



Emails to Algorithms Public Records Challenges in the AI Era

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

I. Introduction

The expansion of the internet and proliferation of electronic applications, messaging platforms, tools, and data have fundamentally altered the public records landscape. As technologies have advanced, public entities and officials across the country have started to integrate messaging applications such as Microsoft Teams, Signal, and WhatsApp to aid in efficient communication in the course of their duties. In recent years, artificial intelligence tools, such as ChatGPT, have also rapidly expanded their capabilities and become widely available to the public. Some public entities and employees have started using this technology to assist with the everyday functions of government. Chat and AI tools are thus starting to impact public entities in many of the same ways that email did in the early 2000s – namely, increasing the amount of data affiliated with public business and creating new categories of records potentially subject to disclosure under the California Public Records Act (“CPRA”).

While public entities have become accustomed to maintaining and producing electronic records like emails, the rapid development of Artificial Intelligence (“AI”) and increased use of encrypted messaging applications pose a variety of new challenges. On the one hand, AI, messaging applications, and workflow technologies present the opportunity for public agencies to increase efficiency and boost their overall operations. On other hand, these technologies can burden public entities by increasing the number of public records subject to disclosure under the CPRA, the amount of data that must be retained by an entity, and the frequency at which public officials must search their personal devices and accounts for public records

This paper explores the novel issues that public entities can expect to face under the CPRA in this new technological era, from determining whether AI prompts constitute public records to best practices with messaging apps like Signal. First, we examine whether and when records generated with AI or other technologies are public records. Second, we analyze the CPRA exemptions that may apply to public records created with AI or other types of applications, including the potential waiver of privileges when utilizing those technologies. Finally, we will discuss the potential benefits of AI technologies for reducing the burden of CPRA requests and give guidance on best practices and policies for managing the potential pitfalls of AI and other applications.

II. Understanding What Constitutes a Public Record

Public entities are subject to the CPRA. Its purpose is to ensure transparency in governance by granting members of the public the right to copy and inspect public records. Thus, any introduction of new technology into operations requires public entities to ask whether that technology creates a new kind of public record subject to disclosure and, if so, how can the public entity revise its retention policies to ensure documents are being properly preserved? With AI transcripts in particular, the volume of data can be substantial which raises serious issues of storage capacity and cost. This makes data storage an increasingly important consideration for public entities and highlights the need for clear, technology-use and deletion policies targeted at

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

reducing the retention of unnecessary data. Such policies are critical to ensure that public entities are not overwhelmed by the cost of retaining and reviewing vast amounts of data.

While record retention and disclosure of public records are separate issues¹, the interplay between record preservation and the policy goals of the CPRA mean that both are important considerations when developing use and retention policies for messaging apps and AI. The amount of data retained and the length of time data is retained determines the pool of potential public records because the content of the record, not its format determines whether something is a public record under the CPRA. This section therefore analyzes the extent to which chat and AI records may be “public records” for purposes of the CPRA. It also discusses appropriate retention practices for those records to ensure adherence with applicable laws while managing data volume.

A. Public Records Under the CPRA

1. AI Interactions as Public Records

AI applications, such as ChatGPT, Microsoft Copilot, Gemini, and Claude, are relatively new tools with rapidly growing capabilities. These AI assistants typically involve a human user: (1) inputting requests for information; (2) requesting a synthesis of information; (3) asking for guidance for how to address a particular issue; or (4) requesting help with complex tasks such as translating or coding (collectively “AI Interactions”). Given the infancy of these technologies and their use in public governance, California courts have not yet addressed whether public employees’ AI Interactions constitute public records under the CPRA. The CPRA broadly defines public records as “any writing containing information relating to the conduct of the public’s business prepared, owned, used, or retained by any state or local agency regardless of physical form or characteristics.”² In light of this definition and the CPRA’s general preference for transparency, it is advisable to consider AI Interactions, which are created and used by an entity, as public records if and when those exchanges relate to public entity business. Further, notes or transcripts from meetings related to public entity business created by AI assistants (such as the automatic note-taker in Zoom or other apps) would also likely qualify as public records.³

¹ Indeed, the California Supreme Court recently confirmed that the CPRA does not impose any standalone retention requirements, as it is solely a law regulating disclosure and access to public records. (*City of Gilroy v. Superior Court of Santa Clara County* (Cal., Jan. 15, 2026, No. S282937) 2026 WL 111709.)

