Leadership-in-Action Report: AI for Air Quality at eGov Foundation

Rishi Dinesh

Georgetown University

Acknowledgement

I would like to thank the Laidlaw Scholars Foundation and the Georgetown University Laidlaw Scholars Program for supporting my research. I would like to thank Colleen Dougherty for her kindness and guidance. I would like to thank my family for their support, without which I would not be able to do this. Finally, I would like to thank my colleagues and peers at eGovernments Foundation in Bangalore for their mentorship, wisdom, and guidance throughout this research process.

Executive Summary

During my Leadership-in-Action (LiA) placement at eGovernments Foundation (eGov) in Bangalore, India, I researched how artificial intelligence (AI) could help reduce air pollution. I wrote a final research brief that would advance eGov's efforts and support an upcoming multi-stakeholder workshop between communities, policymakers, and technologists.

The project produced three concrete outcomes. First, I held four direct conversations with an eGov software engineering team, a former state official, an official at India's AI Centre of Excellence, and a philanthropic funder. Second, I built a stakeholder map of 33 groups across three levels (affected communities, emitters, and authorities), including community organizations and informal-worker advocacy groups that I researched but did not interview. Third, I wrote a brief assessing AI applications across five stages of air quality governance. My manager, Kriti Mittal, and the head of eGov's Policy team, Anumita Raj, reviewed the brief. I revised it in response to their feedback, and it was sent to every workshop participant.

Two caveats frame everything below. The workshop, which aims to bring communities, policymakers, and technologists into one room, has not yet been held – it was delayed until after I left India. Also, my engagement with the communities most exposed to air pollution was through desk research and not conversation, which is a limitation of the work and also one of its findings.

The central conclusion of my final brief at eGov is that AI's value for air quality depends less on the sophistication of the models than on the data systems, institutions, and politics that support them.

Project Evolution

eGov is a Bangalore-based non-profit that works with governments, multilateral organizations, and civil society to expand access to public services through Digital Public Infrastructure (DPI). It is best known for DIGIT, an open-source platform that powers urban services in several Indian states, and it is a founding partner of the National AI Centre of Excellence for Sustainable Cities. That position sits at the intersection of technology, community empowerment, and climate action, which is why I proposed to work there.

My original LiA outline was titled *Empowering Communities and Supporting NGO Capacity Through Responsible AI*. It identified a gap: states are deploying AI in decisions that shape the lives of Dalit, Adivasi, women, and other marginalized communities, while the NGOs that reach those communities are largely absent from the conversation. I planned four objectives: structured conversations with 8–10 NGO leaders and advocates, a scoping memo on AI and sustainability, a community-informed policy brief, and support for a stakeholder event.

Once I was inside the organization, the scope changed. My managers and I recognized that air pollution was an unusually high-impact place to apply that same question, since it is tied directly to health, justice, and equity. The numbers in my brief show why – the average Indian is exposed to 83 $\mu g/m^3$ of fine particulate matter (PM2.5) daily, more than 16 times the World Health Organization's guideline. Long-term exposure caused an estimated 1.67 million deaths in 2019.

Low-income communities, including slum dwellers, migrant workers, small farmers, and Scheduled Castes and Tribes, carry a disproportionate share of that burden. These are the same groups my original outline set out to center.

The pivot also changed the shape of my deliverables. Instead of 8–10 conversations spread across NGO sectors, I held four conversations with institutional and technical stakeholders and complemented them with structured desk research on community and worker organizations. Instead of a separate scoping memo and a longer policy brief, I wrote one brief that does both jobs: it scopes AI applications and sets out the governance conditions they depend on.

Outcomes and Objectives

My outline set four SMART objectives, each tied to a Laidlaw attribute. For each, I set out what I planned, what I delivered, and what the gap between them taught me.

Objective 1: Community Engagement (Empathy and Active Listening)

On this objective, I planned to hold 8–10 structured conversations with NGO leaders, community organizers, and frontline advocates, followed by a written synthesis shared back with participants.

In actuality, I held four direct conversations with an eGov software engineering team, a former state official, an official at the AI Centre of Excellence (a Central government-backed research and policy initiative), and a philanthropic funder. Together they gave me four distinct views of the same problem: how technology is built, how government decisions get made and stalled, how national AI policy is shaped, and how money moves toward solutions. I complemented these conversations with desk research on the groups who live with the consequences. I studied community organizations and advocacy groups representing informal workers, including organizations working with outdoor informal workers, waste pickers, slum dwellers, women who cook with solid fuels, and pregnant women and newborns. I read their published work and programs to understand how they describe the problem and what they ask of policymakers.

Ultimately, I fell short of the original target. The conversations I held were with people closer to institutions than to affected communities. Desk research is a real substitute for learning about those communities' priorities, but it cannot replace hearing from them directly. I treated that gap as a design problem for the next stage: the workshop is built to bring those voices into the room, and the brief was sent to all workshop participants so that anyone it describes, or fails to describe, can respond to it.

