor or citizen	Administrator or citizen
Rejected	Invalid or out of scope, with reason	Administrator
Reopened	Issue persists or recurs after closure	Citizen or administrator

Table 3. Complaint lifecycle states.

The Reopened state is important because it captures recurrence. A complaint reopened several times at the same location signals that the repair or cleanup was inadequate, and the analytics layer uses this signal to flag chronic problem sites.

4.6 Resolution Evidence and Feedback

A field officer cannot move a complaint to Resolved without uploading at least one photograph taken at the location after the work. The framework compares the resolution location with the complaint location and records the time of the upload. The citizen is notified and can view the before-and-after evidence. If the citizen disagrees that the issue has been resolved, the complaint can be reopened with a new photograph, and the reopening is recorded in the history. This closes the feedback loop that conventional channels typically leave open.

4.7 Notifications

At each lifecycle transition the framework notifies the relevant stakeholders through in-application messages and, where configured, SMS or email. Notifications to citizens are brief and state the new status. Notifications to officers include the category, location, and photograph. Overdue complaints trigger escalation messages to the department head according to a configurable service-level target for each category.

5. Duplicate Complaint Identification

5.1 Why Duplicates Matter

When many residents notice the same problem, many of them report it. Treating each report as an independent record inflates workload statistics, may cause multiple field visits for one issue, and obscures the true priority of the problem. At the same time, discarding duplicate reports loses useful information, since the number of people affected is a reasonable indicator of severity. CivicEye therefore merges duplicates into a single parent issue and keeps a count and list of linked reporters.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

5.2 Signals Used

Duplicate identification in CivicEye compares a new complaint against open complaints using up to five signals.

1. **Spatial proximity.** The great-circle distance between the two reported locations is computed using the haversine formula. Complaints farther apart than a category-specific radius are not considered candidates. A smaller radius suits streetlights, which are point objects, while a larger radius suits waste accumulation or road damage that can span a stretch of road.
2. **Time window.** Only complaints that are still open, or were closed within a recent period, are considered, so that a new report about a recurring problem can be linked to the earlier history without being merged into an already resolved issue.
3. **Category agreement.** Complaints of the same category receive a higher score than those of different categories. Different categories at the same location are not treated as duplicates, because a pothole and a broken streetlight at one junction are separate problems.
4. **Textual similarity.** The descriptions are converted into vector representations using a sentence embedding model, and cosine similarity is computed. Because complaints may be written in English, Kannada, Hindi, or mixed language, a multilingual embedding model is preferred where available.
5. **Image similarity (optional).** Feature vectors extracted from the photographs are compared. This signal is the least reliable, since two photographs of the same pothole taken from different angles can look quite different, and so it carries the lowest weight.

5.3 Combined Score

Let a new complaint be n and an existing candidate be c. The duplicate score is a weighted combination of the individual signals:

$$S(n, c) = w_1 \cdot S_{geo} + w_2 \cdot S_{time} + w_3 \cdot S_{cat} + w_4 \cdot S_{text} + w_5 \cdot S_{img}$$

where each component is normalized to the range 0 to 1, the weights sum to 1, and S_{geo} decreases as distance increases. A distance-decay function such as

$$S_{geo} = \max(0, 1 - d / r)$$

where d is the distance and r is the category radius, is a simple option. Category agreement is binary and may be used as a gate: if categories differ, the pair is excluded.

The score is interpreted using two thresholds. If S is above a high threshold, the new complaint is automatically linked to the existing issue and the citizen is informed. If S lies between a lower and the high threshold, the pair is presented to an administrator for confirmation. Below the lower threshold, the complaint is treated as new. Using a review band rather than a single cut-off keeps humans involved in uncertain cases and limits the damage of false merges.

Signals can also conflict, and the design resolves the common cases explicitly. Two complaints thirty meters apart with different categories are separate issues by construction, because the category gate excludes the pair regardless of proximity. Two nearly identical descriptions recorded two months apart may describe a recurring rather than persistent problem; the time-window signal handles this by allowing links to recently closed issues while keeping the link visible in the complaint history. Because no single weighting dominates across all categories, the weights are treated as calibration parameters (Section 5.5) and evaluated at the automatic-link threshold rather than summarized as a single headline score (Section 9.2).

5.4 Recurring Issues and Hotspots

Beyond merging, the framework uses spatial clustering over historical complaints to identify locations that generate repeated reports across time. A density-based method such as DBSCAN groups complaints that lie

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

within a given distance of each other without requiring the number of clusters to be fixed in advance, and it labels isolated points as noise [5]. The resulting clusters are reported as hotspots in the analytics layer, and the reopened-complaint history is used to distinguish locations with many separate issues from locations where the same issue keeps returning.

