

Are the Government's Conversations with AI Accessible Under Public Records Laws and FOIA?

[Steven D Zansberg](#)

Sep 10, 2026 ⌚ 17 min read

Summary

- AI prompts and outputs used by government officials to conduct public business should generally be treated as public or agency records subject to state public records laws and FOIA.
- While courts have not yet resolved how these laws apply to AI chatbot conversations, existing legal principles strongly support disclosure when the records relate to official functions.
- Practical barriers such as poor retention, high costs, and asserted privileges will likely shape future disputes over access.

Nikada via Getty Images

Generative artificial intelligence presents a new frontier in the ongoing dialogue between technology and the law. Time will tell whether, as in the case of other technological advances, generative artificial intelligence will fulfill its promise to revolutionize the way we process information. But AI's novelty does not mean that its use is not subject to longstanding legal principles. . . . ¹

Serving as the mayor of a city or municipality is an extremely demanding job: having wall-to-wall meetings with city councilors, lobbyists, developers, public interest groups, activists, department heads, staffers, and officials from neighboring jurisdictions; attending numerous public events; and responding to direct contacts from constituents. Anyone holding such a 24/7 position, as well as that person's staffers, would welcome assistance from ChatGPT, Claude, Gemini, or other AI chatbots utilizing large language model learning. ²

In 2025, a combined reporting project between Cascade Public Broadcasting and KNKX-TV exposed that city workers, including the mayor's staff, in two Washington cities utilized ChatGPT to write mayoral letters, generate grant applications, draft policy documents, synthesize public comments,

and compose replies to citizens and the media.³ The joint reporting effort revealed alleged bid-rigging performed through AI prompts and responses,⁴ and that half of a letter the mayor of Bellingham sent to state authorities seeking funds for a Native American tribe was actually copied, verbatim, from ChatGPT.⁵ That reporting was enabled by the successful use of the Washington Public Records Act. Similarly, in 2025, the Texas Department of Transportation produced numerous conversations between its employees and ChatGPT in response to a records request under that state's public records law.⁶

But not all public officials have been willing to provide their AI prompts and results to requesters under state and federal public records laws. As recently reported in *Wired* magazine, both the "Department" of Government Efficiency and the Department of Housing and Urban Development denied Democracy Forward's Freedom of Information Act (FOIA) requests for the AI prompts and responses used in those agencies' crafting of official policies; HUD cited the deliberative process privilege (FOIA Exemption 5) as grounds for its withholding decision.⁷ At the state level, the City of Colorado Springs, Colorado, recently refused to provide a local television station copies of the AI chatbot interactions of that city's mayor, citing two exemptions from disclosure in Colorado's Open Records Act: "work product assembled for the benefit of elected officials" and "deliberative process" materials.

Various academics and government accountability groups have recently sounded the alarm over the lack of transparency regarding how the government is using AI to conduct public business, including rulemaking.⁸ No court has yet published a decision applying either FOIA or a state public records act to government officials' "conversations" with AI chatbots. While disputes over records access percolate their way up through the courts, this article examines what the public's right to know "what their government is up to"⁹ is when it uses generative AI to conduct public business.

Determining how state and federal public *records* laws apply to this relatively new technology requires following the same analytic path that courts have used for decades to determine whether records generated by prior communications platforms (e.g., chat rooms, text messages,¹⁰ email, audio/video recordings, Google searches, etc.) were subject to those laws. The first step in that process is to determine whether the particular "records" being sought constitute either "public records" under state law or "agency records" under FOIA. Only if that question yields a positive response does the court need to answer the second question: whether the government has met its burden of demonstrating that one or more statutory exemptions from disclosure of those records applies.

Which Records Are Being Sought?

There are multiple records associated with a government's use of AI technology. As one commentator has aptly stated: "With agencies adopting AI tools that generate probabilistic scores, model outputs, or real-time alerts, we must ask: What parts of an AI system constitute a "record"?"

Is it the input data? The model's decision? The training dataset? The algorithmic weights? Without guidance, lawyers and agencies are left in a legal gray zone.”¹¹

The underlying algorithm(s), filters, biases, and the large language dataset from which responses to prompts are generated all are of profound interest to the public if they are to understand how important policy decisions are made. The results generated by an AI chatbot are determined by the interaction of human prompts, the large language dataset, and the algorithmic programming for processing; accordingly, the public is entitled to know *all* of these data points, just as they are entitled to know the identities of the humans who contribute to the formulation of public policy.

Indeed, several states have passed so-called AI Transparency legislation that mandates the disclosure of the inner workings of AI programs, regardless of whether they are utilized by public officials or nongovernmental private persons.¹² California law requires all police and sheriff's departments to disclose publicly whenever AI has been used to generate any portion of an incident report and to retain the first draft thereof and the audit trail for as long as the official report is retained.¹³

While the inner workings of the particular generative AI bot used by the government is of unquestionable public interest and concern, this article will focus exclusively on records of the prompts submitted by government officials (“inputs”) and the responses that AI provides to those prompts (“outputs”).

Inputs and Outputs Connected to Official Conduct Are “Public Records” Under States’ Laws

The application (or nonapplication) of various records laws to a particular category of record depends, in large part, on how the relevant statute defines what *is* a “public record.” Most state’s public records laws generally define “public records” as (1) any “writing”—usually defined broadly to include any written communication “regardless of physical form or characteristics” (2) that is created, authored, produced, generated by a government employee, or “maintained, or kept” by a government agency (meaning over which it has either physical custody or a right of retrieval and access) *and* (3) whose *contents* bear some demonstrable connection to the conduct of public business.¹⁴ Thus, AI inputs or prompts *unrelated* to any official governmental function (e.g., “purely private” matters, like a personal grocery list or plans for a family member’s birthday party) would *not* be “public records,” even if they are generated by a public official on government-provided communications devices during regular business hours.¹⁵

Because state public records laws are considered “remedial statutes,” courts have tended to provide an expansive, pro-access interpretation to statutory text defining “public records.”¹⁶ As one commentator has put it, regarding Washington State’s public records law:

Courts in Washington have historically taken a wide view of what constitutes a [public] record. For instance, in *Belenski v. Jefferson County* (2015), the court found that automatically generated “internet access logs”—which simply tracked when an agency computer accessed the web—were public records. If a passive log is a public record, a deliberate prompt typed by a public employee to generate government content might be as well. ¹⁷

Indeed, several state courts have held that metadata associated with digitized public records is an intrinsic part of those records, thus requiring government agencies to produce such information under states’ public records laws. ¹⁸ Under these precedents, the content of a government official’s voluntary communications *to* an automated recipient, as well as the content of communications *from* an automated sender to a public official/employee, *for use in his/her conduct of official functions*, would qualify as “public records.”

