ASPEN POLICY ACADEMY
aspen institute

Operational Document

Draft Use Guidance: Generative Artificial Intelligence

By Hande Güven and Jessica Wang

November 2025

This document presents an example version of the City of New York's Use Guidance for Generative Artificial Intelligence, drawing on the City's current Preliminary Use Guidance. This sample guidance was developed by Aspen Policy Academy fellows while participating in the Science and Technology Policy Fellowship. The fellows were asked to recommend updates to the New York City Office of Technology and Innovation's Artificial Intelligence Action Plan. The full project, including a policy brief explaining the fellows' core recommendations, is available here. Please note that all authors' opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

1. Introduction

Generative artificial intelligence (GenAI) technologies are rapidly evolving and becoming increasingly embedded in everyday tools and workflows. They present opportunities for city agencies to improve efficiency, enhance communication, and support public service delivery — for example, by drafting content, summarizing information, or streamlining repetitive tasks.

At the same time, GenAI technologies also raise critical concerns. These concerns relate to accuracy and reliability, data privacy and cybersecurity, algorithmic bias and discrimination, and the potential erosion of public trust if GenAI is used in ways that are opaque, unaccountable, or inappropriate for government contexts.

The City of New York is committed to a responsible, transparent, and equitable approach to emerging technologies. This guidance provides a framework for the safe and effective use of GenAI by city staff, aligning with the city's broader principles for digital governance and public sector innovation. It is intended to help agencies evaluate GenAI tools, understand appropriate use cases, mitigate associated risks, and maintain trust in the city's services and institutions.

2. Purpose

The Office of Technology and Innovation (OTI) is issuing this guidance to support agencies in their exploration of GenAI technologies by providing a framework for ensuring that GenAI use is responsible, transparent, and equitable. In this document, OTI provides both Use Guidance (Section 5) to help agencies understand the appropriate considerations and guardrails for GenAI use and Common Use Cases (Appendix A) to provide clarity on the application of the Use Guidance in commonly expected settings.

3. Authority

OTI is committed to supporting agencies in their exploration of GenAI. The following guidance lays out key considerations for agencies and their personnel with respect to the use of GenAI tools and outlines critical policy requirements for agencies. The guidance is a living document that will be updated on an ongoing basis as GenAI technologies and considerations for use evolve.

4. Roles and Responsibilities

As with all technology products and tools, individuals should access GenAI tools only when such access has been approved by responsible agency personnel and authorized by agency-specific and citywide requirements. The relevant personnel and requirements may vary across GenAI tools but will frequently include the following responsible agency personnel:

- Agency Chief Contracting Officer (ACCO)
- Agency General Counsel
- Agency Chief Information Officer (CIO) or Agency Chief Technology Officer (CTO)
- Agency Chief Information Security Officer (CISO)
- Agency Privacy Officer (APO)
- Agency Algorithmic Tools Liaison
- Business or operational owners

5. Use Guidance

The following framework builds on New York City's AI Principles. It is designed to help city staff evaluate the responsible use of GenAI by providing guiding questions, risk considerations, mitigation recommendations, and example use cases. The Use Guidance for GenAI is structured to move from high-level strategic content at the

beginning to more detailed use cases in Appendix A.

Sections 5.1 through 5.5 contain text that could be pulled into 1-pagers for each of the city's AI Principles, outlining the key risks associated with GenAI use. Each section includes high-level guiding questions to help evaluate risks, explains why these risks matter, offers mitigation strategies, and provides practical examples of appropriate, risky, and prohibited use cases. These questions aim to engage readers in a dialogue about potential risks and pitfalls related to GenAI use in specific contexts while allowing them to formulate answers best suited to their own use cases.

The following table provides an overview of the guiding questions to consider as you evaluate GenAI use. They are designed to help identify and assess associated risks for each principle.

Overview of Guiding Questions

AI Principle 1: Cybersecurity	
Account and Access Security	Are you using a city-managed or enterprise account with appropriate permissions and security controls that follows relevant city or agency policies?
Tool Authorization and Licensing	Are you using a tool that has been reviewed, licensed, and approved for government use — rather than a consumer-facing or unvetted version?
Compliance	Are you compliant with relevant cybersecurity policies, both citywide and agency-specific?
AI Principle 2: Information Privacy	
Data Sensitivity	Does the input contain personal, regulated, or confidential data (e.g., personally identifiable information, health records, financial data)? Are these data appropriate to input into a third-party tool, even with safeguards in place?

Overview of Guiding Questions, Continued

AI Principle 2: Information Privacy, Continued	
Training Data Risk	Could the tool retain or learn from your input data, and have you confirmed whether opt-out or enterprise protections are enabled?
Compliance	Are you compliant with relevant information privacy policies, both citywide and agency-specific?

AI Principle 3: Validity and Reliability	
Audience Scope	Who is the audience for this output — internal or external?
Impact on Decision-making	Will this output influence policies, resource allocation, or services — or is it for general reference or internal brainstorming only?
Error Likelihood	How reliable is the tool, especially when first introduced or used in a new domain?
Error Discovery	How easily can the accuracy of the output be checked, especially by some without subject-matter expertise?