² Gov. Code § 7920.530(a).

³ Public agencies have started to generate policies and guidance for use of AI applications which appropriately categorize information inputs and outputs from AI tools as public records. *See* City of San Jose Generative AI Guidelines, last updated April 24, 2025, <https://www.sanjoseca.gov/home/showpublisheddocument/100095/638255600904303329>; City of Long Beach Generative AI Guidance, <https://longbeach.gov/globalassets/smart-city/media-library/documents/generative-ai-guidance-v1-1>, p. 2.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

Several factors can inform whether AI Interactions relate to the public's business, including: (1) the subject matter of the exchange; (2) the context and purpose of the interaction; and (3) whether an employee was acting in the course of their employment when utilizing the AI (e.g., an employee utilizing ChatGPT to brainstorm policy ideas or create a meeting agenda). These inquiries shed light on whether an AI Interaction falls within the definition of "public record" under the CPRA as they assist with understanding whether a record relates to public business in a "substantive way," which is a fundamental requirement for something to be deemed a public record.⁴ Thus, AI Interactions should be considered public records when the content relates substantively to public business. This is particularly true when AI is used to assist with the creation of other important public documents such as memorandums, staff reports, or contract documents. While an AI Interaction is often not the final product, the CPRA's expansive scope and the public's interest in how important government decisions are made both warrant treating AI Interactions as public records unless the circumstances make clear that a particular AI Interaction is not public (e.g., when AI is clearly used for some personal purpose) or until California courts hold otherwise.

2. Chat Messages as Public Records

California courts have made clear that records from messaging apps can qualify as public records subject to disclosure under the CPRA. These records (such as Microsoft Teams or Slack messages) become public records when they substantively relate to the conduct of public business.⁵ When determining whether private messages relate to public business, courts will look to: (1) the subject matter of the messages; (2) the context or purpose of the messages; (3) whether the author was acting in the scope of public employment; and (4) the audience of the messages.⁶ If a message is found to relate to public business under these factors, it is subject to disclosure, even if the message is stored on a private account or device.⁷

Importantly, any type of chat message can qualify as a public record under this test. Even chats exchanged on ephemeral messaging apps with auto-deletion settings, like Signal, can constitute a public record. In light of general retention requirements and the CPRA, auto-deletion settings on such apps are discouraged. The public may even request screenshots of the settings page of Signal accounts for public officials. While there is no clear authority stating that app settings on private devices are subject to disclosure, courts could reasonably determine that messaging app settings are public records which relate to the conduct of the public's business, in particular because the public has an interest in communications by public officials especially

⁴ *City of San Jose v. Superior Court* (2017) 2 Cal.5th 608, 618 ("[T]o qualify as a public record under CPRA, at a minimum, a writing must relate in some substantive way to the conduct of the public's business.")

⁵ *Id.* at 629.

⁶ *Id.* at 618.

⁷ *Id.* at 614.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

when artful transparency loopholes like Signal exist. All chat-related records substantively related to public business should therefore be considered as potential public records.

3. Metadata as a Public Record

Electronic records, including chat messages and AI transcripts, often have associated metadata. Metadata is the “data about data,” meaning that it is not visible in the text of a record but rather is behind-the-scenes information that provides details and context for the data it describes. Metadata includes information about the author of a document, document creation date, timestamps, device type, document permissions, etc. Metadata is an important consideration when disclosing and retaining electronic records, as it can contain sensitive information, provide crucial detail for locating records, or even house the data necessary for preserving the digital file.

