ChatGPT

Generative AI in the Public Sector: Implications for Work and Collective Bargaining

Executive Summary

- **Generative AI Adoption in Government:** State and local governments are beginning to use generative AI tools (like large language model chatbots and content generators) to streamline operations. Early use cases include drafting routine documents, translating calls (e.g. 911 transcripts), answering public queries via chatbots, and assisting with coding or data analysis. Adoption is uneven: a few jurisdictions are moving quickly with pilot programs and official guidelines, while many others remain cautious or in exploratory phases [1]. Leaders are grappling with how to reap efficiency gains from AI while maintaining accuracy, privacy, and public trust.

- **Shifting Labor Dynamics:** The rise of AI in government work is beginning to alter labor relations. Public sector unions are closely watching AI's impact on jobs and day-to-day work. Many unions are **pushing back against unilateral implementation** of AI-driven changes, insisting that introduction of AI be a **subject of bargaining**. Concerns focus on job security (preventing AI-driven layoffs or deskilling), protection from increased electronic surveillance or algorithmic management, and ensuring employees share in productivity gains rather than being replaced. In some cases, unions have sought **contractual guarantees or side agreements** addressing AI – for example, prohibiting AI from displacing bargaining-unit workers or requiring notice and negotiation before new tech is deployed.

- **AI in Contracts – Emerging Trends:** A growing number of labor agreements now include language on AI and automation. In the private sector, dozens of recent contracts (especially in media, entertainment, and logistics) explicitly limit how AI can be used. **Common provisions** include: bans on using AI to replace workers or cut pay, requirements that AI output be **reviewed by humans** and clearly labeled, commitments to **training and upskilling** workers on new tools, and the formation of **joint labor-management AI committees**. For instance, one union contract states generative AI "*may be used*… as a complementary tool… but *may not be used as a primary tool*" for core work, and that AI use **shall not result in layoffs, position eliminations, or pay reductions**. Even in the public sector, early examples are emerging (e.g. a faculty union's proposed AI committee at a state university), signaling that AI is becoming a standard topic at the bargaining table.

- **Key Considerations for HR and Labor Relations:** HR leaders and labor relations professionals in government need to get ahead of the AI curve. Internally, this means **developing clear policies and guardrails** for employee use of generative AI – covering what uses are permitted, how to protect sensitive data, and what oversight is required. It also means **assessing job impacts**: identifying where AI might *augment* staff work versus where it could displace tasks, and planning for training or role adjustments accordingly. Proactive engagement with unions is crucial: experts recommend involving unions **early** (even forming joint AI working groups) to build understanding, address concerns, and negotiate any necessary agreements before AI tools are rolled out. Ensuring **fairness,**

transparency, and worker input at each step will reduce risk and build trust. With thoughtful strategy, public sector organizations can harness AI for public good while upholding workforce integrity – but the next 12–24 months are critical for putting the right frameworks in place.

(The rest of this briefing provides a detailed analysis of these points, including on-the-ground examples of AI use in government, evolving collective bargaining responses, excerpts from AI-related agreements, and strategic guidance for HR/LR teams.)

Introduction: A Primer on Generative AI in Government

What is Generative AI? Generative Artificial Intelligence refers to a class of AI systems (such as OpenAI's GPT-4, Google's PaLM 2, etc.) that can produce human-like content – text, images, code, and more – in response to prompts. Unlike traditional software that follows explicit instructions, generative AI models are trained on vast datasets and can *generate* new content by predicting likely responses. This capability enables tools like chatbots that can draft emails or answer questions in natural language, image generators that produce graphics from descriptions, and code assistants that suggest programming solutions. The technology grabbed mainstream attention with the late-2022 release of ChatGPT, and has since rapidly advanced in capability and popularity [2].

Why it Matters for Public Sector Work: Generative AI has the potential to significantly improve productivity in government operations. It can automate **non-routine tasks** that previously required substantial staff time – for example, writing first drafts of reports and correspondence, summarizing lengthy documents, translating communications on the fly, or sifting through databases to answer public inquiries. Used well, these tools might help agencies overcome staffing gaps and backlog issues by handling routine elements of work, allowing employees to focus on higher-value services. AI can also assist with analyzing complex data or accelerating processes (imagine an AI summarizing regulations to help determine benefit eligibility faster). In short, generative AI offers *new solutions to perennial public-sector challenges* – doing more with less, improving responsiveness, and innovating in service delivery.

Risks and Challenges: At the same time, AI introduces new risks and uncertainties for government. Generative models are prone to making errors or "hallucinating" false information, so **quality control and human oversight** are essential. Bias is another concern – if the AI's training data has embedded biases, its outputs could unintentionally perpetuate inequity or discrimination. Privacy and security loom large as well: public agencies handle sensitive data, and using external AI platforms raises questions about data exposure. Indeed, San Francisco's newly issued AI guidance explicitly warns city employees *not* to enter confidential information into public AI tools, since those inputs could be seen by the AI providers or leaked publicly [3]. There are also broader ethical questions about transparency (making sure citizens know when they're interacting with an AI) and accountability (who is responsible if an AI makes a faulty decision in a public service context).

Policy Responses Taking Shape: Recognizing these stakes, government organizations are starting to craft **AI usage policies** to guide staff. Many city and county governments have drawn on frameworks from pioneers like Boston and San Jose – as well as federal guidance like NIST's AI Risk Management Framework and the White House's AI Bill of Rights – to establish local guidelines. These policies typically stress *alignment with existing laws* (e.g. nondiscrimination, privacy laws), *risk mitigation* (bias, security, accuracy), *transparency* (disclosing AI use), and *human accountability* for AI-assisted decisions. For example, the City of Seattle's new Generative AI Policy (an "interim" policy launched in late 2023) requires that **any AI-generated**

output be reviewed by a human before use, that AI-generated images or videos be clearly labeled as such, and that employees apply an equity toolkit to evaluate potential biases in any AI application. Seattle's policy also enumerates principles like transparency, fairness, privacy, and security as cornerstones of acceptable AI use. San Francisco issued "preliminary" AI guidelines in early 2024 that *encourage* employees to experiment with AI for tasks like drafting emails or summarizing text, but with strict instructions to **fact-check all AI outputs and disclose when content was AI-assisted**, and to avoid using it for any sensitive or high-stakes decisions [3]. These examples illustrate the balancing act public employers are attempting – encouraging innovation and efficiency gains, while instituting guardrails to prevent mistakes or abuses.

