United States
Consumer Product Safety Commission

Generative AI Use Policy

July 8, 2025

Hengyi Hu
Version 2.0

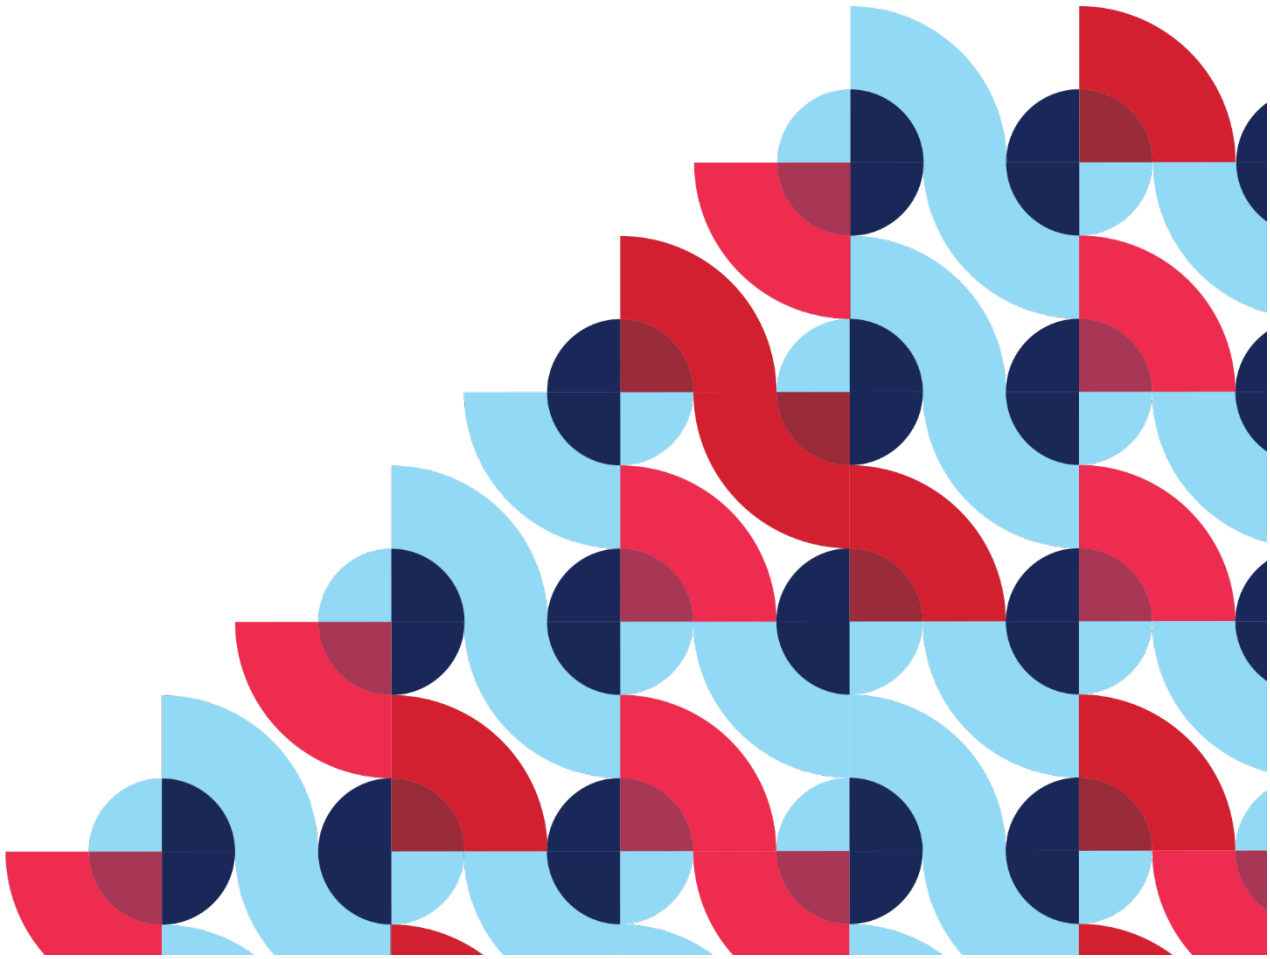

Table of Contents

Review and Updates

Version	Date	Status	Notes
0.1	4/28/25	Complete	EXRR AED review
1.0	5/13/25	Complete	Revised draft for EXIT review
1.1	6/5/25	Complete	Revised draft for Agency Review
2.0	7/8/25	Complete	Final draft after Agency Review

