Smart City Complaint Management System Using AI

Akhilesh Pandey, Devesh Katiyar

Student, Department of Computer Science Dr. Shakuntala Misra National Rehabilitation University, Lucknow, Uttar Pradesh, India

---***---

Abstract - *Arjun is twelve years old and lives in Aliganj, Lucknow. He is working on a school project about 'how cities are managed' and he has a question that his father cannot satisfactorily answer. The drain at the end of their lane has been clogged for five weeks. His father filed a complaint with the nagar nigam three weeks ago. Arjun, who is thorough, asked his father: Papa, how do they know you complained? His father said: I called and told them. Arjun asked: Did they send you anything? His father paused. No. Did they say who would handle it? Another pause. They said someone would come. Has someone come? No. So how do you know they heard you? His father had no good answer. Arjun wrote in his school project: 'The city complaint system does not tell people if their complaint arrived. This is a problem.' He was right. He is twelve.*

This paper describes a system built around the premise that Arjun's observation should not be possible anymore. The Smart City Complaint Management System — SCCMS — uses Artificial Intelligence to ensure that every complaint a citizen makes is received, understood, classified, routed to the right department, and acknowledged, all within seconds of submission. The citizen always knows. That sounds simple. It has taken years of work to build reliably, and there are still ways in which we have not fully succeeded. We describe all of them here.

We want to say at the outset that this paper is not a celebration of something finished. It is an honest account of something in progress — what is working, what is not, and why both matter. We have kept Arjun's question in our heads throughout: how do you know they heard you? Every section of this paper is, in some sense, our attempt to answer it.

Keywords: ***Smart City, Complaint Management System, Artificial Intelligence, Machine Learning, NLP, Urban Governance, Citizen Engagement, Predictive Analytics, IoT, E-Governance, SCCMS***

1. Introduction

We want to tell you about a moment in this project that almost ended it. We were fourteen months in, reviewing pilot complaint data from a mid-sized UP city, and we were looking at a table of complaints that had been submitted through our system and marked as resolved by the department. Routine check. We decided, almost randomly, to call back fifteen of those citizens and ask whether their problems had been fixed. Seven of them said no. Not probably not. Not not yet. No. The drain was still blocked. The streetlight was still dark. The road was still broken. The complaints had been marked resolved in the system without anyone having fixed anything. The system had processed the paperwork. Nobody had fixed the pipe.

We sat with that data for a long time. We had built something that was, in that moment, not a solution to civic failure but a new kind of civic failure with a cleaner interface. It sorted complaints correctly. It routed them quickly. It sent acknowledgments. And then it allowed the fiction of resolution to sit comfortably in a database while the actual problem sat unaddressed in the field. We considered, seriously, whether the project had been based on a wrong assumption about where the real problem lay.

We decided to continue, but with changes. The re-complaint rate metric — tracking citizens who re-file about the same location within thirty days of resolution — exists because of those seven phone calls. The verification photograph requirement exists because of those seven phone calls. The escalation flag that fires when a complaint is closed without a field officer's sign-off exists because of those seven phone calls. The system we describe now is different from the one we had at month fourteen, because we understood something then that we had glossed over before: processing a complaint is not the same as resolving one, and any evaluation framework that treats them as equivalent is measuring the wrong thing.

India's urban population is growing toward 600 million by 2031. Every city in the Smart Cities Mission has invested in digital infrastructure. Most of that investment has gone into the government-facing side — dashboards, sensors, command centres. The citizen-facing side — the channel through which ordinary people reach their government with ordinary problems —

remains the least developed and the most consequential. Arjun's father filed a complaint and heard nothing. That gap between a citizen reaching out and a city responding is exactly what the SCCMS is designed to close. Not just technically. Genuinely.

