AI and Government Workers

Use Cases in Public Administration

By Samantha Shorey

Roosevelt Institute

AI and Government Workers: Use Cases in Public Administration

JULY 10, 2025

By Samantha Shorey

Share f X 🦋 ✉

Table of Contents

Executive Summary

Public administrators are the primary point of contact between constituents and the government, and their work is crucial to the day-to-day functioning of the state. AI technologies have been touted as a way to increase worker productivity and improve customer service in the public sector, particularly in the face of limited funding for state and local governments. However, previous deployments of automated tools and current AI use cases indicate the reality will be more complicated. This report scans the landscape of AI use in the public sector at the state and local level, evaluating its benefits and harms through the examples of chatbots and automated tools that transcribe audio, summarize policies, and determine eligibility for benefits. These examples reveal how AI can make the experience of work more stressful, devalue workers' skills, increase individual responsibility, and decrease decision-making quality. Public sector jobs have been an important source of security for middle-class Americans, especially women of color and Indigenous women, for decades. Without an understanding of what is at stake for government workers, what they need to effectively accomplish their tasks, and how hard they already work to provide crucial citizen services, the deployment of AI technologies—sold as a solution in the public sector—will simply create new problems.

Introduction

Since January, Americans have heard repeated claims from the "Department of Government Efficiency" (DOGE) that the federal government can and will be automated. Agencies have announced "AI-first" strategies following a federal hiring freeze and the buyout, firing, or resignation of 23,000 federal workers. The most immediate effects of this push for AI implementation will be in the places where Americans interact with their government most directly: at the state and local level of public administration. Public administration workers are the contact point between constituents and government, initiating bureaucratic processes and facilitating access to benefits programs. They are the ones who issue marriage licenses and hunting permits. They are who constituents interact with when applying for the Supplemental Nutrition Assistance Program (SNAP) or unemployment compensation, when going to court or to the Department of Motor Vehicles (DMV). AI technologies are already changing these workers' jobs, and therefore how constituents experience critical government operations.

An October 2023 survey from the Bloomberg and Johns Hopkins initiative AI City Connect found that according to mayors, only 2 percent of cities were using AI. A September 2024 Ernst and Young survey of federal, state, and local government workers at the senior level, however, found that more than 50 percent of workers were regularly using AI applications. This discrepancy indicates both a significant increase in AI adoption by public servants in a short time frame, as well as a potential gap between formalized, top-down AI use (where AI is officially approved for an identified purpose) and more ad hoc use by individuals as part of their daily tasks.

Closed-door conversations between public administration stakeholders and Roosevelt researchers surfaced recurring, conflicting themes. On the one hand, representatives from technology companies that are pursuing AI deployment in the public sector sincerely championed AI's ability to improve the experience of employees and customer service for constituents. On the other hand, practitioners

and researchers pointed out that there was little evidence that AI meets the performance requirements necessary to ensure consistent, secure public service. They emphasized that failures in AI systems, such as wrongful benefit denials, aren't just inconveniences but can be life-and-death situations for people who rely upon government programs.

This report considers the existing and potential impact of AI tools, particularly those leveraging generative AI, on government workers in public administration. Bringing together information from local news reports, scholarly articles, publicly available government documents, and conversations in private convenings, we offer a scan of the existing landscape using stories of AI implementation from across the country. As generative AI tools are increasingly installed as an intermediary in the public sector, constituents *and* government workers are becoming end users of AI. Some **good government jobs**—which have been a source of middle-class stability in communities across the country for decades—are already facing elimination. In many cases, however, the jobs are not disappearing but merely getting more complicated.

[F]ailures in AI systems, such as wrongful benefit denials, aren't just inconveniences but can be life-and-death situations for people who rely upon government programs.

Hope (and hype) around offloading repetitive, seemingly simple activities onto machines has impelled technology adoption across many industries. Yet when AI is confronted with the **complexities of the real world**, its **capabilities often fall short**. Workers are left to do much of the work that was supposed to be performed by computers *and* to correct AI's mistakes. More than 75 percent of workers in **a recent survey** reported that, amid expectations of increased productivity from their bosses, AI had made aspects of their job more difficult and "added to their workload in at least one way."

In her book ***Automating Inequality***, Virginia Eubanks details a program in Indiana to modernize the application process for Medicaid and SNAP. In-person eligibility interviews with career caseworkers were replaced with self-serve data entry and conversations with newly hired call-center employees. To initiate their application, applicants were instructed to enter their own information online or through a toll-free number, schedule a follow-up interview, and fax in the necessary documentation. But, at the centralized call center, employees were tasked with complex calls they were too inexperienced to evaluate and millions of pages of documentation they were too overwhelmed by to index. Application denials increased by 50 percent. Beyond the hardship this placed on people seeking assistance, correcting these mistakes fell on the shoulders of remaining state employees, many of whom burned out and left the profession.

To be sure, the *public sector* is much bigger than the field of *public administration*. The US government employs about **19 million people** at the state and local levels, and 2 million federally. Other public sector workers, including **schoolteachers**, **police officers**, and **trash collectors**, are also contending with AI technologies in their work. This report limits its scope to public administrators—workers whose consequential decisions and direct interactions with constituents sustain the day-

to-day-operations of democracy but are often overlooked in discussions of automation in high-stakes settings.

Public administration workers are critical to larger considerations of how AI implementation in government operations will affect the "well-being, safety, and dignity" of human beings across the country. An overburdened government workforce has understandable incentives to use new technologies to improve service and reduce wait times. But when workers have limited control over new technologies and their deployment, new kinds of errors and increased—rather than decreased—workloads emerge.

Background: What Is Public Administration?