In summary, generative AI is quickly moving from tech buzzword to practical reality in government workplaces. The next sections of this report delve into how exactly AI is being used on the ground in the public sector, how labor unions are responding to protect workers, what new contract language is emerging around AI, and what steps HR and labor relations teams should be taking **now** to navigate this new terrain.

Generative AI in State & Local Government: What's Happening on the Ground

Early Use Cases and Pilot Projects: Across the United States, state and local governments have begun deploying generative AI tools in a variety of roles. Most of the current use cases are still **pilot programs or limited experiments** rather than full production systems, but they span a wide range of government functions:

- **Administrative and Drafting Assistance (Tier 1 AI):** Many agencies are exploring AI as a "digital assistant" for routine tasks like writing letters, generating first drafts of reports or press releases, summarizing meeting notes, and creating data charts. These straightforward applications – which one framework labels *Tier 1* AI uses – are popular because they are relatively low-risk and often **low-cost or free**. For instance, employees have used tools like OpenAI's ChatGPT (or its specialized government-cloud version, "ChatGPT Gov") and built-in AI features in Microsoft 365 (Copilot) to draft emails or outline policy documents with greater speed. These tools can help bridge staffing shortages by **boosting individual productivity**, effectively allowing staff to get more done in less time.

- **Public-Facing Chatbots and Translation (Tier 2 AI):** A number of cities are testing AI to improve how they communicate with residents. **Chatbots on municipal websites** can answer frequently asked questions from citizens at any hour, helping people navigate services or find information more quickly. For example, a city might deploy a GPT-powered chatbot to handle questions about trash pickup schedules, permit requirements, or COVID-19 resources. Another compelling use is **real-time translation**: some emergency call centers have looked at AI translation services to instantly translate 911 calls or other service requests in multiple languages. This can be life-saving in communities with diverse languages, ensuring first responders get critical information without delay. These *Tier 2* uses put AI in direct interaction with the public, which is why they come with higher oversight needs – agencies must carefully vet the quality and accuracy of AI responses since there is less human review before the information reaches the user.

- **Data Analysis and Complex Problem-Solving (Tier 3 AI):** The most advanced (and still largely aspirational) uses involve AI analyzing large datasets to assist in policy decisions or operational

efficiency. State governments are interested in whether generative AI and related tools can help, for example, **clear backlogs** in processing applications for public benefits or housing permits. One vision is to have AI read and summarize complicated eligibility rules or scan application files to flag missing information, thereby speeding up approvals in programs like Medicaid or housing assistance. Likewise, housing departments might experiment with AI to evaluate building permit submissions or generate draft affordable housing plans. These *Tier 3* applications could be transformative in addressing chronic bureaucratic logjams, but they are also the furthest from turnkey implementation. They often require custom development, robust data governance, and significant human oversight given the high stakes (e.g. deciding who gets public benefits).

Leaders and Laggards: Progress in adopting generative AI in government is uneven. A few pioneering jurisdictions have leapt forward with formal initiatives: - **State-Level Initiatives:** In September 2023, California's Governor Newsom issued an executive order to encourage state agencies to explore responsible uses of generative AI [1]. By early 2024, New York State had even begun *encouraging* its employees to utilize AI tools in daily work where appropriate [1]. Pennsylvania launched a pilot program for state employees to use ChatGPT for tasks, touting it as the first of its kind in state government [1]. These top-down signals have spurred agencies in those states to identify candidate projects for AI integration.

- **City Governments:** Tech-forward cities such as **San Francisco, Seattle, Boston, New York, and San Jose** have crafted internal AI guidelines and are openly experimenting. San Francisco's guideline (released Dec 2023) not only urges staff to try AI for mundane tasks (with required fact-checking), but also commits to surveying departments on all current and potential AI use cases to build a public inventory of AI in use. Seattle's interim policy (Nov 2023) similarly lays out detailed rules to ensure any use of generative AI is transparent and fair – and that policy will evolve as the city learns from initial trials. These cities are effectively serving as laboratories for how AI might improve municipal operations, from **automating customer service responses to generating draft budgets**. They're also identifying pitfalls and sharing lessons through coalitions (for example, the *GovAI Coalition* is a network of local agencies sharing AI policy templates and best practices).

- **Areas of Rapid Uptake:** Notably, **IT departments and innovation offices** within governments are often the first to play with generative AI, using it to write snippets of code or summarize technical documentation. Some public information officers have used AI to help draft social media posts or press release outlines (with human editing). And as mentioned, call centers and 311 services are actively testing AI chatbots to handle simple queries. These are domains where the technology's current abilities closely match an immediate need.

Conversely, many jurisdictions are **moving more slowly** – often for good reason: - **Resource and Knowledge Gaps:** Smaller cities and counties may lack staff with AI expertise to confidently deploy these tools, or they may not even be fully aware of what AI could do for them. According to research by the National League of Cities, a lot of local governments remain in an "information gathering" stage about AI, unsure how to start. This lack of awareness and capacity naturally slows adoption. - **Legal and Ethical Concerns:** Some public officials are hesitant to green-light AI pilots until clearer rules or laws are in place. They worry about liability if an AI makes a bad decision, or about violating privacy laws inadvertently. Given the public sector's mandate to serve **all** residents fairly, caution is warranted – no city wants to be in the headlines for an AI-related blunder that harms constituents. As a result, certain agencies have hit the pause button on AI until they develop comprehensive ethical guidelines or wait for state-level regulation. - **Workforce Skepticism:** In some cases, employees themselves (and their unions) have raised concerns,

which can slow down management's desire to implement AI. If workers fear that a new AI tool is a step toward automating their job, they may resist participating in pilots or may demand negotiations before any introduction of AI (we will explore this dynamic more in the next section).