Objective 2: AI and Sustainability Scoping (Rigor and Evidence-Based Thinking)

My goal here was to identify three to five intersections between AI and eGov's sustainability work, and assess each for feasibility, community readiness, and governance prerequisites, being as clear about risks as about opportunities.

The brief assesses AI across five stages of air quality governance: mapping airsheds, monitoring through low-cost sensor calibration, forecasting, compliance, and emergency response. Each is grounded in real examples, such as machine learning to identify airshed boundaries, forecasting

models that predict air quality six days ahead, Delhi's AI-enabled DUST 2.0 platform for construction-site compliance, and AI-optimized parametric insurance that pays informal workers when pollution crosses a threshold. For each, the brief ties promise to prerequisites. For example, one in five National Clean Air Programme cities has no continuous monitoring station, and 80% of India's industrial emissions monitors are missing over 1,000 hours of data, so models built on that data inherit its gaps.

My manager and the head of the Policy team recommended that I expand the section on AI skepticism and describe the use cases more clearly. I made both changes in the final draft. Separately, my background in American government policy helped me identify a software package developed by the US Environmental Protection Agency that could be replicated or adapted for an Indian context, which I flagged as an example of learning across governance systems without assuming that any one system's tools transfer unchanged.

Ultimately, this objective matters most to me because AI discussions tend toward optimism. The strongest form of rigor I could offer eGov was a brief that says where AI helps, where it cannot yet, and what has to be true first.

Objective 3: Brief (Translating Knowledge into Action)

Here, I planned to put out a policy brief informed by community conversations and scoping, usable by eGov with its state and national partners.

I delivered a final brief, reviewed and refined as described above, to serve two audiences. It was used to align workshop participants on how to approach the conversation about AI and air quality. It was also shared within eGov to inform how software teams approach the problem, which is where the brief's argument that AI systems should be built on a shared Digital Public Infrastructure (DPI) layer – i.e., reliable data standards, open interfaces, and interoperable registries – is most directly actionable.

Objective 4: Stakeholder Event Support (Collaborative Leadership)

Here, I wanted to support end-to-end planning of a stakeholder event, ensuring civil society is meaningfully represented.

My contribution was the stakeholder mapping that underpins the workshop. I mapped 33 groups in three levels. Level 1 holds 7 groups directly affected by poor air quality, from urban residents and cardiovascular patients to outdoor informal workers, waste pickers, and pregnant women. Level 2 holds 7 groups whose activities emit pollutants, including brick kilns, stubble burning, freight transport, coal power, construction, open waste burning, and mining. Level 3 holds 19 actors who can set standards, enforce them, generate evidence, or fund change: regulators, ministries, scientific institutions, think tanks, donors, and technology organizations. For each group I identified a representative organization or body and a short rationale for why that group belongs in the room.

The workshop has not yet been held, so I cannot report its outcome. The mapping itself, though, makes the original problem visible. Most of Circle 1's representatives are civil society

organizations, while most of Circle 3 are institutions. Putting the three circles in one room is precisely what rarely happens, and it is what the event is designed to change.

Key Findings

The brief's most important finding is that India's air quality problem goes far beyond a simple lack of algorithms – it is more politically, systematically entrenched. India launched the National Clean Air Programme in 2019 to bring 131 cities up to national standards, yet most have recorded only modest improvements or worsening levels. The constraints are incomplete monitoring data, limited institutional capacity, and underdeveloped emergency response. Four findings follow from that.

1. Data is often unavailable or unusable. Many cities lack continuous monitoring stations, and legacy stations that sample only biweekly are used to fill the gaps. Readings from the two main station networks sometimes disagree, and there is no standardized guidance for procuring, siting, or calibrating low-cost sensors. Any AI system trained on this data would reproduce its blind spots, most likely in the neighbourhoods that already have the least monitoring.
2. Data does not flow between institutions. India lacks a unified, interoperable architecture that connects station, industrial, low-cost sensor, and satellite data through common standards and open interfaces – often across multiple levels of government. Building one would require formal multi-level data sharing and shared quality-assurance rules at the Central level.
3. Monitoring is disconnected from response. The brief contrasts Punjab, where satellite-detected stubble burning triggers live text alerts to the state pollution board and prompts field officers to record evidence, with Bihar, where 70% of industrial emissions monitors operate offline and agencies must wait for manually reported data. Outside Delhi-NCR's tiered emergency framework, few cities have automatic triggers that turn a reading into action.
4. Institutions are the binding constraint. State Pollution Control Boards (SPCBs) are broadly understaffed and short on technical expertise, and their directives are often left unimplemented. A forecasting tool that no one is equipped to act on cannot clean the air.

