An AI-Orchestrated Architecture for Responding to FOIA Requests

Jason R. Baron
College of Information
University of Maryland
College Park, MD USA
jrbaron@umd.edu

Devanshu Kejriwal
A. James Clark School of Engineering
University of Maryland
College Park, MD USA
devkej06@umd.edu

ABSTRACT

In the United States, at every stage of the process for handling requests filed under the Freedom of Information Act (FOIA) and corresponding state laws, the process almost entirely involves human interaction and review. The goal of replacing the "as is" model of FOIA processing with an AI model is to eliminate obvious bottlenecks due to labor-intensive actions performed by FOIA staff. In this paper we propose a high-level AI-orchestrated architecture for FOIA processing, in which five AI sub-routines each have their own distinct FOIA task to perform, including: an intake interface, a search protocol, a review protocol for filtering/redacting exempt material, an explanation protocol for the requester to understand what decisions have been made on withholdings, and software that functions to audit the process. At each stage, the process retains a measure of human-in-the-loop review, for the purpose of auditing, quality control, sampling, and process improvements, and other interventions as needed.

KEYWORDS Freedom of Information Act, AI-orchestrated, user experience (UX), chatbot, generative AI, technology assisted review, sensitivity review, FOIA exemptions, agentic AI

Reference: *In: Proceedings of the Artificial Intelligence & Open Government Workshop (AIOG 2026), held in conjunction with the 21st International Conference on Artificial Intelligence and Law (ICAIL 2026), June 8, 2026, Singapore*

1. INTRODUCTION

The process for responding to freedom of information requests for access to documents at all levels of the US government is largely human review-centered. From the initial intake of requests either through an online portal, email address, or U.S. mail, to the collection of putatively responsive documents, through to making redactions and drafting determination letters on what documents (in whole or in part) are being released or withheld, human actors are heavily involved. Searches are still largely conducted by keyword searching to find responsive records, a method that has for over the past 15 years been found by the legal community to be both flawed and inefficient for more complex requests.

At the federal level in the US, in 2025 over 1,200,000 requests were filed pursuant to the Freedom of Information Act (FOIA) [1]. The FOIA provides for agencies to respond to requests within 20 working days, subject to limited exceptions [2]. However, due to the resulting extensive backlogs in processing the expansive volume of requests received, it is not unusual for federal agency FOIA officers to file declarations under oath in court cases that an agency will take years to respond to FOIA requests requiring extensive searches. Reported examples include: a 2017 case involving Federal Bureau of Investigation (FBI) records where the agency stated that it would take them 17 years to respond [3], and in 2021, a FOIA officer of the Food and Drug Administration declared that it would take his agency 55 years to review an estimated 100,000 documents [4].

Employing AI orchestration to cover the series of discrete tasks involved in processing FOIA requests holds out the potential to (i) greatly expedite the process, (ii) improve the user's experience (UX) in interacting with an agency; (iii) provide better explanations as to the decision on releasing or withholding documents (or parts thereof) to the requestor; and (iv) result in more comprehensive responses.

As described below, all but one of the sub-tasks involves AI tools already in place, even if not directed to the FOIA process. From the use of chatbots with generative AI capabilities, to adopting search methods using technology assisted review, to drafting determination letters explaining to the requester why the agency acted as it did, agencies either are capable of or soon could be expected to be capable of adopting all of the above to make the FOIA process more seamless. The one current exception is using an AI tool to perform an adequate sensitivity analysis to filter documents for FOIA exempt material. However, progress is being made in this area as well.

2. RELATED WORK

So far as the authors are aware, to date there has not been research specifically discussing how an AI-orchestrated, multi-stage process, with or without agentic AI capabilities, could be fashioned to improve current FOIA processing. This position paper is intended as a complement a FOIA AI model previously discussed in a non-technical opinion piece by the lead author [5]. More generally, the subject of applying AI tools in US civil discovery has been the subject of recent discussion [6, 7]. With respect to the individual AI sub-tasks described here, some have been studied more than others. Chatbots with embedded generative AI capabilities have been discussed widely [8, 9]. Machine learning (ML) in the form of "technology assisted review" (TAR) as a replacement for keyword searching has been accepted by the community for the past 15 years [10, 11]. Sensitivity analysis studies have been undertaken using a variety of ML classifiers. [12, 33]. The use of machine learning and gen AI tools to perform sensitivity analysis in the form of identifying FOIA exempt material has been studied more recently [12, 13, 14, 15], and To the extent AI orchestration as envisioned here may evolve into a true agentic AI process, there is an additional growing literature on the application of agentic AI to public sector processes, including frameworks for how agentic systems will transform core government functions [27, 28, 29] and discussions of the governance challenges such systems pose for public sector organizations in the absence of human review. [30, 31, 32].