Key Challenges in Implementation: Even for enthusiastic adopters, deploying AI in government comes with challenges that can stall progress. Local officials report difficulty in **vetting AI vendors' claims** – with the AI industry evolving so fast, it's hard to validate which tools actually perform as advertised. Rigorous testing is needed to ensure an AI system is accurate and unbiased, but many agencies don't yet have the capacity (data scientists, testing frameworks, etc.) to do that. Additionally, concerns about **data management** ("Are our data sets even ready or safe to feed into an AI?") and **cybersecurity** (fear of new attack vectors through AI) give some CIOs pause. Lastly, there's the human factor: staff need training to use AI properly, and a portion of employees may be reluctant or uncomfortable using it at all, especially without clear guidance.

In summary, generative AI is **already making inroads in government work**, but in a very patchwork manner. A few trailblazers at the state and city level are demonstrating early wins (and sharing cautionary tales), while many others watch and learn. The coming year or two will likely see broader experimentation as awareness grows and pressure builds to improve efficiency. For public HR professionals, it's critical to understand this landscape – because as AI starts showing up in job functions, it inevitably will show up in labor-management discussions as well. We turn next to how that is unfolding.

Labor Relations in the Age of AI: What Public Sector Unions Are Facing

The introduction of AI into the workplace has not gone unnoticed by labor unions. In the public sector – as in many industries – unions are increasingly attentive to how generative AI and automation might affect jobs, working conditions, and the future of work for their members. Here's how collective bargaining dynamics are beginning to shift:

AI as a Mandatory Subject of Bargaining: Labor unions view the use of AI on the job as something that **must be negotiated**, not unilaterally decided by management. In legal terms, if an employer's use of AI could impact employees' working conditions or job security, it is likely a *mandatory subject of bargaining* (comparable to introducing new equipment or work methods). Unions are therefore asserting their right to be involved from the outset. The NewsGuild, representing media workers, bluntly states that the use of AI to perform bargaining-unit work falls in this category and thus requires negotiation – a stance increasingly backed by union actions. In practice, this means if a city hall or state agency wants to roll out an AI tool that might change how employees do their work (for example, an AI chatbot that handles citizen inquiries previously answered by clerks), the union expects advance notice and a chance to bargain over the decision and its effects.

Rising Worker Concerns: Union members themselves are driving much of this cautious approach. Generative AI burst into public awareness so quickly that it left many workers feeling uncertain or anxious about their future. Surveys and anecdotal reports show public employees have two big categories of concern: - **Job Displacement or Deskilling:** Employees worry that AI could either replace their jobs entirely or deskill their roles. For instance, if an AI system can generate a rough draft of a report in seconds, will management start to question the need for as many analysts or administrative staff? Even if outright layoffs

are not imminent, could the presence of AI erode the skill content of certain jobs, potentially affecting classifications and pay grades? These fears are not abstract – they've been echoed by union members from university faculty to government clerks. A faculty union leader in California noted that while panic isn't widespread, **"job replacement is a valid concern"**, and workers are understandably nervous that AI might make some roles obsolete or reduce the value of their expertise. Public sector unions, remembering past waves of automation, are determined to prevent "technology unemployment" among their ranks. - **Workload and Stress:** Paradoxically, there is also concern that AI could make work more intense rather than easier. If AI boosts productivity, managers might ratchet up performance expectations ("the AI helps you handle 50% more cases, so we can assign you more!") without commensurate pay increases or staffing relief. One union research report pointed out that in past tech transitions, productivity gains often went *"everywhere but the paychecks"* of rank-and-file workers. Unions fear a scenario where AI tools increase surveillance of workers or impose unrealistic metrics. For example, AI might be used to monitor how quickly an employee responds to emails or to analyze call times in a service center, pressuring workers to speed up. Indeed, unions like AFSCME have raised flags about **algorithmic management** – AI-driven systems that could monitor or direct workers in ways that undermine their autonomy and privacy.

Union Demands and Proposals: In response to these concerns, unions are starting to formulate specific demands at the bargaining table: - **Job Protection Clauses:** A top priority is language in contracts that **protects against layoffs or job eliminations due to AI**. Unions want assurances that AI will not be used as a pretext to cut positions or erode labor standards. For instance, in contract talks, a union may propose a clause that "no bargaining unit employee shall lose their job or suffer a reduction in pay as a direct result of the introduction of new AI or automation technology." We are seeing this exact kind of language being won in some sectors (detailed in the next section). The aim is to shift the narrative: AI should be a tool *in employees' hands*, not a replacement for those hands. - **Consultation and Consent:** Beyond just job security, unions are insisting on having a **say in if and how AI is implemented**. This can take the form of contractually mandated notice periods (e.g. the employer must notify the union X days before implementing any new AI system affecting workers) and the creation of **joint labor-management committees or task forces** to evaluate new technologies. The rationale is that workers and their representatives should have input on selecting AI tools, assessing their impact, and establishing policies for their use. One public university's faculty union, for example, has proposed a Letter of Agreement to form a committee with the administration to "investigate emerging issues related to generative AI" on an ongoing basis – essentially a forum to stay ahead of AI developments that could affect faculty work. - **Training and Retraining:** Unions are also advocating for employer commitments to **train workers on new AI tools** and to retrain any workers whose jobs might be significantly changed by AI. If an employee's role is going to evolve into more of an "AI supervisor" (checking AI outputs) rather than manual processing, they should get training to adapt, and possibly even a reclassification or pay adjustment if the skill set changes. Likewise, if certain tasks are automated, the union may seek agreements on reassigning affected employees to other meaningful work rather than simply eliminating their roles. The American Federation of State, County and Municipal Employees (AFSCME) has officially resolved to ensure that "workers whose jobs are impacted by [AI] receive protections, support and retraining as needed" so that they are not left behind by technological change. - **Ethical and Fair Use Policies:** In some negotiations, unions are proposing contract language that sets **ethical guidelines** for AI use. This can include commitments to bias testing (so that AI decisions don't unfairly discriminate among employees or the public) and transparency measures (for example, if AI is used in evaluating employees or making recommendations, those affected must be informed). Unions representing professional employees, like teachers and journalists, have been vocal that any use of AI in content creation must come with clear rules on attribution and quality control, to preserve the integrity of their work. Public sector unions similarly might push for provisions ensuring AI isn't used to make

disciplinary or hiring decisions without human evaluation, to avoid hidden algorithmic biases affecting people's careers.