AI Principle 4: Trust and Transparency	
Output Attribution	Have you clearly marked or disclosed that content was generated or assisted by GenAI when shared externally or in public documents?
GenAI Tool Disclosure	Have you disclosed the use of this GenAI tool in line with any reporting requirements?

Overview of Guiding Questions, Continued

AI Principle 4: Trust and Transparency, Continued	
Human Oversight	Is a human reviewing, editing, or approving the GenAI output before it is finalized or shared?
Ownership and Copyright	Is the human contribution sufficient to satisfy legal or organizational requirements for ownership or reuse?

AI Principle 5: Social Responsibility	
Bias and Equity	Could this tool's outputs reinforce stereotypes, exclude specific populations, or create inequitable impacts in access, language, or outcomes?
Environmental Impact	Is this GenAI use proportionate to the task's importance, considering the energy and infrastructure demands of large models?

5.1 Cybersecurity

Is this GenAI system designed and deployed in a way that protects against cybersecurity threats, including those unique to AI, and does it align with citywide security standards?

Guiding questions to ask

- **Account Access and Security:** Are you using a city-managed or enterprise account with appropriate permissions and security controls that follows relevant city or agency policies?
- **Tool Authorization and Licensing:** Are you using a tool that has been reviewed, licensed, and approved for government use — rather than a consumer-facing or unvetted version?
- **Compliance:** Are you compliant with relevant cybersecurity policies, both citywide and agency-specific?

Why this matters

- GenAI tools not designed for enterprise use often lack essential protections such as encryption, data governance, and administrative controls.
- Use of personal accounts is prohibited by citywide cybersecurity policies.
- Improper account use increases the risk of unauthorized access, data leaks, and other cyberthreats.

What you must do

- Only use preapproved city-managed accounts for all GenAI-related work.
- Comply with citywide cybersecurity policies.
- Confirm that all accounts comply with agency security policies, including multifactor authentication and password complexity requirements.

Example scenarios

- **Appropriate:** Use Copilot Enterprise with your work email on a city-managed browser.
- **Prohibited:** Use a free AI image generator logged in through a personal Gmail account to create visuals for a city project that has not been released publicly.

5.2 Information Privacy

Does the GenAI tool collect, use, or share identifying information, and if so, does it comply with all relevant privacy laws and policies while protecting individual data rights?

Guiding questions to ask

- **Data Sensitivity:** Does the input contain personal, regulated, or confidential data (e.g., personally identifiable information, health records, financial data)? Are these data appropriate to input into a third-party tool, even with safeguards in place?
- **Training Data Risk:** Could the tool retain or learn from your input data, and have you confirmed whether opt-out or enterprise protections are enabled?
- **Compliance:** Have you complied with relevant information privacy policies, both citywide and agency-specific?

Why this matters

- Inputting identifying information or other sensitive or restricted information (e.g., resident data, health records, financial info) into GenAI tools may lead to unauthorized disclosure and retention, re-identification, or violations of city information privacy laws.
    - For instance, some GenAI tools retain input data and use it to train their models, putting confidential or regulated data at risk.
- Some GenAI tools offer personalization in the form of "memory," which allows the tool to save user preferences in order to shape future answers to the user's prompts. This memory may result in the user receiving inaccurate and/or biased answers despite positing well-formed and neutral prompts.
    - For example, a chatbot might remember that a user works for the Department of Education and begin tailoring responses based on the assumption that all queries relate to education policy, which could lead to inaccurate or irrelevant answers if the user later asks about unrelated topics such as citywide procurement or public health initiatives.
- Even nonsensitive data can pose a risk when it is combined with other information or used in unintended contexts.

What you should do

- Follow your agency's internal policies, Citywide Privacy Protection Policies and

Protocols, and the NYC Identifying Information Law when handling any data that could be identifying or sensitive.

- Do not enter sensitive or confidential data into GenAI tools unless the tool has been explicitly approved for that purpose.
- Before using a GenAI tool, check whether input data may be stored or used for training — and disable that feature if possible.
- Use enterprise-licensed tools that support data control and training opt-outs.
- When in doubt, redact or anonymize input data before inputting them into a GenAI tool.

Example scenarios

- **Appropriate:** Use an enterprise GenAI tool with training disabled to summarize a publicly available report.
- **Prohibited:** Paste a resident's medical history into a free-tier chatbot to help write a case summary.
- **Prohibited:** Input an embargoed draft policy into a public GenAI tool that retains data for future training.

5.3 Validity and Reliability

Does this GenAI tool effectively address the specific problem it is intended to solve, and can it perform reliably over time and across different contexts?

Guiding questions to ask

- **Audience Scope:** Who is the audience for this output — internal or external (personnel within your agency, other agencies, or the public)?
- **Impact on Decision-making:** Will this output influence policies, resource allocation, or services — or is it for general reference or internal brainstorming only?
- **Error Likelihood:** How reliable is the tool, especially when first introduced or used in a new domain?
- **Error Discovery:** How easily can the accuracy of the output be checked, especially by some without subject-matter expertise?
- **Error Impact and Recovery:** If an error goes unnoticed, how severe are the potential consequences? How difficult would it be to mitigate or remediate the impact?

