

From Curiosity to Capability

What We Learned From Training
250,000 Public Servants on AI

CONTENTS

Preface	3
Executive Summary	4
Key Findings	5
Recommendations	22
What We Are Learning About Learning	24
Methodology & Course Details	26

This report was created by the
Observatory of Public Sector AI team:

Beth Simone Noveck

Professor at Northeastern University,
Director of the Observatory of
Public Sector AI at the Burnes
Center for Social Change

Anirudh Dinesh

Observatory Infrastructure
Project Manager

Greg Porumbescu

Associate Professor, Rutgers School
of Public Affairs and Administration,
and Director of the Public Interest
Technology and Management
Lab, Rutgers University

Lee Rainie

Director, Imagining the Digital
Future Center, Elon University

Allison Wan

Research Fellow, Burnes Center for
Social Change, Northeastern University

Data infrastructure

Stephan Schmidt

Digital/AI Producer, Burnes Center for
Social Change, Northeastern University

Design and graphics

Mariana Villamizar

Designer, Burnes Center for Social
Change, Northeastern University

Media contact

Bonnie McGilpin

Communications Director, Burnes
Center for Social Change

bmcgilpin@thegovlab.org

We also want to acknowledge
the efforts of the InnovateUS
team in compiling this report:

Agueda Quiroga

Alberto Rodríguez

Anjith Prakash Chathan Kandy

Ayush Khanvilkar

Christine Ine

Dane Gambrell

David Fields

Dhruv Kamalesh Kumar

Elana Banin

Neil Kleiman

Rob Asaro Angelo

PREFACE

Governments across the country are adopting generative artificial intelligence (GenAI).

Yet there is a surprising gap in our understanding of the public sector and its use of GenAI. We know far more about the policies and the platforms than we do about the people expected to use them, especially for public and civic purposes. Still, it is the knowledge, skills, judgment, motivation and confidence of public professionals that determine whether investments in GenAI and other forms of AI translate into stronger government capacity and better public outcomes.

- Are public servants using these tools?
- What prevents adoption?
- Do public professionals work differently with these tools for public benefit?
- What kinds of learning and support help public professionals use AI effectively?
- And what does the growing use of AI by public professionals mean for how governments serve the public?

This report seeks to answer those questions.

Drawing on one of the largest datasets yet assembled on public-sector GenAI learning, the InnovateUS¹ Observatory of Public Sector AI analyzed surveys and learning data from more than 250,000 public-sector professionals across all 50 states and the District of Columbia². These data come from professionals working in state, local, tribal, territorial, and federal government in roles ranging from public safety and human services to education, communications, legal services, technology, and executive leadership.

The findings paint a picture of a public sector that is earlier in its AI journey than many might assume and a workforce in which many do not yet even know about the basics of AI. Despite the extraordinary public attention to the governance of AI, far less is focused on governing with AI.

¹ InnovateUS delivers free at-your-own pace courses and live workshops to public professionals in AI and innovation skills. We have trained over a quarter of a million learners. InnovateUS has partnerships with 60+ states and cities and has upskilled large portions of the workforce in such jurisdictions as Washington, DC, Boston, Ohio, California, New Jersey, Colorado. For more about InnovateUS, see Section 5.

² The findings in this report come from 200,000 anonymous pre- and post-course surveys as well as pre- and post-workshop surveys and interviews with learners and their managers. The data collection period for most of the findings here was between June 4, 2024 and May 31, 2026.

Executive Summary

1 Many public professionals are just getting started with understanding and using GenAI.

For government agencies, the central challenge is not managing AI adoption. It is building the workforce skills required to support responsible adoption in the first place.

3 Leadership knowledge gaps about GenAI are a primary bottleneck to its adoption and effective use.

Leadership capability is one of the most important determinants of whether organizations move from experimentation to effective adoption.

5 Training builds confidence and confidence expands use.

Public professionals with high confidence report intentions to use AI across more than twice as many tasks as those who remain uncertain.

2 The biggest barriers to using GenAI in government are not tied to resistance.

They emerge from uncertainty about how to use it, whether it is relevant, and whether it is permitted.

4 Public professionals primarily use AI for information and communication tasks, such as research, drafting, editing, and summarizing.

Patterns of use vary significantly across roles, suggesting that AI adoption in government is shaped less by the technology itself than by the nature of the work being performed.

6 Public professionals readily identify tasks they want AI to perform.

They are less likely to articulate the public values those tasks are intended to advance. The next phase of public-sector AI adoption will require connecting AI capabilities to public purposes and public values.

Key Findings

FINDING #1

Many public professionals are just getting started with understanding and using generative artificial intelligence (GenAI). For government agencies, the central challenge is not managing AI adoption. It is building the workforce skills required to support responsible adoption in the first place.

The public conversation about artificial intelligence often creates the impression that AI has already become commonplace. Yet among the public professionals participating in InnovateUS learning programs, the reality looks very different. Some 51% of public professionals who took the most basic AI course offered by InnovateUS (“Using Generative AI at Work”) said they had never tried these common tools before the training. Another 14% said they had personally tried such systems but do not use them at work. In other words, nearly two-thirds of these public professionals entered training with little or no practical experience using generative AI.