Union Pushback in Action: We are already seeing instances of unions pushing back in real time against perceived overreach with AI: - In **higher education**, when a dean at Boston University suggested faculty could use AI tools to help manage classes during a grad-student strike, the faculty union (affiliated with SEIU) bristled at the idea. The mere mention of using AI as a strike replacement or stopgap raised *"alarm bells"* for the union, which saw it as a threat to both jobs and academic quality. The incident underscored that even suggestions of AI replacing human roles (in this case, teaching assistants) will draw swift union scrutiny. - In the tech and media industry (which often foreshadows trends for others), unions have not hesitated to strike over AI issues. The recent high-profile strikes by the Writers Guild of America (WGA) and SAG-AFTRA (screen actors) in 2023 prominently featured AI concerns – writers demanded limits on AI-generated scripts and actors sought protections against digital replicas of their likeness. The fact these unions won significant AI provisions through tough bargaining is not lost on public sector unions, who see a precedent to follow. - Even in manufacturing and service industries, unions like the Teamsters (at UPS) and UNITE HERE (hospitality) have gone to the mat to negotiate over automation and AI (e.g. UPS won a ban on drones and driverless vehicles for the contract term). These victories send a message: **unions can successfully negotiate constraints on new technology**. Public sector unions, though dealing with governments rather than private companies, are likely to be similarly emboldened to insist that AI be introduced only in a worker-friendly, negotiated manner.

Challenges for Unions: On the other side of the table, public employers and HR might note that many current union contracts still do *not* mention AI explicitly – this is all very new. Unions are essentially racing to address a moving target, and not all have the technical expertise to know exactly what to bargain for. Some union leaders worry about being overly prescriptive and inadvertently hindering beneficial uses of technology. That's why some are opting for broader protective language (principles and committees) rather than detailed prohibitions that could become outdated. As one higher-ed union activist put it, *specific, detail-oriented conversations* are best – it's a complex topic and not a one-size-fits-all situation. Unions are learning from each other and even looking to other sectors (like the auto workers or entertainment unions) for model language. We can expect that **union coalitions and federations** (e.g. AFL-CIO affiliates) will increasingly share resources on how to tackle AI in bargaining, much as they have with health and safety language in the past.

In summary, the labor relations context around AI in government is evolving quickly. There is a clear **protective stance** from unions: they want a seat at the table for any tech changes, and they prioritize job security, privacy, and fairness. For HR and labor relations professionals, this means any plans to deploy AI will likely require conversation (if not negotiation) with your unions. The next section will highlight some concrete language and agreements that have emerged so far, which can serve as reference points for both union and management representatives navigating this new terrain.

AI in Collective Bargaining Agreements: Emerging Contract Language and Themes

Although we are still in the early days, a variety of **AI-related clauses and agreements** have started to appear in collective bargaining across different sectors. Below is a sampling of notable examples – including policy language from memoranda of understanding (MOUs), side letters, and full contract provisions – that

address AI and automation. These examples, while not all from the public sector, illustrate the themes and protections that are likely to influence public sector bargaining as well.

- **No Replacing Human Creators (Media Industry – *The New Republic*):** Journalists at *The New Republic* magazine, represented by The NewsGuild, won strong contract language regulating AI. The clause specifies that **generative AI "may be used by bargaining unit employees as a complementary tool"** in their editorial work, **"but it may not be used as a primary tool for creation"** of content. In other words, writers can choose to use AI to assist (like suggesting headlines or doing research), but management cannot force them to use AI or use it to bypass them. Crucially, the contract also states that AI **shall not be used to cause layoffs, to fill vacant positions, or to reduce the pay** of Guild-represented staff. This is a clear guarantee that AI won't cost jobs or undermine wages – a precedent-setting protection for content creators.

- **Transparency and Labeling of AI-Generated Content (Digital Media – Ziff Davis):** In a recent agreement at Ziff Davis (a digital media company), the union secured a commitment that if the company uses AI to generate or edit any content that appears alongside employees' work, it **must be clearly identified as "AI-Generated Content."** The contract lays out detailed guidelines for such labeling across text, images, and multimedia. This provision ensures **transparency** to both workers and the audience – readers will know what's written by a human and what had AI involved. It protects the integrity of the workers' contributions and holds the employer accountable for AI use. (This kind of transparency measure could easily translate to public sector communications: imagine a city requiring that any AI-written section of a public report be noted as such.)

- **Mandatory Training on AI Tools (Multiple NewsGuild Contracts):** A common provision in many of the 30+ media contracts with AI language is a requirement that **employees be trained in any AI tools** introduced to the workplace. For example, if a newsroom introduces an AI system to help scan documents for news leads, the employer must provide training on how to use it properly and ethically. The reasoning is twofold: workers should be equipped to use new technology effectively (so they can benefit, not be left behind), and they need to understand the **ethical standards and pitfalls** (to avoid, say, inadvertently spreading an AI-generated error). Some contracts even tie this to discipline – clarifying that employees won't be held accountable for misuse of AI if they weren't trained, and conversely that after training, improper use (like failure to fact-check AI output) could have consequences. This emphasis on training aligns with public sector values of professional development and could appear in government MOUs as well.

- **Advance Notice and No-Layoff Guarantee (Logistics – UPS Teamsters):** In the landmark 2023 contract between UPS and the Teamsters union covering 340,000 delivery and warehouse workers, the union won unprecedented protections against new technology. The contract **prohibits the company from implementing any technology that would result in the layoff of drivers or package handlers for the life of the contract (through 2028)** – including specific mention of **drones, driverless vehicles, and other AI-based innovations**. UPS must also negotiate with the union over any such technology before deployment. This effectively pauses automation-driven job cuts, giving workers security and leverage as these technologies evolve. While this is a private sector example, it's noteworthy for public transit or delivery services that might consider drones or automated vehicles; similar clauses could be sought by public employee unions in those fields.