5.5 Parameter Tuning

The radii, weights, and thresholds are configuration parameters rather than fixed constants, since they depend on the density of the city, the typical size of each issue type, and the quality of location data. The recommended practice is to calibrate them on a labeled sample of historical complaints in which administrators have marked true duplicate pairs, and then to review the false-merge and missed-duplicate rates before enabling automatic linking. The weights can also be learned rather than hand-tuned: given a labeled set of duplicate and non-duplicate pairs, a logistic-regression model over the five normalized signals yields coefficients that can be inspected for plausibility and adopted as the starting weighting, with the review band retained as a safeguard either way.

6. Image-Assisted Issue Classification (Optional Component)

6.1 Purpose and Positioning

Citizens do not always know which category best fits what they see. A blocked stormwater drain may be reported as garbage, and a damaged road edge may be reported as a pothole. Miscategorized complaints are routed to the wrong department and lose time. The optional classification component addresses this by analyzing the uploaded photograph and suggesting a category. The suggestion is displayed to the citizen, who can accept or change it, and the administrator can still correct it during verification. The component is deliberately positioned as an assistant and not as an authority, because classification errors on real-world photographs are unavoidable.

6.2 Model Considerations

The task is a small multi-class image classification problem with classes such as waste, pothole, drainage, streetlight, and an "other or unclear" class. A pretrained convolutional or transformer-based backbone, fine-tuned on a labeled set of civic issue photographs, is the practical approach, because transfer learning reduces the amount of labeled data needed [6]. Where the suggestion must be produced on a phone or on a low-bandwidth connection, a compact architecture in the MobileNet family is a suitable choice [7]. Otherwise the model can run on the server after upload.

For a pilot city, a labeled set on the order of a few thousand photographs distributed across the defined classes, deliberately including night-time and monsoon conditions, is a realistic collection target. Accuracy should be reported per class rather than as a single overall figure, because confusion between visually similar classes, such as pothole and damaged road edge, produces the costliest misroutings; the "other or unclear" class absorbs genuinely ambiguous images and keeps the measured classes honest.

6.3 Confidence Handling

The model outputs a probability for each class. CivicEye shows a suggestion only when the highest probability exceeds a configured confidence threshold. When confidence is low, the interface asks the citizen to choose manually and does not display a guess. This avoids nudging users toward wrong categories. Predictions, user choices, and administrator corrections are logged, which creates a feedback dataset for periodic retraining and for measuring how often the suggestion is accepted.

6.4 Data and Fairness Considerations

Training data should reflect the conditions of the target city: local road surfaces, drain designs, waste containers, night-time images, and monsoon conditions. A model trained mainly on photographs from other regions may misclassify local scenes. Photographs can also contain faces, vehicle number plates, and house

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

fronts. The framework therefore recommends blurring or cropping identifiable people and plates where practical, and restricting access to raw images to authorized staff.

7. Analytics and Decision Support

7.1 Indicators

The analytics layer converts operational records into indicators that help administrators understand both the volume and the quality of service. The principal indicators are summarized in Table 4.

Indicator	Definition	Planning use
Complaint volume	Number of complaints by category, ward, and period	Workforce and budget allocation
Hotspot density	Complaints per unit area, or cluster membership	Targeted inspection and preventive work
Mean and median resolution time	Time from submission to resolution, by category and department	Service-level monitoring
Overdue rate	Share of open complaints past their service-level target	Escalation and staffing
Duplicate ratio	Share of submissions linked to an existing issue	Indicator of severity and visibility of problems
Reopen rate	Share of closed complaints later reopened	Quality of repair and cleanup
First-response time	Time from submission to verification	Administrative efficiency
Citizen confirmation rate	Share of resolutions confirmed by the reporter	Trust and verification of work

Table 4. Principal analytics indicators.

7.2 Operational Definitions

For the indicators to be comparable across wards and periods, each is computed over a reporting window W, normally a calendar month, with its denominator made explicit:

1.0 **First-response time:** the median of ($t_{verification} - t_{submission}$) over complaints verified during W. The median is preferred to the mean because verification delays are heavy-tailed.

1.0 **Resolution time:** the median of ($t_{resolution} - t_{submission}$) over complaints resolved during W. Computing over resolved rather than submitted complaints avoids right-censoring, since complaints submitted late in the window have not yet had time to be resolved.

1.0 **Overdue rate:** the share of open complaints at the end of W whose age exceeds the service-level target of their category.