Unfortunately, no provision of the CPRA explicitly addresses metadata, and there have been no reported California opinions on the extent to which metadata must be disclosed. Other jurisdictions have held that metadata is a public record subject to disclosure if no exemptions apply.⁸ Moreover, the CPRA does provide that electronic records must be produced in native format, assumingly with metadata intact, if so requested (i.e., producing a native Word document that is in possession of the entity instead of producing a PDF printout of the same document).⁹ In absence of direct law on the issue, the prevailing guidance is for public entities to treat requests for metadata the same as all other requests and disclose nonexempt metadata when requested.¹⁰ If a requester does seek metadata, it is important to review any responsive metadata to ensure that private data, system security information, or other sensitive information is redacted prior to release if such material is not responsive to the request or qualifies as exempt under the CPRA.

B. Retention of AI and Messaging Records

1. Retention of AI Interactions

California public entities are required to retain certain categories of records pursuant to statutes and regulations. These requirements often relate to specific record types, such as personnel records. There is no authority directly addressing retention of AI Interactions because, as with all public records, whether AI Interactions should be retained is based on their *content and purpose*, not their format. Thus, if a particular exchange of messages is relevant to a specific facet of the public’s business that has a record series in the public entity’s retention schedule,

⁸ *Lake v. City of Phoenix* (Ariz. 2009) 218 P.3d 1004, 1008; *O’Neill v. City of Shoreline* (Wash. 2010) 240 P.3d 1149, 1154; *Irwin v. Onondaga County* (N.Y. 2010) 895 N.Y.S.2d 262, 268.

⁹ Gov. Code § 7922.570.

¹⁰ See League of California Cities, A Guide to the California Public Records Act, Revised September 2025, p. 14.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

then those AI Interactions should be categorized for retention and saved for the appropriate time period.

To facilitate proper retention and avoid commingling of records, public employees should create work accounts for AI Interactions and use those accounts solely for entity business, not personal matters. Public employees should save any outputs or AI-generated transcripts that relate substantively to an existing record series, or which otherwise have long term value to the public entity's work. The public entity's IT Department should conduct routine purges of AI Interactions that have served their purpose and have not been designated for retention. Since some AI interactions are likely to have longer-term value than informal chat messages, it is advisable to wait six months to two years before purging AI records.

While not directly related to record retention, some public agencies have issued guidance to employees for citing the use of generative AI in documents. For example, the City of Long Beach encourages employees to be transparent about AI use and include citations such as, "This document was drafted with support from (GEN AI SYSTEM) on DD/MM/YYYY. The content was fact-checked by City staff. This information is unmodified."¹¹ As public employees work to gain familiarity with AI and integrate it into their operations, encouraging citations to indicate AI assistance can be beneficial, both for transparency and for knowing when AI Interactions rise to a level of significance that they should be retained. As a general rule, it is recommended to save any AI Interactions related to the drafting of shared documents (i.e., living documents where multiple people can contribute to drafting), particularly if those documents indicate that AI provided assistance. Conversely, AI Interactions in which the user is seeking general information (e.g., how to perform technology tasks, such as converting file formats or translating languages), the AI Interaction likely does not need to be retained unless the user is relying on the output in some other formal documentation that could qualify as a public record. Thus, guidelines for citing AI, instructions on appropriate use cases¹², and retention parameters for particular use cases are critical for informing management of AI data, particularly in this early phase of integrating AI tools.

Overall, AI Interactions should be viewed as another type of information format, like Word documents or emails, that must be retained only when their content relates significantly to public business, an existing record series, or otherwise has long-term value in light of its context and purpose. Clear guidance on AI use for employees is crucial for limiting the data burden of retaining AI Interactions while allowing for transparent engagement with tools that improve efficiency of public entity operations.

¹¹ City of Long Beach Generative AI Guidance, <https://longbeach.gov/globalassets/smart-city/media-library/documents/generative-ai-guidance-v1-1>, p. 6.

