Embracing the Future: Responsible Use of Generative Artificial Intelligence in Local Government Operations

By Nick Cotton-Baez, CIRSA Associate General Counsel

INTRODUCTION

The rapid advancement of artificial intelligence (AI)—a broad field aimed at creating intelligent machines capable for performing tasks typically requiring human intelligence—has transformed various sectors of the economy, and local government is no exception. From automated customer service chatbots to predictive maintenance of infrastructure, AI has the potential to transform the way local governments operate and serve their constituents. As local governments navigate the implementation and use of AI in government operations, it's crucial for municipal officials to understand both the opportunities and risks.

This article focuses on "Generative AI", which is a subset of AI that refers to systems utilizing algorithms and statistical models to process, categorize, analyze, and draw inferences from patterns in the dataset on which the program has been trained (e.g., large language models) to generate statistically probable outputs, such as text, images, and videos, when prompted by users. Popular Generative AI tools and platforms include ChatGPT, Dall-E2, Lensa AI, and Perplexity AI, among many others.

Generative AI platforms and tools offer local governments unprecedented opportunities to enhance public services, optimize job performance, streamline processes, and improve citizen engagement. However, their use gives rise to a variety of risks, and thus their implementation must be approached with caution and foresight.

This article is intended to help CIRSA members identify key risk considerations associated with the use of Generative AI tools in local government—including data privacy, copyright infringement, ethical considerations, and the potential for inaccuracy and bias—in view of local governments' role of remaining at the forefront of innovation while upholding the trust and welfare of the communities they serve.

RISKS

By now, you've likely come across ominous warnings about AI's potential to surpass human intelligence, giving AI systems the ability to undertake actions beyond human control, potentially in malicious ways that may pose existential threats. Maybe you fear AI's potential to increase government surveillance and social manipulation, leading to the erosion of citizen privacy and enabling authoritarian control. While these and other potential outcomes shouldn't be ignored, they're largely out of local governments' control.

That said, there are several known risks surrounding the use of Generative AI that local governments can control and minimize by implementing sound policies and training programs for the appropriate and responsible use of Generative AI. Local governments must familiarize themselves with these risks to inform good policies and practices surrounding the use of Generative AI in local government services and functions.

Risks associated with Generative AI use generally fall into two broad categories: (A) risks associated with user prompts; and (B) risks associated with using AI-generated content.

A. User Prompts

User prompts and AI-generated responses generally are not private, as Generative AI models recycle and learn from previous interactions (i.e., prompts and responses) with users. If used in prompts, security breaches involving Generative AI systems may contain personal information or personally identifying information and other sensitive and confidential information, which could give rise to liability under data protection or other laws.[i] The practice of copying and pasting information from local government records into prompts may increase this risk. Additionally, use of Generative AI may result in the creation of a public record subject to disclosure under the Colorado Open Records Act (CORA). When using Generative AI, local government officials and employees should keep this in mind, particularly if it's in the local government's interest or required by CORA that certain prompts and AI-generated content are kept confidential.

B. Using AI-Generated Content

Many risks associated with using AI-generated content arise from characteristics of the source materials from which the Generative AI system draws to generate its responses to user prompts. Indeed, Generative AI systems are most-often trained on data sourced from the internet, which may produce inaccurate or fictitious responses to prompts, content reflecting the biases of such data, or content containing copyrighted material.

Generative AI is designed to make up information when it does not have an answer, which may lead to inaccurate responses because the data it's using is incomplete or inaccurate. In some cases, Generative AI has even been observed to "hallucinate" website URLs, court cases, and other purported source materials. This is exactly what happened when a Colorado attorney used ChatGPT to generate case citations and case-specific details to support his motion to set aside an unfavorable district court decision entered against his client. The lawyer filed the motion without verifying whether the cases existed, and when the judge noticed the cited cases did not exist, the lawyer lied about his use of ChatGPT blaming the fictitious case citations on a legal intern. The events led to the suspension of the attorney's law license. Lawyers in Texas and New York have also been sanctioned for citing fictitious cases generated by AI chatbots. It easy to see from just these examples in the legal field how use of unverified AI-generated content in any context can significantly damage careers and reputations and, for local governments, public trust. Thus, as discussed more below, your organization's AI policies should require any AI generated content be verified!