Public administrators are part of all local, state, and federal agencies, departments, and offices. They are tasked with the day-to-day work of interpreting and implementing laws and policy, providing services through public programs, and enforcing regulations. For example, SNAP is funded through the multitrillion-dollar federal budget, but it is public administrators in state offices who are responsible for registering applicants, assessing their eligibility, and ensuring that they receive, on average, $187 of "food stamps" each month. Similarly, in most cities, building construction is regulated by local zoning laws, but it is public administrators in permit, license, and inspection offices who review architectural drawings to determine if a proposed structure meets requirements and issue the permit needed to begin the project. In these and other contexts, millions of Americans have their most direct interactions with the government through their interaction with public administrators. Behind the scenes, public administrators also play vital supportive roles—including budgeting and clerical duties—that facilitate government operations.

Public administration is a solid job. On average, salaried public administrators are paid more money than those who do office and administrative support occupations in the private sector. Additionally, they receive substantially more in benefits like paid leave, insurance, and contributions to retirement accounts. On average, public administrators earn a total of about $63 per hour in wage and benefit compensation. Thirty-two percent of public administrators are also part of a union. As the Center for American Progress observes, these qualities have made public sector work an important source of economic security for women and people of color since the 1950s. Today, this is especially true for Black and Indigenous women, who are represented in the public sector workforce at higher rates than they are in the general population. Yet, as the Economic Policy Institute shows, the quality of public sector jobs can vary widely, and in some occupations (such as childcare and facilities maintenance) more than half of workers are paid less than $15 an hour. This disparity is present within public administration as well: An estimated 25.9 percent of administrative support workers are paid less than $15 an hour.

The public sector has recently faced a series of employment challenges. At the onset of the COVID-19 pandemic, more than 700,000 state and local government employees were laid off and an additional 150,000 retired early. As the pandemic waned and the economy recovered, the number of people employed in the public sector consistently trended upward each month. Federal aid from the American Rescue Plan sought to compensate for lost tax revenue and support rehiring in local governments, bolstered state spending on infrastructure projects creating new jobs, and dispersed additional pay to essential public workers. These funds helped local governments recover almost three times faster than they had during the Great Recession.

In 2023, federal, state, and local governments added a combined total of 360,000 (non-postal, noneducation) jobs—including the greatest gains at the local level in 50 years. Yet there were twice as many available state and local job openings as there were hires in November of that year. In states like Hawaii, one in four public sector jobs were unfilled. Public administration roles, including office assistants and eligibility workers, had some of the highest vacancy rates. Government employees have been difficult to recruit due, in part, to "anemic wage growth" post-pandemic compared to the private sector. An analysis of government job postings found that the number of applicants to public sector jobs fell 50 percent between 2021 and 2022.

Despite the strong commitments they feel to their communities, public sector workers report feeling burned out and overwhelmed as they shoulder the work of their missing colleagues. They describe a downward spiral in which understaffing leads to lower job satisfaction, eventual resignations, and more vacancies. According to the *New York Times*, the head of the Division of Human Resources for the state of Idaho expressed a sense of resigned acceptance, reporting that "she was looking for ways to provide services with fewer people, especially for tasks like enrolling people in state benefits," a public administration duty. She added, "there's a lot of jobs that we're going to have to start looking at technology to solve." Inconsistent and declining funding for local governments has made automation especially attractive as a potential solution for closing the labor gap. The city council president of Williamsport, Pennsylvania, echoed this when he suggested that AI may help with public administration duties such as streamlining permitting—maintaining the quality of their municipal services even as the tax revenue that supports operations is shrinking. Here, AI is presented as a tool to alleviate existing staffing and workload problems public sector workers face, which often result from a lack of funding.

Inconsistent and declining funding for local governments has made automation especially attractive as a potential solution for closing the labor gap.

Even so, many AI systems do not independently accomplish a task from beginning to end—they rely upon human users to perform some aspects of the task. In the absence of government employees, these duties would be offloaded onto constituents at the other end of the system. Such is the case for benefits enrollment processes, where automation is partly achieved by requiring people to enter their own information into online forms and scan documents—activities once facilitated by caseworkers. Heteromation, the process that offloads "critical tasks to end users as indispensable mediators," saves costs by moving labor from paid employees to uncompensated users, but it does not improve the customer experience of the public. Constituents already spend a significant amount of time and effort when they apply for permits and public programs—filling out extensive paperwork, attending appointments, and supplying documentation to establish their eligibility. As social policy experts Pamela Herd and Donald Moynihan observe, the costs of understanding and acting in accordance with extensive and complex policies often fall upon constituents. These "administrative burdens" hinder successful applicants and affect socially disadvantaged people (the espoused beneficiaries of public programs) the most.

Current AI systems have specific capabilities that are matched to specific tasks. As such, determining what aspects of a job might be automated requires identifying and assessing the variety of tasks a worker is responsible for. The table below outlines the occupation-specific tasks for four types of public administrators: license and permit clerks, court clerks, municipal clerks, and eligibility interviewers for government programs. These are the nonmanagerial occupations that were (until March 2025) listed as part of the Government & Public Administration "career pathway" identified by the US Department of Labor's Occupational Information Network. The table organizes each occupation's tasks into five activity domains: (1) Communication with the Public; (2) Data Entry, Document Verification, and Payment Processing; (3) Research, Analysis, and Determination; (4) Internal Communication and Operations; and (5) Financial Management. Given that public administrators act as intermediaries between the public and government programs, using AI to automate aspects of their duties will dually impact recipients as well. Thus, the domains are organized from activities that involve a high degree of interaction with the public to those that involve less.