¹² See, e.g., City of Long Beach, Generative AI Use Case Registry, <https://longbeach.gov/smartcity/projects/generative-ai-use-case-registry/>.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

2. Retention of Chat Messages

Like AI Interactions, chats should be analyzed for retention based on their substance and purpose. If chats relate to the public's business and are relevant to an existing record series, they should be designated for retention under the most appropriate category.

It is likely that many chat records will not fall into an existing record series because chats are often used informally and do not have long-lasting significance for the public's business. According to Government Code section 34090, city department heads may destroy records with the approval of the legislative body and consent of the city attorney, without copying the records, after they are no longer required. However, the statute further provides that certain records, including records less than two years old, may not be destroyed. This suggests that all public entity records, presumably including emails and chats, must be maintained for at least two years before they may be destroyed, even if they are no longer of value.

This two-year requirement does not apply to all records. The Government Code requires the Secretary of State to establish guidelines for record retention at the local level.¹³ Those guidelines are contained in the Secretary of State's [Record Management Handbook](#). Under those guidelines, messages which are "transitory," meaning that they are created primarily to communicate informal information, are meant to be short-lived, and have a limited purpose may be destroyed once they have served their purpose.¹⁴ This guidance explicitly applies to emails, but the description can be extended to chat messages, which are equally if not more transitory in nature than emails, given the informality of most chats on platforms like Teams. Accordingly, several local entities characterize chats as "transitory" in their retention policies and permit regular deletion of chats.¹⁵ These policies typically treat all chats as transitory unless the custodian of the chats affirmatively identifies the exchange as something that has more than transitory value. Under these policies, all chats not categorized as eligible for retention are deleted on a routine basis.

While there is a strong argument that policies permitting deletion of emails and chats before two years are legally permissible, there is some dispute about whether these policies comply with Government Code section 34090. In 2019, AB 1184, which would have required emails sent or received regarding public business be retained for a minimum of two years, was

¹³ Gov. Code § 12236.

¹⁴ California Secretary of State, Records Management Handbook, *Chapter 3 – Electronic Records and Storage*, <https://www.sos.ca.gov/archives/records-management-and-appraisal/records-management-handbook/chapter-3>.

¹⁵ See University of California Irvine, [Chat Platform Message Retention Project | Office of Information Technology](#) (destruction of chat messages after one year); City of Fresno, Resolution 2019-122 (instant messages and chats deleted every six months); County of El Dorado, California, Electronic Messaging Retention Policy (Teams chats automatically purged and permanently deleted 29 days from the last modified date).

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

introduced to settle this debate. The bill was vetoed by the Governor. In 2020, AB 2093, which largely mirrored AB 1184, was introduced but ultimately withdrawn. The committee comments for both AB 1184 and AB 2093 mentioned the ambiguity in Government Code section 34090. The comments to AB 2093 also noted that some courts have moved in the direction of applying the existing retention statute to emails. That statement, however, was only supported by citation to a single stipulated order out of San Diego that only applied to the parties in that case.¹⁶

Most recently, in 2022, AB 2370, a bill that would have required state agencies to retain digital records for two years, was introduced. Articles about the legislation noted that “State law already requires that local government do this,” but that state agencies have been allowed to delete informal/transitory emails all along.¹⁷ The bill comments referenced the ambiguity around whether the two-year retention requirement for emails applies to local governments, and expressed concern that creating a rule for state agencies would open the door to eventually applying the same to local agencies.¹⁸ Like the previous two attempts, AB 2370 failed.