Moreover, content produced by Generative AI systems may include copyrighted material, which has formed the basis for many lawsuits brought against Generative AI developers by artists, authors, and other content creators. Generative AI systems are trained using data that has been sourced from the internet, often without regard to copyright or licensing terms. It's often difficult to determine the content used to train a Generative AI system, and the extent to which AI-generated content regurgitates copyrighted material. Some Generative AI systems do not include citations or links to the websites or other sources used to generate responses to user prompts, which may hamper the user from verifying the accuracy and completeness of the generated response and the right to use it.

Additionally, Generative AI systems can reflect the biases (e.g., cultural, political, social, etc.) of the source materials on which the system has been trained, as Generative AI systems cannot compensate for preexisting prejudices, stereotypes, or underrepresented data sets. For example, a Generative AI might pair the term "municipal clerk" with a female pronoun, and the term "mayor" with a male pronoun. The algorithms and statistical models used by Generative AI to parse and process content derived from source materials can also be a source of bias, resulting in decisions that could be discriminatory. Indeed, the use of AI in employment matters, including but not limited to hiring, pay determinations, and performance monitoring, may run the risk of violating Title VII of the federal Civil Rights Act of 1964 or other employment laws.[ii]

While some Generative AI systems implement guard rails to warn against or block offensive, harmful, or unsafe content, others may not. Due to the rapid advancement of Generative AI, even systems with guard rails might allow unwanted content to slip through.

MINIMIZING RISK

With knowledge of the risks associated with Generative AI, and that risks may evolve, local governments should develop and implement policies and training programs to minimize such risks while preserving the benefits of Generative AI to local government services and functions. Policies should contain (1) requirements for responsible user prompts, (2) mechanisms for reviewing AI-generated content for accuracy and bias and ensuring no copyrighted material is used without proper attribution or without obtaining proper rights, (3) provisions for the security of login information and data, (4) procedures for monitoring Generative AI use and policy enforcement, (5) supervisory expectations for employee use of Generative AI, and (6) provisions concerning the policy's interaction with other local government employment policies. By establishing clear expectations for employee use of Generative AI, local governments can foster a culture of accountability and transparency.

A. User Prompts

Because information used in prompts and AI-generated responses are generally not private, local government policies should prohibit employees from submitting sensitive, confidential, or regulated data, or any personal information or personal identifying information (PII) to a Generative AI system. Policies should also warn that use of Generative AI may result in the creation of a public record under the Colorado Open Records Act (CORA). Accordingly, policymakers should consider adopting policies respecting prompts that may result in AI-generated information that the local government desires to keep confidential.

B. Accuracy & Bias

Policies should require users of Generative AI to carefully review AI-generated content for accuracy and bias before using the content in any work product. To that end, when crafting your entity's policies, consider requiring employees to use Generative AI systems that provide links to source materials so that employees may readily evaluate AI-generated information. When available, employees may verify accuracy and identify bias by following links within the Generative AI system to source materials. When source materials are not cited, it is suggested your policies require employees to verify the accuracy and identify bias by independent research.

C. Copyrighted Material

Policies should require careful review of AI-generated content and available source materials to ensure no copyrighted material is published or distributed to citizens or third parties without proper attribution or without obtaining proper rights. As with reviews for accuracy and bias, policymakers should consider requiring employees to use Generative AI systems that provide links to source materials so that employees may readily identify copyrighted material. As to systems that do not provide links, policies should require employees to track down source materials to evaluate AI-generated information for plagiarism concerns and potential copyright infringement. Local governments may wish to take advantage of one of the many internet or software programs that reviews texts for plagiarism to further discourage publication and distribution of work product containing copyrighted material. Some "plagiarism checkers" are available for free. However, be wary of using plagiarism checkers powered by AI, as using these tools carries many of the same risks as the Generative AI systems discussed in this article.

