




AI at CAAs: Issues, Tips, & Sample AI Use Policy

Introduction

Artificial intelligence (AI) is swiftly changing our world.¹ AI technologies are rapidly evolving systems designed to approximate human cognition, act rationally, and perform tasks without significant human oversight.²

Some Community Action Agencies (CAAs) have taken advantage of AI to make their services more accessible to clients, communicate with the public, and boost internal productivity. Others are more cautious about how they can ethically and responsibly step into the world of AI. Whether your CAA is looking to restrict, promote, or simply set expectations with staff with respect to AI use, an AI policy can help manage the potential risks inherent in AI use and help a CAA maintain compliance.

Disclaimer

This sample AI policy was developed by Community Action Program Legal Services, Inc. (“CAPLAW”) and has not been approved by any outside authority, such as the U.S. Department of Health and Human Services. When using this sample to develop, review, or update an AI policy, CAPLAW strongly recommends you consult with legal counsel familiar with your CAA’s operations and the laws of your state.

The contents of this publication are intended to convey general information only and do not constitute legal advice. This publication does not constitute or create an attorney-client relationship. If you need legal advice, please contact CAPLAW or another attorney directly.

How to Use this Resource

This resource is intended to help CAAs understand the legal framework and issues that may arise in drafting and administering an AI policy. There is no one right way to address AI use at a CAA. Rather, this resource and sample policy offers a jumping off point for CAAs. In the sample policy, *Tailoring Tips*, *To Dos*, and [bracketed text] highlight where a CAA needs to make decisions and take actions to customize the policy so it best meets the organization’s needs.

This resource includes a brief background of the legal framework governing AI, which is important to keep in mind when drafting an AI policy and training staff. Regular training on proper AI use and applicable legal requirements will help make any AI policy truly effective.

¹ U.S. Government Accountability Office; [Report to Congressional Addressees: Artificial Intelligence, Federal Efforts Guided by Requirements and Advisory Groups](#) (September 2025)

² U.S. General Services Administration; [Centers of Excellence: AI Guide for Government](#) (Accessed September 2026)

Legal Framework of AI Use

No comprehensive federal law currently governs AI use in the workplace. Instead, existing legal frameworks that predate widespread AI adoption apply to AI and will inform a CAA's AI policy. As AI rapidly evolves, the applicable legal landscape will continue to change. This list of federal laws provides a starting point for CAAs to assess legal risks, incorporate those considerations into AI policies, and build staff training around responsible AI use.

Industry-specific requirements are the rules that, to date, provide the most comprehensive federal framework implicated by AI use. These rules exist for certain categories of data within an industry. Examples include the Health Insurance Portability and Accountability Act (HIPAA)³ in the healthcare industry, the Fair Credit Reporting Act (FCRA) applicable to financial and consumer reports, the Federal Trade Commission Act (FTC Act)⁴ applicable to CAAs with for-profit activities that generate unrelated business income, and the Family Educational Rights and Privacy Act (FERPA) applicable to CAAs that receive U.S. Department of Education funds and informative for CAAs that receive Head Start funds.

Anti-discrimination law generally protects individuals against discrimination on the basis of race, sex, and other classifications under the Equal Protection Clause of the 14th Amendment to the U.S. Constitution. In addition to this general constitutional protection, two major categories of anti-discrimination laws exist: employment and provision of federally funded benefits and services.

The anti-discrimination law that applies to CAAs with respect to their role as employers is Title VII of the Civil Rights Act of 1964 (Title VII). Title VII prohibits employers from making adverse employment decisions based on race, color, religion, sex, and national origin. AI tools used for recruiting and hiring employees, performance management, promotion recommendations, and skill inference present risks of potential bias and discrimination under Title VII. Where AI uses historical data to define successful employees, this data may reflect and amplify past inequities which may result in biased performance evaluations and unequal access to training and development opportunities. CAA staff inputting information on job titles, career gaps, and project assignments should be careful that the AI tool does not unintentionally encode that information as correlated with protected classes. CAAs using AI in recruiting, hiring, and performance management should regularly check on the quality of the input and output data from AI, be transparent about any AI use in this area, and maintain a level of human oversight over decision-making.⁵

3 While the legal requirements under HIPAA are not specific to AI, CAAs covered by HIPAA must have a clear understanding of what constitutes protected health information (PHI), whether PHI is being disclosed to any third-party AI vendors, and if staff are using AI in connection with their work. HIPAA applies to “covered entities” and “business associates.” CAAs who think they are covered by HIPAA or that certain health-related information they collect is covered by HIPAA should reference CAPLAW's [Privacy Rights of CAA Clients FAQ](#). HIPAA governs how PHI may be used or disclosed regardless of whether the PHI is handled by a human, traditional software, or AI tools. A Business Associate Agreement must be in place if a HIPAA-covered entity discloses any PHI to a third-party AI vendor, similar to the requirements regarding the disclosure of PHI to any third party.