The fact that the legislature has failed three times to enact clarifying legislation, the noted existence of public entity policies permitting automatic deletion of email, and the lack of case law strongly suggests that email—and likely chats—can be treated as presumptively transitory and purged through an automatic deletion policy. It is critical, however, that any policy clarifies the responsibilities of the custodian to determine whether a message is transitory or should be retained. Helpful metrics to assist local agency staff in making that determination include: (1) whether the message was sent to communicate informal information, as opposed to formal communications; (2) whether the message loses its value once its immediate purpose is served (i.e., once the recipient receives it); (3) whether the information shared is of only short-term value and does not document official decisions, policies, or financial obligations (e.g., messages sent to plan a meeting are likely transitory while the minutes from that meeting are not transitory); and (4) whether or not the message provides substantive evidence of a transaction, explains a decision, or is required for legal or financial audits. It is also important that the policy make clear that staff should not use the chat feature as a storage device and, instead, should preserve any chats determined to be a public record. Records created by chat can be stored by exporting them via the app’s settings menu (if the option is available) or if necessary, screenshotting the relevant chats and saving the screenshots to an appropriate location for

¹⁶ Cal. Assemb. Comm. on Judiciary, Analysis of AB 2093 (2019-2020 Reg. Sess.), at p. 3 (May 10, 2020).

¹⁷ Los Angeles Times, “Government emails are public records. Deleting them too quickly is not OK” (March 28, 2022), <https://www.latimes.com/opinion/story/2022-03-28/editorial-california-public-records-delete-emails>.

¹⁸ Cal. Assemb. Floor Analysis, Analysis of AB 2370, (2021-2022 Reg. Sess.), at p. 3 (March 23, 2022).

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

retention. It is recommended that public entities provide clear guidance to employees on how and where to save chats that are public records to ensure consistent retention practices.¹⁹

Policies that automatically delete chats (and emails) after a prescribed period of time allow public entities to control costly data proliferation. However, there is a lingering possibility that these policies could be subject to scrutiny and legal challenge under Government Code section 34090.²⁰

3. Retention of Metadata

As explained above, metadata accompanies other electronic records and is not typically considered a standalone record. Thus, the metadata of electronic records, like chats and AI Interactions, should be retained whenever the primary record qualifies for retention. Public entities should work through their IT departments to ensure proper procedures for the preservation of metadata, as there is always a risk that metadata can deteriorate, be lost, or destroyed unintentionally, such as when converting file formats or transferring files between different operating systems.

III. Conducting the Exemption Analysis

Once a public record has been created and is in possession of the public entity, it is subject to potential disclosure under the CPRA unless one of the CPRA's several exemptions apply to support withholding or redacting the document. This section will address guidance for the exemption analysis of AI Interactions, chats, and metadata that are public records. It also discusses the risk of waiving a privilege or exemption when using these technologies.

A. CPRA Exemptions in the Context of AI Interactions

1. Considerations for the Exemption Analysis

AI Interactions that relate to a public entity's business likely qualify as public records. However, the nature of generative AI apps and assistants as brainstorming tools creates unique questions regarding the applicability of established CPRA exemptions. California courts have yet to address this question, but based on statutes and case law, the most likely applicable exemption is the public interest exemption (also known as the "catchall" or "balancing test" exemption).

¹⁹ Public officials should apply any retention policy for chats to messages exchanged on personal devices or accounts if the messages relate to public entity business. If officials use apps with automatic message deletion, then they should opt out of automatic deletion settings or actively preserve public business-related messages before they are deleted.

²⁰ See 2015-2016 Alameda County Grand Jury Final Report, *The Public's Right to Know: Electronic Records Retention and Access* (Grand Jury found automatic email purge after 30-60 days to be an impediment to efficiency and transparency and encouraged cities in Alameda County to ensure email preservation for at least two years).

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

This “catchall” provision does not apply to any specific category of records, but instead allows agencies to justify withholding any record by demonstrating that on the facts of a particular case, the public interest in nondisclosure clearly outweighs the public interest in disclosure (i.e., a balancing test).²¹ This exemption is often invoked in the context of protecting private information, contract negotiations, or where extensive redactions make the remaining information serve little value to potential readers. To the extent that a public entity uses a closed, enterprise AI system²², records generated by AI tools that are a part of the closed system and which contain confidential information can be redacted in accordance with the public interest exemption. However, privacy-based redactions are unlikely to arise in the context of public, non-enterprise AI tools (i.e., open system tools like Claude or ChatGPT which train off of interactions with all users)

– none of which are likely to arise in the context of public, non-enterprise AI tools (i.e., open system tools like Claude which train off of interactions with all users) where inputting private information should be avoided. While the general public interest exemption may lend itself to creative arguments for nondisclosure given its flexibility, any decision to withhold AI records on this ground must be supported with a clear explanation as to weight the balancing of the interests favors nondisclosure. where inputting private information should be avoided. While the general public interest exemption may lend itself to creative arguments for nondisclosure given its flexibility, any decision to withhold AI records on this ground must be supported with a clear explanation as to weight the balancing of the interests favors nondisclosure.