This finding pushes against a common assumption about government’s AI challenges. Much of the public debate focuses on how governments should regulate, govern, or manage AI. But for many organizations, a more immediate challenge


remains: workforce readiness. Awareness of GenAI may be high, yet use remains low, pointing to a familiar knowing-doing gap.

The data also reveal that the knowing-doing gap is uneven across professions. Public professionals working in digital services, information technology, communications, and education were generally more likely to report prior AI use. Those working in public safety, legal services, administrative support, and other operational roles were less likely to report regular use.

Awareness of GenAI may be high, yet use remains low, pointing to a familiar knowing-doing gap.

51% of public sector workers have never tried GenAI. Those who use it employ it for a variety of tasks

Percentage of pre-course survey respondents who report level of using GenAI at work.



Note: Surveys of InnovateUS course takers June 2024–May 2026

Organizations often behave as though their challenge is managing widespread AI adoption. Our findings suggest many governments are facing a more foundational task: helping public

professionals develop the knowledge and judgment required to use these tools responsibly and effectively.

FINDING #2

The biggest barriers to using GenAI in government are not tied to resistance. They emerge from uncertainty about how to use it, whether it is relevant, and whether it is permitted.

Public debate about artificial intelligence often assumes that government workers are resistant to change. Our data suggest a different explanation. Many of these workers are unsure how to use GenAI tools. Even more notably, they do not know what is allowed, what works, or why it matters. Together, these three barriers account for most non-use: 28% of non-users said they did not think AI applies to their work; 23% said they did not know how to use GenAI; and 19% said they did not know if they are allowed to use AI in their jobs.

None of the three leading barriers reflects active opposition to AI. Instead, they point to the notion that workers struggle to understand how a new technology fits within existing responsibilities, professional norms and organizational rules. Uncertainty extends to AI's value as well: before the course, more than a third of respondents (37%) said they did not know whether AI would make their work easier or harder.

The findings suggest that many governments may be misdiagnosing their AI adoption challenge.

Organizations often focus their attention on procurement, governance frameworks and technical infrastructure. These investments are important.

Yet for many public professionals, the more immediate questions are practical ones:

- How does this apply to my job?
- Can I use it?
- What can I use it for?
- What can I trust it to do?
- What are the risks?

It is important to note that here, too, perceived barriers are not evenly distributed across professions. For some public professionals, especially those in legal, public safety, and operational roles, the central question is relevance. They are less likely to see obvious opportunities to apply AI in their day-to-day work. Among non-users in public safety, 43% said AI does not apply to their work, compared with 32% overall.

KEY FINDINGS

For others, especially those already comfortable with technology, the question is permission. Many report uncertainty about whether organizational policies allow them to use AI tools at all. Among non-users in digital roles, 28% were unsure whether they were allowed, compared with 18% overall. The absence of a policy can function as a policy. When expectations about GenAI use are ambiguous, uncertainty itself becomes a barrier to experimentation and learning. This points to an immediate need to update existing usage protocols so the public workforce has clear, practical guidance, rather than leaving them to navigate uncertainty on their own.


These findings point toward the conclusion that governments seeking to accelerate responsible AI

adoption should focus as much on the clarity of their vision –and permissions– as on technology. Public professionals need practical guidance about what AI is, how it can support their work, and what constitutes acceptable use.

When expectations about GenAI use are ambiguous, uncertainty itself becomes a barrier to experimentation and learning.

Non-users of GenAI cite three key barriers to GenAI adoption: knowledge, skills gaps, perceived irrelevance

Percentage of pre-course survey respondents who report these as primary reasons for non-use.



Note: Surveys of InnovateUS course takers June 2024–May 2026. Percentage for the access option is based on the 37,907 respondents surveyed after it was added in January 2026.

FINDING #3

Leadership knowledge gaps about GenAI are a primary bottleneck to its adoption and effective use. Leadership capability is one of the most important determinants of whether organizations move from experimentation to effective adoption.

When public professionals report uncertainty about AI, they frequently look to supervisors and managers for guidance. Yet many leaders are navigating the same uncertainties as the people they supervise. Our findings suggest that significant AI use for public benefit could be unlocked through expanded leadership learning. Among managers and leaders participating in InnovateUS learning programs, nearly a quarter (24%) of those not using AI said the primary reason was not knowing how to use it. Another 22% said they did not know they were allowed to use it, and 19% said they do not believe it applies to their work. That adds up to a meaningful share of managers who are not yet in a position to guide their staff on GenAI. Managers are also more likely than non-users overall to see AI as relevant to their work: 19% said it does not apply, compared with 32% of all non-users. What holds them back is not knowing how to use it or whether it is allowed.

Much of the current discussion about AI training focuses on frontline employees. The assumption is that leaders already understand the technology and that the primary challenge is teaching workers how to use it. But many managers are still developing their own understanding of AI's capabilities, limitations, risks, and potential applications. As

a result, uncertainty at the top can cascade throughout an organization.