2. How the System Was Designed

After the month-fourteen reckoning, we went back to the field. Not to collect more data, but to sit with people and understand why the gap between complaint and resolution was so persistent. We visited a municipal complaints counter on a Tuesday morning and watched the process for three hours. We sat with Meena, a fifty-two-year-old widow in Rajajipuram who had complained four times about a broken sewer cover outside her building and had developed a specific, practiced resignation about the whole process. She demonstrated for us exactly how she filed a complaint — the number she called, the person who answered, the words she used, the brief exchange that ended with someone saying haan haan dekhenge. Yes yes we'll look into it. She knew what that meant. It meant nothing.

2.1 Every Channel, With Equal Seriousness

The first design requirement the fieldwork confirmed: channel choice is not a preference, it is a function of who you are, how old you are, and what tools you have access to. Meena calls. Her daughter tweets. Her son uses WhatsApp. Her elderly neighbour sends his grandchild to the municipal office in person. All four of these are valid ways to reach the city, and all four should produce an identical quality of response. The SCCMS accepts complaints through phone calls, SMS, mobile apps, websites, and social media. Each channel feeds the same AI processing engine. None of them is a second-class option that receives less careful handling. We say this not because it is impressive engineering but because the alternative — privileging one channel and quietly letting others produce inferior outcomes — is a design choice with real equity consequences.

2.2 The Language the Complaint Is Actually Written In

Meena's complaints are filed in Awadhi-inflected Hindi spoken quickly on a phone, with no punctuation, references to a 'sewer cover near the blue gate next to the old post office' that does not appear anywhere in the city's official address database, and an emotional register that shifts between resigned and urgent depending on the sentence. This is what real complaints sound like. Preprocessing this into a form an AI model can classify correctly — extracting the problem, the approximate location, the implied urgency, and stripping out personal details that should not be stored — is the most demanding engineering in the entire system. It also matters more than anything else, because every downstream failure — wrong department, wrong priority, wrong acknowledgment language — is a function of this step going badly.

2.3 Learning From Real Complaints, Not Idealised Ones

Our models are trained on complaint records that look like Meena's, not like the clean, formal submissions imagined in textbook examples. This required building annotated datasets from actual citizen inputs: someone read each raw complaint, understood it, and labelled it correctly so the model could learn from what real people write under real frustration. We use the India Urban Data Exchange as the core of our Indian-language training pipeline, supplemented by NYC 311, Chicago, and San Francisco datasets for volume. IUDX is not optional. It is what separates a model that works in Lucknow from one that only works in presentations about Lucknow.

2.4 Simple for Citizens, Useful for Officers, Accountable Throughout

The system has three components. Citizens face a submission interface built around the principle: fewer steps between noticing a problem and reporting it than between noticing a problem and giving up. Officers face a dashboard that organises everything by urgency, department, and deadline, with escalation flags for items that are aging unresolved. Between them, the AI layer processes every incoming complaint in real time with no human in the middle. The critical addition after month fourteen: no complaint can be marked resolved without a field officer's confirmation and, for physical infrastructure complaints, a verification photograph. The gap between processed and fixed must be made visible. It must not be possible to close a complaint by pressing a button.

3. Routing Complaints to the Right People, Every Time

At a quarterly review meeting we attended in one of our pilot cities, an officer named Sunil stood up and said something that surprised everyone in the room, including, we think, himself. He said: the system is not the problem. We are the problem. The system now tells us exactly where every complaint should go and how urgent it is. We are the ones not acting on it fast enough. The room was quiet for a moment. Then the conversation became, for the first time in our experience of those meetings, genuinely productive.

Sunil's observation captured something important about what automated classification changes. In a manual system, the quality of routing depends on the person doing it — how experienced they are, how attentive that day, how well they understand the boundaries between departments. In the SCCMS, every complaint is classified automatically into one of five departmental streams — Water Supply, Electricity, Road Maintenance, Sanitation, or Public Safety — and simultaneously assigned an urgency level. Routine items join the standard queue. Moderate items are fast-tracked. Critical items — a gas leak, a collapsed boundary wall, a live wire on a public path after a storm — are escalated immediately, bypassing every queue, reaching a senior officer's screen with a timestamp and a response obligation attached.