1. Purpose

This Generative Artificial Intelligence (GenAI) Usage Policy will establish a comprehensive framework for the ethical, secure, and effective use of GenAI tools and services at the U.S. Consumer Product Safety Commission (CPSC). This policy is designed to guide CPSC personnel in responsibly integrating GenAI systems into agency workflows, while safeguarding the public trust, individual rights, and the agency's mission to protect consumers from hazardous products.

This policy addresses challenges presented by GenAI, including the potential for output inaccuracy, hallucinations, data leakage, and social or algorithmic bias. It provides baseline requirements and best practices for ensuring that outputs generated by Artificial Intelligence (AI) align with CPSC's scientific standards, legal mandates, and ethical values.

Furthermore, the policy seeks to:

- Promote efficiency and innovation through the strategic use of GenAI tools
- Prevent misuse of GenAI that may cause harm to individuals, erode public confidence, or disrupt CPSC operations
- Encourage transparency, traceability, and human oversight in the development and use of AI systems
- Mitigate risks related to bias, misinformation, and lack of accountability in generative content
- Ensure that GenAI applications reflect the Commission's core values of scientific integrity, fairness, accuracy, and consumer safety

By establishing this policy, CPSC aims to both comply with federal directives and model excellence in the responsible and mission-aligned use of GenAI technologies.

2. Scope

This policy applies to all individuals affiliated with CPSC who access or use GenAI systems in the course of their official duties. This includes, but is not limited to:

- Federal employees (permanent, temporary, or term-limited)
- On-site and remote contractors or consultants
- Interns, fellows, and agency volunteers
- Other individuals granted access to CPSC systems or networks under formal agreements

The policy governs GenAI usage across all work environments, including CPSC offices, telework settings, and any remote or hybrid arrangements. This policy encompasses both:

- In-tenant GenAI tools: These are AI-enabled services integrated within CPSC's enterprise systems and managed under its IT security protocols. Examples include Microsoft 365 Copilot, Copilot Chat, and other AI services within CPSC's secure cloud tenant. These tools operate under enterprise governance controls, including identity management and data retention policies.
- Publicly Available GenAI tools: These are commercial, web-based AI platforms accessible over the internet that operate on shared infrastructure not controlled by the CPSC. Examples include ChatGPT (OpenAI), Microsoft Copilot (personal), Google Gemini, Anthropic Claude, xAI, Perplexity, and other services. These tools reside outside of government security baselines and present higher risks related to data handling, privacy, and reliability.

This policy will establish guidelines to ensure that 1) all GenAI use complies with federal regulations and CPSC standards, 2) promote responsible experimentation with these tools while maintaining trust in CPSC's work, and 3) prevent improper sharing or exposure of sensitive or confidential information, whether intentional or accidental. This policy is informed by and in alignment with the following directives and guidance:

- Executive Order 14179, *"Removing Barriers to American Leadership in Artificial Intelligence"* (2025)
- OMB Memorandum M-25-21, *"Accelerating Federal Use of AI through Innovation, Governance, and Public Trust"* (2025)
- OMB Memorandum M-25-22, *"Driving Efficient Acquisition of Artificial Intelligence in Government"* (2025)
- CPSC Privacy and Security Notice[1], including Personally Identifiable Information (PII)
- CPSC Directive 1801: Epidemiologic Data Sharing Under the CPSA, including Section 6(b) Information Disclosure[2]
- Consumer Product Safety Act (CPSA), 15 U.S.C. § 2051, et seq.
- Consumer Product Safety Improvement Act of 2008, Pub. Law No. 110-314
- CPSC Information Quality Guidelines