When leaders are uncertain, employees often become more cautious. When leaders cannot distinguish between appropriate and inappropriate uses of AI, professionals receive mixed signals about experimentation. When leaders lack confidence in their own understanding of the technology, organizational adoption slows.


This dynamic helps explain one of the most striking findings from the previous section: Fully a fifth (19%) of non-users report uncertainty about whether AI use is permitted.

The challenge is particularly important because research on organizational change finds that managerial support is among the strongest predictors of successful technology adoption. (See Section 1, Literature Review) Employees rarely adopt new tools simply because the technology

Significant AI use for public benefit could be unlocked through expanded leadership learning.

Barriers to adopting GenAI differ by professional role

Percentage of pre-course survey respondents not using GenAI selecting each barrier



By percentage, the color scheme of this chart



Note: Surveys of InnovateUS course takers June 2024–May 2026. Role figures come from a later survey version (37,386 respondents), so overall percentages differ from the chart on page 8.

exists. They adopt them when leaders create an environment in which learning, experimentation, and responsible use are expected and supported.

This suggests that governments may need to rethink how they sequence AI capacity building. Investing in leadership capability creates multipliers. A manager who understands AI can influence dozens of professionals. An executive team that understands AI can shape the culture and expectations of an entire agency.

When leaders lack confidence in their own understanding of the technology, **organizational adoption slows.**

FINDING #4

Public professionals primarily use GenAI for information and communication tasks, such as research, drafting, editing, and summarizing. Patterns of use vary significantly across professions, suggesting that AI adoption in government is shaped less by the technology itself than by the nature of the work being performed.

Much of the public discussion about artificial intelligence focuses on autonomous agents, advanced analytics, and fully automated workflows. Our findings suggest that, in government today, AI is being used much more pragmatically. Public professionals primarily rely on generative AI to support information and communication tasks that are central to their daily work.

Among respondents who use GenAI, editing content is the most common application, with 70% reporting that they use AI for this purpose at least monthly. Content creation follows closely

(70%), along with conducting research (51%) and summarizing information (57%). More specialized applications, including organizing data, translating text, simulating conversations, and writing code, are used less frequently. Note that because these questions were asked only of respondents already using GenAI, these figures describe how existing users apply AI rather than overall adoption across the workforce. Taken together, these findings suggest that GenAI is being integrated first where it can reduce the cognitive burden of knowledge work. Rather than replacing professional judgment, public professionals are using AI to draft, edit, research, summarize, organize information, and communicate more efficiently.

The aggregate figures, however, mask important differences across occupations within the public sector. Educators are among the broadest and most likely users of GenAI, reporting especially high use for creating content (78%) and editing content (73%), but also above-average use for summarization, translation, and simulated conversations.

Communications professionals concentrate their use on drafting, editing, and summarization, reflecting the communication-intensive nature of their work. Digital professionals are the heaviest users of AI-assisted coding. Legal staff report the lowest use for most tasks, with one notable exception: 62% use AI for research, above the 51% average. These differences suggest that there is no single model of AI adoption across government. Instead, public professionals are adapting AI to the specific demands of their occupations. The technology is not reshaping work in the same way everywhere; rather, existing workflows are shaping how AI is used.

This finding has important implications for workforce development. Much of today's AI training treats AI literacy as a generic capability applicable across occupations. Our results suggest a different approach. While foundational AI skills remain important, public professionals are likely to benefit most from learning experiences that demonstrate how AI can support the specific tasks associated with their roles.

The findings also point toward an important research agenda. Much of the existing evidence on AI adoption has focused on occupations such as software development, where improvements can be measured in terms of coding speed or productivity. Those studies have generated valuable insights, but they capture only a small portion of government work. Public-sector AI is being applied across a much broader range of information-intensive activities that underpin policy development, service delivery, regulation, communication, and administration. As adoption matures, research should move beyond measuring overall AI use to understanding how different professions

integrate AI into their work, how those practices change organizational capacity, and whether they ultimately improve public outcomes.

Public professionals are likely to benefit most from learning experiences that **demonstrate** how AI can support the **specific tasks** associated with their roles.

Different public professionals use GenAI for different purposes

Among pre-course respondents already using GenAI at work, percentage using it for each task at least monthly