3. AN AI MODEL ARCHITECTURE

The goal of replacing the "as is" model of FOIA processing with an AI-orchestrated model is to eliminate obvious bottlenecks caused by human review, where in the typical agency the FOIA office is underfunded and lacks the human capacity to efficiently process incoming requests and current backlogs.

Figures 1 provides a schematic overview of the proposed multi-stage AI model. As shown on Figure 1, the following five AI sub-routines each have their own distinct FOIA task to perform, including: an front-stage intake interface that will constitute the user's experience (U/X) with the agency throughout the FOIA decision making process (AI Stage 1); and on the back-end, a collection, AI search, and preservation protocol (AI Stage 2); a review protocol for filtering/redacting sensitivities consisting of FOIA exempt material (AI Stage 3); a decision and explanation protocol (AI Stage 4); and AI software that audits the entirety of the process (AI Stage 5).

As described below, only in the simplest requests are the stages actually carried out in linear fashion. Rather, in the case of larger and more complex requests, we believe that the orchestration model is best carried out iteratively, best conceived of as a series of "loops," as set out in Figure 2, involving the requester and AI to narrow and refine a request en route to the requester receiving a final determination with respect to documents released and documents withheld (in whole or in part).

AI Stage 1 consists of an initial agency FOIA Chatbot set up as the point of contact for the requester to initiate requests and engage in a subsequent dialogue with the agency throughout the FOIA process. The chatbot would rely on the latest version of a large language model (LLM) for the purpose of holding a dialogue with FOIA requesters. Intake would consist of a request for records on a given topic, likely triggering queries from the chatbot to better understand the request being made and resolve ambiguities in language. Loop 1 in figure 2 depicts the chatbot reaching consensus with a requester in setting the scope of clear parameters (e.g., names or positions of likely custodians or organizational components; dates of record creation).

The FOIA chatbot would serve as the continuing interface with requesters for the purposes of providing interim and final determinations on release and withholdings, as well as responding to questions and additional input. This process could include the identification and execution of any necessary consultations with or referrals to other agencies with notice to submitters of requests.

At AI Stage 2, putatively responsive records would be collected, searched, and preserved on agency networks and in agency repositories. Depending on the volume and complexity of the incoming request, the AI software would be tasked to perform either keyword searches using Boolean operators, or machine learning in the form of TAR. As depicted in Loop 2, the FOIA Chatbot would inform a requester of the results of the search, for the purpose of further narrowing or clarifying search parameters in cases where search results have turned up a voluminous amount of records.

Stage 3 AI orchestration would consist of performing a sensitivity analysis once the narrowed parameters of the document set have been agreed to by all parties. As conceived here, this consists of using machine learning classifiers to segregate out FOIA exempt material. In the U.S., aside from collections involving national security records on stand-alone networks, the vast majority of exempt material falls into one of the following categories: (i) Exemption 5, consisting of documents that in whole or in part contain attorney-client communications, attorney work product, or material subject to the deliberative process privilege; (ii) Exemption 6, consisting of personal information, including but not limited to numerical forms of PII; (iii) Exemption 4, involving trade secrets and other forms of confidential information; and (iv) Exemption 7, requiring redactions to protect law enforcement investigations.

At AI Stage 4, a generative AI application would be employed to explain in narrative form the results of the entirety of any review conducted, to deliver to the FOIA chatbot for purposes of communicating to the requester the result of the FOIA search and review. The AI-generated narrative would enhance the explanation process by following a best practices model for the formatting of determination letters as advocated by the FOIA Advisory Committee to the US Archivist [16]. This would consist of a full explanation of how records were searched, what AI processes were used, why certain documents were withheld in full or in part, and providing additional information pertinent to the requester including appeal rights.