[1] https://www.cpsc.gov/About-CPSC/Policies-Statements-and-Directives/Privacy-Policy
[2] https://www.cpsc.gov/Business--Manufacturing/section-6b-information-disclosure

3. GenAI Use Policy

3.1. Appropriate Use of GenAI

The appropriate use of GenAI at CPSC must be in accordance with the directives and guidance listed in the Scope section, as well as CPSC's internal policies on securing safety data, transparency, and responsible technology adoption. This section outlines behavioral and procedural expectations for all users of GenAI systems at CPSC, with clear distinctions between in-tenant tools (those integrated into CPSC's Azure enterprise environment) and publicly available tools, i.e., those operating outside CPSC's tenant and control.

3.1.1. In-Tenant GenAI Tools (Microsoft 365 Copilot, Copilot Chat)

In-tenant GenAI tools are those that are deployed within CPSC's secure IT environment and managed under the agency's cybersecurity, data governance, and identity access policies. These tools comply with federal mandates, including data protection requirements and usage monitoring.

When using in-tenant GenAI tools, users must:

1. Ensure the tool is used only for appropriate, mission-aligned purposes, such as internal productivity, summarization, or drafting support for non-sensitive tasks.
2. Consult with supervisors and stakeholders regularly to ensure AI-generated outputs align with agency priorities and responsibilities.
3. Critically review and edit GenAI outputs before using it in any formal capacity.
4. Comply with CPSC's Information Quality Guidelines when using GenAI to support analytical or communications work that is influential following the definitions in the Information Quality Act.

3.1.2. Publicly Available GenAI Tools (e.g., ChatGPT, Google Gemini, xAI)

Publicly available GenAI tools are third-party platforms that operate outside of CPSC's IT infrastructure. These tools may transmit data across the public internet, store user input indefinitely, or use submitted content to retrain underlying models. Therefore, use of these tools must be treated with greater caution and risk awareness. Users of publicly available GenAI systems must adhere to CPSC's Privacy and Security Notice[3].

When using Publicly Available GenAI Tools, users must:

[3] https://www.cpsc.gov/About-CPSC/Policies-Statements-and-Directives/Privacy-Policy

1. Not share or input any proprietary or confidential CPSC data (e.g., PII, CPSA Section 6(b), pre-decisional information) with external GenAI systems.
2. Not share any other private or proprietary CPSC information such as confidential internal communications, rules, or pre-decisional information with external GenAI systems.
3. Report any confidential data breaches or incidents involving AI systems to their supervisor, the CPSC Chief Information Security Officer.
4. Also follow guidance for In-Tenant GenAI Tools, above.

To reduce the risk of inappropriately disclosing information, CPSC staff should limit GenAI use to in-tenant Microsoft Copilot/Copilot Chat solutions whenever possible. Publicly available GenAI tools should only be used for general knowledge searches, researching publicly available information, and summarizing publicly available websites, papers, websites, etc.

3.2. Identifying and Mitigating Bias in AI-generated Content

In alignment with Executive Order 14179 and OMB Memorandum M-25-21, all CPSC personnel must proactively identify and mitigate potential biases in GenAI outputs, especially when such outputs may inform decision-making, public safety communications, or analyses that affect consumers or regulated entities. Bias in AI can stem from skewed data, model limitations, or contextual misunderstandings. To ensure fairness, accuracy, and alignment with CPSC's values, GenAI users must adhere to the following practices:

1. Establish familiarity with common types of bias in AI systems, including:
    a. Data bias: When training data overrepresents or underrepresents certain populations or characteristics.
    b. Algorithmic bias: When model design leads to unintended, skewed, or inequitable outcomes.
    c. Confirmation bias: When users selectively accept AI outputs that align with pre-existing beliefs or expectations.
2. Stay current with internal training or resources on AI and AI accountability.
3. Review AI-generated content critically, especially in contexts that involve:
    a. Product safety and public health safety communications
    b. Scientific or technical information
    c. Legal, regulatory, or enforcement actions
    d. Public health safety and consumer guidance
4. Seek feedback from a diverse set of colleagues and supervisors to ensure outputs are free from embedded assumptions or unintended implications.
5. Implement safeguards and controls in workflows where GenAI is used to align with risk management guidance from OMB M-25-21, particularly for safety-impacting use cases.