	Edit content	Create content	Summarize text	Conduct research	Sort multiple texts	Analyze or sort data	Translate or simplify text	Simulate conversations	Write or analyze code	Total respondents
Overall	70%	70%	57%	51%	38%	38%	36%	31%	27%	21,286
Accounting and finance	68	69	50	49	33	34	29	26	25	823
Administrative support	72	73	56	45	36	36	33	32	25	1,511
Communications and marketing	77	78	65	50	40	37	38	30	23	453
Compliance and regulation	73	68	56	51	32	32	29	24	19	980
Digital	66	63	60	56	45	43	41	33	69	252
Education	73	78	59	51	42	42	48	37	26	5,355
Engineering and technical	66	64	59	60	43	44	35	31	45	636
Human resources	75	75	53	50	34	34	30	30	19	672
Information and communications technology	67	67	60	62	41	42	34	32	55	1,117
Information and knowledge management	73	70	64	57	45	47	39	38	42	401
Legal and parliamentary	55	52	46	62	28	28	27	23	19	482
Managerial/leadership	78	78	62	52	38	36	28	30	19	1,927
Monitoring and audit	70	71	53	53	35	37	33	34	26	361
Project and program management	79	77	65	53	40	39	28	27	20	1,396
Public safety	59	59	52	47	37	37	37	34	30	1,955
Research	68	57	56	61	33	40	26	22	57	493
Service delivery	64	66	51	42	30	29	29	27	19	773
Strategic policy	78	70	65	67	38	36	34	26	23	307
Other role	59	58	47	44	29	29	30	24	21	1,392

By percentage, the color scheme of this chart



Note: Surveys of InnovateUS course takers June 2024–May 2026

FINDING #5

Training builds confidence and confidence expands use. Public professionals with high confidence report intentions to use AI across more than twice as many tasks as those who remain uncertain.

General debates about AI training often focus on knowledge. Our findings suggest that confidence may be even more important. Across InnovateUS learning programs, participants consistently reported increases in both knowledge and confidence following training.

Those reporting a lot of knowledge grew from 13% of respondents before the course to 45% after taking the course. Furthermore, 31% of respondents before taking the course said they had no knowledge of AI and that number shrank to 5% after the course. At the same time, the share of respondents reporting they were very confident about using GenAI before the course was 27% and it shot up to 44% after taking the course. Those reporting they were somewhat confident grew from 23% before the course to 34% after taking it. And those reporting no confidence shrank from 9% pre-course to just 2% after taking the course. These before-and-after figures come from different groups of respondents, not the same people tracked over time.

In the case of more than 22,000 participants who took the “Using Generative AI at Work” course, we were able to identify specific individuals and compare their pre-course and post-course


answers. The findings confirm what we saw in the data aggregating the tens of thousands course takers.

Training increases knowledge across the board in significant ways and the course very clearly bolstered their confidence. A related finding emerged from our survey questions about the level of knowledge gained by different groups after taking the course. The knowledge gains were the most pronounced among those who said before the course that they had no knowledge of GenAI. That group’s reported knowledge gain was more than double those who reported some knowledge or a lot of knowledge before taking our course. The confidence level of course takers was particularly likely to improve among those who reported

Those reporting a lot of knowledge grew from 13% of respondents before taking the course to 45% after taking the course. The level of reported knowledge grew, as well.


Training increases GenAI knowledge...

Self-reported knowledge of GenAI before and after the course (%), cross-sectional samples



...and confidence

Self-reported confidence using GenAI before and after the course (%), cross-sectional samples



Note: Surveys of InnovateUS course takers June 2024–May 2026

having no confidence in using GenAI before taking the InnovateUS course. Again, the gains were substantially larger for this group than for those who began with higher levels of confidence.

One of the most striking findings emerges when confidence is examined alongside intended use. Public professionals who reported high confidence after training also reported intentions to use GenAI across a much broader range of tasks than those who remained uncertain.

Respondents with high confidence reported plans to use AI for an average of 3.8 different tasks.

Those with low confidence reported plans to use AI for only 1.5 tasks. Public professionals with high confidence do not simply use AI more frequently. They imagine more possibilities for how AI might support their work.

These gains are significant in their own right. Public professionals who better understand AI are better positioned to evaluate its strengths and weaknesses, identify appropriate uses, and recognize potential risks. In addition, higher levels of both knowledge and confidence are associated with greater GenAI usage and more

In effect, **confidence functions as a bridge between awareness and adoption.**

positive perceptions of GenAI’s impact on work. This suggests that changing attitudes towards training – especially among those who do not have high levels of knowledge about GenAI – will likely lead to broader behavior shifts as well.

In effect, confidence functions as a bridge between awareness and adoption. Many public professionals begin their AI journey uncertain about how these tools work, whether they are useful, or whether they are permitted. Training helps reduce that uncertainty. As confidence grows, professionals

become more willing to explore and learn how AI can support different aspects of their work.

Importantly, confidence should not be confused with blind trust. Effective AI use requires skepticism, verification, and professional judgment. The goal of training is not to encourage uncritical adoption. It is to help public professionals develop the knowledge and confidence necessary to use AI responsibly.

The findings suggest that governments should evaluate AI learning programs by assessing:

- ➔ What participants know at the end of a course
- ➔ Whether participants leave feeling capable of applying what they have learned
- ➔ Whether they end up applying those skills

Higher levels of confidence expand intended use

Average number of tasks respondents said they were “more likely to do when using generative AI,” by post-course confidence level



Note: Surveys of InnovateUS course takers June 2024–May 2026

SPOTLIGHT

Building AI capability at scale in Washington D.C.

Washington, D.C., mandated an AI training requirement for all government employees and the contractors that work with them. According to DC's Chief Technology Officer Stephen Miller: "We don't give someone access to our fleet systems without making sure they have a responsible driving record."