4 The FTC Act prohibits unfair or deceptive acts or practices in or affecting commerce, but the FTC's jurisdiction is limited to entities “organized to carry on business for its own profit or that of its members”. [15 U.S.C. §§ 44-45](#). Most CAAs, as 501(c)(3) charitable nonprofits, likely fall outside this definition and therefore outside FTC Act jurisdiction. However, CAAs that operate revenue-generating programs or maintain for-profit subsidiaries could potentially fall within the FTC Act's reach for those activities.

5 A number of states have adopted laws that regulate employers' use of AI-assisted hiring and employment tools and often require employers to provide notice regarding AI use for employment-related purposes. In Illinois, for example, an employer cannot use AI to discriminate intentionally or unintentionally against any protected class covered by Illinois state law when used to recruit, hire, and promote employees, as well as when used to renew employment, select employees for trainings or apprenticeships, discharge, or discipline staff. The law also requires employers to inform applicants or employees about AI use. See, e.g., Manatt; [AI-Assisted Hiring Faces a New Compliance Landscape in 2026](#) (December 2025).

The anti-discrimination law that applies to CAAs as federal grant recipients is Title VI of the Civil Rights Act of 1964 (Title VI). Title VI prohibits discrimination based on race, color, or national origin in programs or activities that receive federal financial assistance, meaning that CAAs are prohibited from discriminating in the provision of their federally funded services based on such classifications. Where AI is used to evaluate eligibility for CAA services and programs, CAAs should similarly exercise caution that the AI tool does not reflect or amplify inequities or unequal access to services.

Reasonable accommodations may be required in the workplace under the Americans with Disabilities Act (ADA), a federal law that protects individuals with disabilities from discrimination and requires employers to engage in an interactive process when a qualified individual with a disability requests an accommodation.

AI tools might unintentionally violate the ADA by screening out qualified individuals with disabilities in the application process or employees on the job. The Equal Employment Opportunity Commission (EEOC) recommends employers adopt the practice of providing proactive disclosures about how an AI tool evaluates applicants or employees.⁶ For example, if employers inform their employees that they will be assessed partly on the basis of keyboard usage, employees with a vision impairment would know to request an alternative means of measuring productivity as a reasonable accommodation. CAAs who wish to limit the use of AI tools, such as AI-enabled glasses or AI notetakers, should consider that the reasonable accommodation framework may require employers to permit their use.⁷

Data privacy laws at the state and federal levels govern a CAA's responsibility to protect certain personal information, such as a client's name, address, date of birth, phone number, email address, income, Social Security number, government-issued ID number, or bank account number, among other data. Various funding source and program rules may define protected client information differently. Funding source requirements are found in the laws and regulations that govern the programs, and/or referenced in the grant agreements and contracts between awarding agencies and CAAs. The Office of Management and Budget's (OMB) Uniform Administrative Requirements, Cost Principles, and Audit Requirements for Federal Awards (Uniform Guidance) includes definitions of "personally identifiable information" and "protected personally identifiable information" that may apply in the absence of funding source rules.⁸

A CAA should obtain client consent before disclosing or sharing any client-specific information, both internally and with third parties, through AI tools. A release of information form typically authorizes the CAA to use and disclose client information for specified purposes. CAAs may ask clients to sign such a form at the time information is collected. The form should specify what information is being collected, how it will be used, to whom it will be disclosed, and whether it will be shared internally across programs and departments or with external parties. For further information regarding client rights to privacy, reference CAPLAW's [Privacy Rights of CAA Clients FAQ](#).