Why this matters

- GenAI tools are not uniformly reliable across all tasks. Their performance depends on factors such as how they are prompted, the specific use case, the context in which they are applied, and the quality and recency of the training data and underlying model.
- GenAI tools are pr to certain systemic risks related to the validity and reliability of their outputs. The following are some relevant terms that you may hear:
  - **Hallucination:** A model creates a plausible-sounding answer that is factually incorrect.
  - **Context rot:** A model gets confused when a conversation runs long.
  - **Sycophancy bias:** Models tend to output answers that aim to please the user, leading to confirmation bias.
- Outputs can appear polished and convincing even when they are incorrect, outdated, or misleading.
- Errors that go undetected can lead to confusion, reputational harm, inequitable outcomes, or operational inefficiencies.
- Tools may degrade in performance over time, especially when they are used in rapidly evolving subject areas.

What you should do

- Clearly define the purpose of the GenAI tool before use and ensure that its capabilities align with the task at hand.
- Rigorously test GenAI tools in your domain before wider use.
- Devise a strategy for validating outputs and incorporate human-in-the-loop review from experts where relevant.
- Always confirm outputs using trusted sources or expert validation.
- Understand the limitations of the tool you are using. If you know the limitations of a particular tool, help others understand those limitations and document them for other users.
- As a user, understand the limitations of the tool you are using, including the potential for inaccuracy or other errors. When possible, help others (including coworkers, staff, or the audience) understand the limitations associated with GenAI tools by documenting them in your work or through internal conversations.
- Monitor the performance of GenAI tools over time and reevaluate their reliability as use cases evolve and as models are updated.

Example scenarios

- **Appropriate:** Use a GenAI tool to draft a first-pass summary of public meeting notes for internal review, followed by staff revision and validation before release.
- **Risky:** Use GenAI to draft talking points for a press conference on a policy announcement without subject-matter review.
- **Risky:** Use GenAI to filter and recommend a list of candidates who will progress to the interview stage for a job position.

5.4 Trust and Transparency

Would the public understand how this GenAI tool is being used, and have you taken steps to ensure transparency, explainability, and the option for human oversight or appeal where appropriate?

Guiding questions to ask

- **Output Attribution:** Have you clearly marked or disclosed that content was generated or assisted by GenAI when shared externally or in public documents?

- **GenAI Tool Disclosure:** Have you disclosed use of this GenAI tool in line with any reporting requirements, such as Local Law 35?
- **Human Oversight:** Is a human reviewing, editing, or approving the GenAI output before it is finalized or shared?
- **Ownership and Copyright:** Is the human contribution sufficient to satisfy legal or organizational requirements for ownership or reuse?

Why this matters

- Transparency is essential to building public trust in government services—especially when new technologies like GenAI are involved.
- Without clear communication, the public may not understand how GenAI-created content is produced, where human judgment was applied, or how to question GenAI-assisted outputs.
- Failure to disclose GenAI use can lead to public confusion, diminished credibility, and legal risks — especially if outputs appear authoritative but contain errors or were produced without adequate oversight.
- Lack of attribution or unclear authorship can raise questions about intellectual property, accountability, or the legitimacy of content.

What should you do

- Disclose when content is generated or assisted by GenAI.
- Follow agency protocols for reporting the use of GenAI tools, including pilot or experimental use cases.
- Ensure that human review and accountability are in place, especially for content that affects public rights, services, or perceptions.
- Understand how copyright laws and city policies apply to GenAI-assisted content, and ensure adequate human contribution where required.

Example scenarios

- **Appropriate:** Note in a public newsletter that GenAI was used to generate initial drafts, with final edits by a city staff member.
- **Appropriate:** Submit your use of a GenAI summarization tool to your agency's AI inventory or other tracking process.

- **Risky:** Publish GenAI-created FAQs on a public website without staff review, attribution, or disclosure.
- **Risky:** Release a policy explainer written by GenAI without clarifying authorship or verifying its legal accuracy.

5.5 Social Responsibility

Could the use of this GenAI system result in unfair or inequitable outcomes, and have you taken steps to identify and mitigate bias and harm — especially to vulnerable communities?

Guiding questions to ask

- **Bias and Equity:** Could this tool's outputs reinforce stereotypes, exclude specific populations, or create inequitable impacts in access, language, or outcomes?
- **Environmental Impact:** Is this GenAI use proportionate to the task's importance, considering the energy and infrastructure demands of large models?

Why this matters

- GenAI tools are trained on vast datasets that may reflect historical biases, social inequalities, or cultural imbalances, introducing the risk that outputs will replicate or amplify those harms at scale.
- Without proper safeguards, GenAI can produce exclusionary language, reinforce harmful narratives, or disadvantage communities that have been historically underserved or misrepresented.
- GenAI systems — especially large-scale models — can be resource-intensive, consuming significant energy and computing power. While these impacts may be less visible, they contribute to environmental burdens that public agencies should consider, particularly when the task could be accomplished through simpler means.
- A socially responsible approach to GenAI requires active efforts to test for bias, engage diverse perspectives, minimize unnecessary energy use, and prioritize equity and inclusion throughout the tool's design and deployment.