For the purpose of providing a window of transparency to requesters on the results of a review and redaction process at Stage 3, we conceive of Stages 3 and 4 as optimally being carried out as part of a sampling process, with a Loop 3 back to the requester after a small sample of responsive documents have been reviewed. The sample could be sent after Stage 3 without an explanation in detail, but ideally would include an interim determination letter with an explanation of withholdings or redactions undertaken. It would be up to the agency to set rules in place on the size of the sample and how many samples would be undertaken before moving on to a full process of reviewing remaining documents and providing a final determination letter, pursuant to Stages 3 and 4. After the final determination letter is provided by the FOIA chatbot to the requester, the current requirements for filing a further administrative appeal would apply. If an appeal is taken, as shown in Loop 4 there would be a further review at Stage 3 and a second determination letter sent at Stage 4.

Finally, we provide for an AI Stage 5, where the AI orchestration role would entail oversight of the process employed in AI Stages 1 through 4, through the accumulation of metrics. At each stage the process anticipates some measure of human-in-the-loop review, for the purpose of auditing, quality control, sampling, and process improvements, as well as requested by the individual filing the FOIA request. However, we anticipate that only in unusual cases would there be a need for eyes-on review on a document-by-document level prior to a final decision being rendered.

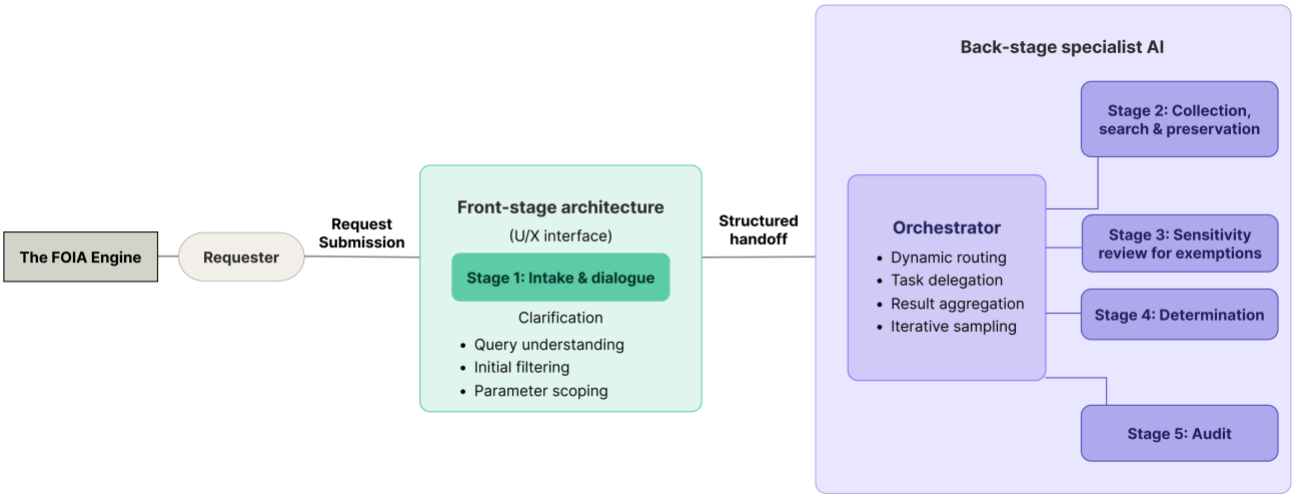

Fig. 1. An AI-Orchestrated Architecture

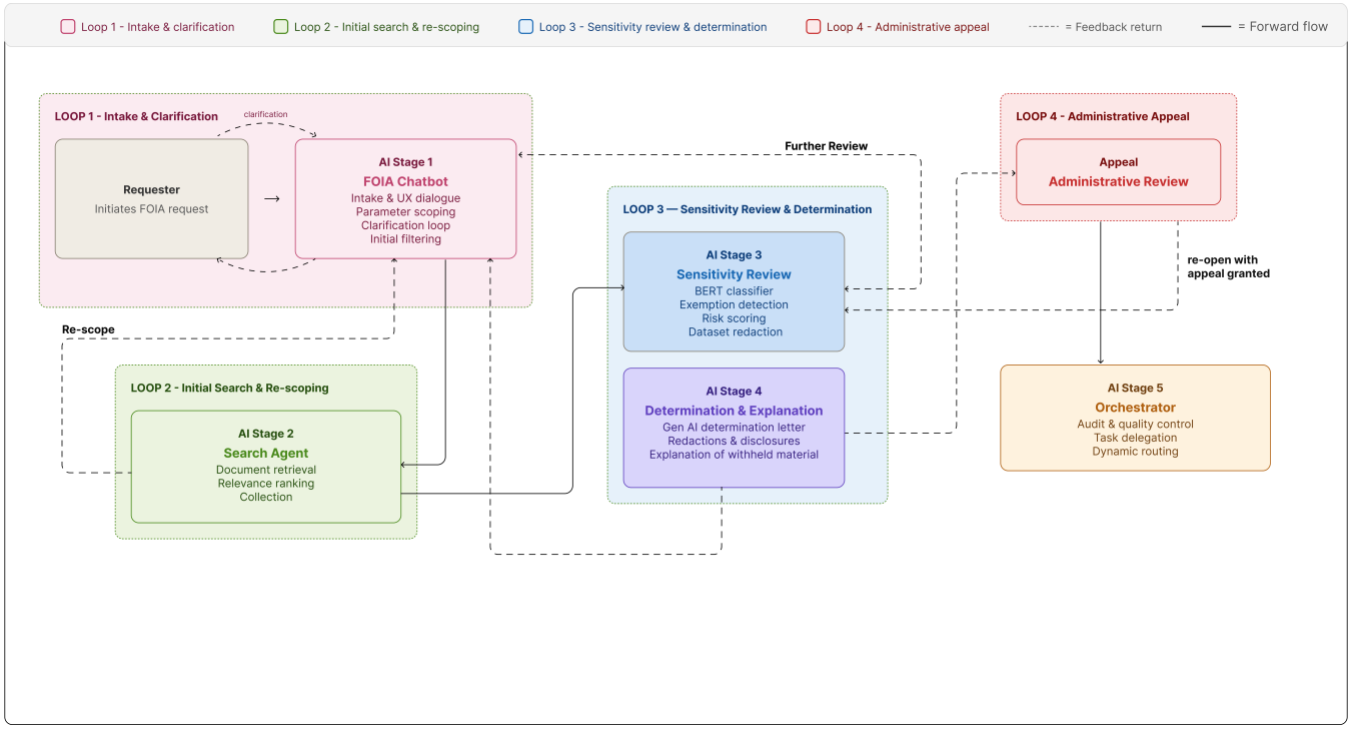

Fig. 2. Iterative Loops in the Review Process

3.3 Software Components.

Subject to limitations described in section 4, the above tasks can be largely deployed using existing AI software. Table 1 illustrates representative software applications presently available to accomplish a seamless AI process. The list is not intended to be exhaustive and we expect will be superseded in small or large part in short order.

Table 1. Representative software for multi-agent FOIA processing.

Agent	Technology	Role in FOIA Processing
Stage 1: Intake & U/X Dialogue	Large Language Model API (e.g., GPT-5, Claude)	Powers natural language understanding and dialogue. Enables the chatbot to converse with the requester, interpret free-text FOIA requests, and ask clarifying questions in natural language.