⁶ EEOC; [Visual Disabilities in the Workplace and the Americans with Disabilities Act](#), Question 16 (July 2023)

⁷ AI-enabled glasses raise other legal concerns when used in the workplace, as seen under the following sections in this resource: Attorney Client Privilege, IP Law, Labor Laws, and Recording Laws. For more information on how AI can provide accommodations and remove obstacles for workers with disabilities, see Job Accommodation Network; [How Artificial Intelligence Impacts Workplace Accommodation](#) (Accessed September 2026).

⁸ [2 CFR 200.1 Personally Identifiable Information \(PII\); 2 CFR 200.1 Protected Personally Identifiable Information \(Protected PII\)](#)

Notice and disclosure requirements may obligate CAAs to inform clients, employees, or regulators that the CAA is using AI in decision-making or operations. Several states have enacted mandatory disclosure rules around AI use in the employment context. For example, Utah’s Artificial Intelligence Policy Act requires disclosure when generative AI is used in regulated occupations or in interactions where a person could reasonably believe they are communicating with a human, which may be relevant to CAAs using AI chatbots or similar tools in client-facing service delivery.⁹ CAAs should assess whether their AI use requires new or amended privacy notices or disclosures under state law, particularly regarding automated decision-making and opt-out rights.

Whether CAAs must disclose AI use pursuant to their receipt of federal or other grants depends largely on the contract and grant requirements as well as how AI is used. In many cases, no affirmative disclosure obligation exists merely because a federal grantee uses AI to assist with internal productivity. CAAs should carefully review their federal grants and contracts to understand if current or future AI use may materially affect the grant or contract performance or if AI use and disclosure is otherwise specifically addressed.

Attorney client privilege protects certain communications from discovery in a lawsuit if they are (1) between a client and their attorney, (2) intentionally kept confidential, and (3) for the purpose of obtaining or giving legal advice. The work-product doctrine similarly protects certain materials prepared in anticipation of litigation from discovery, such as legal notes or strategies. If a third party is present in conversations with an attorney or documents are not prepared in anticipation of litigation, these protections may not apply.

Inputting confidential or litigation-related information into an AI tool raises concerns under both doctrines. Because an AI chatbot is not an attorney, communication with an AI tool is generally not a communication made to counsel for the purpose of obtaining legal advice and therefore may not qualify for attorney-client privilege. Additionally, many consumer-facing AI tools collect and retain user data, which may be treated as a disclosure to a third party that undermines the confidentiality required for privilege to attach.

Whether the work-product doctrine protects AI-assisted litigation materials is an emerging and fact-specific question. While courts remain split on the issue, they generally do not treat every third-party disclosure (including, potentially, inputs into an AI tool) as an automatic waiver of work-product protection.¹⁰ Some courts have held that a represented party who uses an AI tool on their own initiative, without any direction or involvement from their attorney, cannot claim work-product protection over the resulting materials. Other courts have reached the opposite conclusion, holding that the work-product doctrine broadly protects materials prepared by or for a party or its representative (not merely by attorneys) if prepared in anticipation of litigation.¹¹

⁹ Utah State Legislature; [S.B. 149 Artificial Intelligence Amendments](#)

¹⁰ See *United States v. AT&T Co.*, 642 F.2d 1285, 1299 (D.C. Cir. 1980); *United States v. Sanmina Corp.*, 968 F.3d 1107, slip op. at 21–23 (9th Cir. 2020)

¹¹ Compare *Bice v. Robb*, 511 F. App’x 108, 110 (2d Cir. 2013), with *Wultz v. Bank of China Ltd.*, 304 F.R.D. 384, 394–95 (S.D.N.Y. 2015), and *Hertzberg v. Veneman*, 273 F. Supp. 2d 67, 76–77 (D.D.C. 2003). Those courts which recognize work-product doctrine protections in this area note that this approach is consistent with the text of Federal Rule of Civil Procedure 26 which establishes the framework for disclosures in litigation, specifically during discovery.

Since the application of privilege and work-product protections to AI is still developing, CAAs should exercise caution. Asking an AI tool a legal question or inputting confidential information related to a legal matter into an AI chatbot may put privilege and work-product protections at risk. Similarly, AI-enabled glasses or notetakers used in conversations with an attorney may compromise privilege. CAAs should ensure staff understand these risks and refrain from inputting privileged, confidential, or litigation-related information into AI tools without guidance from counsel.