What you should do

- **Bias and Equity**
  - Consider how GenAI outputs may affect different populations, especially those who face language, accessibility, or cultural barriers.
  - Avoid using GenAI in ways that could replicate or reinforce harmful stereotypes, especially in public-facing content or decision-support contexts.
  - When feasible, consult impacted communities or subject-matter experts to identify risks and mitigation strategies.
- **Environmental Impact**
  - Avoid invoking a model for tasks that could be completed with traditional tools (e.g., search engines, templates, spreadsheets). Only use GenAI for tasks that benefit from scale, creativity, or automation.
  - Write prompts to be as specific as possible and include all possible context to reduce unnecessary re-generation or editing cycles. This includes specifying the tone, length, and format of the desired response in your prompt. For example, “always summarize in bullet points,” “CSV document with 3 columns that say ...,” “mention key words and sources if any.”
  - Consider using smaller models that are faster and less resource-intensive when available for lower-risk or internal tasks.
    - Small language models (SLMs) are lighter-weight GenAI tools that use less compute and resources than large language models (LLMs).
    - SLMs are designed to perform well for simpler tasks and can be fine-tuned for particular domains.
    - A number of open-source SLMs exist that can provide a cost-free or low-cost alternative with a smaller footprint for organizations looking to customize their GenAI use.
  - Store and reuse high-quality GenAI outputs rather than re-generating them every time.
  - Select vendors with green cloud commitments. Prioritize vendors with transparent energy use, data center efficiency metrics, and carbon offsetting or renewable energy commitments.
  - Incorporate sustainability into procurement criteria. Evaluate tools not only for price and performance, but also for energy impact and emissions transparency.

Example scenarios

- **Appropriate:** Run a small pilot of a GenAI-assisted intake tool for city services and gather feedback from diverse residents before scaling.
- **Risky:** Use GenAI to auto-generate service descriptions without testing for inclusive language or accessibility across reading levels.
- **Risky:** Use a large model to look up basic information that could easily be found through less resource-intensive tools.

6. Definitions

Generative artificial intelligence (GenAI) — Any AI system whose primary function is to generate content, which can take the form of code, text, images, and more.

Generative AI tools — The integration of generative AI models into a variety of software or browser applications, including word processors, email, calendars, and chatbots, which may be run locally or by an application programming interface (API).

AI-generated content — Any content produced by GenAI.

Appendix A: Common Use Cases

A.1 Document summarization

Document summarization refers to the use of GenAI tools to create short summaries or overviews of longer materials such as reports, meeting transcripts, policy documents, or emails. This use case can improve productivity, streamline communication, and help staff quickly extract key information from dense or lengthy content.

- Many city agencies work with long reports, stakeholder memos, or public meeting transcripts.
- Summarization can reduce reading time and surface relevant insights.
- This task is often low risk when used internally and with public or nonsensitive documents.

Key risks to consider

- **Misrepresentation or omission of key content:** Summaries can appear fluent and well structured while omitting important nuances or misrepresenting the source. This is especially risky if outputs are used to inform decisions or public messaging.
- **Data sensitivity:** Summarizing documents that contain nonpublic city data, draft policy, or nonpublic data using a GenAI tool that retains or learns from input can pose privacy or security risks.
- **Output attribution:** Failing to disclose GenAI assistance in creating summaries can erode public trust — especially if outputs are distributed externally or used in public-facing documentation.
- **Bias and equity:** Summaries may reflect biases present in the source material or model. For example, they may overlook underrepresented voices in transcripts or deprioritize equity-related themes in policy documents.

Recommended practices

- Use summarization tools with public documents or internally distributed materials that do not contain sensitive or regulated data.
- Always review and validate GenAI-created summaries, especially if they will be shared outside your team or agency.
- When summarizing long transcripts (e.g., public hearings), consider chunking the document and summarizing section by section to more easily verify the accuracy of the summary.
- Avoid uploading confidential memos, drafts, or PII-containing documents into GenAI tools—unless the tool is enterprise approved for that data type.

- Disclose the use of GenAI if summaries are shared externally or incorporated into formal documentation.

A.2 Content drafting

Content drafting refers to the use of GenAI tools to generate written materials such as emails, policy memos, web content, newsletters, talking points, FAQs, and internal documentation. The goal is often to speed up writing, generate first drafts, or rephrase existing content in a different tone or format.

- GenAI tools can accelerate the writing process, especially for routine or templated content.
- Drafting is often the first step in a multistage content workflow, where human review and editing follows.
- This use case is widely applicable across roles — from communications to operations to policy teams.