- **Comprehensive Tech Impact Provisions (Hospitality – Las Vegas Culinary Union):** In late 2023, UNITE HERE Local 226 (the Culinary Union) in Las Vegas negotiated new contracts for tens of thousands of casino hotel workers that addressed automation and AI head-on. These contracts **strengthened existing technology protections** by requiring: **advance notice** to the union before any new technology is introduced, **free training and upskilling** opportunities for workers affected by new tech, **severance pay and transfer options** for any workers displaced, and **privacy safeguards** for employees being monitored by AI systems. For example, if a hotel wanted to introduce a robot to deliver room service (potentially impacting room servers), they would have to notify the union well in advance, discuss alternatives, train staff to work with or operate the new system, and ensure that no one simply loses their job without compensation or another opportunity. Such a holistic approach could serve as a model for public sector agreements, especially in health care or administrative offices where automation might change how work is done. It's about creating a **"technology buffer"** – a set of guarantees that soften any negative impact on workers.

- **Joint AI Committees and Expert Input (Auto Industry – UAW & Ford):** The United Auto Workers' 2023 contract with Ford Motor Company includes detailed Letters of Understanding addressing future tech like AI. One letter establishes a **joint union-management committee to research AI technology for its implications on worker safety and operations**. The company must share information about new AI-driven processes and engage with the union through this committee. Additionally, Ford agreed to provide **advance notice** of any planned technological changes (such as automation in assembly processes) and to invest in training for workers to be able to fill the high-tech jobs created by these changes. Essentially, the union secured a voice in *how* AI is adopted and assurances that workers will benefit (through training) rather than be left behind. In the public sector, we are likely to see similar structures – for instance, a state DOT might form a labor-management committee on AI if considering automated traffic monitoring systems, ensuring workers can raise concerns and adapt.

- **Regular Consultations on AI (Tech Sector – New York Times Guild):** When the tech employees at The New York Times (members of the NewsGuild) won their contract in 2023 after a strike, one of the provisions was the creation of a **Generative AI consultation committee**. The contract language calls for a joint committee to **"discuss the potential impact of Generative Artificial Intelligence,"** meeting at least twice a year at the union's request. This committee doesn't outright limit management's use of AI, but it guarantees a structured, ongoing conversation where the union can get information and provide input about AI initiatives. It's a recognition that AI is evolving, so a static one-time agreement might not suffice – ongoing dialogue is needed. A similar approach could be very useful in government settings: for example, a city could agree with its unions to hold semiannual meetings on workplace technology, including AI updates, so that both sides stay aligned and any brewing issues are discussed before they become crises.

- **Letters of Agreement on AI (Public Higher Ed – Oregon State University):** In early 2024 bargaining, the faculty union at Oregon State University (UAOSU) proposed a forward-looking Letter of Agreement on Artificial Intelligence. Rather than trying to codify exact rules (since academia is still figuring out how AI will be used in teaching/research), they proposed forming a **joint committee of union and administration representatives to meet at least semiannually and "investigate emerging issues related to generative AI."** The union explained that AI issues are *"too nebulous to build a prescriptive article around"* at present, but they wanted a mechanism in place to **"work collaboratively to be prepared"** as AI impacts arise during the contract period. This proactive

approach – essentially an insurance policy that the parties will handle AI developments together in good faith – could be a template for other public sector contracts that are coming up for negotiation now, while AI usage is still in flux.

Themes Across Agreements: From these examples, a few **common themes** emerge: - **Protecting Jobs and Wages:** Virtually all AI-related contract language seeks to ensure that technology is used to assist workers, not replace them. "No layoff due to AI" clauses and commitments to reassign or compensate any displaced workers are central. - **Worker Involvement and Oversight:** Whether via formal committees or required consultations, unions are securing a role in monitoring and guiding AI implementation. This reflects a demand for *transparency* – the employer must share what AI is being used and why – and for *worker voice* in how it's deployed. - **Training and Adjustment:** Many agreements insist on employer-provided training on new tools and in some cases extended support (like retraining or severance) if jobs fundamentally change. This ensures workers can adapt and ideally thrive alongside AI. - **Ethical and Transparent Use:** Contract clauses often address the *manner* in which AI is used – requiring things like human review of AI decisions, clear labeling of AI-generated materials, and respect for privacy and intellectual property. These not only protect workers but also uphold quality standards in the work itself (important in public service contexts). - **Flexibility for Future Tech:** Given the rapid pace of change, several unions have chosen mechanisms (like MOUs or side letters) that allow revisiting the topic or updating terms as AI evolves. This is a pragmatic recognition that what might be adequate policy today could need strengthening tomorrow.

For public sector HR and labor relations practitioners, these emerging contract provisions are instructive. They offer a menu of potential approaches – from hard-and-fast prohibitions to collaborative processes – that can be tailored to the public service environment. The **implication** is that AI language is likely to appear in your next round of contract negotiations if it hasn't already. Being prepared to negotiate on these points (and understanding what similar employers have agreed to) will put you ahead of the curve.

What HR and Labor Relations Teams Should Be Doing Now

As generative AI continues to gain ground in government workplaces, public sector HR and labor relations (LR) professionals have a pivotal role in shaping a responsible and constructive path forward. This section outlines strategic guidance – essentially a checklist of questions and actions – to help HR and LR teams navigate the next 12–24 months. The overarching goals are to **assess risk**, **build internal alignment**, engage unions proactively, and ensure that AI adoption happens with fairness and foresight.

1. Educate and Build Internal Knowledge: First and foremost, HR should ensure that key leaders and stakeholders understand what generative AI is (and isn't). This might involve: - **Internal briefings or workshops** on AI basics, tailored to government contexts. A "brief primer" for managers can dispel myths and clarify capabilities. The aim is to prevent both unrealistic expectations ("AI will solve everything overnight") and undue fear ("AI will immediately replace everyone"). - Keeping abreast of **policy guidance** from trusted sources. Organizations like the National Association of Counties (NACo), National League of Cities (NLC), and groups like the *GovAI Coalition* are publishing templates and frameworks for AI governance. HR can digest these and pick out relevant pieces for their agency. - **Monitoring legislation and regulation**: With federal and state rules around AI emerging, someone in HR/LR should track developments (e.g. EEOC guidelines on AI in hiring, or state laws requiring AI impact assessments). Compliance with any new mandates will be part of HR's responsibility.