Intellectual property (IP) law is comprised of a set of rights to exclude others from making, copying, or using certain intangible creations of the human mind. Typically, IP law covers three forms of legal protection, each which protect a different type of creation: patents, copyrights, and trademarks. The Uniform Guidance also requires recipients and subrecipients to take certain actions when federal funds are used to pay for intangible property, i.e., property with no physical existence, such as trademarks, copyrights, and trade secrets.

Copyright grants creators of original works of authorship exclusive rights in their creative work, such as literary, graphic, and musical works. CAAs may routinely create copyrightable content, including educational materials, program guides, client resources, grant reports, and outreach materials. When CAAs use generative AI to help create such content, two copyright concerns arise: (1) whether the AI-assisted work will receive full copyright protection and (2) whether the AI output might infringe third-party copyrights. Courts have not, to date, recognized copyright in works that lack a human author, including AI-generated works.¹² However, works created by humans with the assistance of AI may be entitled to copyright protection depending on the nature of human involvement in the creative process. When AI determines the expressive elements of its output, that generated material is not considered the product of human authorship and is not protected by copyright. Some courts have found that having a human provide prompts to AI alone does not provide sufficient human control to make users of an AI system the authors of the output.¹³

For CAAs using AI to draft content, this means ensuring active human participation in the creative process and documenting that creative involvement. CAAs should also be aware that AI outputs may infringe copyrights in works used to train the AI tool. Copyright owners may be able to establish infringement if the AI program had access to their works and created substantially similar content.

Trademarks prevent businesses from misrepresenting the source of goods or services. Generally, a business may use any word, name, symbol, or device to identify particular goods or services. Generative AI can create logos, brand names, and slogans which may unintentionally create content that infringes a third party's trademark.

Questions regarding ownership and IP protection of AI-generated work product remain subject to regulatory attention and ongoing litigation. CAAs should be careful to ensure any AI-generated or

¹² Congressional Research Service; [Generative Artificial Intelligence and Copyright Law](#) (Updated July 2025)

¹³ U.S. Copyright Office; [Copyright and Artificial Intelligence, Part 2: Copyrightability](#) (January 2025)

assisted deliverables under federal grants and contracts are accurate and do not unintentionally incorporate protected third-party content. CAAs should also ensure that where AI is used in communications, CAA branding remains consistent so as not to negatively impact CAA trademarks (e.g., ensure AI does not distort a trademarked logo).

The right of publicity is an IP law concept reflected in most state laws. Right of publicity laws generally create a private right of action for the unauthorized commercial use of another person's name, image, or likeness or other aspects of their identity.¹⁴ Certain AI use raises concerns under right of publicity laws, such as using AI to imitate voices in song or deepfakes, which are AI-generated or altered content that can include real people's name, image, or likeness. While copyright laws do not prohibit voice imitations, right of publicity laws may protect and provide remedies for these AI voice and likeness misappropriations.

Trade secrets protect information that the owner makes reasonable efforts to keep secret and that derives its value from not being generally known to others. Both federal and state law provide trade secret protection. Output data from AI tools may be protectable as a trade secret if the information is sufficiently secret and not generally known. While trade secrets may offer a level of protection for certain AI outputs, AI also poses a risk to trade secret protection. A donor list that is protected, kept confidential, and has economic value is an example of a possible CAA trade secret.¹⁵

AI use can risk disclosure of trade secrets to third parties. To protect trade secrets, employers should educate employees on what constitute sensitive data and track the inputting of potentially sensitive data into AI tools and prompts. AI-enabled glasses, for example, may inadvertently capture trade secrets, internal communications, and client data and share such information with third-party AI vendors. These devices often connect to third-party AI vendors, meaning that information captured by the device may be transmitted to and stored by those vendors.

Labor laws, including the National Labor Relations Act (NLRA), protect employees' rights to engage in concerted activity concerning wages, hours, and working conditions. CAAs must ensure that the integration of AI in the workplace does not infringe upon employees' protected rights under the NLRA. When audio and video are recorded through AI-enabled glasses, for example, this recording of workplace conditions may qualify as protected activity, especially if it documents safety concerns, harassment, discrimination, or union-related activity. CAAs should ensure their AI policies are narrowly tailored and clearly tied to legitimate business justifications as not to infringe on employees' rights under the NLRA.