Public Administration Responsibilities

Job Type	Activity Domain	Tasks
License and Permit Clerks (driving, marriage, fishing, hunting, construction)	Communication with the Public	Answer questions or provide advice to the public regarding licensing policies and procedures.
	Data Entry, Document Verification, and Payment Processing	Question applicants to obtain required information; evaluate information on applications to verify completeness and accuracy; verify the authenticity of identification documents; code information on license applications for entry into computers; collect fees.
	Research, Analysis, and Determination	Perform record checks on past or current licensees; issue permits and licenses.
	Internal Communication and Operations	Perform general office duties such as typing or proofreading correspondence, distributing and filing official forms, and scheduling appointments.
Court Clerks	Communication with the Public	Answer inquiries from the public regarding judicial procedures, court appearances, trial dates, adjournments, outstanding warrants, summonses, subpoenas, witness fees, or fines; instruct parties about timing of court appearances.
	Data Entry, Document Verification, and Payment Processing	Examine legal documents submitted to courts for adherence to laws or court procedures.
	Research, Analysis, and Determination	Search files and contact witnesses, attorneys, or litigants to obtain information.
	Internal Communication and Operations	Prepare dockets and calendar of cases to be called; prepare documents recording the outcomes of court proceedings; prepare and issue orders of the court such as probation orders, release documentation, sentencing information and summonses; record case depositions, court orders and arrangements made for payment of court fees; perform administrative tasks such as answering telephone calls, filing court documents, and maintaining office supplies.
Municipal Clerks (city recorders)	Communication with the Public	Maintain office-tracking systems of correspondence or follow-up actions; issue public notifications of all official activities or meetings; respond to requests for information from the public, other municipalities, state officials, or legislative offices.
	Research, Analysis, and Determination	Research information in the municipal archives upon request by public officials or private citizens.
	Internal Communication and Operations	Prepare ordinances, resolutions, or proclamations so that they can be executed, recorded, archived, and distributed; prepare meeting agendas; manage the direct filing, maintenance, safekeeping, and digitization of municipal documents; record and edit the minutes of meetings and distribute them to appropriate officials or staff members.
	Financial Management	Perform budgeting duties, such as assisting in budget preparation or expenditure review; record and maintain fiscal records and accounts; perform contract administration duties assisting with bid openings or awarding of contracts.
	Other	Participate in municipal election administration by preparing and distributing ballots, training election officers, and tabulating results.

Eligibility Interviewers for Government Programs (welfare, unemployment insurance, social security, public housing)	Communication with the Public	Interpret and explain information such as eligibility requirements, application details, payment methods, and applicant's legal rights; answer applicants questions about benefits and claims procedures; refer applicants to other agencies for assistance or job openings; schedule benefits claimants for adjudication interviews to address questions of eligibility.
	Data Entry, Document Verification, and Payment Processing	Assist applicants to complete forms; interview applicants for public assistance; check employer and other references to verify answers and obtain additional information; monitor the payments of benefits throughout the duration of a claim.
	Research, Analysis, and Determination	Conduct annual, interim, and special housing reviews and home visits to ensure compliance with regulations; interview recipients at specified intervals to certify their continued eligibility; compile, record, and evaluate personal and financial data to verify and determine eligibility status; initiate procedures to grant, modify, deny, or terminate assistance.
	Internal Communication and Operations	Provide relevant information gathered during the interview to social workers; keep records of assigned cases and prepare reports.

Get the data • Created with Datawrapper

AI in Public Administration: Key Tasks and Use Cases

This section catalogues how AI is currently being used in two activity domains within public administration work at the state and local level: Communication with the Public and Research, Analysis, and Determination. These domains, both involving the exchange of information between workers and constituents, can and do overlap with one another but typically have different intended outcomes. Drawing on documented use cases and concerns around AI in adjacent sectors, we then anticipate the potential impact of generative AI on government workers.

What Is AI?

AI or "artificial intelligence" is a label that is broadly applied to a variety of computer-based systems that accomplish tasks automatically in a way that simulates human-like thinking. AI systems can make "**predictions, recommendations, or decisions**" through "**recognizing patterns**, learning from experiences, drawing conclusions, or taking action." AI relies on computer **algorithms**, which automatically carry out a "logical series of steps" to organize and process data. **Machine learning** algorithms continuously refine this logical series of steps based on exposure to additional data. Yet AI also includes more simple "expert systems" that have been available for decades. Expert systems concretize a set of **rules and procedures** followed by human practitioners into software programs so they can be carried out automatically. At times, the terms AI, machine learning, and algorithms are used interchangeably. However, not all algorithms are machine learning algorithms, and AI is more than machine learning.

Generative AI is a class of AI technologies that uses large amounts of data and advanced computing techniques to produce text or images. For generative AI tools that produce text—like OpenAI's chatGPT, Google's Gemini, and Microsoft's Copilot—the data that informs this process is also text, which has been collected from sources such as **publicly available websites**. This data is used as the basis for a "large language model," which uses machine learning to identify patterns within the collection of existing text and then combines it to create "new" text that resembles the patterns. The emergence of user-friendly and widely available generative AI products in 2023 heightened the fervor around AI in public sector work. Users can interact

with these technologies through actions that mimic natural conversation: typing questions and requests into a chat box where they receive a humanlike response. Generative AI's intuitive interface has made AI less difficult to integrate into existing processes and more accessible to average workers. Importantly, **artificial _general_ intelligence** (AGI)—a hypothetical single AI technology that can do all human actions—**does not currently exist**.

Communication with the Public

a. Chatbots

Chatbots are the most obvious and commonly proposed way of using generative AI in government agencies. Currently, public administrators answer many recurring questions from the public—providing information about how to apply for licenses, permits, or benefits. Their responses may not be especially customized to the people asking the question, but are rather a method of communicating procedural information. This information typically exists on websites or in printed resources, but it can be difficult for members of the public to find independently. By rapidly answering basic questions and automatically pointing people toward relevant documents, AI could improve the experience of residents and reduce the amount of repetitive communication required of government workers.