Bias mitigation is a shared responsibility, and the agency's use of AI must reflect CPSC's expectations of integrity, impartiality, and consumer protection. When in doubt, seek input from legal, ethics, or policy officials to ensure AI content is appropriately reviewed and responsibly applied.

3.3. Accuracy of AI-generated Outputs

CPSC is responsible for producing information and decisions that are scientifically sound, legally defensible, and aligned with its mission to protect the public. All outputs generated or supported by AI, including GenAI, must be thoroughly reviewed and validated for accuracy, relevance, and appropriateness before being used in official work product, communications, or decision-making.

3.3.1. Accuracy Verification of GenAI Outputs

All AI-generated outputs, including text, summaries, code, visualizations, and recommendations, must be reviewed by qualified CPSC staff before they are used in any formal or external capacity. Users must not treat GenAI output as authoritative or factual without evaluating its content. Human oversight and critical reviews are crucial when the output relates to product safety communications, scientific or technical information, regulatory and enforcement actions, and consumer guidance.

To verify and ensure the accuracy of AI-generated content, users must:

1. Cross-reference information with official, reliable sources (e.g., peer-reviewed studies, CPSC databases, regulatory documents, government statistics).
2. Fact-check citations and ensure reference materials are correctly attributed.
3. Test and validate outputs, especially for data analysis, predictive modeling, or generative code.
4. Involve subject matter experts (SMEs) when outputs pertain to specialized knowledge domains such as science, engineering, or public policy.

CPSC personnel are ultimately responsible for the content they submit, publish, or act upon — regardless of whether that content was AI-assisted. Users must:

1. Take ownership of AI-assisted work products.
2. Avoid overreliance on AI systems, especially in critical judgments.
3. Identify and exclude any inaccuracies, hallucinations, or misleading outputs.

3.3.2. High-Impact Use Cases

CPSC follows the guidance from OMB M-25-21, which characterizes high-impact use cases as those where "output serves as a principal basis for decisions or actions that have a legal, material, binding, or significant effect on rights or safety." For CPSC, this applies to many potential applications in Safety Operations, including hazard pattern identification and evaluation, automation of data coding/analysis, evaluating the impacts of CPSC actions on safety, targeting products or shipments for inspection and other key activities. Other CPSC high-impact applications include interpreting statutes, regulations, or agency policy documents, and providing public-facing content such as notices to consumers and language translations.

For high-impact use cases, CPSC management will determine if the risks outweigh the operational benefits. Those that are approved will follow a separate guidance titled "Implementing Risk Management Practices for High-Impact AI," which implements the OMB M-25-21 guidance, that will document the procedures, responsibilities, and controls required to identify, assess, mitigate, and monitor risks.

3.4. Transparency

CPSC is committed to ensuring that the public, stakeholders, and internal users understand when, how, and why AI is being used. Transparency in using GenAI is critical when it plays a substantive role in shaping outputs that may influence external communications, consumer understanding, regulatory activities, or safety-related determinations.

In cases where AI substantively contributes to the generation of materials such as summaries, policy drafts, reports, or outreach content, the role of AI must be documented and, when appropriate, disclosed to the intended audience, including through annual publication of AI Use Cases.

Additionally, at the team level, GenAI users will:

1. Develop best practices and Standard Operation Procedures at the team level on how to disclose the use of AI systems. Create and enable processes to check and double-check AI generated content.
2. Keep track of GenAI usage, including purpose, inputs, outputs, and any actions taken based on the AI-generated results.
3. Maintain awareness of the GenAI Use Policy and other guidelines and best practices such as the NIST AI RMF[4].

[4] https://www.nist.gov/itl/ai-risk-management-framework

4. Participate in training sessions and workshops to stay updated on AI-related developments, risks, and mitigation strategies.

