Democratic AI:

Of the People, by the people, for the people

Max Michaels | Kenzo Fujisue

In *The Wizard of Oz*, the most unsettling fact is not that the Wizard is an imposter. It is that almost no one asks what lies behind the curtain until the damage is already done. Authority is accepted because it sounds confident, looks impressive, and arrives wrapped in the machinery of certainty. That is how an illusion of intelligence becomes a structure of power.

That is a useful way to think about AI in public life. The danger is not that machines suddenly become sovereign. It is that institutions begin to exercise power through them while making democratic responsibility harder to see. AI becomes the curtain behind which authority can move faster, speak more confidently, and explain less.

AI will not present itself as a new regime of power. It will arrive as assistance: a chatbot that answers questions, a system that organizes cases, a tool that processes forms, a dashboard that helps officials prioritize. Each function may seem modest on its own. Together, they create a new way for institutions to see, sort, recommend, and act.

That shift matters because once AI enters welfare offices, classrooms, clinics, courtrooms, tax departments, newsrooms, and the phones in people's hands, it ceases to be merely administrative software. It becomes part of the machinery of public authority. It influences what officials notice, what citizens are told, what cases move faster, and which priorities take precedence. In that sense, AI gives organizations a new way to make consequential decisions without always taking visible responsibility for them.

That is why the core issue is not efficiency alone, but artificial responsibility: how AI allows institutions to act while making responsibility harder to find. The governance problem begins when a consequential outcome is produced by a chain of systems, vendors, prompts, models, and approvals, yet no one person or office can clearly explain who decided what, under whose authority, and with what right of appeal.

This problem becomes clearest not in theory, but in administration. Imagine a housing agency that uses an AI system to decide which families receive emergency assistance first. One mother is pushed down the queue because the system labels her case low risk. When she asks why, the vendor says the model only made suggestions. The agency says the tool merely provided recommendations. A caseworker signed the result but did not meaningfully alter it. In the end, the mother is left appealing a decision that seems to have involved everyone and no one at once.

What failed in that case is not simply the model. What failed is the chain of accountability around it. Governments have long used tools that ordinary citizens cannot fully inspect, from tax rules to risk scores. But AI thickens the fog. It adds outsourced infrastructure, proprietary models, adaptive outputs, and weak appeal paths. Human review alone is not enough if the human reviewer lacks the time, authority, or information to reverse the outcome.

The next phase of public-sector AI will make this challenge even sharper. These systems will not only answer questions. They will search records, complete forms, draft notices, manage cases, initiate workflows, and recommend actions. As soon as AI participates in action rather than advice alone, the relevant question changes. It is no longer just whether the answer is accurate. It is whether the workflow is lawful, reviewable, accountable, and reversible.

That creates a second danger beyond technical error: cognitive surrender. When an AI system sounds composed, fluent, and official, citizens may stop questioning it. Public servants may also defer to outputs they do not fully understand because the recommendation came from the system. In both cases, the institution begins to borrow confidence from the machine without earning legitimacy through explanation.

And that is why AI does not lighten the democratic burden of government; it increases it. Drafting, sorting, and summarizing may happen faster. But judgment, justification, and appeal do not accelerate merely because the software does. Accountability remains slow work. A polished dashboard or a clean model score does not reduce the obligation to explain who acted, why they acted, and how that action can be challenged.

The task for government therefore shifts from simple adoption to constitutional discipline. Machines can accelerate recommendations, but institutions must still verify, justify, and own the outcomes. That is not a flaw in democracy. It is one of its essential protections. Speed is useful. But speed without accountable authority is only a faster route to opaque power.

If that is the risk, then the response cannot be vague. Every public AI system should have a defined purpose, a registered owner, documented data sources, clear human decision rights, meaningful appeal mechanisms, and the ability to halt or reverse high-impact actions. Contracts must preserve access to documentation, cybersecurity review, audit trails, and intervention rights. Public officials must remain responsible even when software is deeply embedded in the workflow.

Some argue that such requirements will slow innovation. But innovation without accountability is not progress; it is risk transfer. A denial that arrives faster is still a denial. A more elegant dashboard is not the same as a fairer state. A chatbot is not a substitute for citizenship. The question is not whether government can automate more. It is whether it can do so without concealing responsibility behind systems that citizens cannot effectively challenge.

None of this means AI has no public value. It can reduce wait times, widen access, translate information, flag fraud, and support overburdened employees. But democracies should not measure success by the number of tools deployed or the labor hours removed. They should measure success by whether services improve, errors fall, schools work better, communities become safer, public decisions grow more legible, and trust is strengthened rather than eroded.

Ultimately, a democracy should never be satisfied with asking only what the system decided. It must also ask who authorized it, under what authority, on what evidence, and who remains answerable for the result. The democratic test is not whether a system can produce a suggestion. It is whether responsibility can still be located when that recommendation shapes a public decision. What matters is not whether AI allows democracies to govern more efficiently, but whether public responsibility survives its use. In a trustworthy democracy, responsibility cannot be artificial.