Basic chatbots have been used in customer service situations like these for at least a decade. **Chatbots** are software programs specifically designed for conversational interaction, where a human user types a query and the chatbot automatically delivers a response. Basic chatbots use non-generative AI to identify keywords in a user's question. These keywords are tied to a prewritten answer stored in a database. Describing the chatbot design process in detail, the state of **Georgia** explains that it developed 30 question-and-answer pairs that addressed the "vast majority of constituent questions." A similar bot in **Texas** has provided prewritten responses to common questions posed by the public more than 21 million times. The COVID-19 pandemic impelled many state governments to **create chatbots**, as unemployment and benefits offices were inundated with calls and questions during a period of widespread layoffs. Basic chatbots provide consistent and accurate information because their responses are prewritten by content experts, but they can be frustratingly limited. They struggle to interpret the various ways humans describe problems and can only provide a set number of responses.

Currently, government chatbots do not give responses specific to an individual user.
Yet personalization is one of the primary value propositions of generative AI, and
so it is likely that governments will explore how to leverage these capabilities to
improve automated functions. An **opinion essay in *Nature*** from researchers at the
Alan Turing Institute enthusiastically speculates that government might tailor
public services "according to your personal details and previous interactions with
the state." One way this may be achieved is through adapting chatbot
communication to individuals based on identity or demographic characteristics
that make them members of particular groups—a marketing practice known as
targeted messaging. For example, one of the duties of eligibility interviewers is
referring applicants to other relevant agencies who may be able to provide
additional assistance. A person who is identified as "low-income" when answering
questions about unemployment insurance, for example, could also be
recommended for an application to SNAP. In fact, a **2024 Florida bill** attempted to
create a generative AI–powered platform that would match recipients of
unemployment benefits to job opportunities based on an individual skills and
interests assessment. Matching and recommendation systems rely upon classifying
individuals and thus come along with important questions about how potential
members of this group are identified, the accuracy of the classification
mechanisms, and the **assumptions** that are associated with group membership
and hard-coded into software programs.

Public administrators also undertake highly individualized forms of
communication, such as informing people about the timing of their court
appearances or outstanding fines. When chatbots are used to facilitate these tasks,
they typically link to **secondary services** that allow users to look up details about
their case online. To fully integrate these forms of communication, chatbots would
need access to both the publicly available information they currently rely upon *and*
databases of personal information. At the federal level, this kind of access is
limited by the **Privacy Act of 1974**, which prohibits the disclosure of data held by
federal agencies without written consent. As legal expert **Kevin Frazier argues**,
"entering [personally identifiable information] into an AI system likely qualifies as
'disclosure.'" Further, the Privacy Act **prohibits the combination of data** from
separate government agencies into a single centralized database. But, at the state

and municipal level, privacy protection for government data varies. Storing data related to an individual's identity comes with inherent risks around **unauthorized use**. This can occur when organizations use personal data in ways not originally intended by the individual, or when outside parties (such as hackers) gain access to data through security breaches. If governments use products from large technology companies that retain the data they put into AI systems, they forfeit some of their power to protect constituents from these potential harms.

Impact on Workers: Fast-Paced and Frustrated

Chatbots are currently designed with the expectation that there will be interactions that are too complex to be handled by technology. In these cases, members of the public are theoretically "handed off" to human government workers who help answer their questions. **Interviews** with public administrators indicate that chatbots—which automate the repetitive parts of communication—effectively reduce the overall length of calls and focus worker's efforts on more specialized cases. Although framed as a benefit, workflows that make employees work faster on the most complicated cases intensify the experience of work and may contribute to **job stress**.

Further, public administrator's conversations with members of the public will typically begin after they have already been failed by chatbots. This means that the interpersonal interactions that workers do have will be characterized by consistent frustration from users and an increased responsibility for diffusing tension. Public administrators will also be responsible for addressing problems that arise from either errors made by the system or errors made by users. Automated **"self-service" technologies** fundamentally change the relationship between workers and customers. Rather than working in *service to* program recipients, public administrators become *supervisors of* program recipients who perform the former duties of employees (but without pay or training). Public sector employees strongly perceive their work as having a **positive social impact**. By reducing their ability to see their work as a substantive contribution to the well-being of others, AI may decrease the motivation and meaningfulness public administrators derive from their jobs.

b. Translation and Transcription

Language translation is an already widespread use of AI technology. Early forms of automated translation have been available **since the mid 2010s**, powered by machine learning models not dissimilar from the ones marketed today. Currently, an estimated **5 million low-income people** interact with AI-translated text across the public sector. Title VI of the **Civil Rights Act** requires that people who have **limited English proficiency** have "equal access to federally funded programs" through language interpretation and translation. For public administrators who determine eligibility for benefits programs, **TechTonic Justice** founder Kevin De Liban explains that this would include translating "program rules, key forms, written notices about a person's application or eligibility, all interaction with agency staff, and administrative appeals processes." While a **March 2025 executive order** from the Trump White House has designated English as the official language of the United States, it remains unclear how this might affect federal agencies in practice. Further, many state and local governments have their own **language access statutes**.

AI can be used to translate static text (in documents and websites) and to translate dynamic interactions in automated chatbots. When OpenAI released the

government-facing version of ChatGPT, it highlighted how the program had been used to translate Minnesota state agencies' forms and website landing pages into Spanish, Hmong, and Somali. Similarly, Google highlighted its partnership with Dearborn, Michigan, translating digital documents for permit applications into Arabic and Spanish. The Miami court system's chatbot "SANDI" offers a more dynamic conversational example. SANDI uses AI to help residents find and complete forms for "non-complex family court matters" on its websites—and can interact with users in both English and Spanish.