Key risks to consider

- **Inaccurate or misleading content:** GenAI-created text can sound polished but contain factual inaccuracies, outdated information, or incorrect assumptions. This is particularly risky for content related to policy, law, or public services.
- **Overreliance on GenAI without human review:** Publishing or distributing GenAI-drafted content without thorough human editing may lead to reputational harm, misinformation, or missed errors, especially when shared with the public.
- **Use with sensitive or nonpublic inputs:** Drafting based on documents that contain PII, internal deliberations, or embargoed material in a non-enterprise GenAI tool may violate privacy policies or inadvertently expose confidential information.
- **Undisclosed GenAI assistance in public communications:** If GenAI-created text is presented as fully human authored without disclosure, it may undermine public trust, especially in contexts involving policy positions, resident communications, or leadership messaging.
- **Cultural or equity-related tone risks:** Without review, GenAI may use language that lacks cultural sensitivity, misrepresents community perspectives, or excludes certain groups in tone or phrasing.
- **Cultural or equity-related tone risks:** Without review, GenAI may use language that lacks cultural sensitivity, misrepresents community perspectives, or excludes certain groups in tone or phrasing.

Recommended practices

- Use GenAI to create first drafts, then refine through human review and subject-matter expertise.
- Avoid using GenAI to generate final versions of public-facing or legally significant content without validation.
- Use enterprise-approved tools for drafting, especially when inputs involve internal memos, drafts, or sensitive context.
- Edit for tone, clarity, cultural awareness, and legal accuracy — especially when addressing diverse audiences or publishing externally.
- Disclose GenAI involvement in content creation if relevant to the audience or public context.

A.3 Meeting notetaking

GenAI tools can be used to assist with meeting notetaking by summarizing transcripts, identifying action items, or producing readable minutes from audio recordings, video calls, or live conversations. Some tools also offer real-time transcription and analysis features.

- City staff often participate in meetings, workshops, and public forums for which accurate notes and summaries are needed.
- GenAI can streamline the documentation process, reduce manual effort, and ensure that key takeaways and decisions are captured consistently.
- These tools are especially useful for meetings that are long or have many participants.

Key risks to consider

- **Incorrect or misleading summaries:** GenAI tools may omit important context, misattribute statements, or misrepresent tone, especially in nuanced or multi-speaker discussions.
- **Failure to disclose AI involvement in meeting documentation:** If GenAI-created notes are presented as fully human authored, others may assume a level of editorial review or context awareness that is not present.
- **Inappropriate use of sensitive or nonpublic meeting content:** Uploading transcripts of confidential meetings into GenAI tools that retain input data may expose sensitive conversations or violate privacy protections.
- **Overreliance on GenAI without human review:** Using GenAI-created notes without verifying accuracy can result in miscommunication or missed follow-up items, especially if notes are shared widely or become part of the official record.

Recommended practices

- Use GenAI to generate draft summaries or action items, but ensure that they are reviewed by a meeting participant before distribution.
- Clearly label GenAI-created notes if shared externally or formally archived.
- Avoid using GenAI tools to process confidential, sensitive, or legally privileged meetings unless the tool is enterprise-approved and privacy-compliant.
- Choose tools with clear data retention policies, and disable training or storage features when possible.
- Consider accessibility and inclusion: Ensure that notes reflect all participants' contributions and that transcripts are usable for those with hearing or language needs.

A.4 Translation

Translation refers to the use of GenAI tools to convert written content from one language to another. It may be used to increase accessibility, support multilingual engagement, or prepare draft materials in multiple languages before professional review.

- New York City agencies frequently serve linguistically diverse communities.
- GenAI offers a fast, low-cost way to generate draft translations of public information, internal documents, or service-related content.
- GenAI can support multilingual outreach, especially in early drafting or prototyping stages.

Key risks to consider

- **Inaccurate or culturally inappropriate translations:** GenAI tools may mistranslate terms, overlook cultural nuances, or produce overly literal translations that are confusing or misleading — especially in technical or community-specific contexts.
- **Lack of human review or community validation:** Unreviewed translations may introduce unintended meaning or tone, which could misinform the public or lead to distrust, especially in public-facing service communications.
- **Exposure of sensitive or internal content:** Translating nonpublic documents (e.g., case files, internal memos) in GenAI tools not approved for sensitive data may risk disclosure or unauthorized model training.

Recommended practices

- Always engage human review from a native speaker — especially for public-facing content.

- Engage bilingual staff, community partners, or professional translators to verify accuracy, tone, and cultural nuance.
- Avoid using GenAI to translate confidential, legal, or sensitive content unless the tool is enterprise approved for such data.
- Clearly disclose GenAI use when translations are shared externally or published online, and label them as drafts when appropriate.
- Be especially cautious when translating time-sensitive or high-impact messages, such as emergency alerts, legal notices, or eligibility requirements.

A.5 Media (image and video) production

This use case involves employing GenAI tools to create visual media – including illustrations, infographics, stock-style images, avatars, or synthetic videos – to support communications, presentations, outreach materials, or prototypes.

- Many city agencies need visuals for digital content, public education, program promotion, or internal presentations.
- GenAI tools make it easy to create custom visuals quickly and at low cost, even without graphic design expertise.
- Visual GenAI is often used for brainstorming, prototyping, or supplementing communication workflows.