Although no analogy is necessary to prove this point, consider a public official, in preparing various government reports or policy recommendations, inputting numeric “raw data” into a computer program (or a calculator) and then obtaining various results/analyses therefrom (e.g., sums, averages, means, standard deviation, etc.); no one would question that both the data inputted into and the tabulations produced by that program would constitute public records. That this analogy involves only numbers, not words, is a distinction without a difference. Indeed, that AI chatbots (unlike calculators) are notoriously susceptible to factual errors, and outright “hallucinations,” makes the need for public access to both the inputs and outputs all the more critical.

Indeed, California’s Special Districts Association has notified its members that “AI-generated content used in agency business will likely fall within C[alifornia] P[ublic] R[ecords] A[ct]’s broad definition of ‘public records.’ ... Agencies must be prepared not only for compliance, but for *public scrutiny of how AI shapes government communications*.” ¹⁹

Similarly, writing in Coates’ Canons, the North Carolina local government law blog, Kristi Nickodem concludes that “[i]f a state or local government official or employee is using a generative AI tool (e.g., ChatGPT, Gemini, Claude, Copilot) to carry out the duties of their role, then both the prompts the individual gives to the AI tool (records ‘made’ in connection with the transaction of public business) and the outputs the individual receives from the AI tool (records ‘received’ in connection with the transaction of public business) would be subject to North Carolina’s public records law.”

²⁰ This conclusion applies, with equal force, to numerous—if not all—other states’ public records laws.

Inputs and Outputs Connected to Agency Business Are “Agency Records” Under FOIA

The federal Freedom of Information Act (FOIA) applies only to “agency records” of Executive Branch agencies. ²¹ As in state public records laws, “records” are defined broadly to include “any

information” in “any format, including an electronic format.” ²² As is true under states’ public records laws, records whose content is not germane to any official functions or activities of that agency do not constitute “agency records.” ²³ But the Supreme Court has made clear that the FOIA was intended to provide public access to “the information available to an agency in its decision-making processes,” including, by way of example, “studies, trade journal reports, and other *materials produced outside the agencies* both by private and governmental organizations,” which agency personnel use “[i]n performing their official duties.” ²⁴

Under the above precedents, records generated by federal agency employees in the scope of their duties (i.e., inputting prompts into an AI chatbot) and the records received from the chatbot and kept by that employee for use in conducting that agency’s public business constitute “agency records” subject to FOIA. ²⁵

The Federal Government Is “All In” on Using AI to Conduct the Public’s Business

On April 3, 2025, Director of the Office of Management and Budget Russell T. Vought issued Memorandum M-25-21, mandating all federal “agencies *must* adopt a forward-leaning and pro-innovation approach that takes advantage of [AI] technology to help shape the future of government operations. . . . Agencies are *directed to accelerate* the Federal use of AI.” ²⁶ Section 4 of that memo imposes additional monitoring and evaluation of pilot testing of AI that performs “high impact” functions: “AI is considered high-impact when its output serves as a principal basis for decisions or actions that have a legal, material, binding, or significant effect on rights or safety.” ²⁷

A handful of concrete examples exemplify the general trend. ²⁸ In July 2025, the Environmental Protection Agency publicly touted that it had, as of that date, documented 103 instances of agency employees’ “AI use or planning,” and its administrator declared, “At EPA, we’re embedding innovation into our systems, not resisting it. From *AI tools that speed up environmental reviews* to smarter monitoring and emissions tracking” ²⁹ In September 2025, the Federal Communications Commission (FCC) publicly proclaimed its “commitment to accelerate Federal use of AI.” ³⁰ Indeed, in its first publicly released “AI Use Cases” in February 2026, the FCC reported that it had licensed Microsoft Copilot, for between 101–1,000 agency users, for the purpose of “[g]enerating first drafts of documents, briefing, or communication materials using AI.” ³¹ In February 2026, the Department of Justice reported 315 instances of AI use in 2025, including, controversially, “using AI to predict rates of recidivism for current inmates and sex offenders.” ³² And in July 2026, the Department of Veterans Affairs announced that “VA GPT is an on-network generative AI chat interface that employees are using to assist with basic administrative tasks (drafting emails, summarizing documents, summarizing meeting notes, etc.)[, ... with] over 95,000 users onboarded.” ³³

Under existing caselaw, all “writings” in agencies’ possession, custody, or control that document, or are utilized in conducting official agency business, constitute “agency records” subject to FOIA. And an executive order entered at the end of the first Trump administration dictates that “Agencies *shall* be transparent in disclosing relevant information regarding their use of AI to appropriate stakeholders, including the Congress and *the public*, to the extent practicable and *in accordance with applicable laws and policies*.” ³⁴

Logistical Hurdles to Accessing Public Records and Agency Records

Merely because a document or record is covered by a state’s public records law or FOIA does *not* mean you can obtain a copy of it. A group of researchers, led by Alex Spangher at Stanford University, is engaged in a systematic survey of both state and federal agencies and their use of AI chatbots in performing public business. ³⁵ In April 2026, the group submitted 24 FOIA requests to various federal agencies and 81 requests to state and local governments across the nation under states’ public records laws. Those requests sought records showing how widespread their use of generative AI is, the policies they have in place governing such use, and records of the actual AI chatbot conversations of public employees.