AI can also be used to transcribe spoken language. Often called "speech-to-text," this functionality is a familiar part of many smartphones that allows users to dictate statements that then appear in written form. AI is marketed to a wide variety of industries as a note-taking tool to capture and summarize conversations in meetings among employees. But, this capability is especially pertinent to the public sector, where open meetings laws often require that certain committee or city council meetings are recorded and accessible to the public. Further, the Americans with Disabilities Act (ADA) requires that communication from state and local governments (such as public meetings) is captioned for people who are deaf or hard of hearing. Currently, municipal clerks are responsible for recording and editing meeting minutes and distributing them to appropriate offices. While AI can be used to create transcriptions of speech automatically, when active, it records overheard speech indiscriminately. It lacks the human understanding of the delicate social negotiations that clerks perform as they determine what is private or a matter of public record.

> ### TRANSLATING CITY COUNCIL MEETINGS IN THE BAY AREA
>
> Several cities in northern California are piloting a tool that pairs AI transcription with AI translation on handheld devices that provide real-time captioning in almost 50 languages. A San Jose city clerk emphasized that AI translation had expanded access to public meetings beyond the three languages previously supported with human interpreters (English, Spanish, and Vietnamese) and also reduced yearly costs. But in nearby San Francisco, city officials have declined to use these real-time translation tools, even though they've been adopted by other cities in their region. They argue that mistake-prone technologies are at odds with the fundamental goal of translation: to convey a message in such a way that it accurately reflects the intentions of the speaker and is understandable to the recipient. Though these AI programs are informally referred to as performing "interpretation" (the live translation of speech) they are really a series of technical hand-offs. The programs convert speech into text in the first language, then translate the written text into a second language. Like in the slumber party game of "telephone," each step in the process increases the potential for error.

Impact on Workers: A Multilingual Mess

All the potential benefits and challenges present in the automated communication of chatbots are amplified by AI translation. Currently, public administrators must satisfy the vast informational needs of residents, *and* do so in multiple languages. Without AI, this is achieved through hiring multilingual public administrators or through paying human translators. For static text in forms and documents, translation can occur once until the original text is updated. For dynamic text—such as digital correspondence or in-person conversation—translation requires constant labor. Unmet translation needs are a barrier for community members'

participation in local government and hinder access to basic services like the DMV. Rather than spending time requesting and waiting for a human translator, AI translation could help more people in linguistically diverse places and help them faster.

Public administration requires that translators (whether they be human or machine) are attuned to particular aspects of language. Applying for permits or benefits involves specialized terminology that can be difficult for AI to accurately translate. Nuanced differences between terms, such as being "fired" vs. "laid off," for example, can impact determinations for unemployment insurance. Similarly, discussions about city planning involve terminology for locations tied to histories and geographies that are culturally specific rather than literal. To achieve some of the specificity required for public programs, AI companies propose that existing large language models can be fine-tuned with additional data. For example, a collaboration including the New Jersey Department of Labor has produced a set of training data to translate Spanish words specific to unemployment programs. Yet, beyond gaps in vocabulary, generative AI has more inherent operational problems. As machine translation expert Amelash Teka Hadgu and AI ethicist Timnit Gebru outline in *Scientific American*, using generative AI to transcribe spoken language can introduce a type of error called "hallucinations" that weren't part of more traditional automated systems. Hallucinations occur because generative AI tools attempt to predict strings of words based on their previous, recurrent association in training data. This results in not only words that are incorrectly transcribed but entire sentences that are totally fabricated.

Reviewing and correcting AI translation creates "an extra layer of work" for human translators who have to "clean up the mess" of the tools' mistakes. Further, when public administrators don't speak the language they are translating text into, they are unable to evaluate or verify AI outputs for accuracy, and mistakes go undetected. Research shows that generative AI users who feel less confident in their own ability to complete a task are also less likely to think critically and steer AI's responses. This is to say that users with a low level of language proficiency may more willingly accept AI outputs or hesitate to correct them. Even as AI companies pursue technical fixes for recurrent errors by fine-tuning models to disciplinary or cultural contexts, multilingual people are needed to identify terminology that the technology routinely gets wrong and to gather information about the proper alternative.

The unfounded belief that AI can automate all aspects of translation devalues the skills of multilingual public administrators in the present, who remain just as important to serving constituents in linguistically diverse places.

Perhaps more than any other area of application, translation highlights a persistent dynamic between AI and labor. Our cultural fears about AI's impact on the future of work are often about obsolescence—the fear that the technology will perform so well and so cheaply that humans cannot compete in labor markets and they will be replaced by machines. Yet, presently, government decisions to downsize are preceding the technological capabilities that would impel this progression. Obsolescence requires AI technologies that can perform the collection of tasks that make up an occupation *fully, automatically,* and *with a reasonable amount of*

precision. Turning to a function like AI translation, we can clearly see where this assumption falls short. Certainly, AI can facilitate basic translation work for common languages in government offices. But effective public administration still requires the labor of multilingual people who can correct the errors made by AI, perform more complex activities beyond the capabilities of AI, and facilitate translation for languages that are not represented in AI systems. Yet, the unfounded belief that AI can automate all aspects of translation devalues the skills of multilingual public administrators in the present, who remain just as important to serving constituents in linguistically diverse places.

Research, Analysis, and Determination

a. Summarization

The practice of public administration does not just involve carrying out a set of predetermined procedures. Rather, it requires "substantial discretion" from public administrators who interpret and enact government policies in singular situations. In his seminal work *Street-Level Bureaucracy*, Michael Lipsky documents that public sector workers (who are largely conceptualized as lower-level employees) regularly make consequential decisions about how to apply rules and interpret cases. Digital technology, which initially acted as a support to human decision-making, now increasingly acts as a leader in this process—taking public administrators from the street-level to the screen-level.