	spaCy (NER)	Named entity recognition system that extracts structured parameters such as names of custodians, agencies, locations, and date ranges from unstructured user input.
	Redis	Session cache that stores conversation history, ensuring contextual continuity across multiple dialogue turns.
Stage 2: Collection, Search & Preservation	Vector databases (Pgvector, FAISS, Milvus, Pinecone, Weaviate, Qdrant, Chroma)	Semantic search systems that retrieve conceptually similar documents using vector embeddings, even when keyword overlap is minimal. Open-source options include FAISS and Chroma; managed cloud options include Pinecone, Weaviate, and Qdrant. Selection will depend on deployment environment and scale.
	Relational databases (PostgreSQL, Microsoft SQL Server, Oracle)	Primary relational database storing document metadata, retrieval results, and ranking scores for downstream processing. PostgreSQL is the open-source baseline; Microsoft SQL Server and Oracle are the major proprietary alternatives widely deployed in government environments.
	Preservation layer	Ensures collected records are preserved in place and protected from alteration or deletion for the duration of the FOIA process.
Stage 3: Sensitivity Review	BERT Classifier (Legal-BERT fine-tuned) [25]	Machine learning classifier that evaluates whether document segments fall under FOIA exemption categories, enabling automated sensitivity detection.
	Triton Inference Server	GPU-accelerated inference engine that enables scalable, high-throughput execution of the BERT classifier.
	PII detection Regex + LLM-based approaches	Regex (rule-based pattern matching) provides a first-pass method to detect structured PII such as SSNs, phone numbers, and passport numbers.
Stage 4: Determinations	Large Language Model API + RAG (LlamaIndex) [26]	Generates FOIA determination letters with explanations for withheld material, ensuring responses are consistent with FOIA case law and prior decisions.
Stage 5: Audit	Go / Asynq (Orchestration Layer)	Coordinates multi-agent workflow execution, manages task routing, and escalates cases to human reviewers as requested or needed.
	PostgreSQL (Append-only Audit Log)	Immutable logging system that records all decisions, confidence scores, and provides for human overrides for compliance, transparency, and legal auditability.