The group’s preliminary findings have cataloged deployments of AI by 53 federal agencies, all 50 states, and 77 local governments. The group also has assembled the first corpus of government AI chatbot logs, comprising 3,216 chat threads, between 2023 and 2026. The study’s authors report that the two most frequently encountered barriers to accessing AI chatbot conversations are (1) the agencies’ claim/assertion that they have “no records responsive to the request,” and (2) the exorbitant costs being assessed as the condition for providing the responsive records. The state of Colorado, for example, asked the researchers for \$9 million to produce records responsive to the group’s requests. ³⁶

Both of these barriers—(1) inadequate records retention and (2) prohibitive costs imposed on requesters for search and retrieval—have thwarted public access to digitized and cloud-based records, even before the advent of generative AI. ³⁷ While solutions to these two significant roadblocks to public access are beyond the scope of this article, advocates for the public’s right to know must focus their efforts on requiring AI-related records retention and challenging exorbitant fees that pose an insurmountable barrier to accessing those records.

Notably, on point (1) above, the U.S. Army has issued an AI Compliance Memorandum that expressly requires all system owners to “ensure that all user interactions, *including prompts and AI-generated content*, are properly identified, *retained* and secured as official records” and to “provide access to AI prompts and content classified as government records and maintain them to support timely FOIA responses.” ³⁸ Similarly, the Department of Homeland Security has contracted with OpenAI to process and maintain all communications with the department as federal records that may be subject to FOIA requests. ³⁹

Statutory Exemptions Most Likely to Be Asserted in Response to Requests to Access AI Chatbot Conversations with Public Employees

Records subject to either state public records laws or the federal FOIA are subject to various exemptions from disclosure. Among them are personnel files, investigatory files of criminal law enforcement, trade secrets, confidential commercial data, privileged records, and various other expressly itemized exceptions.⁴⁰

This article will not survey all the statutory exemptions that governments may assert as grounds to withhold AI chatbot conversations. Instead, it will focus on the statutory exemptions most likely to be asserted as grounds for nondisclosure.

And, as noted above, both the City of Colorado Springs and the “Department” of Government Efficiency have denied public access to public officials’ AI prompts and responses, asserting they are subject to two related state-law exemptions: “work product ... assembled for the benefit of elected officials” (an exclusion from Colorado’s definition of “public records”⁴¹) and the deliberative process privilege. How might those related statutory exemptions fare when asserted as grounds for withholding government employees’ AI chatbot prompts and responses under state and federal records laws?

Government offices assert that when an elected official, or other government employee, jots down preliminary thoughts on a paper pad or (using a laptop, phone, or other digital device) prepares a draft of a speech or a proposed policy, such writings constitute “work product assembled for the benefit of an elected official” or deliberative process material. Assuming, for present purposes, that is correct, why should the transmittal of those drafts to an AI chatbot (or to Google, Meta, Verizon, or Microsoft) transform those writings into something not protected by those exemptions? If an AI chatbot performs the same functions as a human staffer or outside consultant, e.g., proposing revisions or actually generating the “first draft” of an agency policy, how are such writings any different from those generated by a human, in terms of privilege? To respond to these questions, we must consider the underlying purposes of those related privileges/exemptions.

Are Humans Required at Both Ends of the “Conversation”?

A creature of the common law, the deliberative process privilege serves to shield from public view pre-decisional communications among government employees (and/or outside consultants), particularly recommendations, advice, and subjective opinions of subordinates to higher-ranking superiors—so that the person transmitting such views will not be “chilled” from offering frank and candid advice.⁴² The privilege “serves to assure that subordinates within an agency will feel free to provide the decisionmaker with their uninhibited opinions and recommendations without fear of later being subject to public ridicule or criticism.”⁴³

But does this rationale make any sense when the recommendations or “advice” come not from a human being, but from an inanimate machine or computer program? AI chatbots are programmed to generate responses to prompts automatically, regardless of the identity or rank/position of the human who submits them, the nature of the prompt, or the purpose to which the response is to be used. Thus, the principal rationale for the so-called privilege—to avoid “chilling” the transmission of candid recommendations, opinions, or advice from subordinates to their superiors—seems completely inapposite to information provided by a computer program.⁴⁴

As John Davisson, deputy director of enforcement at the Electronic Privacy Information Center, has stated: “If there is a deliberative process, occurring *between human beings* that work for [a federal] agency, where the material is both pre-decisional and leading up to some sort of decision, that might be a justifiable assertion of deliberative process privilege there,” but “AI systems, computers are not entitled to candor.”⁴⁵ In other words, being notified that its responses to prompts *will be disclosed to the public* is not going to “chill” *a computer program* from processing the prompt(s) and generating responses as dictated by its programming.

However, federal courts applying Exemption 5 have not required two human participants sharing views, opinions, or thoughts on contemplated agency action. For example, two courts have held that the exemption applies to computer programming used to process selected data input by a government employee or a retained outside expert.⁴⁶

Indeed, Exemption 5 even has been applied to a government official’s own writing, *not shared with any third party*, generated for his/her own benefit, i.e., in preparing a “working draft” or notes for future consideration in formulating agency policy.⁴⁷ Why is that? Because in addition to the “avoid chilling candid communications” rationale, the deliberative process privilege also serves “to protect against premature disclosure of proposed policies before they have been finally formulated or adopted; and to protect against confusing the issues and misleading the public by dissemination of documents suggesting reasons and rationales for a course of action which were not in fact the ultimate reasons for the agency’s action.”⁴⁸

So, at least some courts have held that two humans exchanging communications is not a prerequisite for asserting that AI inputs and outputs are subject to the deliberative process privilege. Accordingly, whether the invocation of that privilege will be sustained must turn on the particular factual context of each case: How “deliberative” in nature are the AI chatbot communications, and how relevant is that content to an actual *policy-making* decision?⁴⁹ Of course, pre-decisional materials that form the basis of actual agency action (and are incorporated into that policy or decision) are not subject to the deliberative process privilege;⁵⁰ but it may be difficult, without an audit trail, to determine which portions of AI chatbot outputs were so incorporated.

Does Conversing with an AI Chatbot Waive All Claims of Privilege?

For purposes of this discussion only, let's assume that certain prompts submitted to an AI chatbot, in a particular factual context, would ordinarily constitute "deliberative process" material. Like all others, the deliberative process privilege can be waived through voluntary disclosure to third parties. So, does the "sharing" or disclosure of those prompts to an AI chatbot waive that privilege?