Serving as a connective point between government policy and constituent needs, public administrators must collect and analyze information from both these sources in order to make the decisions that are part of their daily work. On the government end, early digital technology supported public administrators by storing policy information, but such documentation is extensive, complex, and changing. IBM describes laborious contemporary work routines where public administrators are left to "click through endless screens, hunting and pecking for information stored in different databases." Alternatively, generative AI is characterized by consultancies as being capable of leading the information-seeking process. It can be used to "to quickly look up and summarize key file data from relevant regulations, policies, and procedures . . . to provide guidance and policy suggestions at the point of decision-making." Industry experts posit that using AI in this way not only makes finding information faster for workers, but facilitates consistency across cases by ensuring that workers have equal awareness of the same information.

In the working-facing tech examples, AI does not independently perform the core task (granting, modifying, or terminating assistance for an applicant). It **provides public administrators** with the policy "data and evidence they need to take action." The **New Hampshire Department of Employment Security** is also exploring the use of generative AI to collect and summarize evidence on the other end of this dynamic, from *constituents*. In contested unemployment claims, people who have recently lost their jobs (and their former employers) will interact with chatbots tasked with gathering facts and "packaging" them into summaries for human reviewers. The New Hampshire chatbot seeks to automate the interviews and reference checks that are currently conducted by public administrators to assess initial and continued eligibility. These conversations may include sensitive information, financial data, and follow-up visits to an applicant's home. The text-based communication produced through chatbot interactions would also create a foundation for generative AI in other steps of the benefits application process, such as preparing reports for record-keeping and sharing with social workers.

Impact on Workers: Responsibility for Risks

A **comparison test** conducted by Amazon for the Australian government found that humans were better than AI at summarizing regulatory documents in every criteria that they measured. The researchers anticipate that generative AI summaries may make document analysis *more* labor-intensive for workers, requiring that they refer to original documents to check facts, gain context, and find missed information. If workers fail to engage in these types of oversight, summaries—such as those included in the White House's **"Make Our Children Healthy Again" report**—may include misstated source information or cite sources that do not exist. Errors not only undermine the confidence of readers but can **harm workers**, who "may suffer a loss of time, resources, and even reputation."

To avoid these potential pitfalls, **city** governments have written generative AI **guidelines** emphasizing that AI is **imperfect** and that oversight is **essential**—instructing government workers to thoroughly assess the accuracy and fairness of

AI outputs. But as information studies scholar **Ben Green** argues, the emphasis on human oversight in AI policy ultimately places responsibility for technology on the shoulders of workers, not the companies that produce them or the institutions that deploy them.

b. Decision-Making

Automation has already had far-reaching impacts on the recipients of public assistance programs. Computer algorithms have made **eligibility decisions** about almost all of the 92 million low-income people in the United States when they enroll in programs such as Medicaid, Medicare, and SNAP. For example, at least 25 state **unemployment insurance agencies** have begun using **facial recognition** technologies to perform identity verification required to qualify for benefits. Agencies using this technology automate procedures once performed by public administrators with the goal of making them faster and easier for constituents, while also making fraudulent activity more difficult. Facial recognition technologies are a red flag for civil rights experts for many reasons, including that their accuracy varies widely based on **gender and race**. One **MIT study** found that basic identification of dark-skinned women using facial recognition technology is less accurate than a random guess. A system that works smoothly for some folks and could mire others in denials and paperwork is discriminatory. Despite the risks of using such a tool to decide whether people in poverty get their essential needs met, the 2020 report for the Administrative Conference of the United States found no systematic **government-wide protocols** to assess bias in these tools.

For more than a decade, automated technologies have been used in benefits determination processes at the Social Security Administration. The **Quick Disability Determination** (QDD) program uses a predictive algorithm to identify claims where "a favorable disability determination is highly likely and medical evidence is readily available" and prioritizes them for expedited human review. In **Social Security case adjudication** (which occurs when a claimant appeals an initial disability determination), AI is used to review and improve the quality of decision documents that staffers draft. In both these examples, automated technologies are assistive tools used by government workers. These applications are recognized by experts as some of the **most successful** in government. Yet, as the **National Academy of Social Insurance** (NASI) observes, they are not without their weaknesses. The QDD program privileges cases involving commonly occurring medical conditions. To improve quality and timeliness for cases with more complex diagnoses that aren't well-developed in AI systems, such as mental health, the government will need to solve systemic problems like understaffing. NASI also suggests adding additional guardrails to ensure accuracy and to avoid biased decisions, given the historical biases that may be present in Social Security Administration data.

ADJUDICATING UNEMPLOYMENT CLAIMS IN NEVADA

Nevada is extending the transcription and summarization capabilities of generative AI a step further by incorporating automated decision-making. The state is developing a tool intended to make initial determinations for a backlog of 40,000 appeals cases in the state's unemployment court. It will transcribe virtual hearings and summarize evidence to create a written document that includes "its own analysis of whether a person's unemployment claims should be approved, denied, or modified." The generative AI model used for the analysis will be trained on Nevada unemployment law documents and a collection of previous cases.

Still, government workers are not obsolete. They become the users of "recommender systems," acting as human reviewers who approve or override AI's recommendation. This relationship is muddled by the fact that, because it retrieves relevant laws that workers may not have been previously aware of, the AI is "somewhat of a training tool as well," as described by one Nevada official. Legal experts also warn that automated systems designed to speed up determinations can cause human reviewers to feel pressure to make their decisions quickly as well. Should a reviewer hastily agree with an ultimately incorrect AI recommendation, it is unclear if it can be legally overturned.