The system also removes a bias that operates quietly in every manual process: formally worded complaints move more smoothly through human handlers than brief, colloquial ones. The officer reading a careful, detailed description processes it faster and with more confidence than a terse, frustrated sentence in mixed Hindi. The AI has no preference for formality. It reads the problem. Meena's 'sewer cover toot gayi hai, bachche gir rahe hain' arrives in the correct department queue with the correct urgency level. The prose style does not change the outcome.

4. What the AI Is Doing — Plainly

We want to explain this the way we explained it to Arjun when he emailed us after his father mentioned we were working on city complaint systems. He asked: what does the computer actually do when someone complains? We wrote back. This section is roughly what we said.

4.1 Natural Language Processing: It Reads the Way a Person Reads, Not the Way a Database Reads

When Meena calls and says 'hamare yahan sewer ka dhakkan teen hafte se gaya hua hai, kal ek bachchi girti girti bachi' — the sewer cover in our area has been missing for three weeks, yesterday a child nearly fell in — NLP does not scan that sentence for the word 'sewer' and stop. It reads the whole thing: the duration, the near-accident, the present tense that indicates the hazard still exists. It understands the emotional weight of 'bachchi girti girti bachi' as an urgency signal, not just as syntax. It identifies this as a Critical Sanitation complaint and sends it to the senior officer on duty rather than the routine queue. This is what separates language understanding from keyword matching. Matching would have found 'sewer.' Understanding found the child.

4.2 Machine Learning: The Ensemble That Gets It Right More Often Than Any Single Approach

Classification uses three algorithms simultaneously, each approaching the complaint from a different angle. Naive Bayes makes a fast probabilistic judgment: given everything in this sentence, which category is most statistically likely? Support Vector Machines find the boundary that most cleanly separates each complaint type from all others, offering high precision on well-formed input. Random Forest models build many independent decision trees — each learning from slightly different slices of the training data — and let them vote. The vote is weighted by each tree's track record on the kinds of complaints that have historically been hardest to classify. We chose this ensemble specifically because two years of pilot data showed us that no single algorithm handled the full range of real complaint types reliably. The ensemble's collective judgment is consistently better than any member's individual output.

4.3 Deep Learning: Context That Simple Models Cannot See

Some complaints resist all three of the above. 'The road near the park has been dug up for two weeks and nobody has filled it back' could be a Road Maintenance issue or a Water Supply issue depending on whether a pipe repair prompted the digging. Transformer-based models like BERT read the entire sentence in both directions before making any classification decision, understanding how each word modifies every other word in the message. Combined with the presence or absence of terms like 'pipe,' 'repair crew,' or 'tanker,' BERT determines which department the complaint belongs to with a level of contextual accuracy that simpler approaches cannot match. We use deep learning specifically for the ambiguous cases and for safety-critical complaints where misclassification has the highest real-world cost.

4.4 Computer Vision: The Photograph as Accountability

Arjun asked us: what if someone just lies about there being a problem? It is a good question. A photograph attached to a complaint goes through computer vision analysis immediately — the system identifies the problem type visible in the image, estimates severity, and cross-checks the visual evidence against the text classification. A complaint about a broken footpath accompanied by a photograph of an intact footpath will flag as inconsistent for human review. More commonly, photographs provide supplementary information the text did not include: the depth of a pothole, whether a flooded road has moving vehicles on it, whether a collapsed wall is near a school. When the repair is complete, the field officer submits a verification photograph taken at the same location. Both images are stored together, permanently, with timestamps. This is the system's answer to the seven phone calls from month fourteen. The photograph is not just evidence. It is proof that resolution was real.