¹⁴ Congressional Research Service; [Artificial Intelligence Prompts Renewed Consideration of a Federal Right of Publicity](#) (Updated January 2024)

¹⁵ The Uniform Defend Trade Secrets Act of 2016, [18 U.S.C. § 1831 et seq.](#), defines trade secrets broadly as "all forms and types of financial, business, scientific, technical, economic, or engineering information, including patterns, plans, compilations, program devices, formulas, designs, prototypes, methods, techniques, processes, procedures, programs, or codes, whether tangible or intangible, and whether or how stored, compiled, or memorialized physically, electronically, graphically, photographically, or in writing if (A) the owner thereof has taken reasonable measures to keep such information secret; and (B) the information derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable through proper means by, another person who can obtain economic value from the disclosure or use of the information."

Recording laws raise concerns in two-party consent states. At the federal level, the Electronic Communications Privacy Act and Wiretap Act operates on a “one-party consent” basis, meaning it is legal to record a conversation if at least one person in the conversation is aware of and has consented to that recording. Many states, however, require “two-party consent” which means that every participant in a private conversation must be notified and consent before any recording.

When employees use AI to record meetings or other private workplace conversations, this poses a risk in states with two-party consent laws. CAAs using AI recording and transcription tools such as AI notetakers or AI-enabled glasses should provide advance notice in meetings in which such tools are used and should use active consent prompts that require others in the meeting to consent before recording begins.

Contract law at the state level governs a CAA’s relationship with AI vendors. CAAs should review all contracts with AI vendors, including reviewing the terms and conditions associated with free AI platforms. CAAs should ensure employees are trained not to use AI outside of the CAA’s policy and not to input any sensitive or work-related data into their own personal AI accounts.

While CAAs should review any AI vendor contract in its entirety, certain provisions warrant special consideration. Most importantly, CAAs should understand data ownership, sharing, and use rights. The contract should clearly state that the CAA is the sole and exclusive owner of the data uploaded to the AI tool and should specifically address whether the vendor may use, aggregate, or share the CAA’s data for purposes such as training AI models. The contract should also clarify how the vendor will protect the security and confidentiality of data inputs, including the definition of a data breach, notification requirements, and reimbursement obligations if the vendor is at fault.

Other important contract provisions address insurance, limitation of liability, and indemnification. The contract should address the insurance carried by each party, such as a cyber-liability insurance policy that protects from losses such as data and privacy breaches, and ask the vendor to maintain insurance coverage. A limitation of liability clause limits the amount of exposure an AI vendor will have in lawsuits and claims. A CAA should understand if potential damages owed by an AI vendor are capped and when the cap applies. The AI vendor should agree to indemnify and/or defend a CAA from certain claims or lawsuits by third parties resulting from a breach of the vendor’s obligations. For further information on reviewing key vendor contract provisions, see CAPLAW’s [Is Your Head in the Cloud? Contemplating Cloud Computing for CAAs](#).

Effective AI vendor management requires both initial due diligence and ongoing oversight. Before engaging an AI vendor, CAAs should vet the vendor’s data practices, security measures, and third-party disclosure policies. After engagement, CAAs should periodically review vendor performance, monitor for changes to terms of service or data practices, and reassess whether the vendor continues to meet the organization’s needs and risk tolerance.

This material is supported by Grant Number 90ET0505-03 from the ACF Office of Community Services, Community Services Block Grant within the Administration for Children and Families, a division of the U.S. Department of Health and Human Services. Neither the Administration for Children and Families nor any of its components operate, control, are responsible for, or necessarily endorse this material (including, without limitation, its content and any services or tools provided). The opinions, findings, conclusions, and recommendations expressed are those of the author(s) and do not necessarily reflect the views of the Administration for Children and Families and the ACF Office of Community Services, Community Services Block Grant.

The contents of this resource are intended to convey general information only and do not constitute legal advice. Any communication through this resource or through CAPLAW’s website does not constitute or create an attorney-client relationship. If you need legal advice, please contact CAPLAW or another attorney directly.

Sample CAA Policy

[Name of Community Action Agency]

AI Use Policy

[Community Action Agency] (“CAA”) recognizes that artificial intelligence (“AI”) provides unique opportunities for CAA and Staff to increase productivity and foster innovation. To assist Staff in making responsible decisions about the use of AI and to protect CAA’s programs and clients, CAA has established the following guidelines for the responsible, ethical, and legally compliant use of AI.