Individuals who are impacted by automatic determinations are left with little recourse to understand—or, if necessary, correct—these decisions. For example, state governments in Idaho and Arkansas used algorithms to calculate the number of home-care hours covered by Medicaid. These algorithms made drastic and unexplained cuts to disabled peoples' services. Although the courts ultimately ruled that these programs were illegal, the cases highlight the barriers constituents face when seeking to hold these systems accountable for their harms. In Idaho, the Department of Health and Welfare developed its software "in house" but initially refused to provide explanatory information about the technology, claiming it was a "trade secret." In Arkansas, the Department of Human Services relied on an algorithm built by an external nonprofit. Still, the lawsuit was filed against the state, *not* the nonprofit company that created the technology. In both cases, "the plaintiffs still have very little understanding of exactly how and why the system had reduced their benefits." When states or technology companies hide information about how policies are encoded into decision-making systems, they make it impossible for members of the public to identify problems and register complaints. As legal scholar Danielle Citron argues, this violates many states' administrative law requirements and the democratic process as a whole.

> When states or technology companies hide information about how policies are encoded into decision-making systems, they make it impossible for members of the public to identify problems and register complaints.

Impact on Workers: Less Communication, Lost Context

Government officials increasingly emphasize the importance of human "referees" who oversee AI-powered decision-making processes. They present AI as automating the communication, research, and initial analysis that inform decisions, which are ultimately made by government workers. But AI-aided decision-making is not a straightforward process. In their research overview, public administration experts Saar Alon-Barkat and Madalina Busuioc identify two main cognitive biases that can impact the quality of human decisions when AI is involved in other sectors. First, users tend to defer to algorithmic recommendations, even when the recommendations conflict with outside information or raise concerns requiring more research. Second, users are more likely to defer to an algorithmic recommendation when it aligns with cultural stereotypes. Taken together, AI has the potential to act as a risky shortcut: helping people make decisions faster, but less thoroughly or fairly.

Recipients of public programs and practitioners also view human workers as essential to good decision-making—particularly because of their ability to understand context. Workers gather contextual information through their conversations with multiple people involved in a case, supporting their ability to evaluate the quality of AI outputs. But if communicative interactions between constituents and public administrators are automated, public administrators are stripped of the opportunity to gain understanding specific to an individual's daily experiences, personal history, or culture. Thus, even if workers can override an AI recommendation, the use of AI reorganizes the system in such a way that workers can't gain the information they need to make high-quality decisions.

AI in Local Government: Policy Overview

Proponents of implementing AI in government work make two primary arguments: It will improve the delivery of government programs and reduce the workload of government employees. Alex Foard, the executive director of research and collaboration for New York City's Office of Technology and Innovation, summarized this dual objective when he stated: "AI has the potential to make the *government run better* and provide 8 million people with *easier access to services and benefits*" (emphasis added). Government offices are "constantly torn" between the public's conflicting demands for their operations to be more responsive to constituents *and* less labor-intensive. The use cases we have outlined in this report are evidence of experimental deployments in state and local governments to remedy this paradox through the perceived efficiency of AI.

Pro-consumer arguments for AI implementation are part of larger efforts to improve the "customer experience" of government programs. Long wait times in government offices and stalled paperwork in application backlogs are familiar parts of constituents' experience with public administration. More than half of city mayors reported that they were exploring the use of generative AI for constituent-service assistance in 2024. Beginning with an executive order from the Biden administration, members of the public have been reenvisioned as "customers" who should receive quality service like they do from commercial organizations. Public administration duties were a central focus for this order. Agencies were directed to remove requirements for in-person eligibility interviews and on-paper forms during enrollment for SNAP, Social Security, and unemployment benefits. It also emphasized the need for electronic permitting. Here, digitization is a mechanism for saving individuals' time, and AI is increasingly part of this charge.

> Government offices are "constantly torn" between the public's conflicting demands for their operations to be more responsive to constituents *and* less labor-intensive. The use cases we have outlined in this report are evidence of experimental deployments in state and local governments to remedy this paradox through the perceived efficiency of AI.

An advisory council established by **Tennessee HB 2325** is focused specifically on AI in public administration and is intended to create an action plan that supports "public employees in the efficient and effective delivery of customer service." **Virginia HB 6001** similarly allocated funds to pursue the integration of AI in "day-to-day activities" at the state DMV to "improve the overall customer experience." Beyond the successes, failed legislation also reveals areas where the government is exploring application of AI and associated concerns. State senators in California are grappling with how AI is being used in **direct communication** with the public, and during **benefit enrollment** processes in particular.

Anti-bureaucratic arguments for AI implementation are characterized by the desire (which is not unique to government work) to increase the productivity of employees' labor. Repetitive tasks that are tedious and time-consuming for workers are well-matched to automation, which performs best when a task is predictable and consistent. A blog post from the **New America Foundation** observes "one of the most significant advantages technology offers is its ability to automate repetitive tasks" that are part of bureaucratic processes. If AI can alleviate these duties, there are fewer duties for government workers to perform. This could have real benefits. Government workers are currently overextended. Offices are understaffed, and the demand for services is great. Delegating tasks to AI may relieve some of this burden, improving the quality of life for government workers and potentially even the quality of their work, as they have more time to devote to more complex tasks. But rather than justifying reductions in workforce, current use cases underscore the necessity for government workers who will continue to work alongside AI.

State and local governments are also at the forefront of attempts to manage the use of artificial intelligence in government projects. They are "implementing policies at a quicker, more direct, and more impactful level than their national counterparts," researchers at New York University's **GovLab** observe. Despite several **executive orders** over the last **five years**, there is no **comprehensive federal guidance** about the government's use of AI. **Municipal governments** have sought to fill this void with rules and regulations that define their own terms for appropriate AI use. State lawmakers proposed nearly **150 bills** regarding the government's use of AI in 2024. Most of the enacted bills—in states including West Virginia, Indiana, and Delaware—create general task forces to study the use of AI by state agencies.

> [R]ather than justifying reductions in workforce, current use cases underscore the necessity for government workers who will continue to work alongside AI.