By 2026, more than 32,000 District employees (or 86% of the workforce) had enrolled in the Responsible AI for Public Professionals course. The Observatory collected data from these learners, resulting in an early picture of how public-sector AI adoption develops at scale.

The patterns closely resemble those observed across the national sample. Most employees

GOVERNMENT OF THE
DISTRICT OF COLUMBIA
MURIEL BOWSER, MAYOR

octo
OFFICE OF COUNCIL ON TECHNOLOGY

innovate(us)

District of Columbia AI Learning Hub

The District of Columbia has partnered with InnovateUS to provide asynchronous and free AI innovation training programs for state workers.

InnovateUS offers learning on AI, data, digital, and other innovation skills to help public sector professionals better serve their constituents and deliver more effective, equitable, and engaged policies and services.



START WITH THE BASICS: WHAT YOU'LL LEARN IN THE RESPONSIBLE AI FOR PUBLIC PROFESSIONALS COURSE SERIES

Responsible AI for Public Professionals is a two-part course series that showcases the practical ways you can use Generative Artificial Intelligence (GenAI) in your work as a public professional.

In Part 1: Using GenAI in Government, you will learn practical applications of GenAI to enhance your daily government work and gain skills to identify opportunities where AI can improve public service delivery.

In Part 2: Understanding Risks and Building Resilience, you

District of Columbia AI Learning Hub. Website by InnovateUS.

had little or no prior experience with GenAI. The principal obstacles to adoption are organizational rather than technical. The most frequently cited reasons were uncertainty about whether AI was relevant to their work, uncertainty about whether its use was permitted, and a lack of knowledge about how to use it. Even among employees in digital roles, many remained unsure about agency policies governing AI use.

Participation in the training was associated with substantial increases in employee confidence and intended use. Following the course, the share of employees reporting that they were very confident using generative AI increased, while the share reporting no use declined. The largest improvements occurred among employees who entered the program with the lowest levels of confidence. This pattern suggests that structured training can broaden AI capability across the workforce rather than primarily benefiting early adopters.

Patterns of use also vary across occupations. Daily use is highest among employees working in digital services, human resources, engineering, and communications, while adoption is lower in public safety, legal, and frontline service roles. These differences likely reflect variation in job tasks, organizational policies, and the availability of immediately applicable use cases.

Employees who report using generative AI most commonly apply it to drafting and editing documents, conducting research, and summarizing information. Coding is comparatively uncommon. This distribution is consistent with findings elsewhere in this report showing that generative AI is being incorporated first into information processing and communication tasks that are common across many government occupations.

The early results indicate that large-scale AI adoption requires more than access to AI tools. Employees need opportunities to develop practical skills, understand organizational expectations,

and identify where AI can contribute to their daily work. Washington, DC's experience illustrates how a coordinated training strategy can strengthen those conditions across an entire public workforce.

FINDING #6

Public professionals readily identify tasks they want GenAI to perform. They are less likely to articulate the public values those tasks are intended to advance. The next phase of public-sector AI adoption will require connecting AI capabilities to public purposes and public values.

To better understand how public professionals are using AI, InnovateUS asks participants about specific tasks they perform such as research, drafting, editing, summarization, data analysis, coding, and other day-to-day responsibilities. These activities matter because they contribute to larger goals like delivering services, improving responsiveness, increasing accessibility, strengthening public trust, reducing administrative burden, or helping residents solve problems.

Still, the surveys show there is a gap between the tasks public professionals say they perform with GenAI and the limited way these workers cite the broader democratic purposes their work serves. When asked open-ended questions about the way these workers hope to use AI in the future, they most frequently described activities like drafting reports, creating presentations, conducting research, analyzing documents, editing communications, and summarizing information. In other words, they list many of the job activities they are already performing with AI.

Fewer open-ended responses from these public professionals cite outcomes such as improving service delivery, enhancing democratic accountability, bolstering the transparency and openness

of government operations, increasing accessibility, reducing wait times, strengthening public participation, improving decision-making, or creating better experiences for citizens.

This finding should not be interpreted as a criticism of public professionals. In many ways, it reflects the current state of AI adoption. Most organizations remain in the early stages of learning how AI can support individual productivity. It is natural that professionals focus first on the tasks immediately in front of them. Before people can rethink systems, they must first understand tools.

The surveys show there is a gap between the tasks public professionals say they perform with GenAI and the limited way these workers cite the broader democratic purposes their work serves.

Yet the distinction between tasks and outcomes is important. Governments do not exist to draft reports, summarize documents, or edit communications. They perform those tasks to advance public purposes and create public value. Whether a particular AI use improves efficiency or effectiveness depends on how workers connect tool use to those broader goals.

A public professional who uses AI to draft a document more quickly may save time. A public professional who uses AI to make services easier to access, decisions easier to understand, or the government more responsive may create public value. As governments move from experimentation to adoption, the challenge will increasingly be measuring the impact of new ways of working with AI, helping public professionals connect technological capabilities to changes in how their organizations serve the public.