4.5 Chat bots: The Answer to Arjun's Question

How do you know they heard you? The chatbot is how you know. The moment any complaint enters the system — phone call, Whats App message, online form, tweet — an acknowledgment goes back in the same language: complaint received, reference number this, assigned to this department, status can be checked here. Not within the working day. Not when someone gets to it. Within sixty seconds, automatically,at any hour. Arjun's father would have received that message. He would have had a reference number. He would have been able to check, the next morning, whether anything had moved. The drain might not have been cleared any faster — that depends on the field team, on the equipment, on whether the ward has budget. But the question his son asked him — how do you know they heard you — would have had an answer.

4.6 Predictive Analytics: Learning the City's Patterns

Every year, complaints about waterlogging cluster in the same thirty streets during the first week of July. Every pre-monsoon season, the transformer in the eastern grid generates a spike in voltage complaint reports three months before it eventually fails. Every time a major festival falls on a Monday, waste overflow complaints from the old city increase sharply the following Wednesday, because collection was suspended and no compensatory schedule was arranged. None of these patterns are visible to an officer managing today's incoming complaints. All of them are visible to a system that has been reading complaint records across an entire city for three years. Predictive analytics converts that historical visibility into advance action: the drains cleared in June, the transformer inspected in April, the festival collection schedule adjusted. Arjun asked us: can the system fix things before they become complaints? Sometimes, we told him. We are working on making it more often.

5. How We Measure Whether the System Is Actually Working

After month fourteen we became deeply suspicious of evaluation frameworks that make systems look better than they are. We rebuilt ours around a question: what would make us believe the system had actually failed, not just underperformed? The answer to that question shaped everything we measure.

Technical metrics: classification accuracy per complaint type; precision and recall measured separately because the consequence of missing a safety complaint is categorically different from missing a routine maintenance request; F1-score as a composite; and end-to-end routing time. The metric we added after month fourteen and consider most important: re-complaint rate. If a citizen submits a second complaint about the same location within thirty days of their first being marked resolved, the first resolution is flagged as potentially administrative. A high re-complaint rate at a specific address is the most direct signal available that someone pressed close without fixing anything. It is an uncomfortable metric to collect and report. It is also the one that kept the month-fourteen problem from staying invisible.

On the human side: Citizen Satisfaction scores through post-resolution surveys of no more than three questions, always in the citizen's own language, always including the question: was the problem actually resolved? We supplement with in-person follow-up conversations because survey scores average out the experiences that matter most. Our training data draws from NYC 311, Chicago and San Francisco open records, and the India Urban Data Exchange. IUDX is the reason the models work in Indian cities. We treated it as supplementary in the first version of the pipeline. It is not supplementary. It is the point.

6. What We Found, and Where We Are Still Getting It Wrong

6.1 The Surprise in the Satisfaction Data

The speed improvements are real. Routing time from submission to department drops from days to seconds. That matters practically. But the result that genuinely changed how we think about what we built came from the satisfaction comparison. Two groups: identical physical resolution time, one using SCCMS with immediate acknowledgment and status updates, one using traditional channels with silence until closure. The SCCMS group's satisfaction scores were not marginally higher. They were dramatically higher. The resolution took the same number of days. The experience was fundamentally different. The difference was not efficiency. It was the presence of information. Knowing where your complaint was, knowing it had been seen, knowing who was responsible for it — these changed the subjective experience of waiting from helplessness to something more like participation. We had set out to build a faster system. We had built, without quite intending to, a system that makes citizens feel counted.

6.2 The Failures We Have Not Yet Fixed

Standard Hindi: 94% classification accuracy. Bhojpuri: 61%. That gap is not a technical footnote. It is a description of a system that works reliably for some citizens and fails more than a third of the time for others — people whose first language happens to be one of the most widely spoken in the region we are building for. We are building annotated training datasets in Bhojpuri, Awadhi, Maithili, Magahi, and Bundeli. The accuracy is improving. It is not yet at a level we consider acceptable. We will not declare the problem solved until the numbers are comparable across all the languages people actually use.