One federal court has applied the privileges routinely asserted as grounds for withholding "documents" from discovery in litigation to AI chatbot inputs and outputs, finding that neither attorney-client privilege nor the work product doctrine applies,⁵¹ while two other federal courts have ruled that such conversations by pro se litigants can qualify for protection under the work product doctrine.⁵²

In March 2026, Judge Jed Rakoff, of the Southern District of New York, found that an AI user's prompts to the AI chatbot Claude (by Anthropic) were not subject to either of those doctrines.⁵³ The party asserting those privileges, Bradley Heppner, was a criminal defendant, charged with securities fraud and other crimes. He claimed he had used Claude to prepare outlines of defense strategies for his attorneys, although it was undisputed that his attorneys had not requested him to do so. Judge Rakoff categorically denied Heppner's assertion of privilege with respect to "the AI documents." Not only is Claude "not an attorney," and therefore incapable of providing legal advice, but more importantly, the AI communications were not "confidential" because

the written privacy policy to which users of Claude consent provides that Anthropic collects data on both users' "inputs" and Claude's "outputs," that it uses such data to "train" Claude, and that Anthropic reserves the right to disclose such data to a host of "third parties," including governmental regulatory authorities. ... The policy clearly puts Claude's users on notice that Anthropic ... may "disclose personal data to third parties in connection with claims, disputes[,] or litigation." ... [Accordingly,] Heppner could have had no "reasonable expectation of confidentiality in his communications" with Claude.⁵⁴

Under Judge Rakoff's ruling in *Heppner*, it would appear that sharing information with an AI chatbot waives any claim to confidentiality, or at least it *might*. *Heppner*'s ruling is limited to that case's particular facts. It is far from clear that the same would be true if the operative facts were different. Suppose, for example, that the government employee's prompts were input into an AI chatbot that did *not* utilize user prompts to train it, and the chatbot was contractually bound to not disclose either inputs or outputs to any third party. In the dispute mentioned above between a television station and the City of Colorado Springs, the City maintains that the two AI chatbots the mayor used to help him draft speeches and correspondence are not trained on user prompts and are contractually bound not to share any information it receives or provides with any third party. Thus, the City asserts, no waiver of any privilege has occurred.

Recall, too, that two federal courts have distinguished *Heppner*, in a markedly different procedural and factual context, and have held that *pro se* litigants submitting prompts to AI chatbots to aid

them in preparing court filings did not waive their claim to the “attorney work product” doctrine applicable in civil discovery disputes. ⁵⁵

So, here too, the outcome of future lawsuits challenging denials under state and federal records law of AI chatbot conversations of government employees is likely to turn on the particular circumstances under which those writings are generated and shared.

Conclusion

Agencies and officials at all levels of government have widely embraced generative artificial intelligence in conducting public business. ⁵⁶ AI is the latest in a long line of technological innovations that facilitate human communications: the written word, printing press, telegraph, telephone, radio, audio/video recording, email, text messaging, web-based chatrooms, Zoom. As Judge Rakoff has stated, while technology advances at breakneck speed, legal principles remain unchanged. Public records laws, including FOIA, were enacted to enable the citizenry to monitor the work of, and thereby hold accountable, *everyone in government*, i.e., those constitutionally bound to serve the interests of We the People. ⁵⁷

AI chatbots are designed by humans and are “trained” using massive amounts of human-generated text, images, and audio. Unlike all the other communications media listed above, generative artificial intelligence independently *creates content* and, in doing so, based on its programming, inevitably introduces certain biases, distortions, and outright factual errors into the conversations in which it participates. ⁵⁸ This justifies an *even greater public interest* in knowing (1) which AI chatbot(s) are being utilized to formulate public policy and otherwise conduct public business, (2) who is using AI chatbots, (3) what algorithms and datasets are utilized by those chatbots, (4) what are the actual prompts that government officials submitted to those chatbots, (5) what are the responses generated by the chatbots, and (6) how are those responses thereafter put to use in conducting public business.

Practically all of these data points are automatically documented in digitized records made, maintained, or kept by government offices that utilize AI chatbots. And all of those writings are quite appropriately subject to public scrutiny under state and federal public records laws. AI use by government has quickly become the rule, not the exception, and (as is true outside government) will soon be practically ubiquitous. Without access to those writings, the public will not be able to understand what its government is up to.

Endnotes

1. *United States v. Heppner*, 820 F. Supp. 3d 292, 299 (S.D.N.Y. 2026).
2. According to one source, the 11 “best” AI Chatbots in 2026 are (1) ChatGPT, <https://chatgpt.com/>; (2) Claude, <https://claude.ai/>; (3) Perplexity, <https://www.perplexity.ai/>; (4) Gemini, <https://gemini.google.com/>; (5) Copilot, <https://copilot.microsoft.com/>; (6) DeepSeek,