SPOTLIGHT


Public Voice

PUBLICVOICE.INNOVATE-US.ORG


Public Voice is an experimental evaluation tool developed by the Observatory of Public Sector AI to better understand whether AI training leads to meaningful changes in how public professionals work. The project was created to solve what we came to call the “six-word problem.” When asked open-ended questions at the end of an InnovateUS course, participants’ responses averaged just six words—too little detail to understand how they were actually applying what they had learned.

Public Voice addresses this challenge by combining short voice responses with AI-generated follow-up questions that encourage learners to provide concrete examples of how they are using AI in their work. Rather than collecting generic reactions to a course, the tool captures evidence of specific tasks, workflow changes, and remaining barriers to adoption.


Following successful pilots in spring 2026, Public Voice is now integrated into the Observatory’s evaluation process across InnovateUS learning programs. The richer qualitative data it generates complements traditional surveys, enabling researchers to move beyond measuring learner satisfaction toward understanding whether AI training is producing measurable changes in professional practice. We anticipate having richer insights to share in our second report.



Text input



Voice input + Follow-up questions



Summary of main themes

Recommendations

Our findings point toward smarter, more effective ways to introduce and cultivate the use of GenAI by public professionals for public purposes. These recommendations are based on consistent patterns across a very large sample. We have not yet tested them in practice, and we plan to do that with our partners.

1 Treat workforce capability as public infrastructure

Governments invest billions in digital systems and AI technologies. Yet those investments create public value only if public professionals have the knowledge, judgment, and confidence to use them effectively and responsibly. Workforce capability should be treated as a core component of government modernization, not as an optional professional development activity. Without sustained investment in people, even the most advanced technology cannot deliver better government.

2 Move from AI literacy to AI for public purpose

The next phase of AI learning goes beyond helping public professionals understand AI or identify potential use cases. It must help them apply AI to advance public purposes: improving service delivery, increasing accessibility and responsiveness, strengthening decision-making, and expanding meaningful public participation.

3 Build organizational capability along with individual skills

Training is necessary, but it is not sufficient. Public professionals also need access to appropriate tools, clear governance, supportive leadership, and opportunities to apply what they learn in their daily work. Governments should approach AI adoption as an organizational transformation opportunity, creating the conditions that enable responsible and sustained use.

4 Invest in leaders as force multipliers

Managers and supervisors determine whether new skills translate into new ways of working. Leaders who understand AI can model responsible use, establish clear expectations, identify high-value opportunities, and create a culture that supports learning and experimentation. Building leadership capability is one of the fastest and most scalable ways to increase organizational capability.

5 Measure public problem solving with AI

The success of AI should be measured by the public value it creates for governments, public professionals, and residents. Evaluation should emphasize outcomes such as higher-quality services, greater accessibility and responsiveness, improved efficiency and institutional effectiveness, stronger public and employee voice, and more effective solutions to public problems.

What We Are Learning About Learning

Our extensive engagement with public professionals also gives us fresh insights into the learning process itself. Among the revelations:

1 Never forget that learning is **social**

Create opportunities for practitioners to share experiences, exchange examples, and learn from one another. Public professionals often trust the experiences of colleagues facing similar challenges.

2 Teach **new ways** of working

Focus on how to integrate AI responsibly into work. Go beyond how to do today's tasks faster. Teach how to work differently.

3 Prioritize **confidence**, not just knowledge

Effective learning programs should focus on helping participants feel capable of applying what they have learned on the job.

4 Train leaders alongside staff

Leadership learning is a force multiplier because it leads to culture change.

5 Tailor learning to professional roles

Public professionals use AI differently depending on the nature of their work. Develop instruction based on roles and problems to be solved.

6 Focus on public value

Teach how to use technology to benefit residents, improve governance and deepen democracy.

7 Learning is necessary but not sufficient

Training can increase knowledge and confidence, but sustained adoption also depends on access to tools, clear guidance, supportive leadership, and a culture that permits applying learning in practice.

Methodology & Course Details

From Curiosity to Capability

What We Learned From Training
250,000 Public Servants on AI

About InnovateUS

Who we are

InnovateUS is a no-cost, philanthropically funded peer-to-peer learning initiative. We provide live and at-your-own-pace learning on data, digital, innovation, and AI skills, for public professionals by public professionals.

Our hypothesis

Investment in public sector professionals produces measurable improvements in government performance and better outcomes for residents.

Our mission

To improve lives by enabling public service professionals to deliver more effective and equitable policies and services.

Our reach

More than 250,000 learners across all 50 states, the District of Columbia, and 80 countries. We work directly with 60+ state and city partners, whose leaders comprise our advisory board, supported by 270+ leading faculty and 300+ completed workshops.

What we offer

Peer-to-peer learning through courses, workshops, and coaching, designed for and by public professionals. Self-paced courses build foundational knowledge, live workshops reinforce and apply skills, and coaching programs support hands-on implementation. All programs are no-cost, include a certificate of completion, and run on our open Moodle platform or a partner learning management system. Our website is available in seven languages.