1.0 **Duplicate ratio:** the share of submissions in W that were linked to an existing parent issue, whether automatically or after review.

1.0 **Reopen rate:** the share of complaints closed in W, or in a trailing three-month window, that were subsequently reopened. The trailing window reduces censoring for issues that reopen slowly.

1.0 **Citizen confirmation rate:** the share of resolutions made visible to the reporter for at least seven days that the reporter confirmed, treating silence as non-confirmation so that the measure stays conservative.

1.0 **Hotspot density:** complaints per square kilometer per category per W, or membership in a DBSCAN cluster computed over a trailing six-month window.

Two presentation rules apply to all indicators. Indicators based on fewer than thirty events in W are displayed with a minimum-volume warning rather than as precise figures, and every rate is shown alongside its absolute count so that small denominators cannot masquerade as stable percentages.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

7.3 Views for Different Users

Administrators see a city-level dashboard with a map of open and recent complaints, filters by category and ward, and trend charts. Department heads see a view limited to their own queue, with aging and overdue indicators and per-officer workload. Planners can export aggregated data for periodic reviews. Citizens may be shown a public, privacy-protected view of the status of issues in their neighborhood, which supports transparency and discourages repeated reporting of known issues.

7.4 From Reporting to Planning

The value of analytics lies in the decisions it supports. A persistent cluster of garbage complaints at a market junction may indicate that collection frequency or the number of bins is inadequate, which is a planning decision and not a cleanup task. A concentration of drainage complaints before and during the rainy season can justify advance desilting in the affected wards. A high reopen rate for road repairs in one zone may point to a problem with materials or workmanship. In each case, the framework shifts the response from reactive case handling toward preventive planning.

8. System Design and Implementation Considerations

8.1 Architecture

CivicEye can be implemented as a conventional three-tier application. Citizens use a mobile application or responsive web client. A backend application server exposes secured interfaces for complaint submission, workflow actions, and analytics queries. A relational database with spatial extensions stores complaints, users, assignments, and history, while an object store holds photographs. The intelligence layer runs as a separate service so that duplicate detection and classification models can be updated without changing the workflow code. A notification service sends SMS, email, and in-application messages.

8.2 Core Data Model

The data model is built around the complaint and its history. The main entities are listed in Table 5. Storing every state change in StatusHistory, rather than overwriting a status field, provides a complete audit trail and is the basis for calculating resolution times and delays.

Entity	Key attributes
User	Identifier, role, contact details, ward or department
Complaint	Identifier, category, description, latitude, longitude, ward, status, timestamps, parent issue identifier
Media	Identifier, complaint identifier, file reference, type (complaint or resolution), capture time
Assignment	Complaint identifier, department, officer, assigned time
StatusHistory	Complaint identifier, previous state, new state, actor, time, remarks
DuplicateLink	Child complaint, parent complaint, score, decision (automatic or reviewed)
Department	Identifier, name, categories handled, areas covered, service-level targets

Table 5. Core data model entities.

8.3 Location Handling

Location data is captured from the device's positioning service and shown on a map for the citizen to confirm or adjust. Because device positioning can be inaccurate in dense urban areas, the framework stores the reported accuracy and allows a manual pin correction. Each complaint is also mapped to an administrative boundary such as a ward using a point-in-polygon test, so that routing and reporting follow the municipality's own organization. Spatial indexes allow nearby-issue queries and duplicate candidate searches to run quickly.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

8.4 Security and Privacy

A civic platform handles personal data and must protect it. The framework applies role-based access control, encrypted communication, and hashed credentials. Contact details of citizens are visible only to staff who need them. Complaint photographs are checked for file type and size, and metadata that is not needed is removed. Public views display issues without the name or contact details of the reporter. Retention rules specify how long images and personal data are kept. Because the framework handles personal data of residents, implementers should align their deployment with applicable Indian data-protection law and municipal policy.

8.5 Preventing Misuse

Open reporting systems can be abused through false, repeated, or malicious complaints. CivicEye reduces this risk through account verification, rate limits on submissions, the mandatory photograph and location, verification by an administrator before routing, and a record of rejected complaints per user. Abuse-prevention measures should be proportionate so that they do not discourage genuine reporting, especially by users who are new to digital services.

8.6 Accessibility and Inclusion

A framework intended for a diverse city must serve users who differ in language, literacy, and device. The interface should support local languages, use icons for categories, and work on low-end phones and weak connections, for example by compressing images and queuing submissions for upload when connectivity returns. Residents without smartphones can be supported through assisted submission at ward offices or through helpline operators who enter complaints into the same system, so that all channels feed one record.