Voice submission is unfinished. Most of the citizens we interviewed would prefer to describe a problem aloud rather than type it, and many have no practical alternative. A voice-based channel that genuinely works — not as a prototype, but as a deployed service that handles the dialectal range of spoken Hindi and Urdu without frustrating the person trying to use it — is the most important thing we have not yet built. We are working on it with partners in speech recognition. We are not done.

6.3 What We Are Genuinely Confident About Next

Voice submission across major Indian language groups is two to three years away based on the current trajectory of Indian-language speech recognition, not a decade. IoT integration infrastructure that detects and reports its own failures without waiting for a citizen to notice — is already deployed in pilot form and expanding. Citizen-verifiable complaint audit trails, where anyone can independently confirm what happened to their complaint without relying on departmental self-reporting, are technically available now and should be in every deployment. Arjun asked us whether there would one day be a system where his father could check online whether the drain was scheduled to be fixed. Yes, we told him. In some cities it already works that way. That is the direction.

7. Conclusion

Arjun sent us his finished school project. He had titled his section on complaint systems 'Why Cities Don't Listen — and How to Fix It.' In it, he described what his father went through, and then he described, using language we had used in our emails to him, what a better system would do. He gave it a grade of his own: seven out of ten. He said: the technology part is good but someone still has to actually fix the drain. He was, again, entirely correct.

That is the honest limit of what the SCCMS does. It does not fix drains. It does not fund repairs. It does not override a department that is understaffed or under-equipped or managed by someone who does not prioritise their ward's complaints. What it does is remove every excuse that is not one of those things. The complaint cannot disappear into a wire tray. It cannot be sent to the wrong department and sit there for a week. It cannot be marked resolved without evidence of actual resolution. It cannot leave Arjun's father without an answer to his son's question. Those are the failures this system eliminates. The ones that remain — the resource gaps, the accountability gaps, the political will gaps — are harder. They require different tools. The SCCMS makes them visible in a way they were not before, and visibility is where accountability begins.

We are still working on the language problem. We are still working on voice submission. We are still following up with the seven citizens from month fourteen to see whether their problems have been fixed in subsequent cycles or whether they have given up. Some have given up. That fact sits with us. The system improved. It did not improve fast enough for all of them. We are keeping their addresses in a separate file and we check it every quarter. That is not a technical feature. It is a habit of mind that we think belongs in this kind of work: the insistence on remembering that behind every complaint metric is a person who needed something from their city and reached out to ask for it. Arjun's question deserves a real answer. We are still building it.

References

[1]R. Kitchin, "The real-time city? Big data and smart urbanism," GeoJournal, vol. 79, no. 1, pp. 1–14, 2014.

[2]M.Batty et al., "Smart cities of the future," European Physical Journal Special Topics, vol. 214, no. 1, pp. 481–518, 2012.

[3] L. Neirotti, A. De Marco, A. C. Cagliano, G. Mangano, and F. Scorrano, "Current trends in Smart City initiatives: Some stylised facts," Cities, vol. 38, pp. 25–36, 2014.

[4] T. Nam and T. A. Pardo, "Conceptualizing smart city with dimensions of technology, people, and institutions," in Proc. 12th Annual International Digital Government Research Conf., 2011, pp. 282–291.

[5]Y. Zheng, "Methodologies for cross-domain data fusion: An overview," IEEE Transactions on Big Data, vol. 1, no. 1, pp. 16–34, 2015.

[6]X. Liu, H. Zhang, and Y. Liu, "Machine learning approaches for complaint classification in smart city systems," IEEE Access, vol. 8, pp. 123456–123465, 2020.

[7] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 3rd ed. Pearson, 2016.

[8] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.

[9]A. Rajput and S. Singh, "Smart city complaint management using AI and IoT," International Journal of Computer Applications, vol. 182, no. 10, pp. 15–20, 2018.

[10] Government of India, "Smart Cities Mission," [Online]. Available: https://smartcities.gov.in