3.5. Prohibited Uses of GenAI

CPSC is committed to the safe, ethical, and legally compliant use of AI technologies. This policy explicitly prohibits the misuse of AI in any form and establishes safeguards to prevent unintended, harmful, or unauthorized applications. In accordance with federal directives, the following are explicitly prohibited:

1. Using AI to make decisions that affect individuals' rights or safety without human oversight – e.g., eligibility for programs, legal determinations, product risk assessments, etc.
2. Entering private or sensitive information into external, publicly available GenAI tools. This includes PII, enforcement actions, manufacturer-submitted data, trade secrets, or information labeled "For Official Use Only."
3. Using AI to generate or disseminate misleading or deceptive content.
4. Allowing AI to operate autonomously to impact public health without appropriate safeguards.
5. Bypassing established guidance or governance.
6. Failing to report dangerous or unethical AI behavior.

3.6. Risk Management Practices

A separate guidance titled "Implementing Risk Management Practices for High-Impact AI" will document the procedures, responsibilities, and controls required to identify, assess, mitigate, and monitor risks associated with AI systems that are designated as rights-impacting or safety-impacting.

4. AI Governance and Oversight

4.1. AI Governance Board

To ensure effective oversight and responsible use of GenAI across CPSC, an AI Governance Board (AGB) will be established as a central component of CPSC's AI and Data governance framework. The AGB will be chaired by the Chief AI Officer (CAIO) in evaluating AI use cases, monitoring risk mitigation strategies, and promoting transparency, accountability, and ensuring continued compliance with applicable federal guidance. The AGB will serve as a forum for reviewing AI implementations, advising high-impact or high-risk use cases or systems, and coordinating input from legal, privacy, security, and scientific perspectives. The structure,

membership, roles, and procedures of the AGB will be detailed in a separate AI Governance Board Charter and published on CPSC.gov.

4.2. Stakeholders and Oversight

CPSC will maintain a structured and transparent approach to governing the use of GenAI. Oversight responsibilities are distributed across CPSC's leadership and governance bodies to ensure accountability, alignment with mission objectives, and the responsible use of GenAI technologies. The following offices and entities serve as key stakeholders in the governance of AI systems at CPSC:

1. Executive Leadership (OEX): The Office of the Executive Director (OEX) provides strategic direction for AI use across the agency. OEX is responsible for approving the AI Use Policy and any future amendments. In coordination with governance stakeholders, OEX promotes an agency-wide culture of safe, secure, and trustworthy AI adoption, in alignment with federal AI guidance.
2. Chief AI Officer (CAIO): The CAIO provides leadership in the implementation and oversight of GenAI systems. The CAIO advises on responsible use practices, leads agency coordination on AI risk assessments, and ensures that rights- and safety-impacting use cases are identified, documented, and governed in compliance with federal requirements.
3. The Chief Data and Information Officer (CIO): The CIO provides leadership in the acquisition and implementation of GenAI systems. They are responsible for acquiring Microsoft Copilot licenses and establishing a shared Azure AI space, along with IT support for GenAI and open-source tools.
4. Chief Information Security Officer (CISO): The CISO provides critical oversight of data security and privacy risks associated with AI use. They are responsible for evaluating and responding to incidents involving unauthorized disclosure or improper sharing of data with external GenAI systems.
5. Scientific Integrity Official (SIO): The SIO is the primary Agency-level contact for questions regarding scientific integrity and oversees the process for responding to allegations of violations including cases where GenAI tools may have inappropriately influenced scientific publications, analyses, or research communications. The SIO ensures that AI use upholds CPSC's standards for scientific integrity, objectivity, and reproducibility.
6. Office of the General Counsel : Legal advisors provide guidance on intellectual property, regulatory compliance, privacy law, and acceptable use of AI. They also assess legal risks associated with external vendor systems and AI-generated content.
7. Office of Risk Reduction: Technical staff provide guidance on integration of AI tools into safety mission, including context, utility and risk.

4.3. Training and Awareness

The CAIO and AI Governance Board will deliver periodic training sessions and briefings on:

1. Ethical use of AI and GenAI
2. Risk identification and mitigation
3. Transparency and accountability obligations
4. Real-world use cases and lessons learned across government
5. Other training and presentations, such as effective prompt engineering