2. Develop Clear AI Usage Policies: If your government or department hasn't done so yet, HR should lead (or collaborate in) the creation of a **generative AI policy for employees**. This policy should define how staff may use AI on the job, with an emphasis on **safety, ethics, and consistency**. Key questions to address: - **Permissible vs. Prohibited Uses:** What tasks are employees allowed to use AI for? (e.g. drafting routine memos, yes; making final policy decisions, no.) And what uses are off-limits? For instance, many policies forbid using AI to generate official communications without review, or using AI for any form of decision that impacts individual rights (like approving benefits or evaluating employee performance) without human oversight. Lay these out clearly. - **Data Privacy and Security:** What types of information can be input into AI systems? It's wise to prohibit entering any confidential, personal, or sensitive data into external AI tools [4]. If using an approved internal AI system, outline how data will be handled. Employees should be made aware that prompts to tools like ChatGPT might be stored by the provider. The policy might require use of only **city-approved AI platforms** that meet security standards. - **Human Oversight and Accountability:** The policy must stress that employees remain responsible for the outputs. For example, Seattle's policy mandates human review of all AI-generated content before it's used officially. Define who must sign off on AI-generated work and in what contexts. The idea is to prevent blind reliance on AI and ensure a human is accountable for any final product. - **Ethical and Equity Checks:** Include guidelines to mitigate bias or ethical issues. Some governments require an "equity assessment" if an AI tool is being adopted, to examine potential disparate impacts. The policy can encourage (or require) employees to consider biases in AI outputs and correct them. Align this with your existing diversity and equity commitments. - **Transparency:** Make it a norm to disclose AI involvement. Whether internally (employees tagging a document as having AI assistance) or externally (letting the public know when a chatbot is automated), transparency builds trust. In fact, several cities and counties are committing to **public transparency about AI uses**, even putting inventories of AI systems online. - **Training and Support:** State that the organization will provide training on approved AI tools. This not only helps employees use them effectively, but also signals that misuse will be addressed through training rather than immediate discipline. (A carrot-before-stick approach.) - **Updating Mechanism:** Technology will change – note that the policy will be revisited periodically. An AI policy shouldn't be a dead document; consider establishing an AI governance committee or designating a point person to update guidelines as needed.

Having a solid policy framework in place not only guides employees, it also serves as a **starting point for discussions with unions**, who will want to ensure the policy is fair. In fact, inviting union input while drafting the policy can be a trust-building measure.

3. Proactively Engage Unions and Employees: Don't wait for a union to raise a grievance or demand bargaining over AI – **engage them early**. Some approaches: - **Joint Briefings or Workshops:** Host informational sessions about AI for union representatives and employees. Show them the kind of AI tools being considered, their limitations, and get feedback on concerns. This demystifies AI and reduces fear of the unknown. - **Form a Joint Committee or Working Group:** As seen in several examples, a labor-management AI committee can be invaluable. Consider establishing one now (even outside of contract negotiations) to discuss any planned AI pilots. For instance, if IT wants to try an AI chatbot, sit with union reps in a working group to outline how it will work, what jobs or workflows it touches, and how you'll evaluate it. Make it clear the committee's purpose is collaborative problem-solving, not just a formality. A standing committee can also monitor implementation of your AI policy and suggest improvements. - **Communicate Intent and Benefits:** When introducing an AI tool, frame it clearly: how will it help employees and the public? Emphasize that it's there to *assist*, not replace. For example, if deploying an AI to draft routine emails, explain to staff that the goal is to free up their time for more complex customer service calls that only humans can handle. By articulating a vision where AI handles drudgery and employees

handle the nuanced work, you align with the union's interest in preserving meaningful jobs. And be honest about what AI can't do – this helps set realistic expectations and shows you're not viewing it as a magic bullet. - **Negotiate Ground Rules if Needed:** In some cases, particularly if an AI implementation could tangibly affect workloads or job duties, it might be appropriate to reach a side agreement or MOU with the union up front. For example, you could agree that *no employee will be laid off or face a loss of hours due to this AI pilot, and the union will be provided results/data from the pilot after 6 months to evaluate its impact.* Such an agreement creates accountability and eases union anxiety, allowing the pilot to proceed under mutually understood conditions.

Engaging unions early not only heads off adversarial reactions, it often improves the end result – employees on the front line might spot issues with an AI tool that managers or vendors didn't, saving headaches down the road. As one expert put it, **worker participation improves AI outcomes** because the people doing the work can help ensure the technology actually fits their needs and doesn't inadvertently create more problems.

4. Assess Job Impact and Plan for Workforce Changes: HR should take the lead in analyzing how AI might reshape various jobs and in developing plans to address those changes positively. Key steps include: - **Task Mapping:** Look at the major job classifications in your workforce and identify tasks within those jobs that AI could plausibly handle or assist with. This doesn't mean those jobs go away – it means those tasks might be done differently. For example, if "data entry" is a big part of certain clerical roles, note that AI could automate some data transcription. If "researching case law" is something staff attorneys do, note AI's capabilities in quickly summarizing legal documents. By mapping this out, you can foresee which roles will be most affected. - **Classification and Duty Updates:** If AI takes over some tasks, employees' **primary duties might shift**. You may need to update job descriptions or even reclassify certain positions. For instance, an "Office Assistant" whose job was 50% typing dictated letters might, with AI, spend that time doing customer outreach instead – that moves the role towards a different skill set. HR should be proactive in considering whether such changes warrant reclassification (possibly to a higher grade if the remaining work is more complex) or simply retraining. In any case, involve the union in these discussions to avoid classification grievances later. Many union contracts have provisions requiring bargaining if job duties substantially change – be mindful of those. - **Training and Upskilling Programs:** Once you identify where AI can augment work, plan how to **train staff for that transition**. This could mean digital literacy training, AI tool workshops, or even partnering with local community colleges or universities for short courses. Remember, as AFSCME's resolution emphasizes, workers should have "access to opportunities that AI creates through education and training". Budget and plan for this now, as it may require new training funds or adjusting professional development priorities. - **No-Layoff Strategies:** Decide on your stance regarding workforce size if AI improves efficiency. A recommended approach (and one likely to ease union fears) is to commit that you will achieve any AI-related efficiencies **through attrition or reassignment, not through layoffs**. For example, if an AI system in a permitting office means each staffer can handle 30% more permits, you might gradually reduce headcount by not backfilling some vacancies over time – as opposed to pink-slipping 30% of the staff immediately. Communicate this philosophy clearly. It aligns with the **"no displacement without redeployment"** theme unions advocate. If there is absolutely no alternative to a position being eliminated, be prepared with generous measures (offer transfers elsewhere in the organization, or retraining for a different role, or at last resort, severance packages). Ideally, though, emphasize how you will use AI to *augment* your workforce's capacity to serve the public better, not to cut budgets by cutting people.