8.7 Integration

Many municipalities already operate helplines, grievance portals, and asset registers. CivicEye is designed to coexist with these through standard interfaces: complaints arriving through other channels can be entered or imported with the same structure, and resolved records can be shared back. This approach avoids the creation of a parallel silo and allows gradual adoption.

8.8 Scalability and Performance

The framework is expected to operate at municipal scale, on the order of tens of thousands of active complaints per year for a large city, which is modest by web-application standards; the operation that grows superlinearly if implemented naively is the duplicate-candidate search. Three design choices bound this cost. First, candidates are pre-filtered by the category-specific radius and the open-or-recently-closed time window before any text or image comparison, so expensive embedding comparisons run only on a small candidate set. Second, the spatial index supports bounded nearest-neighbor queries within a ward, so the search cost is proportional to complaints in the neighborhood rather than the whole city. Third, sentence embeddings are computed once at submission and stored, so comparison is a vector distance rather than a model inference. Under these constraints a single commodity application server is expected to serve a mid-sized municipality, but capacity should be validated in the pilot (Section 9.3) rather than assumed.

9. Evaluation Methodology

This paper presents the framework design. A full empirical evaluation requires a pilot deployment and labeled data, and results from such a study are left for subsequent work. This section specifies how the evaluation should be conducted so that the claims of the framework can be tested in a reproducible manner.

9.1 Research Questions

1.0 RQ1: Does structured capture with photograph and location reduce the proportion of complaints that need clarification before action?

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

1.0 RQ2: How accurately does the multi-signal approach identify duplicate complaints, and how many false merges does it produce?

1.0 RQ3: Does evidence-based closure and citizen feedback change the rate of reopened complaints and the level of citizen confirmation?

1.0 RQ4: How accurate and how useful is the optional image suggestion in helping citizens choose the right category?

1.0 RQ5: Does the analytics layer help administrators identify hotspots and plan resources more effectively than existing records?

9.2 Metrics

For duplicate identification, a set of complaint pairs labeled by administrators serves as ground truth. Precision, recall, and the F1 score are computed for different threshold settings, with particular attention to precision at the automatic-link threshold, since false merges hide genuine issues. For image classification, accuracy, per-class precision and recall, and a confusion matrix are reported on a held-out test set that consists of real user photographs, together with the suggestion acceptance rate in live use. For workflow performance, the metrics of Section 7.1, namely first-response time, resolution time, overdue rate, and reopen rate, are compared before and during the pilot. For user experience, a structured questionnaire such as the System Usability Scale can be given to citizens and staff, and the completion time of the reporting task is recorded.

9.3 Study Design

A recommended design is a pilot in one or two municipal wards over a period of at least three months, with a baseline period of equal length using existing channels for comparison. Where possible, a neighboring ward with similar characteristics that continues with existing channels can serve as a comparison group. Data are collected from system logs and, for the labeled evaluation sets, by administrator annotation. Results should be reported with confidence intervals and with an account of factors that may confound the comparison, such as seasonal rainfall or festival periods that affect waste generation.

9.4 Threats to Validity

Several threats should be considered. Adopters of a reporting application may not represent all residents, so complaint patterns can reflect who uses the platform as much as where problems exist. A rise in complaint counts after launch may reflect easier reporting and not worse conditions. Administrators who know they are being observed may change their behavior. Image models evaluated on photographs from one season may not generalize to another. The evaluation should acknowledge these limits and, where possible, use complementary data such as field inspections.

10. Discussion

10.1 Expected Benefits

The framework is expected to benefit three groups. Citizens gain a simple way to report, receive visible updates, and see evidence that work was done. Field and department staff receive complaints that already include a location and photograph and are relieved of repeated visits for the same issue. Administrators obtain a consistent record of complaint activity that supports performance monitoring and planning. For the municipality as a whole, the combination of transparency and traceability can strengthen trust, which is a precondition for continued citizen participation.

10.2 Relation to Smart City Goals

Smart city initiatives emphasize the use of data and digital services to improve urban management. CivicEye contributes by turning everyday citizen observations into structured geographic data and by

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

providing a feedback channel in which residents are partners in monitoring the city. Its modular design allows later integration with other urban data sources, for example sensors on waste bins or street lighting controllers, so that citizen reports and automated sensing complement each other.