Key risks to consider

- **Misinformation or visual inaccuracy:** GenAI-created media may appear realistic while containing factual inaccuracies, anachronisms, or misleading imagery (e.g., nonexistent landmarks, distorted maps, fictitious people).
- **Undisclosed use of GenAI-created imagery:** Sharing GenAI-created images without disclosure may erode public trust — especially in contexts where authenticity matters (e.g., documentation, historical references, public records).
- **Cultural insensitivity or exclusion:** Some GenAI tools may underrepresent or stereotype certain identities, communities, or body types — especially when used without specific guidance on inclusivity.
- **Use of copyrighted material in training data:** GenAI tools trained on internet images may replicate copyrighted styles or content. This can create legal or reputational risks if used in official government materials without proper review.
- **Sensitive or confidential prompts:** Prompts referencing confidential projects, internal policies, or identifiable individuals should not be entered into GenAI tools that retain or learn from inputs.

- **Impersonation of public officials:** GenAI-created media can impersonate public officials or their likeness.

Recommended practices

- Use GenAI tools for prototyping, brainstorming, or supplemental visuals only.
- Disclose that images or videos were created with GenAI if they are used in public-facing materials or communications.
- Review media for accuracy, representational fairness, and inclusivity, especially if depicting people, neighborhoods, or services.
- Do not generate images in the likeness of public officials, figures, or other real people.
- Avoid uploading sensitive prompts that could reveal private, draft, or internal information.
- Consult legal or procurement guidance before publishing GenAI-created content that may raise copyright or intellectual property concerns.

A.6 Code generation

GenAI tools can assist with creating code snippets, scripts, configuration files, or entire software components based on written prompts. These tools are often used for automating tasks, building prototypes, learning new frameworks, or debugging.

- City agencies may have internal staff who write or maintain code to support digital services, data analysis, or automation.
- GenAI can help technical staff work more efficiently, especially when solving common problems or learning new tools.

Key risks to consider

- **Security vulnerabilities or flawed logic:** GenAI-created code may include insecure practices (e.g., hardcoded secrets, lack of input validation) or subtle logic errors that are hard to detect but can lead to real-world vulnerabilities.
- **Use of unlicensed or copyrighted code:** Some GenAI tools may generate code derived from open-source repositories without respecting license terms, raising potential legal or compliance issues for public sector use.
- **Deployment without proper review or testing:** Relying on GenAI-created code without adequate review, testing, or documentation may result in unstable systems or long-term maintenance issues — especially if ownership or authorship is unclear.
- **Use with sensitive data or credentials:** Prompting GenAI with environment variables, API keys, or internal architecture details can expose sensitive information

if the tool stores input or is not enterprise-approved.

Recommended practices

- Always review, test, and validate GenAI-created code before using it in operational systems.
- Follow agency security standards – including code review and vulnerability scanning – before deploying GenAI-assisted code.
- Avoid pasting secrets, credentials, or internal architecture details into GenAI tools unless using secure, enterprise-approved environments.
- Review output for licensing conflicts and ensure it complies with city policy and open-source use guidelines.
- Document where and how GenAI was used in the development process, especially for transparency and future maintenance.

If you'd like to learn more, see the full project, including a policy brief explaining the fellows' core recommendations, at aspenpolicyacademy.org/project/nyc-genai-guidance-2025.