4. THE PRESENT-DAY STATE OF MODEL COMPLETENESS

The above-described AI-orchestrated architecture contains features that range from easy to difficult to implement given the present state of AI technologies. The following consist of "pain points" that need to be worked out in considering deployment of the described model.

(1) Rule-based chatbots already are in present-day use by government agencies in dealing with members of the public [17, 18]. Incorporating generative AI features into chatbots is not difficult; however, ensuring that agency chatbots stay within reliable boundaries in the queries they ask and the language they use is still not without risk. There will be a need for human supervision of any FOIA chatbot to ensure that conversations remain within the bounds of what members of the public would expect when dealing with a human agency representative.

(2) The use of TAR methods for purposes of searching for responsive records is still an esoteric domain for the vast majority of FOIA professionals. It is conceded that the larger majority of FOIA requests will continue to be small or modest in size,

necessitating only that agency staff perform keyword searching to meet their obligations in performing an "adequate search." Nevertheless, as the size of electronic record repositories increase —most notably, so-called "Capstone email archives" alone amounting to tens or hundreds of millions of emails preserved [19]—no question exists that FOIA searches in response to even narrowly tailored requests will grow in size to a point where TAR methods are reasonable to carry out.

(3) Using AI tools to filter for sensitivities is still in the early stages of development. Rule-based pattern matching using regular expressions ("regex") remains the most common approach for finding structured forms of PII, such as Social Security numbers, phone numbers, and passport numbers, in records under review for FOIA [20]. While regex works for structured and predictable PII, other contextual, unstructured or non-standard forms of sensitive personal information are commonly found in government records. More work is required to explore other methods including applying LLM-based methods and other traditional NLP techniques that are more suitable for detecting sensitive information in unstructured text.

Methods for filtering for other types of exempt material beyond personal information are less developed. Recent work has had some success (F1 scores over 70%) in filtering parts of documents that are subject to the deliberative process privilege, as mentioned in section 2 [13, 14]. There are no known "one size fits all" AI tools currently on the market that attempt to filter for *all* US FOIA exemptions.

(4) The state of generative AI software's ability to draft FOIA determinations has been the subject of only one research paper to date [15]. That 2023 research, using ChatGPT-3.5, found that the quality of explanations in AI-generated narratives concerning why documents in whole or in part were withheld were on a par with at least many present-day agency determination letters, which themselves often display striking deficiencies. Further research is needed to determine how well the present-day versions of ChatGPT, Claude, Perplexity and other generative AI applications perform under similar tests.

(5) One element of the US FOIA statute is the requirement that an agency apply a "foreseeable harm" standard when deciding whether the Exemption 5 "deliberative process privilege" exemption is applicable. Courts have held that an agency is required to establish "foreseeable harm" with sufficient, concrete justification to continue to withhold such material, even if it is properly categorized as pre-decisional and deliberative [21]. At present, the authors are unaware of any attempt having been made to train an AI application to recognize "foreseeable harm," although it is conceivable to imagine training AI tools based on past agency determinations with respect to similar material. At bottom, this is the kind of contextual determination that humans still hold greater expertise in than algorithmic methods.

5. THE AI "BLACK BOX," AGENTIC AI, AND HUMAN-IN-THE-LOOP REVIEW

The authors are aware that a great deal of public skepticism exists with respect to "handing over" to AI making final decisions with respect to government services and benefits. And specifically to the FOIA, some have expressed deep concerns that without some substantial level of oversight—including perhaps by external actors to an agency's workforce—the government may have reasons to bias FOIA determinations when the algorithmic methods are known and controlled only by agencies themselves [22]. As Capaldi [23] has stated:

> The role of AI in such sensitive, high-stakes determinations raises critical concerns about transparency and fairness; there is a growing tension between the efficiency AI provides and the transparency required by FOIA. As the system evolves, the need for accountability in AI-driven decisions grows more urgent, especially as human oversight diminishes.