These findings lead to the brief's central argument: AI for air quality should be built on shared digital public infrastructure (DPI), not deployed as stand-alone tools, and it should be introduced where the supporting data and institutions are ready. The equity dimension follows. The same communities that bear the heaviest exposure are the least likely to appear in the data and in the rooms where these systems are designed, which is why the stakeholder map and the workshop matter as much as the technology.

Project Challenges

My outline named four projected challenges: limited community access and trust, data sensitivity, fast-moving policy, and the danger of imposing a Western AI paradigm. Some materialized, some did not, and the real challenges were not all the ones I predicted.

Outreach

This was the largest gap between plan and outcome. Frontline organizations are stretched thin and reaching them in the time available was harder than my outline assumed. I responded by being clear about what I had done, using desk research to represent these groups on their own terms, and designing the stakeholder map so that the people most affected sit at its center and in the workshop's invitation list.

Scope

Moving from NGO adoption of AI to air quality meant starting a new technical subject, with its own science, institutions, and vocabulary, during the placement. I handled it by anchoring the work in the governance question I already cared about, which is who benefits from and who is excluded from AI systems, and by learning the subject through the conversations I held and a structured review of the landscape.

Timelines

The workshop was delayed until after I left India, which meant I could not see my stakeholder mapping tested. I responded by making the mapping self-explanatory, with a rationale for each group, so that eGov's team can use it without me.

Imposing Western Views

This was the risk I took most seriously. The brief's emphasis on prerequisites, the expanded section on AI skepticism my reviewers asked for, and my caution about the EPA software package all come from this concern. Despite my concern (or maybe because of it), I felt that I successfully avoided imposing my frameworks on eGov's research. Instead, I treated AI as one possible tool among several, and a tool whose success depends on local data, local institutions, and local consent. I adopted eGov's approaches, finding solutions through "samaaj, sarkaar, bazaar" – society, government, and the market.

Lessons for Myself

I came to eGov planning to listen to communities. What I learned is that listening at the scale I intended requires relationships that are built before a project starts, not within it. If I did this again, I would begin identifying and approaching community organizations in the first week, and I would ask eGov's partner network for introductions to the groups on my stakeholder map upon its completion.

Second, writing a brief for an organization that wants to act made me more disciplined about separating what AI has been shown to do from what it is hoped to do. My reviewers' request to expand the skepticism section was the most useful feedback I received, because it pushed me to make the brief's caution as specific as its optimism.

Third, studying American government policy gave me a concrete advantage: I could recognize an EPA-developed software package as a candidate for adaptation. The same background taught

me to ask what would have to change for it to work in India, including its data, institutions, and incentives, before suggesting it should.

Coming from the US also gave me an interesting view of how societies structure relationships. I noticed the contrast between communal and individualist habits, such as how decisions are shared, how responsibility is felt, and how we relate to institutions and markets. That contrast is directly relevant to technology. An air quality alert that assumes individuals will protect themselves, for instance, reflects a different idea of responsibility from one that assumes households, workplaces, and communities will respond together. It reinforced my conviction that technology for communities has to be designed with local norms in mind.

Recommendations for eGov

Based on this work, I offer the following recommendations to eGov. They follow from the brief's findings and from the limits of my own engagement. eGov should:

1. Use the stakeholder map to ensure that affected communities and the organizations that represent them (informal-worker groups, community health and housing organizations, and waste-picker groups) take part as participants, not just as subjects of discussion.
2. Begin engaging with the groups I mapped but did not interview. Start with the organizations already in eGov's partner network.
3. Select applications for pilots where monitoring data, institutional capacity, and a response mechanism are already in place, and document honestly where they are not.

Alignment with the SDGs and Laidlaw values

1. SDG 11 (Sustainable Cities and Communities): air quality is a core urban sustainability issue, and my work centered participation in it by mapping who should be at the table before solutions are designed.
2. SDG 13 (Climate Action): many sources of PM2.5, including coal power, transport, and open burning, also drive climate change, so the same institutional and data foundations support both goals.
3. SDG 16 (Peace, Justice and Strong Institutions): the brief's finding that institutional capacity is the binding constraint, and its emphasis on inclusive decision-making, speak directly to accountable and participatory governance. Because of the health burden on low-income communities, the project also connects to SDG 3 (Good Health and Well-being).

Laidlaw Attributes

Each objective drew on one: empathy and active listening in the stakeholder conversations and the choice to map affected communities, rigor and evidence-based thinking in the assessment of AI applications and the expanded skepticism section, translating knowledge into action in a brief that two audiences have used, and collaborative leadership in the stakeholder mapping that prepares the workshop. Across all four, I tried to act as a global citizen by listening before proposing, by being honest about what I had and had not done, and – as eGov itself strongly

believes in – by treating technology as something to be shaped with communities and not delivered to them.