5. Evaluate and Mitigate Risks (Bias, Accuracy, Legal Compliance): HR, along with legal and IT teams, should institute a process for vetting AI tools for potential risks: - **Bias and Fairness Audits:** Before deploying an AI tool, ask: could its output be biased or discriminatory? If using AI in hiring or promotion decisions (e.g., screening resumes or analyzing interview videos), be extremely cautious – algorithmic bias in hiring is a hot-button legal issue. Many jurisdictions (like New York City) have laws requiring bias audits of AI hiring tools. Even in service delivery, ensure the AI is tested for equitable performance across languages, demographics, etc. Insist on vendors providing bias audit results, or engage an independent expert to evaluate the tool. - **Accuracy and Quality Control:** Determine acceptable accuracy rates and error types for an AI in a given task. For instance, a chatbot giving out general info might be tolerated with a low error rate if answers are double-checked, but an AI drafting a council bill must be nearly perfect. Put in place a **review workflow** – who checks AI outputs and how. Perhaps a rotating team of experienced staff reviews a sample of AI-generated material each week to catch any drifts in quality. - **Legal and Policy Compliance:** Ensure any AI use complies with existing laws – e.g., **ADA (Americans with Disabilities Act)** if an AI tool is used in HR (does it unfairly screen out people with disabilities?), or public records laws (AI-generated documents might be subject to FOIA/public records requests, so you need to preserve them). Also, does using a cloud AI service comply with data residency or privacy policies your government has? Run these checks by legal counsel. - **Surveillance and Privacy:** If AI is used to monitor employees (even indirectly, like analyzing their computer usage for productivity), be very careful. Not only can this damage trust and morale, it may conflict with contract language or labor law (the NLRB's General Counsel has indicated that pervasive electronic surveillance of employees can violate labor rights in some cases). Always ask: is this use of AI truly necessary, and have we weighed its impact on employee privacy? It might be better to avoid certain "big brother" uses of AI entirely, or at least negotiate them with the union before implementation.

By conducting these risk evaluations, HR can develop a **set of guardrails** that satisfy both management objectives and union concerns. In fact, many of the local government AI policies being published explicitly highlight risk mitigation strategies focusing on bias, privacy, and security – HR should incorporate those into labor discussions.

6. Foster a Culture of Collaboration and Continuous Improvement: Finally, approach AI as an opportunity for **labor-management collaboration** rather than conflict. The best outcomes will come when both sides work together to integrate technology in a way that improves public services *and* job quality. To that end: - **Internal Alignment:** Ensure that HR, IT, legal, and department heads are all on the same page about the AI strategy. Mixed messages will create confusion. HR should articulate a unified position: we support trying AI to improve service, but we are committed to doing so responsibly and with our employees' involvement. - **Public Communication:** Consider jointly communicating (management and union together) to the public or governing boards about AI initiatives. For example, a city might issue a report or press release noting that it is piloting an AI tool **with input from its labor unions** and that the goal is to enhance services while safeguarding workers. This shows a united front and builds credibility that the implementation is well thought out. - **Stay Flexible and Open to Feedback:** Encourage employees to report issues or suggest improvements with any AI tools they use. Maybe even set up a dedicated feedback channel. Front-line workers will notice problems (or hacks and workarounds) that higher-ups won't. Use that feedback to adjust policies or provide additional training. Iteration is key – treat initial deployments as learning experiences, and be willing to pause or roll back something if it's not working as intended. - **Ethical Leadership:** HR and LR professionals should champion the ethical use of AI, keeping in mind public service values. If a proposed AI application doesn't sit right (say, an algorithm that decides who gets inspected by a city agency, which might unfairly target certain neighborhoods), speak up and reevaluate. It's easier to earn

union and employee trust if they see management being *thoughtful and values-driven* about tech adoption, not just chasing the latest trend or budget cut.

By considering these steps, HR and labor relations teams can turn a potentially divisive issue into a collaborative modernization effort. The presence of AI in the workplace is inevitable, but whether it becomes a source of labor conflict or a jointly managed innovation is largely up to how we approach it now. Many unions have signaled they are **not against technology per se** – they just want to ensure it's implemented in a way that is fair and transparent. As one union leader noted, *"we can be in a place where we're creative and look to other industries"* for solutions, learning how to adapt without panic. With honest dialogue and careful planning, public sector HR can lead the way in modeling that kind of high-road approach to AI integration.

Conclusion: Looking Ahead to an AI-Augmented Public Workforce

Generative AI is poised to be a game-changer in how government services are delivered and how public employees do their work. The next 12 to 24 months will be a formative period in which many agencies move from small-scale experiments to broader adoption of these tools. Alongside that technological shift, we will see a parallel evolution in labor relations and workforce management. Public HR leaders and labor professionals who proactively engage with these changes will put their organizations in the best position to harness AI's benefits while avoiding its pitfalls.

What's Next on the Technology Front? We can expect AI capabilities to continue advancing rapidly. Vendors are likely to offer more **government-tailored AI solutions** – for example, AI systems pre-trained on legal codes for use by city attorneys, or chatbots specifically designed for DMV services. AI will also become more embedded in the software public agencies already use (Microsoft's Copilot integration into Office products is one example, and similar features will appear in case management systems, HR systems, etc.). This means AI tools might "show up" in workers' daily routines even without a big initiative – simply through routine software updates. HR needs to stay ahead of software procurement and updates, ensuring that any embedded AI features are covered by your policies and training.