- <https://www.deepseek.com/>; (7) LLaMA, <https://www.llama.com/>; (8) Tidio, <https://www.tidio.com/>; (9) Intercom, <https://www.intercom.com/>; (10) ProProfs, <https://www.proprofschat.com/>; and (11) Knock AI, <https://www.knock-ai.com/>. *Best AI Chatbots of 2026: Tested & Ranked for Research, Productivity, Customer Support, and More*, KNOCK.AI, <https://www.knock-ai.com/blog/best-ai-chatbots> (last visited Aug. 31, 2026).
3. See Nate Sanford, *Washington City Officials Are Using ChatGPT for Government Work*, CASCADE PBS (Aug. 26, 2025), <https://www.cascadepbs.org/news/2025/08/wa-city-officials-are-using-chatgpt-to-write-government-documents/>; Nate Sanford, *Washington Cities' AI Policies Play Catch-up as Officials Embrace New Tech*, CASCADE PBS (Aug. 27, 2025), <https://www.cascadepbs.org/news/2025/08/washington-cities-ai-policies-play-catch-up-as-officials-embrace-new-tech/>.
 4. [Nate Sanford](#), *Bellingham Staffer Asked ChatGPT to "Exclude" Vendor from City Contract*, KNKX (Jan. 5, 2026), <https://www.knkx.org/government/2026-01-05/city-of-bellingham-chatgpt-ai-contract-vendor>.
 5. Sanford, *Washington City Officials Are Using ChatGPT for Government Work*, *supra* note 3.
 6. See, e.g., William Snider, Microsoft Pilot, "I'm doing the safety moment for our meeting this morning. I'd like to do it on crowd safety. Can you give me some talking points?" (July 28, 2025), https://cdn.muckrock.com/foia_files/2026/02/02/I_m_doing_the_safety_moment_for_our_meeting_this_m.html (last visited Aug. 31, 2026); Aaron Lenz, Microsoft Copilot, "compose an email that lets the recipient know we are asking them to confirm project status so that we can determine if their project should be migrated to a new system or finish in the existing system" (Nov. 17, 2025), https://cdn.muckrock.com/foia_files/2026/02/02/compose_an_email_that_lets_the_recipient_know_we_a.html (last visited Aug. 31, 2026).
 7. Vittoria Elliott, *DOGE Used AI for Housing Policy. The Government Won't Say How*, WIRED (Jul 14, 2026), <https://www.wired.com/story/doge-deployed-ai-housing-policy/>.
 8. See Sean Michael Hewhouse, *AI's Growing Role in Rulemaking Raises New Transparency Questions*, Gov. EXEC. (July 27, 2026), <https://www.govexec.com/technology/2026/07/ai-growing-role-rulemaking-raises-new-transparency-questions/415037/>; see also Nathan E. Sanders & Bruce Schneier, *AI Use by the US Government Is Ballooning. And the Lack of Transparency Is Troubling*, THE GUARDIAN (June 15, 2026) (pointing to the French Digital Republic Act passed in 2016, which "requires all algorithms used to automate government administrative decisions to be subject to public records requests, to be appealable to a human reviewer, and to have mandatory notification of the use of automation to those affected by the decisions"), <https://www.theguardian.com/commentisfree/2026/jun/15/ai-use-by-the-us-government-is-ballooning-and-the-lack-of-transparency-is-troubling>; Avèro Advisers, *AI Governance in Government: Trust Requires Transparency* (Mar. 28, 2025), <https://averoadvisors.com/insights/ai-governance-in-government-trust-requires-transparency/>.
 9. U.S. DOJ v. Reporters Comm. for Freedom of Press, 489 U.S. 749, 773 (1989) (stating that the underlying rationale of the FOIA is to effectuate "the citizens' right to be informed about 'what their government is up to'" (internal citation omitted); see also *infra* note 59.
 10. See, e.g., Steve Zansberg, *Cloud-Based Public Records Pose New Challenges for Access*, 31 COMM. L., no. 1, Winter 2015, at 12, /content/dam/aba/publications/communications_lawyer/january2015/cl_win15_v31n1.pdf; Helen Vera, "Regardless of Physical Form": Legal and Practical Considerations Regarding the Application of State

Open-Records Laws to Public Business Conducted by Text Message, 32 COMM. L., no. 4, Spring 2017, at 24, /content/dam/aba/publications/communications_lawyer/cl_v32_4_spring17_v1.pdf.

11. David Staples, *Public Records Laws Aren't Ready for the Age of AI—And That's a Problem*, MEDIUM (June 26, 2025), <https://medium.com/@dstaples836/public-records-laws-arent-ready-for-the-age-of-ai-and-that-s-a-problem-c14efe0ba671>.
12. Malihe Alikhani & Aidan T. Kane, *What Is California's AI Safety Law?*, BROOKINGS (Dec. 23, 2025), <https://www.brookings.edu/articles/what-is-californias-ai-safety-law/>; Kim Phan, Chris Willis & Taylor Gess, *Colorado Rewrites Its AI Law*, CONSUMER FIN. SERVS. L. MONITOR (May 13, 2026), <https://www.consumerfinancialserviceslawmonitor.com/2026/05/colorado-rewrites-its-ai-law/>; *Summary: Artificial Intelligence 2025 Legislation*, NAT'L CONF. OF STATE LEGISLATURES, <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation> (last updated July 10, 2025).

One high-profile controversy surrounds the Trump administration's decision to adopt Elon Musk's AI chatbot, Grok, for use by federal agencies, including the Department of Defense. *See, e.g.*, Konstantin Toropin & David Klepper, Associated Press, *Pentagon Embraces Musk's Grok AI Chatbot as It Draws Global Outcry*, PBS NEWS (Jan. 13, 2026), <https://www.pbs.org/newshour/world/pentagon-embraces-musks-grok-ai-chatbot-as-it-draws-global-outcry>. Public Citizen recently sought OMB's agency records concerning this decision. *See* Letter from Pub. Citizen to Off. of Mgmt. & Budget, Freedom of Information Act Request—Final OMB Determinations, Implementation Records, and Interagency Coordination Regarding xAI's (July 14, 2026), <https://www.citizen.org/article/foia-request-seeks-omb-records-on-grok-ais-use-by-the-federal-government/>.

Others have raised concerns about government using AI to formulate policies affecting people's substantive rights. *See, e.g.*, The Learning Studio, *The Dark Side of AI in Government Explained*, YOUTUBE (Aug. 4, 2026), <https://youtu.be/2hOfCbOnhyY?si=hmxXqsdU8u9hH-vv>; *Government Use of AI*, ELEC. PRIV. INFO. CTR. (Oct. 28, 2025), <https://epic.org/issues/ai/government-use-of-ai/>.