D. Security

Policies should require compliance with all applicable federal and state data protection laws, including those pertaining to data security and acceptable computing. Policies pertaining to data privacy and security should be reviewed by your municipal attorney or an attorney specializing in cyber security law, as associated laws and regulations are complex and filled with legal and industry jargon.

Moreover, Generative AI users should be made responsible for safeguarding their own system login information and should require adherence to specific security requirements such as the use of strong passwords, frequent password changes, multi-factor authentication (if available), prohibiting credential sharing, and reporting of unauthorized access events and security incidents. Policymakers should carefully consider which email addresses and servers employees should be required to use to create Generative AI accounts, as hackers gaining access to Generative AI accounts created using work credentials may be able to use the same credentials to access to local government systems.

E. Supervisory Expectations

Policymakers should also consider whether employees must obtain permission from, or at least inform, supervisors when employees use Generative AI in connection with the formation of opinions, conclusions, or other information integrated into employee work product. Policies might also benefit from listing prohibited and acceptable employee uses of Generative AI. Moreover, policies should ensure citation to the sources underlying the AI-generated information used in work product rather than to the Generative AI system used to generate it, thus enabling supervisors to independently verify the information used in the work product.

F. Interaction with Other Employment Policies

Policymakers should also consider the interaction between policies related to the use of Generative AI and other existing employment policies, and potential conflicts that may arise due to overlap. If a specific use of Generative AI could implicate other employment or information technology policies of the local government, then the Generative AI policy should contain a statement as to which policy governs in the event of a conflict.

G. Monitoring & Enforcement

Policymakers should also consider reserving the right to monitor and audit employee use of Generative AI to ensure compliance with its Generative AI policy and to protect the entity's data, reputation, and legal and government interests. Moreover, policymakers should consider consequences for violations of the local government's Generative AI policy. Some policymakers may prefer to confine consequences to the loss of the privilege of using Generative AI for work-related purposes, or computer privileges in general. Others may wish to subject violators to disciplinary action authorized under the local government's general personnel policies, even up to termination.

CONCLUSION

The integration of Generative AI into local government operations presents both significant opportunities and considerable challenges. As discussed in this article, the potential for enhanced public services, improved efficiency, and increased citizen engagement is tempered by critical risks, including risks related to data privacy, the propagation of biases, and copyright infringement.

While this article outlines known risks associated with Generative AI use and offers tips for mitigating them, Generative AI is a new and rapidly evolving field, and thus the potential policy impacts and risks to local government organizations are not fully known. Accordingly, local governments must keep a pulse on the changing landscape of opportunities and risks associated with Generative AI.

Ultimately, the successful adoption of Generative AI will depend on a balanced approach that prioritizes innovation without compromising faith and trust in government systems and information, or public welfare. By striking the right balance, CIRSA members can position themselves at the forefront of innovation while safeguarding the interests of their organizations and the communities they serve.

If you have questions about this article or would like samples of AI use policies for employees, contact CIRSA's Associate General Counsel Nick Cotton-Baez at **nickc@cirsa.org**.

Note: This article is intended for general information purposes only and is not intended or to be construed as legal advice on any specific issue. Readers should consult with their entity's own counsel for legal advice on specific issues.

---

i. For example, Colorado local governments are subject to the data privacy requirements of House Bill 18-1128, codified at C.R.S. Section 24-73-101, et seq. Among other provisions, these statutes require a government entity that maintains, owns, or licenses personal identifying information (PII) to implement and maintain reasonable security procedures for the protection of PII. You can read this **CIRSA article** to learn more about these data privacy requirements.
ii. The EEOC has published guidance on how the use of AI in employee assessment and selection can potentially violate Title VII and the ADA. The two sets of guidance can be viewed **here** and **here**. In addition, effective February 2026, Colorado employers will face new requirements and potential liabilities related to the use of AI in employment decisions under Senate Bill 24-205. The application of this legislation to public entities is not clear and a discussion of this Bill is beyond the scope this article.

Publication Date: 12/31/2024