One benefit of the municipal approach is the proximity of local governments to their constituents, which facilitates access to community-level feedback that can be meaningfully reflected in policy. For example, in San Diego a coalition of 30 organizations succeeded in passing a city council ordinance against automated technologies (such as license plate scanners and gunshot detectors) that are used by the police department for **surveillance**. Activists argued that these technologies not only violated community members' privacy but were technologically inept. An early installation of a similar **gunshot detection technology** in San José, California, was only correct about half of the time. Data evaluating this installation was available because the city of San José has pioneered vendor agreements that require technology companies to provide **performance metrics** to the city. The city's **GovAI coalition** brings together members of city governments nationwide to "promote responsible and purposeful use of AI in the public sector." Participating cities, including **Boston**, **Boise**, **Salt Lake City**, and **New York City**, have also created generative AI guidelines for government workers to inform them of the risks and responsibilities of using this tool.

Lawmakers in 10 states are attempting to delineate how AI can be used to streamline government decision-making processes. Many of these bills seek to require "impact assessments" that report on an AI product's performance quality and fairness. In New York, lawmakers passed the LOADinG Act ("Legislative Oversight in Automated Decision-Making in Government Act"), which **forbids state agencies** from implementing new automated decision-making systems in public assistance programs. The most significant support for government workers comes in a 2025 executive order from the **Pennsylvania governor's office**—announced in partnership with SEIU, the union representing state benefits and social services workers. Among the eight core principles is a **commitment** to "prioritize the input, expertise and experience of the employees," which is expected to manifest in the establishment of a collaborative group that will evaluate AI use from the perspective of workers.

Conclusion

At the federal level, executive orders on AI from both the Biden and Trump administrations are reflective of an ongoing negotiation around the role that government workers will play in establishing, sustaining, and overseeing the operation of AI. During Trump's **first presidency**, the White House sought to accelerate government workers' use of AI at the federal level and build public trust in AI technologies. Through establishing the **AI Community of Practice**, the General Services Administration sought to connect government workers so they could share tips and areas of application across agencies. Their set of **guiding principles** emphasized that use of AI in government should be "purposeful and performance-driven," "responsible and traceable," and "regularly monitored." The Biden administration's **executive order** on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence declaratively stated that ensuring responsible use of AI "start[s] with people, our Nation's greatest asset." Aligned with this, the 2024 **Bipartisan House Task Force on Artificial Intelligence** report highlighted the need for creating an AI-ready government workforce.

Yet increasingly, allies and senior officials with the second Trump administration use AI to justify firing public sector workers. Georgia Congressman **Rick McCormick**, for example, used AI to explain downsizing decisions at the Centers for Disease Control and Prevention, claiming that the 1,200 workers who were fired had done jobs that were "duplicitous" (likely meaning "duplicative") of AI. Here, unproven AI technologies are presented as a replacement for human labor. Under the guidance

of tech billionaire Elon Musk, **DOGE piloted** several AI technologies intended to automate the communication and decision-making processes of government employees. These technologies are blunt instruments that act in bulk, dismissing carefulness as inefficiency. Yet a Silicon Valley executive who formerly worked at DOGE reflected that his time there **challenged his perception** that employees are a source of government waste: He found himself "surrounded by people who 'love their jobs,' who came to the government with a sense of mission driving their work."

[T]he public sector should not be a testing ground for tools that haven't been evaluated, tested, and established as truly beneficial.

While this report has focused on local-level public administration workers, the faulty logic behind DOGE is mirrored in cases across the country. Without an understanding of who government workers are, what they need to effectively accomplish their tasks, and how hard they already work to provide crucial citizen services, the deployment of AI technologies—sold as a solution in the public sector—will simply create new problems. When asked to answer for the current limitations of AI, pro-innovation policymakers and tech enthusiasts often respond that these flaws are merely temporary and that the technology will improve. But, as **Aiha Ngyuen** argues, the public sector should not be a testing ground for tools that haven't been evaluated, tested, and established as truly beneficial. AI technologies make workers' interactions with constituents more intense, devalue the skills of multilingual workers, increase workers' individual risk and responsibility, and reduce workers' opportunities to gain the contextual information necessary to make high-quality decisions. These impacts cannot be dismissed as inconsequential while we wait for them to go away.

<div style="text-align:center;">

DOWNLOAD PDF

</div>

▼ Suggested Citation

▼ Acknowledgments

Related Resources

Democratically Deploying AI Means Letting Labor Lead

APRIL 29, 2025
By Sharon Block, Michelle Miller

The Risks of Generative AI Agents to Financial Services

SEPTEMBER 26, 2024
By Todd Phillips

Promoting Innovation Ecosystems in Antitrust: A Framework for Antitrust Analysis Applied to Emerging AI Technologies

JUNE 27, 2024
By Ketan Ahuja

Uber for Nursing: How an AI-Powered Gig Model Is Threatening Health Care

DECEMBER 17, 2024
By Katie J. Wells, Funda Ustek Spilda

AUTHOR

Samantha Shorey

FELLOW, AI

Samantha Shorey is an assistant professor of communication at the University of Pittsburgh and a field researcher who studies automated technologies in the workplace. Her recent research as a co—principal investigator of the National Science Foundation—funded project "The Transformation of Essential Work" focused on the adoption and adaptation of AI technologies by low-wage workers during the COVID-19 pandemic. Shorey will be a panel member of Stanford University's 100 Year Study of Artificial Intelligence (AI100) during the 2026 convening.

Before joining the University of Pittsburgh, Shorey was an assistant professor at the University of Texas at Austin, where she was a member of UT's ethical AI research initiative, Good Systems. She earned her PhD in communication from the University of Washington, where she conducted research in the department of human centered design and engineering. In 2019, she was a fellow at the Smithsonian Museum's Lemelson Center for the Study of Invention and Innovation. Her master's degree is from the University of Massachusetts Amherst.

READ MORE ▶

TAGS: Labor, Worker Power