However, in addition to the general balancing test, there are several specific manifestations of the public interest exemption that might apply to AI Interactions.²³ The iterations of this exemption that are most likely to apply in the AI context are the deliberative process exemption and drafts exemptions, which are discussed below. Exemptions which relate to the privacy of information itself (rather than the privacy of a process or exchange) are less

²¹ Gov. Code, § 7922.000; *Long Beach Police Officers Assn. v. City of Long Beach* (2014) 59 Cal.4th 59, 67.

²² An enterprise closed AI system is a proprietary AI platform designed for, and deployed within, an organization and which restricts access to its underlying models, training data, and codebase. Unlike public AI tools (such as public ChatGPT), closed systems ensure that organization data, proprietary information, and user queries never leave the organization’s secure, firewalled environment, providing security and confidentiality.

²³ The nature of the record and surrounding circumstances guide the invocation of the public interest exemption. As such, the exemption is used in a variety of contexts and there have been specific outgrowths of the exemption both in statute (e.g., drafts under Gov. Code, § 7927.500) and in caselaw (e.g., deliberative process under *Times Mirror Company v. Superior Court* (1991) 53 Cal.3d 1325) which address specific types of records while still ultimately requiring the balancing test analysis to properly invoke the exemption for any particular record.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

likely to apply to AI records because inputting private information into an AI program itself arguably exposes that information to the public, depending on the AI tool.

(a) Deliberative Process Exemption

The deliberative process privilege can be invoked to withhold records from disclosure in response to PRA requests in order to protect “frank discussion of legal or policy matters” where disclosure could greatly hamper governmental efficiency.²⁴ Specifically, “predecisional” documents that are generated before the adoption of a policy or decision are the type of document typically protected by deliberative process.²⁵ Conversely, communications “made after the decision and designed to explain it” are not privileged.²⁶ Further, purely factual and investigative information will generally not be protected as part of the deliberative process unless “actually related to the process by which policies are formulated” or “inextricably intertwined” with “policy-making processes.”²⁷ Given the focus on protecting the decision-making process, the deliberative process exemption is most often applicable to protect predecisional discussions of an entity’s high level decision-makers, such as councilmembers, managers, and department heads. Overall, the key question in every case is “whether the disclosure of materials would expose an agency’s decision making process in such a way as to discourage candid discussion within the agency and thereby undermine the agency’s ability to perform its functions.”²⁸ However, even if disclosure may impact the decision making process, if the impact is minimal, then disclosure may be still be required if it is justified by the public’s interest in receiving the information at issue.²⁹

The purpose and scope of the deliberative process exemption make it appear a worthy candidate for AI Interactions when they are used to develop a policy or help the public agency come to a decision. On one hand, AI Interactions can be viewed as a reflection of a public employee’s deliberative process, albeit through communication with a computer assistant rather than another employee. The release of employee AI Interactions could thus arguably impede an employee’s willingness to use these tools, to the detriment of efficiency. On the other hand, the lack of person-to-person interaction involved with AI Interactions raises the potential that courts will not view such records as worthy of the same protections as candid discussions among employees when deciding important policy matters and/or determine that the public interest in disclosure outweighs the interest in non-disclosure. Indeed, the deliberative process exemption was developed to safeguard free *discussion*, so expanding the exemption to apply to AI Interactions would involve the novel argument that candid deliberation worthy of protection can

²⁴ *Times Mirror Company v. Superior Court* (1991) 53 Cal.3d 1325, 1340.