Scope

This CAA AI Use Policy (“Policy”) applies to employees, board members, volunteers, contractors, and other personnel (collectively, “Staff”) of CAA. The Policy applies to all AI use by Staff in connection with CAA operations, programs, services, communications, grant administration, data management, and internal processes (“CAA Work”), whether AI tools are CAA-provided or Staff-selected and whether a CAA-provided or Staff personal device is used.

This Policy is intended to guide CAA practice and training; it does not create a contract of employment or alter at-will employment status.

Staff’s use of AI requires Staff to exercise certain responsibilities described in this Policy and in other applicable CAA policies. Any conduct that adversely affects the CAA’s programs or clients or Staff’s performance may result in disciplinary action up to and including termination.

This Policy is not to be applied or interpreted in a manner that interferes with any activities protected by state or federal law, including the National Labor Relations Act.

Definitions

For purposes of this Policy:

Artificial Intelligence (AI): Technologies, systems, or tools that perform tasks typically requiring human intelligence, including but not limited to predictive analytics, machine learning, natural language processing, generative AI, large language models, image or audio generation, and automated decision systems.

TAILORING TIP:

Definitions and language should be informed by state AI laws and relevant grant or contract provisions, including the Uniform Guidance.

AI Tool: AI systems that create new content, such as text, images, audio, video, or code, based on prompts.

AI Output: Summary, analysis, content, documentation, or recommendation provided by an AI tool.

Personally Identifiable Information (PII): Any information that identifies, relates to, describes, or could reasonably be linked to an individual, including but not limited to names, addresses, contact details, and identification numbers.

Confidential Information: Nonpublic information from or about CAA, including client or employee data such as financial, health, education, or employment information, program data, financial information, and partner or vendor confidential information.

AI Use

Every use of AI at CAA must incorporate qualified human Staff when making decisions that affect CAA's clients, programs, and communities. AI is a tool, not a decision-maker. Staff remain accountable for decisions, actions, and outcomes that use or rely on AI. AI must be used in accordance with and in furtherance of CAA's mission and values.

AI use must comply with this Policy and all other existing CAA policies. Staff who are uncertain whether a proposed AI use complies with applicable law or this Policy should consult with their **Supervisor** before proceeding.



TO DO: Develop approval criteria or a decision framework for supervisors, which could potentially be tied to the risk classifications described below. Consider designating a specific role (e.g., IT, Compliance, or a designated AI coordinator) to handle or support approval decisions for higher-risk use cases. Lack of criteria creates two risks:

- 1) Inconsistent application: Without criteria, similarly situated Staff may receive different answers depending on which Supervisor they ask, undermining the Policy's fairness and predictability.
- 2) Uninformed approvals: Supervisors may not have the legal or technical background to evaluate AI-related risks. Without guidance, they may approve uses that present compliance concerns or deny beneficial low-risk uses out of excessive caution.

TAILORING TIP: Consider listing approved AI Tools in an exhibit that can be updated from time to time without updating the whole Policy.

Staff may use only the following AI Tools for CAA Work:

- [ChatGPT; Claude; Microsoft Copilot; Google Gemini; enterprise AI operating within secure environments.]

TAILORING TIP: Review each AI Tool and vendor's capabilities in coordination with what the CAA's ideal use cases are for AI use. Consider reviewing the contracts of several vendors to understand which will provide the best protections for the CAA.

Staff wishing to use AI Tools for CAA Work should discuss the parameters of their use with [their direct supervisor or a staff member responsible for AI use oversight] (the "Supervisor"). Staff must inform the Supervisor [verbally/in writing] how the AI Tool will be used. Supervisors must approve or deny requests within [number] days.

TO DO: Properly train supervisors on the legal risks associated with AI use as well as the Policy's permitted and prohibited use cases.

Generally acceptable uses of AI Tools include, but are not limited to:

- [Select any example permitted use cases from Exhibit A, or add a CAA's own use cases.]

All AI Output must be carefully reviewed by a human for accuracy before relying on it for CAA Work. If a reliable source cannot be found to verify AI Output, that information cannot be used for CAA Work. When used as a resource in final work product, AI Tools must be properly cited and documented in [CAA internal database and/or the CAA Work itself (e.g., in footnotes)].