13. CAL. PENAL CODE § 13663 (2026). For other states' mandatory AI use and disclosure laws, *see* LEXIS's "Artificial Intelligence State Law Survey," <https://bit.ly/4ww7Lkx> (last visited Sept. 6, 2026; subscription required).
14. *See, e.g.*, *State Freedom of Information Laws*, NAT'L FREEDOM OF INFO. COAL., <https://www.nfoic.org/state-freedom-of-information-laws/> (last visited Aug. 31, 2026); *Resource: Public Records Law and State Legislatures*, NAT'L CONF. OF STATE LEGISLATURES, <https://www.ncsl.org/center-for-legislative-strengthening/public-records-law-and-state-legislatures> (last updated Apr. 16, 2025); Andrea G. Nadel, Annotation, *What Are "Records" of Agency Which Must Be Made Available Under State Freedom of Information Act*, 27 A.L.R.4th 680 (Supp. 2026); Cristina Abello, *Access to Electronic Communications*, REPORTERS COMM. FOR FREEDOM OF THE PRESS (2009), <https://web.archive.org/web/20220901054416/https://www.rcfp.org/wp-content/uploads/imported/ELECCOMM.pdf>; Marjorie A. Shields, Annotation, *Disclosure of Electronic Data Under State Public Records and Freedom of Information Acts*, 54 A.L.R.6th 653 (Supp. 2020).
15. *See, e.g.*, *Denv. Publ'g Co. v. Bd. of Cnty. Comm'rs of Arapahoe Cnty.*, Colo., 121 P.3d 190 (Colo. 2005) (holding that sexually explicit text messages exchanged between two government employees on county-provided devices, while on duty, were not "public records" because the content of those messages did not bear "a demonstrable connection" to the discharge of public functions); *Wick Commc'ns Co. v. Montrose Cnty. Bd. of Cnty. Comm'rs*, 81 P.3d 360 (Colo. 2003) (holding that when a public employee

makes only passing reference to his or her professional conduct in the entries of a private diary, that diary is “made, maintained, [and] kept” purely in that employee’s private capacity).

16. See, e.g., *Energy Transfer v. Friedman*, 265 A.3d 421, 428 (Pa. 2021) (“As ‘remedial legislation to facilitate government transparency and accountability,’ the R[ight] T[o] K[now] L[aw] is ‘construed to maximize access to public records’ in an agency’s possession.”) (internal citation omitted); *Wyo. Dep’t of Transp. v. Int’l Union of Operating Eng’rs Local Union 800*, 908 P.2d 970, 973 (Wyo. 1995) (in light of “[t]he remedial purpose of the public records act,” “we construe the act liberally in favor of disclosure, interpreting all exemptions narrowly”).
17. Anna Nollan, Stephanie Berntsen & Alicia Lowe, *Navigating the Digital Tide: Generative AI and the Washington Public Records Act*, SCHWABE (Mar. 9, 2026), <https://www.schwabe.com/publication/navigating-the-digital-tide-generative-ai-and-the-washington-public-records-act/>.
18. See, e.g., Peter Kozinets, *Access to Metadata in Public Records: Ensuring Open Government in the Information Age*, 27 COMM. L., no. 2, Summer 2010, at 1, <https://www.steptoelaw.com/a/web/2465/4095.pdf>; Reporters Comm. for Freedom of the Press, *Requesting Metadata with Electronic Records*, THE NEWS MEDIA & THE LAW, Spring 2011, at 17 <https://www.rcfp.org/journals/the-news-media-and-the-law-spring-2011/requesting-metadata-electro/>.
19. Kristin Withrow, *Public Records Act Meets AI: What Recent Requests Reveal About CHATGPT Use in Government*, CAL. SPECIAL DIST. ASS’N (May 12, 2026), <https://www.cstda.net/blogs/kristin-withrow1/2026/05/11/public-records-act-meets-ai> (membership required).
20. Kristi Nikodem, *The Intersection of Artificial Intelligence and the Public Records Act*, COATES CANONS (Mar. 11, 2026), <https://canons.sog.unc.edu/blog/2026/03/11/the-intersection-of-artificial-intelligence-and-the-public-records-act/>.
21. 5 U.S.C. § 552(a)(3)(a) (2026); see also *id.* § 551(1) (defining “agency”).
22. *Id.* § 552(f)(2).
23. *U.S. DOJ v. Tax Analysts*, 492 U.S. 136, 145 (1989) (citing *Kissinger v. Reporters Comm. for Freedom of the Press*, 445 U.S. 136, 157 (1980)).
24. *Id.* at 144.
25. Foreign governments have codified the right of public access to government AI usage records. See, e.g., *Freedom of Information (FOI) and Artificial Intelligence*, INFO. COMM’R’S OFF. (last visited May 6, 2026) (“If staff at a public authority use AI for work purposes, the information generated will be subject to FOIA along with the prompts used.”); Loi 2016-1321 due 7 octobre 2016 pour une République numérique [Law. No. 2016-1321 of October 7, 2016 on the Law for a Digital Republic], JOURNAL OFFICIEL DE LA RÉPUBLIQUE FRANÇAISE [OFFICIAL JOURNAL OF THE FRENCH REPUBLIC], October 10, 2016; Adriane Louyer, *Transparency Obligations for Algorithms in the Public Sector*, HOUDART & ASSOCS. (Oct. 3, 2022), <https://www.houdart.org/les-obligations-de-transparence-des-algorithmes-dans-le-secteur-public/>.
26. Memorandum from Russell T. Vought, Off. of Mgmt. & Budget, Exec. Off. of the President, Memorandum M-25-21, Accelerating Federal Use of AI Through Innovation, Governance, and Public Trust at 1–2 (Apr. 3, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of->

AI-through-Innovation-Governance-and-Public-Trust.pdf (emphasis added). On that same date, OMB released another memo, M-25-22, which provides further guidance on the processes by which federal agencies are to acquire and begin utilization of AI in performing agency functions. Memorandum from Russell T. Vought, Off. of Mgmt. & Budget, Exec. Off. of the President, Memorandum M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (Apr. 3, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf>.