The evolution from the above-described AI-orchestrated workflow to one that incorporates elements of true autonomous decision making in the form of agentic AI should take into account these concerns.

Examples of agentic AI decision making could arise in the area of a chatbot making decisions on whether to refer requests to other agencies for conferral; AI software choosing which of a variety of search methodologies might be employed in finding responsive documents, depending on the nature, complexity, and voluminous of the requests; and AI making decisions on whether to offer requesters further opportunities to refine requests based on the review of interim samples.

A core design requirement is that the FOIA chatbot keeps a requester continuously informed and updated with respect to the agency's FOIA processing. To the extent an agentic AI process is employed to manage the administrative task of executing multi-step workflows with the iterative loops described, doing so will lead to more efficient processing of requests. The additional subtasks performed by AI tools contemplate the option of human review, to the extent an agency wishes to embed an appeal mechanism where requesters are asking for a "real" agency employee to intervene. There should always be a mechanism (embodied in AI Stage 5) for human review on a "system" level in performing continuous evaluation of the quality of the process, with audits, controls, samples, and data analysis, as resources permit.

"Explainability" can be further enforced at the database level through preserved logs representing an immutable audit event stream of all actions taken during the processing of a given FOIA request, as illustrated in Figure 3. The diagram depicts how the AI System, comprising the Orchestrator and AI Stage 1 through software, create structured log events to an append-only logging layer at every step of FOIA processing. Each event record captures the request identifier and timestamp, the specific agent and model version responsible for the action, snapshots of both input and output, and any human override decisions along with their stated rationale. These events can be written permanently to a PostgreSQL database configured with append-only logic, meaning no record can be modified or deleted after it is created. The result is a fully auditable, tamper-resistant chain of custody for every decision made during the lifecycle of a FOIA request, providing both internal oversight and a legally defensible record of the AI system's conduct.

We wish to also note that the "as is" model of human review in FOIA is not without its own considerable "black box" elements, given the degree to which FOIA requesters routinely experience delays due to backlogs, frustrations when attempting to engage in a dialogue with agencies, and material gaps in the explanations received in determination letters. What an AI-orchestrated

process offers in the alternative is open dialogue with agencies at every key stage in the process, coupled with user-friendly explanations with respect to each decision made, all in an expedited fashion.

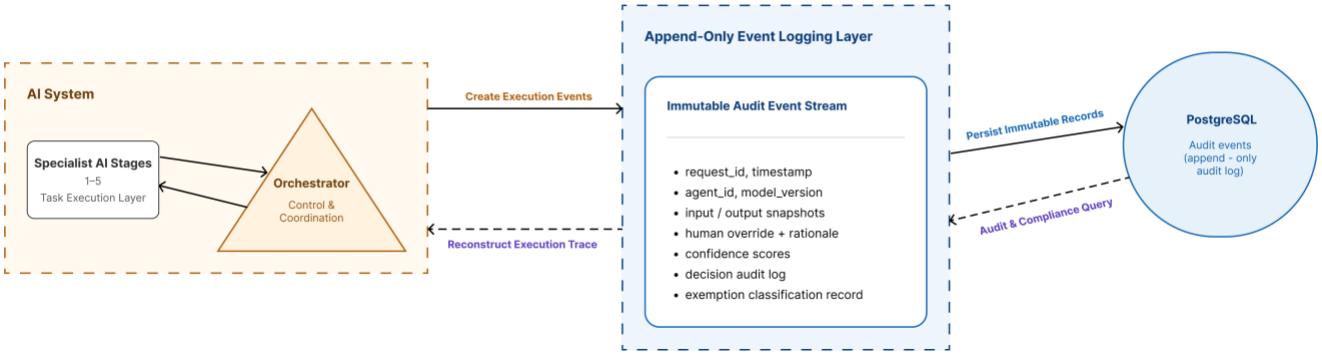

Fig. 3. Immutable Audit Trail

6. CONCLUSION

This paper has had the modest aim of sketching out a "first draft" in modeling what an AI-orchestrated process might look like for FOIA processing. Without direct access to a governmental agency's record repository and its FOIA staff, the authors are not presently in a position to directly perform a full-scale research experiment involving linking all of the discrete elements referenced in the proposed model. Perhaps conducting such research is best left to agency personnel themselves.