A New Normal for Collective Bargaining: On the labor relations front, it's almost certain that **AI clauses will become standard in public sector contracts** over the coming years. Just as today most contracts have language on past practices, health and safety, and nondiscrimination, tomorrow's agreements may routinely address technology and AI. We may see statewide frameworks (for example, a public employee union federation negotiating general AI principles applicable across agencies) or local unions including AI side letters during mid-contract periods if something arises suddenly. The themes will likely mirror what we outlined: no layoffs without negotiation, training guarantees, consultation committees, and protecting rights and privacy. Labor relations professionals should prepare for these conversations – researching what other jurisdictions are agreeing to, and considering what management can agree to that both protects employees and preserves management's ability to innovate. It's a delicate balance, but many private sector deals have found win-win language, suggesting that compromise is achievable.

Opportunities for Collaboration: Rather than viewing AI as a tug-of-war, progressive public employers and unions might find that it's an opportunity for **collaboration to improve public services**. For instance, a union and management could jointly pilot an AI in an understaffed area (say, an AI assistant to help process veterans' benefit applications faster) with an agreement that if it works, the saved time will be reinvested into outreach to veterans, and workers involved will get training to manage the AI and do the outreach. In

other words, productivity gains could be channeled into *expanding or improving services*, not just budget cuts. This kind of approach can align the mission of serving the public with the workforce's interests – a true "high road" use of technology.

Managing the Risks – A Continued Priority: That said, HR must remain vigilant about the risks. Issues like AI bias or a poorly thought-out implementation can have real human consequences – from eroding public trust to inadvertently discriminating against employees or residents. The public sector has an obligation to use AI in a way that is **transparent, accountable, and equitable**. We might see more oversight from legislators or auditors on this; for example, city councils could require reports on AI use or states might mandate impact assessments. By proactively adopting good governance practices now (as many cities are doing with AI inventories and ethical use policies), agencies can stay ahead of regulatory pressure and demonstrate to the public that they are using AI responsibly.

The Human Element – Irreplaceable: One reassuring insight that has emerged from early adopters is that AI, for all its marvels, highlights the value of human judgment and empathy more than ever. AI can crank out information and even make basic decisions, but humans are needed to interpret nuance, make ethical choices, and connect with people. Public service is fundamentally about people helping people – AI is a tool that can assist, but not replace, that human mission. As AFSCME's resolution underscored, having workers *"at the table"* in designing and implementing AI is crucial to its success. When employees are involved, AI is more likely to be used in ways that actually make work better and services better.

Confidence and Preparedness: By now, armed with an understanding of how generative AI is unfolding in government and labor contexts, HR and labor relations professionals should feel more confident stepping into this fast-moving space. There will undoubtedly be surprises and challenges ahead – technologies we can't yet foresee and scenarios that test our policies. But the core principles remain steady: **fairness, transparency, collaboration, and a focus on public service values**. If we keep those principles at the heart of our approach to AI, we can navigate the uncertainties effectively.

In conclusion, generative AI is neither a magic fix for all government problems nor an unstoppable threat to public jobs – it is a powerful new set of tools that, like any tool, can have positive or negative effects depending on how we handle it. The responsibility lies with public sector leaders, HR professionals, and unions **together** to guide AI integration in a way that **augments human capability, protects workers, and enhances the services we deliver to our communities**. By learning from what's happening on the ground now, anticipating the implications for bargaining, and asking the right questions, we can approach this technological evolution with both integrity and ingenuity.

References:

- Urban Institute – Christina Prinvil et al., *"A New Approach to Helping Local Governments Navigate Generative AI,"* June 2, 2025. (Overview of tiered AI use cases in local government and recommendations for AI policies.)

- San Francisco Chronicle – Chase DiFeliciantonio, *"S.F. finally created AI rules for city workers. Here's what they say,"* Jan. 9, 2024 3 1 . (Summary of San Francisco's generative AI guidelines for city staff, emphasizing fact-checking, disclosure, and cautious use.)

- GeekWire – Taylor Soper, *"Seattle releases a generative AI policy to set guardrails around new tech for city employees,"* Nov. 2, 2023. (Details of Seattle's AI policy requiring human review of AI outputs, equity toolkits, and principles aligning with federal guidance.)

- Center for Democracy & Technology – Maddy Dwyer and Quinn Anex-Ries, *"AI in Local Government: How Counties & Cities Are Advancing AI Governance,"* April 15, 2025. (Analysis of AI policies across 21 local governments, highlighting common trends like risk mitigation, transparency, and alignment with existing laws.)

- Inside Higher Ed – Susan D'Agostino, *"Faculty unions seek AI guidelines in contracts and policies,"* April 18, 2024. (Exploration of how faculty unions are reacting to AI in academia, raising concerns about job security, autonomy, and pushing for contract language on AI.)

- AFSCME (American Federation of State, County and Municipal Employees) – *Resolution No. 7: "Artificial Intelligence,"* 46th International Convention, 2024. (AFSCME's official stance advocating for worker involvement in AI use, equitable sharing of AI benefits, training for affected workers, and strong safeguards through collective bargaining.)

- Center for American Progress – Aurelia Glass, *"Unions Give Workers a Voice Over How AI Affects Their Jobs,"* May 16, 2024. (Report on how unions across industries have negotiated AI-related provisions to protect jobs, limit surveillance, and shape AI deployment.)

- The NewsGuild-CWA – Danielle Newsome, *"Guild members are winning strong protections from employer-pushed AI,"* May 12, 2025. (Summary of AI language in 30+ media industry contracts, including examples of protections on use of AI, job security guarantees, labeling requirements, and joint committees.)

- Washington Center for Equitable Growth – Virginia Doellgast and Nell Geiser, *"Boosting U.S. worker power and voice in the AI-enabled workplace,"* Feb. 19, 2025. (Policy brief comparing German and U.S. approaches to worker voice in AI adoption, with examples from recent union contracts like WGA, SAG-AFTRA, UPS, and others.)

- United Academics of Oregon State University (UAOSU) – Bargaining Update, *"Generative AI, Grievances & Termination…,"* Apr. 16, 2024. (Union bargaining update describing a proposed Letter of Agreement to form a joint committee on generative AI issues, as an approach to handle AI developments during the contract term.)

---

1 3 4 SF finally made AI rules for city workers. Here's what they say
https://www.sfchronicle.com/sf/article/s-f-lays-ai-guidelines-city-workers-here-s-say-18597152.php

2 Faculty unions seek AI guidelines in contracts and policies
https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2024/04/18/faculty-unions-seek-ai-guidelines-contracts