How we work

Our content is evidence-based, developed collaboratively with public leaders and expert advisors, and rigorously reviewed for bias and usefulness. It is vendor-neutral, with examples spanning tools like ChatGPT, Claude, Gemini, and Copilot, and openly licensed under Creative Commons (CC BY-SA 4.0).

We are non-commercial and non-partisan. We collect only the minimal data needed for access, never share it with third parties, and never sell learner data. All our content is vendor- and platform-neutral.

Courses

Overview

For each InnovateUS course, learners are invited to participate in short, voluntary, anonymous surveys to learn about how the public sector thinks about and uses AI. Pre-course surveys measuring baseline attitudes and behaviors are given at the beginning of the course, and post-course surveys measuring learning outcomes are given immediately after completion, within the final module of the course. Currently, we do not conduct any follow-up surveys but plan to do so in the future. As surveys are administered through a link within the course, only individuals taking the courses through our government partners are able to complete the surveys.

Data collection began in June of 2024 and is ongoing. For this report, we include all responses completed by May 31, 2026, among those participants taking the “Using GenAI at Work” course. Survey responses collected for other courses are not reported here, but we include overall numbers in the table below for reference. The final sample size in this report consists of 106,300 pre-course respondents and 67,002 post-course respondents. Over this period, multiple survey versions were administered with variation in the survey questions asked. As a result, specific analyses in this report often draw on subsets of the full time range or sample. The exact sample sizes are specified alongside the relevant findings.

Measurement standardization

Since we aggregate across multiple survey versions, some questions were asked with slightly different response options. To ensure comparability, we harmonized response scales in select cases and recoded open-ended responses into existing categories. Self-reported generative AI knowledge, which was measured on both 3- and 5-point scales, was standardized to a 3-point scale. Course recommendation ratings, collected on both 5- and 10-point scales, were standardized to a 10-point scale.

For a professional role, respondents could select “Other” and provide a write-in response. Using keyword matching, we recoded approximately 2,300 of these responses into existing categories, about 1,000 of them into “Public safety,” which was not included as a response option in earlier versions of the survey.

Data limitations and potential sources of bias

Because both survey participation and, in some cases, course enrollment are voluntary, there is potential for self-selection at multiple stages. First, if course enrollment is voluntary or required for accessing generative AI tools, there is the likelihood that the learner population may be more knowledgeable, experienced, or motivated to learn about generative AI compared to the

broader public sector population. Secondly, there is self-selection into taking the surveys with significant attrition moving from the pre-course to post-course survey.

With the absence of demographic data, we cannot mitigate these selection effects through common practices such as survey weighting. However, we can compare respondents who report being required to take courses by their supervisor against those who do not to get a sense of self-selection into course enrollment. Likewise, for pre-course respondents with a learner ID, we are able to compare whether there is self-selection moving from the pre-course to post-course surveys.

When comparing respondents who were required to take the courses to those who are voluntarily participating, voluntary participants report higher levels of pre-course knowledge and usage, and are more likely to believe generative AI will make their work easier.

When comparing respondents who only completed the pre-course survey against those who completed both survey phases, those completing both phases report similar pre-course usage and are slightly more likely to view generative AI as making their work easier.

Taken together, these comparisons suggest that selection into both course enrollment and survey completion is present but modest. Voluntary participants appear somewhat more engaged and optimistic about generative AI, while respondents who complete both survey waves differ slightly in baseline usage and perceptions. These differences should be kept in mind when interpreting results, as estimates may overstate baseline familiarity and perceived benefits relative to the broader public sector population.

Finally, while our data provides valuable insights into generative AI attitudes and behaviors among the public sector workforce, its opt-in nature means

it should not be interpreted as representative, and margins of error cannot be calculated. In addition, all measures are self-reported and collected at the beginning and end of each course. They capture knowledge, confidence, and intended use, not observed behavior or outcomes, and relationships such as the link between confidence and intended use are associations, not evidence of cause and effect.

Measure (all from Pre-course survey)	Required	Voluntary
Avg. Knowledge	1.7	1.9
Avg. Usage frequency	1.6	2.3
Work impact – Easier (%)	40.7	47.7
Work impact – Harder (%)	2.7	2.6
Work impact – Neither (%)	14.6	13.6
Work impact – Don’t know (%)	41.9	36.1
Total respondents	17,047	89,253

* Data covers just those taking “Using Generative AI at Work” course

Measure	Pre-Course Only	Completed both phases
Avg. Confidence	3	3
Avg. Usage frequency	2.5	2.5
Work impact – Easier (%)	46	48.1
Work impact – Harder (%)	2.9	2.6
Work impact – Neither (%)	15.7	16.1
Work impact – Don't know (%)	35.1	33.2
Total respondents	4,313	19,640

* Data covers just those taking “Using Generative AI at Work” course

Workshops

Overview

Workshops are live, practical, and skill-oriented learning sessions, typically lasting 60 to 90 minutes. Led primarily by practitioners and public sector professionals, they combine expert and practitioner insights with real-world examples and opportunities for peer-to-peer learning. Sessions typically begin with a presentation or, when multiple speakers are involved, a facilitated conversation, followed by 15 to 20 minutes of participant Q&A and discussion. This live, interactive format enables participants to engage directly with speakers and learn from one another by sharing experiences, questions, and challenges from their own work.