The present need for AI substitutes to increasingly archaic, cumbersome, and monumentally frustrating human-review centered processes used by governmental employees only gets greater with each passing day. NARA currently holds over 600 million White House emails from the past 30 years, constituting billions of pages [24]. The over 400 agency components of the Executive branch in the US and their corresponding email and total repositories of electronic records vastly dwarf even the accumulation of White House email to date. No requester deserves being informed by an agency that it will take years or even decades to receive a response to a legitimate FOIA request made in good faith.

Our proposed model, elements of which are largely in place with AI improvements on the horizon, is an effort to solve a systemic problem in the administration of freedom of information laws. In the interest of government accountability, the government's commitment to investing in the resources to accomplish an AI-orchestrated FOIA process remains worthy of serious consideration.

AI USE DISCLOSURE

Generative AI tools were used to assist with formatting and summaries of the table and figures. All outputs were reviewed and verified by the authors.

ACKNOWLEDGEMENTS

We wish to thank Eric Caplan, David Cooper, Deborah Moore, colleagues at the Center for Archival Futures, and anonymous reviewers for helpful comments and suggestions.

REFERENCES

[1] 2024 Annual FOIA Report Summary, U.S. Dep't of Justice (2025). https://www.justice.gov/oip/media/1398111/dl?inline.
[2] 5 U.S.C. § 552(a)(6)(A)(i) & (ii)(I); *id.*, § 552(a)(6)(B).
[3] J. Gerstein 2017. Judge Balks at FBI's 17-year FOIA timeline for FOIA request. Politico. https://www.politico.com/blogs/under-the-radar/2017/07/29/judge-balks-fbi-foia-timeline-17-years-241127.
[4] J. Greene 2021. Wait what? FDA wants 55 years to process FOIA requests over vaccine data. Reuters. https://www.congress.gov/117/meeting/house/114270/documents/HHRG-117-GO24-20211201-SD014.pdf.
[5] J.R. Baron, E. Wilczek 2026. FOIA's Future: Agentic AI's Potential to Transform the FOIA Requester eXperience. In: SunshineFest 2026 Conference, Washington, D.C. https://sunshineweek.org/wp-content/uploads/2026/03/AI-Baron-and-Wilczek.pdf.
[6] D. Austin 2026. The Promise and Limitations of Agentic AI: Artificial Intelligence Best Practices. E-discovery Today https://ediscoverytoday.com/2026/02/04/the-promise-and-limitations-of-agentic-ai-artificial-intelligence-best-practices/
[7] G. Foster 2026. The Rise of AI Agents: The Next Big Thing in E-discovery. Alvarez & Marsal. https://www.alvarezandmarsal.com/thought-leadership/the-rise-of-ai-agents-the-next-big-thing-in-ediscovery.
[8] M. Kim 2024. Unveiling the e-Servicescape of ChatGPT: Exploring User Psychology and Engagement in AI-Powered Chatbot Experiences.14 Behavioral Science 558. https://doi.org/10.3390/bs14070558.
[9] J. Ng 2026. A structured framework for effective and responsible generative artificial intelligence chatbot prompt engineering throughout the scientific process: a comprehensive guide for the health and medical researcher. Front. Artif. Intell., Sec. Medicine and Public Health. https://doi.org/10.3389/frai.2026.1745928.