TAILORING TIP: Time frames may depend on staff availability and capabilities as well as the role leadership wants AI to play in day-to-day operations. Longer review times may lead to slower AI adoption and fewer productivity benefits, while protecting compliance by ensuring fulsome review of the request.

Staff wishing to use AI Tools for any use cases not explicitly permitted in this Policy must request permission to do so from the Supervisor. The use of AI Tools for the following purposes is prohibited unless prior [written] authorization is received from the Supervisor. This is not an exhaustive list. CAA may add to this list at any time or reject the use of AI Tools at its discretion.

- [Select any example prohibited use cases from Exhibit A, or add a CAA's own use cases.]

TAILORING TIP: CAAs may decide not to require approval before AI is used but may require Staff to inform Supervisors after-the-fact when AI was used to complete a work product.

Staff may not use AI:

- In any way that Staff know or should know infringes or is likely to infringe any patent, trademark, trade secret, copyright, or other intellectual property or other rights of any person.
- For automated decision making, including profiling, unless Staff make all required disclosures, provide any required right to opt-out, ensure human review of the decision, and otherwise comply with any applicable laws.
- To generate or provide legal, financial, health, medical, tax, or other professional advice or decisions without review by a licensed and qualified professional and disclosure of such AI use.

Staff must meaningfully review any AI Output before taking action on matters that may materially affect access to or delivery of essential services, benefit

eligibility determinations, client outcomes, employment decisions, or matters presenting significant financial, legal, or reputational risk to the CAA. Staff may not integrate any AI Tool with internal CAA software without written permission from the Supervisor.

Training

All Staff must complete any AI training provided by CAA before using AI Tools for CAA Work. Supervisors must complete additional training before approving any AI use. Failure to complete the required training may be treated as a violation of this Policy.

Confidentiality and Privacy

Staff may not input Confidential Information, PII, or other High-Risk Information into AI Tools except as permitted under the risk classification framework below. Staff may not represent AI Output as being their own original work and must disclose AI assistance in accordance with CAA guidelines and any applicable legal or other compliance requirements.

Before entering any information into an AI Tool, Staff must classify that information as low-, medium-, or high-risk. High-Risk Information may only be entered into CAA-approved enterprise AI Tools that have been specifically authorized for such use, and only with prior written approval from Supervisor. Staff who are uncertain how to classify information should consult with their Supervisor before proceeding. [Staff who would like to use AI in connection with Medium-Risk Information must also receive prior written approval from Supervisor.] Low-Risk Information may generally be input into AI Tools in accordance with the Policy.

High-Risk Information

- Confidential Information related to Staff, clients, and donors, including but not limited to financial, immigration or legal status, health and mental health information
- PII of Staff, clients, and donors
- Data protected by specific laws, such as HIPAA
- Income, benefits, or financial details of Staff and CAA
- Proprietary information, trade secrets, or other CAA intellectual property

Medium-Risk Information

- Aggregated program statistics without identifying information
- Anonymized examples
- Internal meeting notes without any Confidential Information, PII, or other information referencing Staff, clients, or donors

TO DO: Develop specific disclosure guidelines. Key questions for CAAs to address include:

- (1) Internal disclosure: Must Staff inform Supervisors when AI was used to develop CAA Work? Should AI use be logged or tracked?
- (2) External disclosure: When must AI assistance be disclosed to clients, funders, or the public? Some grant applications and funder agreements may require disclosure of AI use.
- (3) Format: How should AI use be disclosed? Without clearer guidance, Staff may be uncertain what level of disclosure is expected, which could lead to either over-disclosure (creating administrative burden) or under-disclosure (creating compliance risk).

TAILORING TIP: This language reflects current market practice. A blanket prohibition with respect to High-Risk Information remains a valid conservative choice, but may limit beneficial AI use cases that could be safely enabled with enterprise-grade tools.

The tool-dependent approach reflected in the language distinguishes between:

- (1) Consumer-grade AI tools (e.g., free ChatGPT, public Claude) that may use inputs for model training, lack enterprise security controls, and generally should not receive high-risk information regardless of approval.
- (2) Enterprise AI tools with appropriate protections (e.g., Microsoft 365 Copilot, enterprise versions of ChatGPT/Claude with data processing agreements, no-training commitments, SOC 2 compliance, and BAAs for HIPAA-covered data) that may be appropriate for certain High-Risk Information with proper controls and approval.