10.3 Limitations

The framework has limitations that should be stated plainly. First, a digital system cannot repair roads or collect waste, and its benefits depend on the capacity and willingness of departments to act on the complaints it organizes. Second, reporting is unevenly distributed across neighborhoods, and wards with higher digital access may appear to have more problems than wards with fewer reporters, which can bias resource allocation if analytics are read without caution. Third, duplicate identification can fail in both directions, merging distinct issues or missing true duplicates, and requires calibration and human review. Fourth, image classification is sensitive to the conditions of the training data. Fifth, location accuracy in dense areas can be limited. Finally, this paper presents a design and evaluation plan, and the effectiveness of the framework in practice remains to be demonstrated through a pilot.

10.4 Ethical and Governance Considerations

Civic data can reveal patterns about neighborhoods and individuals. Transparent policies should state what is collected, who can see it, and how long it is kept. Analytics used for resource allocation should be audited for equity so that quiet wards are not neglected merely because they report less. Where automated components such as duplicate linking or image suggestions are used, the system should remain explainable to administrators and allow human override. Public performance data should be presented honestly, including overdue and reopened complaints, because transparency that shows only favorable numbers undermines its purpose.

10.5 Adoption and Departmental Incentives

Transparency mechanisms succeed or fail on the incentives of the staff they expose. A public reopen rate is unwelcome news for a department whose repairs fail early, and a department head who controls data entry can achieve good-looking numbers by closing complaints without work. Three design responses follow. First, closure is performed by the supervisor rather than the field officer, with citizen confirmation as a parallel check, so that no single actor controls the record. Second, the analytics layer should be introduced to departments as a workload-justification tool, showing where effort is consumed by duplicates and repeat failures, before any public dashboard is released. Third, phased transparency, in which internal performance views precede public ones, gives departments time to improve processes rather than numbers. These are governance arrangements rather than software features, but the framework's credibility depends on them.

11. Conclusion and Future Work

This paper presented CivicEye, a digital framework for reporting, managing, and analyzing urban civic issues. The framework combines structured citizen complaints that contain text, photographs, and location with a verification and routing process, a defined lifecycle with evidence-based closure, a multi-signal approach to duplicate identification, an optional image-based categorization aid, and an analytics layer that supports planning. By treating the complaint as a traceable record from submission to closure, the framework aims to improve transparency, reduce redundant effort, and give authorities evidence for allocating resources.

Several directions are open for future work. The first is a pilot deployment in a municipal ward to measure the indicators and answer the research questions in Section 9. The second is the development of a labeled dataset of civic issue photographs and complaint texts in local languages, which would benefit the wider research community. The third is the extension of duplicate identification to use multilingual embeddings and learned weights in place of hand-tuned parameters. The fourth is predictive analytics that forecast seasonal complaint peaks, such as drainage issues before the monsoon, so that preventive action can be

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

planned. The fifth is integration with sensor data and municipal asset registers. Finally, studies of citizen engagement and equity of reporting across neighborhoods would help ensure that the benefits of such systems reach all residents.

Acknowledgment

The author thanks the faculty and colleagues of SVR College of Commerce & Management Studies for their guidance and support.

References

[1] S. R. Arnstein, "A ladder of citizen participation," Journal of the American Institute of Planners, vol. 35, no. 4, pp. 216-224, 1969.

[2] T. Nam and T. A. Pardo, "Conceptualizing smart city with dimensions of technology, people, and institutions," in Proc. 12th Annual International Digital Government Research Conference (dg.o), 2011, pp. 282-291.

[3] M. F. Goodchild, "Citizens as sensors: the world of volunteered geography," GeoJournal, vol. 69, no. 4, pp. 211-221, 2007.

[4] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in Proc. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2019, pp. 3982-3992.

[5] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in Proc. 2nd International Conference on Knowledge Discovery and Data Mining (KDD), 1996, pp. 226-231.

[6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.

[7] A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, 2017.

[8] mySociety, "FixMyStreet: report, manage and fix local problems," https://www.fixmystreet.com and platform documentation at https://www.fixmystreet.org (accessed October 2026).

[9] SeeClickFix, municipal service-request reporting platform, https://seeclickfix.com; see, e.g., City of Albany, NY, "See Click Fix," https://www.albanyny.gov/2443/See-Click-Fix (accessed October 2026).

[10] Ministry of Housing and Urban Affairs, Government of India, "Swachhata App and Swachh City Portal," https://www.swachh.city (accessed October 2026).

[11] "CivicSync: Local Civic Complaint Management System," International Journal of Creative Research Thoughts (IJCRT), 2026. Available: https://ijcrt.org/papers/IJCRT2604638.pdf

[12] M. Shakunthala, "Crowdsourced Civic Issue Reporting and Resolution System," IJEETR, 2026. Available: https://www.ijeetr.com/index.php/ijeetr/article/download/527/491

(This work is licensed under a Creative Commons Attribution 4.0 International License.)