27. Memorandum M-25-21, *supra* note 26, at 14. Federal agencies' enthusiastic embrace of AI predated the current presidential administration. *See, e.g.,* Billy Mitchell, *DHS Releases Commercial Generative AI Guidance and Is Experimenting with Building Its Own Models*, FEDSCOOP (Nov. 30, 2023), <https://fedscoop.com/dhs-releases-commercial-generative-ai-guidance-and-is-experimenting-with-building-its-own-models/>.
28. *See infra* note 56; *see also* Valerie Wirtschafter, *Assessing the State of AI Adoption Across the Federal Government*, BROOKINGS (Apr. 15, 2026), <https://www.brookings.edu/articles/assessing-the-state-of-ai-adoption-across-the-federal-government>.
29. *See* Ellie Borst, *EPA Use of Artificial Intelligence Skyrockets*, E&E NEWS/POLITICO (July 30, 2025) (also stating "EPA is planning to integrate artificial intelligence into processes at the agency."), <https://www.eenews.net/articles/epa-use-of-artificial-intelligence-skyrockets/>.
30. *See* U.S. FED. COMM'NS COMM., ARTIFICIAL INTELLIGENCE COMPLIANCE PLAN FOR OMB MEMORANDUM M-25-21 (Sept. 2025), <https://www.fcc.gov/sites/default/files/fcc-artificial-intelligence-compliance-plan.pdf>.
31. *AI Use Cases*, FED. COMM'NS COMM., <https://www.fcc.gov/sites/default/files/ai-use-cases.xlsx> (last visited Aug. 31, 2026).
32. Miranda Nazzaro, *DOJ Ramps up AI for Legal Work, Crime Predictions, Surveillance, Inventory Shows*, FEDSCOOP (Feb. 11, 2026), <https://fedscoop.com/justice-department-artificial-intelligence-ai-surveillance-inventory-predictive-technology-algorithm-bias/>.
33. *AI Use Case Inventory*, U.S. DEP'T OF VETERANS AFFS., <https://department.va.gov/ai/ai-use-case-inventory/> (last visited July 6, 2026).
34. Exec. Order No. 13,960, Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government, 85 Fed. Reg. 78,939, 78,940 (Dec. 8, 2020), <https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government> (emphasis added).
35. *Stanford AI-Usage FOIA*, MUCKROCK, <https://www.muckrock.com/organization/stanford-ai-usage-foia/> (last visited Aug. 31, 2026).
36. Alexander Spangher, *Public Records Request—Colorado Generative AI Records*, MUCKROCK, <https://www.muckrock.com/foi/colorado-127/public-records-request-colorado-generative-ai-records-208935/> (last visited Aug. 31, 2026).
37. *See, e.g.,* Zansberg, *supra* note 10.

38. Jane Edwards, *Army CIO Issues Guidance on AI Compliance with Records Management*, FOIA, EXECUTIVEGOV (Sept. 3, 2025), <https://www.executivegov.com/articles/army-garciga-application-owners-ai-records-mgmt> (emphasis added).
39. See Rebecca Heilweil, *Generative AI Could Raise Questions for Federal Records Laws*, FEDSCOOP (Apr. 22, 2024) (“an amendment to OpenAI’s terms of service uploaded to the DHS website established that outputs from the model are considered federal records, along with referencing freedom of information laws”), <https://fedscoop.com/generative-ai-could-raise-questions-for-federal-records-laws/>.
40. See 5 U.S.C. § 552(b) (2026) (listing nine statutory exemptions from disclosure).
41. Colorado courts have held that “[a]dvice’ is commonly understood to mean ‘the way in which one regards something,’ and is synonymous with *a person’s* view or opinion.” *City of Fort Morgan v. E. Colo. Publ’g Co.*, 240 P.3d 481, 486 (Colo. App. 2010) (internal citation omitted) (emphasis added); see also *id.* at 487 (citing legal dictionary definition of “advice” as “[g]uidance *offered by one person ... to another*”) (internal citation omitted) (emphasis added). Under these precedents, the responses from an AI platform are not “advice.”
42. *Machado Amadis v. U.S. Dep’t of State*, 971 F.3d 364, 370 (D.C. Cir. 2020) (holding that “recommendations from subordinates to superiors lie at the core of the deliberative-process privilege”).
43. *Coastal States Gas Corp. v. Dep’t of Energy*, 617 F.2d 854, 866 (D.C. Cir. 1980); see also *Kidd v. DOJ*, 362 F. Supp. 2d 291, 296 (D.D.C. 2005) (protecting documents on the basis that disclosure would “inhibit drafters from freely exchanging ideas, language choices, and comments in drafting documents”) (internal citation omitted); *Lewis-Bey v. U.S. DOJ*, 595 F. Supp. 2d 120, 133 (D.D.C. 2009) (protecting documents whose release “would have the effect of inhibiting the free flow of recommendations and opinions”) (internal citation omitted); *Reliant Energy Power Generation Inc. v. FERC*, 520 F. Supp. 2d 194, 205 (D.D.C. 2007) (“Disclosure of internal communications ... can hamper the candid exchange of views ...”).
44. See, e.g., *Petroleum Info. Corp. v. U.S. Dep’t of Interior*, 976 F.2d 1429, 1435 (D.C. Cir. 1992) (finding “release of materials that do not embody agency judgments—for example, materials relating to standard or routine computations or measurements over which the agency has no significant discretion—is unlikely to diminish officials’ candor or otherwise injure the quality of agency decisions”).
45. Elliott, *supra* note 7 (emphasis added).
46. *Cleary, Gottlieb, Steen & Hamilton v. Dep’t of Health & Hum. Servs.*, [844 F. Supp. 770, 782–83 \(D.D.C. 1993\)](#) (holding that “the computer software involved with the Philen Study is within the protected sphere of the deliberative process privilege, because ... the privilege encompasses the decision-making process behind the culling and selection of relevant facts”; the programs thus “reflect[ed] their creator’s mental processes” and revealed her “scientific deliberations and opinions”); *Goodrich Corp. v. U.S. EPA*, 593 F. Supp. 2d 184, 187–89 (D.D.C. 2009) (holding that an incomplete computer modeling program was protected by the deliberative process privilege because the model was “still in development,” and thus was akin to a “draft”); but see *NRDC v. U.S. EPA*, 954 F.3d 150, 158–59 (2d Cir. 2020) (rejecting government’s assertion of the deliberative process privilege where “the core model provides only highly abstract insights into EPA deliberations, which are too far removed from any policy judgments to render the core model deliberative,” and “it contains algorithms instructing the model on how to process the input data. While [it] may in some respects reflect internal agency deliberations, disclosure of its analytical tools cannot reasonably be anticipated to impair the quality of agency decisionmaking.”).