CAAs should work with legal counsel and IT to identify which tools, if any, meet its security threshold for High-Risk Information.

- Grant application drafts with general descriptions

Low-Risk Information

- Publicly available information, research, and data
- General program descriptions
- Marketing and communications templates
- Training content without any client references

TO DO: Review related policies, such as acceptable use policies, IT security policies, data privacy policies, records retention policies, social media policies, and codes of conduct, to identify provisions that apply to AI use. Consider explicitly cross-referencing those policies here or incorporating AI-specific guidance into existing policies to maintain consistent compliance across all areas of Staff conduct.

AI use policies are most effective when they expressly integrate with an organization's existing governance framework. Staff may not realize that their acceptable use policy, IT security policy, data privacy policy, records retention policy, or code of conduct already impose requirements that apply to AI use.

Violations of this Policy

Any violation of this Policy will result in disciplinary action, up to and including termination of employment.

Staff should report any violations of this Policy or potential data privacy breaches pursuant to this Policy to their Supervisor. If the violation or potential breach involves the Supervisor, CAA staff should report to the [Executive Director or Board Chair]. CAA prohibits any form of discipline, reprisal, intimidation, or retaliation for reporting a violation of this Policy.

CAA Policies Remain in Effect

Staff using AI Tools must continue to abide by all other applicable CAA policies. Failure to follow CAA policies may result in termination of Staff's AI use and discipline, up to and including termination of employment.

The [DEPARTMENT NAME/IT/Human Resources/Compliance] is responsible for the administration of this Policy. If you have any questions regarding this Policy or questions about using AI Tools in connection with CAA Work that are not addressed in this Policy, please contact [DEPARTMENT NAME/CONTACT PERSON].

Additional Terms

The policy has been adopted voluntarily by CAA and is not intended to give rise to contractual rights or obligations. CAA may modify, add to, terminate, or deviate from this policy at any time and from time to time, without prior notice.

Effective Date and Revision History

Policy Effective Date:

Version:

Revision History		
Version	Date	Description of Changes
1.0	[]	Initial Version

TAILORING TIP: Consider requiring Staff to sign an Acknowledgement of Receipt that confirms they received, read, and understood the Policy. An acknowledgement serves two purposes: (1) it reinforces that Staff are on notice of the Policy's requirements and (2) it provides documentation that may be useful if disciplinary action becomes necessary. Some organizations include the acknowledgement as an exhibit to the policy while others handle it through onboarding or an annual compliance certification process.

TAILORING TIP: Tracking version history is a best practice for policies that may be updated more frequently, such as where technology and legal requirements are continuing to evolve.

Exhibit A – Sample Use Cases of AI

The below represents several examples of AI use cases at CAAs. CAAs should use these examples to develop their own list of use cases and may do so through a board committee, staff survey, or otherwise. CAAs should also revisit the list of use cases after a set period of time, such as every 6 months or 1 year, to keep the Policy relevant as AI Tools evolve quickly.

Sample Permitted Use Cases

- General knowledge questions meant to enhance understanding on a topic
- Summarize online research
- Brainstorm ideas related to projects
- Iterate/expand upon existing brainstorms related to projects
- Create formulas for Excel spreadsheets or similar programs
- Develop or debug code
- Draft or revise [internal/external] email communications for [spelling, grammar, and/or substantive content]
- Draft social media posts

Sample Prohibited Use Cases

- Notetaking during [internal/external] meetings
- Use of AI Output in final work product without human review and verification
- Submission of AI Output to funders or clients without disclosure of AI assistance
- Reliance on AI Output legal, medical, or financial advice without review by a qualified professional
- Submission of organizational data into an AI Tool
- Employment decisions, including recruitment, hiring, retention, promotions, transfers, performance monitoring, discipline, demotion, or terminations
- Use of AI-enabled glasses absent approval
- Inputting PII of Staff, clients, or prospective clients or applicants, including for eligibility determinations, into an AI Tool
- Any uses Staff know or should know infringes or is likely to infringe any patent, trademark, trade secret, copyright, or other intellectual property or other rights of any person

[Download a word version of the policy](#)