Workshops are organized into thematic series that explore a topic from multiple perspectives and angles, helping participants build practical knowledge and skills across interconnected sessions. They have been offered both in person and through live online webinars. To date, InnovateUS has hosted 308 workshops featuring 289 speakers and reaching 38,800 total participants.

We capture baseline attitudes and behaviors using voluntary and anonymous pre-workshop surveys administered at the start of each session. Post-workshop attitudes are measured through a voluntary, anonymous survey distributed immediately afterward. While workshops are open to the general public, the majority of attendees (approximately 70%) are public sector workers.

Data collection for this report began on April 10, 2025, and is ongoing. This report includes data collected through May 31, 2026. During this period, survey instruments varied: some questions were workshop-specific, while others captured similar concepts with minor wording differences. Where possible, we standardize comparable questions across workshops for reporting. As a result of this variation, sample sizes differ by question, ranging from several hundred to several thousand respondents; exact sample sizes are reported where relevant. Since individuals may attend multiple workshops, some respondents may appear more than once in the data.

Lastly, because workshop attendance and survey participation are voluntary, the sample is subject to self-selection.

Key Facts & Figures

The Observatory of Public Sector AI draws on one of the largest data-sets yet assembled on workers and their AI learning and adoption.

Learning Programs



More than
70,000
completions of the
flagship “Using Generative
AI at Work” course



308
live workshops
delivered



More than
38,500
workshop participants



289
expert speakers

Research Dataset



106,300
pre-course survey
respondents



67,002
post-course survey
respondents



22,303 matched pre- and post-course responses
allowing individual-level analysis

Who Participates

The public professionals participating in InnovateUS learning programs span the job categories below. The table maps the InnovateUS job categories to the closest occupational groups in the U.S. Bureau of Labor Statistics Standard Occupational Classification and the federal Office of Personnel Management occupational series.

Note:

Categories reflect respondents' self-described work functions and correspond approximately, rather than exactly, to SOC occupations and OPM series. Respondents who selected "Other" and described their role in open text were assigned to a listed category when the description clearly matched one (for example, "economist" to researchers and data analysts). Descriptions matching no category, or more than one, remain classified as Other. The final row was not a survey category: health, social service, and facilities roles are not separately captured by the survey's answer options and were identified through open-text responses.

The survey questionnaires we use in our courses can be found at:

[INNOVATE-US.ORG/RESEARCH/QUESTIONNAIRES](https://innovate-us.org/research/questionnaires)



InnovateUS job category	U.S. Bureau of Labor Statistics Standard Occupational Classification	Federal OPM occupational series (U.S. Office of Personnel Management)
Program managers	General and operations managers; project management specialists	Program management; management and program analysis
Policy analysts	Management analysts; political scientists	Management and program analysis
Attorneys	Lawyers; paralegals and legal assistants	General attorney; paralegal specialist
Human resources professionals	Human resources managers; human resources specialists	Human resources management
Budget and finance staff	Budget analysts; accountants and auditors; purchasing agents	Budget analysis; accounting; contracting
Educators	Elementary, secondary, and postsecondary teachers; instructional coordinators	Education program; training instruction
Public safety professionals	Police officers; firefighters; correctional officers; emergency management directors	Police; fire protection; correctional officer
Administrative professionals	Executive assistants; secretaries and administrative assistants; office support staff	Secretary; miscellaneous administration and program
Technology and digital services professionals	Software developers; web and digital interface designers; computer systems analysts; information security analysts	Information technology management

InnovateUS job category	U.S. Bureau of Labor Statistics Standard Occupational Classification	Federal OPM occupational series (U.S. Office of Personnel Management)
Agency leaders and managers	Chief executives; general and operations managers	Program management; Senior Executive Service
Compliance and regulatory professionals	Compliance officers; occupational health and safety specialists; licensing examiners and inspectors	General inspection, investigation, enforcement, and compliance
Service delivery and case management professionals	Social and human service assistants; eligibility interviewers, government programs; customer service representatives	Social services; contact representative
Engineers and technical staff	Civil engineers; mechanical engineers; electrical engineers; engineering technicians	General engineering; engineering technical
Researchers and data analysts	Data scientists; statisticians; economists; survey researchers	Data science; statistician; economist
Health, social service, and facilities professionals (identified through open-text responses)	Registered nurses; community health workers; social workers; building and grounds maintenance workers	Nurse; social work; facilities management



@innovateus-govlab



www.innovate-us.org



hello@innovate-us.org

innovate(us)

InnovateUS provides no-cost, at-your-own pace, and live learning on data, digital, innovation, and AI skills for public service professionals. Governed by our government partners, free content and courses can be hosted on your learning management system.