[10] M.R. Grossman & G.V. Cormack 2011. Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review. Rich. J.L. & Tech. 17:3, art. 5. https://scholarship.richmond.edu/cgi/viewcontent.cgi?article=1344&context=jolt.
[11] Da Silva Moore v. Public Groupe, 287 F.R.D. 182, 193 (S.D.N.Y. 2012) (Peck, M. J.), *aff'd,* 2012 WL 1446534 (S.D.N.Y. Apr. 26, 2012).
[12] G. McDonald 2019. A Framework for Technology-Assisted Review: Using Sensitivity Classification to Prioritise Documents for Review, Ph.D Dissertation, University of Glasgow. http://theses.gla.ac.uk/41076.
[13] Karl Branting, B. Drown, C. Giannella, et al. 2023. Automated Detection of Sensitive Content in Government Records, Artif. Intell. and Law 31. https://link.springer.com/content/pdf/10.1007/s10506-023-09383-6.pdf.
[14] J.R. Baron, M. Sayed, D. Oard 2022. Providing More Efficient Access to Government Records: A Use Case Involving Application of Machine Learning to Improve FOIA Review for the Deliberative Process Privilege. J. on Computing and Cultural Heritage 1, art. 5. https://dl.acm.org/doi/abs/10.1145/3481045.
[15] J.R. Baron, N. Rollings, D. Oard 2023. Using ChatGPT for the FOIA Exemption 5 Deliberative Process Privilege. In: Proceedings of ICAIL 2023, AI and Intelligent Assistance Workshop, June 2023, Braga Portugal. https://ceur-ws.org/Vol-3423/paper4.pdf.
[16] Final Report of the 2022-2024 FOIA Advisory Committee to the Archivist, Recommendation 2023-01. National Archives and Records Administration. https://www.archives.gov/files/ogis/documents/finalreport.6.17.24.pdf.
[17] Internal Revenue Service (n.d.). Chatbot accessibility guide. https://www.irs.gov/help/irs-chatbot-accessibility-guide;
[18] U.S. Dep't of Veterans Affairs (n.d.). VA Chatbot. https://www.va.gov/contact-us/virtual-agent/.
[19] National Archives and Records Administration 2014. The Capstone Approach and Capstone GRS. https://www.archives.gov/files/records-mgmt/email-management/final-capstone-white-paper.pdf.
[20] Relativity (n.d.). Personal Information Detectors. https://help.relativity.com/RelativityOne/Content/Relativity/aiR_for_Data_Breach_Response/Personal_Information_detectors.htm
[21] Reporters Committee for the Freedom of the Press et al. v. FBI, 3 F.4th 350 (D.C. Cir. 2021).
[22] L. Kamb 2023. Some U.S. government agencies are testing out AI to help fulfill public records requests, NBC News. https://www.nbcnews.com/news/us-news/federal-agencies-testing-ai-foia-concerns-rcna97313.
[23] R. Capaldi 2025. FOIA and the Use of AI in Government: Freedom of Information or an Empty Promise of Openness? J. Nat'l Ass'n Admin. L. Judiciary 45:2. https://digitalcommons.pepperdine.edu/naalj/vol45/iss2/2.
[24] D. Hoffman 2024. A digital tsunami is coming. The National Archives is in trouble. Washington Post. https://www.washingtonpost.com/opinions/2024/09/20/national-archives-troubles-digital/.
[25] J. Devlin, M-W. Chang, K. Lee, K. Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT 2019): 4171 – https://aclanthology.org/N19-1423/.
[26] P. Lewis, E. Perez, A. Piktus, et al. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In: Advances in Neural Information Processing Systems 33 (NeurIPS 2020), Vancouver, Canada. https://proceedings.neurips.cc/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf.
[27] C. Schmitz, J. Rystrøm, J. Batzner 2025. Oversight Structures for Agentic AI in Public-Sector Organizations. In: Proceedings of the 1st Workshop for Research on Agent Language Models (REALM 2025), 298. Association for Computational Linguistics. Vienna, Austria. https://aclanthology.org/2025.realm-1.21/.
[28] V. Almeida 2026. Governing AI Agents with Democratic 'Algorithmic Institutions.' Harvard Berkman Klein Center. https://www.techpolicy.press/governing-ai-agents-with-democratic-algorithmic-institutions/
[29] M. Abou Ali, F. Dornaika, J. Charafeddine 2026. Agentic AI: a comprehensive survey of architectures, applications, and future directions. Artif. Intel.l Rev 59, 11. https://doi.org/10.1007/s10462-025-11422-4.
[30] S. Passi 2025. Agentic AI has a Human Oversight Problem. http://dx.doi.org/10.2139/ssrn.5529058.
[31] L. Zhu, Q. Lu, Q., M. Ding, *et al.* 2026. Designing meaningful human oversight in AI. AI Ethics 6, 286. https://doi.org/10.1007/s43681-026-01147-7
[32] I. Sprongl 2026. Governing Agentic AI: Designing Authority in Systems that Act Without You. http://dx.doi.org/10.2139/ssrn.6221439.
[33] M. Sayed, B. Mallekav, D. Oard 2022. Comparing Intrinsic and Extrinsic Evaluation of Sensitivity Classification. In: Advances in Information Retrieval: 44th European Conference on IR Research (ECIR 2022), Stavanger, Norway. https://doi.org/10.1007/978-3-030-99739-7_25.