47. *See, e.g.,* Touarsi v. U.S. DOJ, 78 F. Supp. 3d 332, 346 (D.D.C. 2015) (“An individual government officer’s notes may constitute deliberative documents if they are used to assist the decisionmaking process. *See, e.g.,* Baker & Hostetler LLP v. U.S. Dep’t of Commerce, 473 F.3d 312, 321 (D.C. Cir. 2006) (‘notes taken by government officials often fall within the deliberative process privilege’).”).
48. *Coastal States Gas Corp.*, 617 F.2d at 866.
49. *See, e.g.,* Citizens for Resp. & Ethics in Wash. v. U.S. Dep’t of Homeland Sec., 648 F. Supp. 2d 152, 160 (D.D.C. 2009) (requiring release of portion of emails not discussing policy); People for the Am. Way Found. v. Nat’l Park Serv., 503 F. Supp. 2d 284, 301–02 (D.D.C. 2007) (refusing to allow agency to withhold document containing “predecisional guidance relating to upcoming events” because agency had not shown connection to “any type of governmental policy formulation or decision”); Sun-Sentinel Co. v. U.S. Dep’t of Homeland Sec., 431 F. Supp. 2d 1258, 1277–78 (S.D. Fla. 2006) (declining to protect email communications containing advice to agency director because they contained recommendations on press relations, not on matters relating to agency’s “mission”), *aff’d sub nom.* News-Press v. U.S. Dep’t of Homeland Sec., 489 F.3d 1173 (11th Cir. 2007).
50. *NLRB v. Sears, Roebuck & Co.*, 421 U.S. 132, 161 (1975).
51. *See infra* note 55. *See also* Kelso Lorne St. Jacques Anderson, *Communicating with AI Jeopardizes Privilege Claim*, 51 LITIG. NEWS, no. 4, Summer 2026, /content/aba-cms-dotorg/en/groups/litigation/resources/litigation-news/2026/summer/communicating-ai-jeopardizes-privilege-claim/.
.html
52. *Warner v. Gilbarco, Inc.*, 820 F. Supp. 3d 629, 636–37 (E.D. Mich. 2026) (“[T]o the extent Defendants argue that Plaintiff waived the work-product protection by using ChatGPT, the work-product waiver has to be a waiver *to an adversary* or in a way likely to get in an adversary’s hand. ... And ChatGPT (and other generative AI programs) are *tools, not persons*, even if they may have administrators somewhere in the background.”) (internal citations omitted) (emphasis in original); *Morgan v. V2X, Inc.*, No. 25-cv-01991-SKC-MDB, 2026 U.S. Dist. LEXIS 67939, at *10–11 (D. Colo. Mar. 30, 2026) (“It is true that AI systems like ChatGPT, Claude, Gemini, and others widely available to the public, collect user data for training and other purposes. But in this Court’s estimation, that does not eliminate all expectations of privacy or automatically waive protections.”).
53. *United States v. Heppner*, 820 F. Supp. 3d 292 (S.D.N.Y. 2026).
54. *Id.* at 296–97. Judge Rakoff also rejected Heppner’s claim that his prompts to Claude, and that chatbot’s responses thereto, were protected “work product” of his counsel because neither was “prepared by or at the behest of counsel, nor did they reflect defense counsel’s strategy.” *Id.* at 298.
55. *See supra* note 54.
56. As of April 13, 2026, OMB reports 3,611 individually reported AI uses among 46 federal agencies, including 445 “high impact” cases. *2025 Federal Agency AI Use Case Inventory*, GITHUB, <https://github.com/ombegov/2025-Federal-Agency-AI-Use-Case-Inventory/blob/main/README.md> (last visited Aug. 31, 2026). One in five federal agencies reported using generative AI to respond to FOIA requests as of October 2025. *See* Weslan Hansen, *NARA: 1 in 5 Federal Agencies Turn to AI for FOIA Requests*, MERITALK (Oct. 6, 2025), <https://www.meritalk.com/articles/nara-1-in-5-federal-agencies-turn-to->

ai-for-foia-requests/; Ronald L. Capaldi, *FOIA and the Use of AI in Government: Freedom of Information or an Empty Promise of Openness?*, 45 J. NAT'L ASS'N ADMIN. L. JUDICIARY 25, 27–28 (2025).

57. See, e.g., *Fann v. Kemp*, 515 P.3d 1275, 1280 n.1 (Ariz. 2022) (“The core purpose of the public records law is ‘to allow the public access to official records and other government information so that the public may monitor the performance of government officials and their employees.’”) (citation omitted); *Reno Newspapers, Inc. v. Gibbons*, 266 P.3d 623, 626 (Nev. 2011) (“[T]he purpose of the [public records act] is to further the democratic ideal of an accountable government by ensuring that public records are broadly accessible,” which “promote[s] government transparency and accountability.”); *State ex rel. Gannett Satellite Info. Network v. Petro*, 685 N.E.2d 1223, 1227 (Ohio 1997) (“The purpose of Ohio’s Public Records Act ... is to expose government activity to public scrutiny, which is absolutely essential to the proper working of a democracy.”).
58. See, e.g., Cal. St. Bar Standing Comm. on Pro. Resp. & Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* at 4 (2026) (“AI systems are known to produce outputs that may be false, misleading, fabricated, incomplete, or biased, and reliance on such outputs without appropriate scrutiny and review is inconsistent with a lawyer’s professional obligations.”), <https://www.calbar.ca.gov/sites/default/files/portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf>; Conor Murray, *Why AI “Hallucinations” Are Worse Than Ever*, FORBES (May 6, 2025), <https://www.forbes.com/sites/conormurray/2025/05/06/why-ai-hallucinations-are-worse-than-ever/> (noting that “[q]uestions asked outside of the data the AI model knows can lead to the bot responding with incorrect information,” and reasoning models like ChatGPT and OpenAI are “designed to maximize the chance of giving an answer, meaning the bot will be more likely to give an incorrect response than admit it doesn’t know something”).

Author



Steven D Zansberg
Zansberg Beylkin LLC

Steven D. Zansberg, is a founding partner of Zansberg Beylkin LLC in Denver. Mr. Zansberg represents news media organizations in seeking access to judicial

proceedings, records and open meetings,
libel defense, asserting the...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association |

https://www.americanbar.org/groups/communications_law/resources/magazine/2026-fall/ai-accessible-under-public-records-laws-foia/