²⁵ *Id.* at 1341.

²⁶ *Id.*

²⁷ *Id.* at 1342.

²⁸ *Id.* (internal citation omitted.)

²⁹ *See State v. Superior Court of Los Angeles County* (2024) 101 Cal.App.5th at 235.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

occur between one person and an AI tool, not just between decisionmakers. The unknown future of AI and its implementation in public entity operations may ultimately guide how courts decide to apply the deliberative process privilege to AI Interactions,³⁰ but until then, the scope of the CPRA and competing arguments for application of the deliberative process exemption suggest that public entities should be sparing with their invocation of the exemption for AI records. If it is invoked at all, public entities should have well-articulated grounds for why the entity's functions would be harmed by release of such a deliberative record and understand that the decision may be subject to legal challenge.

(b) Preliminary Drafts Exemption

The CPRA also exempts from disclosure “[p]reliminary drafts, notes, or interagency or intra-agency memorandums that are not retained by the public agency in the ordinary course of business, if the public interest in withholding those records clearly outweighs the public interest in disclosure.”³¹ This “drafts” exemption extends some protection from public scrutiny for writings which relate to pending actions or decisions. To qualify for the “drafts” exemption, the record must be a preliminary draft, note, or memorandum; it must not be retained in the ordinary course of business; and the public interest in withholding the record must clearly outweigh the public interest in disclosure.³²

AI Interactions that relate to, or constitute a part of, a preliminary draft may qualify for the drafts exemption. Like the deliberative process exemption, some may argue that AI Interactions do not implicate freedom of expression concerns sufficient to warrant invocation of the drafts exemption. However, the drafts exemption may be more appropriate and cause less transparency concerns than other exemptions because preliminary drafts, even those generated by AI, are eventually released as final versions. However, there is also an argument that the public has an interest in the information underlying public documents, particularly when that information is derived from AI assistants which can hallucinate, contain biases, or provide incomplete or untrustworthy information. One way to address this public interest while still preserving the ability to invoke the drafts exemption is to include citations in public documents that make clear when AI was used in the creation of a document (as discussed in Section II(B)(1) above). This way, the use of AI is not a secret, the public can scrutinize the document, and the public entity will have greater incentive to verify the information before releasing it. Thus, it is recommended that public entities include a statement in documents that were created with AI

³⁰ Indeed, courts may ask whether we have reached a point that AI is necessary to the deliberative process, such that deprivation of a ‘private audience’ with an AI assistant would genuinely undermine public entity operations. Or, courts could favor the opposing position that decision-making occurred without AI before and thus allowing public insight into public entity AI Interactions does nothing to interrupt the true avenues for deliberation.

³¹ Gov. Code, § 7927.500.

³² *Citizens for a Better Environment v. Department of Food and Agriculture* (1985) 171 Cal.App.3d 704, 711–712; Gov. Code, § 7927.500.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

assistance. Additionally, any use of the drafts exemption should be based on well-articulated reasons that serve the goals of protecting the drafting process.

While there are arguments which may support applying exemptions to AI Interactions, it is important to reiterate that there is no express authority on the issue. Thus, it is recommended that AI-use policies be implemented, which provide guidance about appropriate uses of AI when conducting public entity work and discourage the input of any information that is not suitable for public release into public, open universe AI tools.

2. Managing the Risk of Privilege Waiver

Because AI assistants often learn from interactions and data from a vast number of users, some AI technologies are not contained to “closed universes” and thus do not guarantee that the information inputted will remain private. Indeed, AI applications often share data and the prevailing guidance is to avoid inputting any private or sensitive information into an AI operating in an “open universe”, even due to the risk of data breaches alone.³³ Public entities often deal with private information, including information about building and technology security, private consumer information, and confidential material submitted by vendors and constituents. **Accordingly, public employees should never input private, sensitive, confidential, or privileged information into an AI assistant that is not a part of a secure, closed universe.**

For those entities which have licensed with technology vendors to establish AI integrations which operate in a secure, closed and private system, there is a much lower risk of breaches which could release private information. With closed system AI tools, entities should coordinate with Information Technology staff and consultants to ensure that the technology is set up according to best practices and contains adequate security measures. Then, those entities should establish clear policies for when certain private or confidential information may be used in the enterprise AI system.