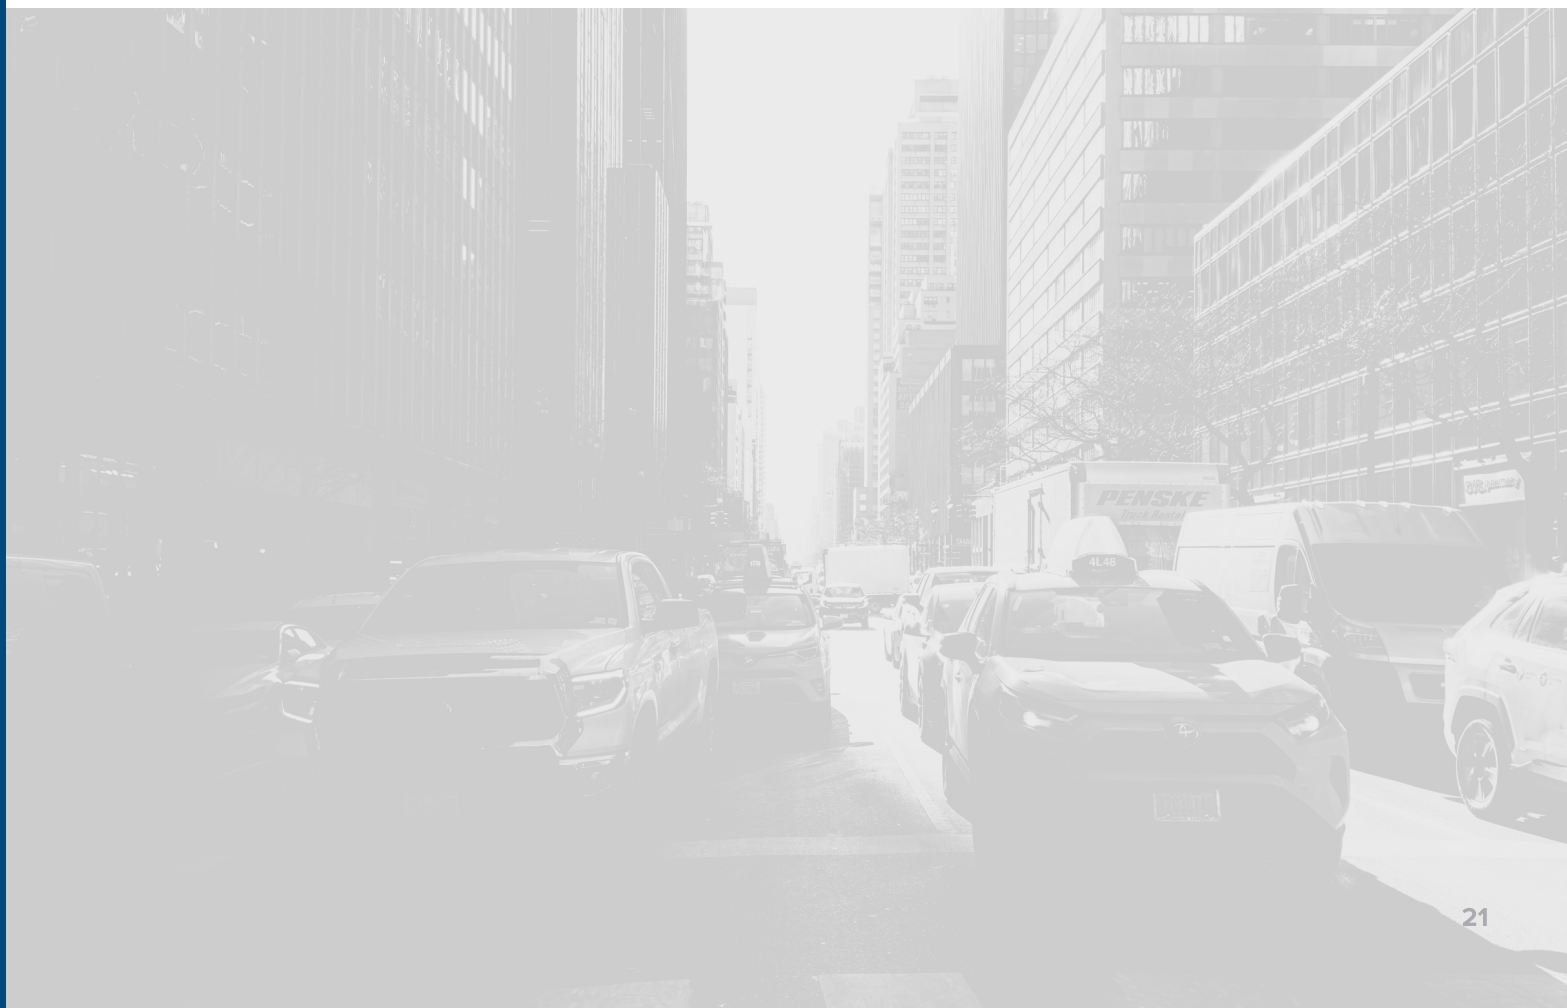

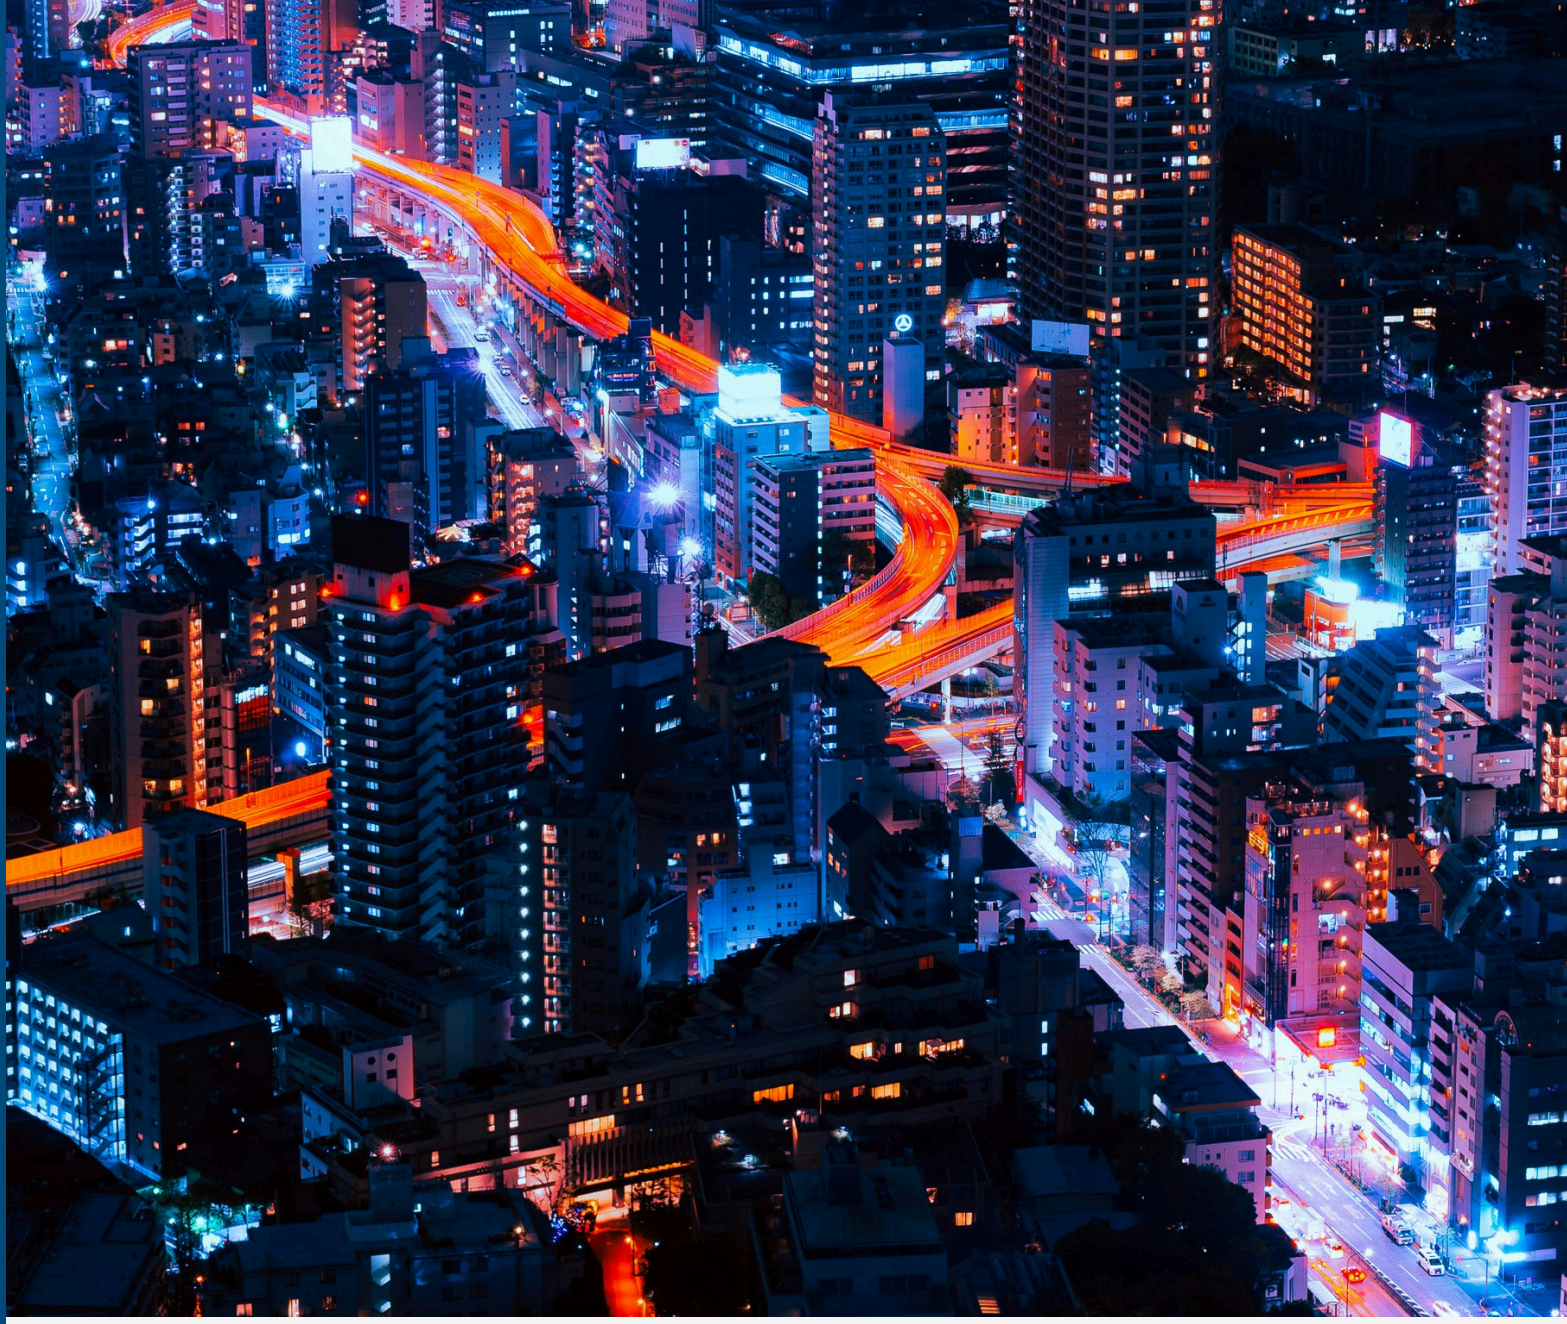

About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors – from tech to the environment, science to civic engagement – with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.