Additionally, AI assistants that are not a part of closed systems pose the additional risk that public employees may inadvertently waive a privilege, including the attorney-client privilege, by uploading confidential documents or information into the AI. This is in part because AI assistants are not presumed to be a private space – on the contrary, many AI models collect and learn from any information they receive. Thus, giving an AI assistant private information could be viewed as akin to sending private information to a third party, reducing a public entity’s ability to claim a legitimate privacy interest in the material. Indeed, a federal court in New York recently held that a criminal defendant’s interactions with a public AI platform (Claude) were not

³³ See Stanford University Human-Centered Artificial Intelligence, “Be Careful What You Tell Your AI Chatbot” (Oct. 15, 2025), <https://hai.stanford.edu/news/be-careful-what-you-tell-your-ai-chatbot>; City of Long Beach Generative AI Guidance, <https://longbeach.gov/globalassets/smart-city/media-library/documents/generative-ai-guidance-v1-1>, p. 2.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

protected by the attorney-client privilege or the work-product doctrine.³⁴ This risk is particularly relevant now, as courts have yet to weigh in on whether and when AI Interactions may be exempt.

Some organizations have started to use AI notetaking technology to transcribe oral conversations during meetings. If an AI transcription is used for a meeting between public employees, the record will likely qualify as a public record and withholding the entire transcript could be difficult to justify in light of transparency considerations under the CPRA. Further, if a third-party AI application is used, it may constitute a waiver of privilege. Thus, AI transcription technologies should be used cautiously, if at all. Indeed, meeting conversations can include stray comments, speculation, or even sensitive information that will require redaction. Further, transcripts generated by AI can misunderstand context or fail to accurately reflect the tone of statements, including sarcastic or joking comments. Those statements may not qualify for redaction on their own, and run the risk of creating false impressions about what was discussed in a meeting. Thus, notetaking technologies should generally be avoided when public employees are engaging in predecisional deliberations or discussing privileged or sensitive information. Transcription technologies are likely more appropriate for meetings with constituents or third-parties. If a public entity chooses to allow the use of AI notetaking, it is recommended that a policy be established which requires that a human verify and revise the transcription as appropriate to accurately reflect the discussion.³⁵

B. CPRA Exemptions in the Context of Chat Records

Chat records are akin to email records in terms of the exemption analysis. That is, emails and chats generally have the potential to qualify for a variety of different exemptions including deliberative process (e.g., predecisional discussions of a new policy), drafts (e.g., exchange of draft language of a press release), medical and personnel records (e.g., human resources emails about private personnel records or information), attorney-client privilege (e.g., city attorney advice), and general privacy (e.g., contact information and other private data about constituents). Thus, chat records should be reviewed for potential exemptions depending on their context and the type of information exchanged.

Importantly, chats or text messages often involve a longer exchange which may contain both exempt and non-exempt information. In those instances, the record should be released with

³⁴ *United States v. Heppner*, 25-cr-00503-JSR (S.D.N.Y.)

³⁵ Companies have even started to develop AI assistants which purposely include obviously false or obscure information in automated notes or reports that the human reviewing must remove during the verification process, adding a layer of accuracy and assurance that the document has been reviewed by a human. (See Philip Lukens, “Automated report writing: Benefits and risks for police” (March 1, 2025), <https://www.police1.com/artificial-intelligence/automated-report-writing-benefits-and-risks-for-police>.) As technologies develop, these types of verification systems will help entities capitalize on the efficiency of AI tools while ensuring continued human oversight.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

redaction of the exempt information whenever possible, rather than nondisclosure of the entire record.³⁶ Chats and texts also often contain both responsive and nonresponsive information within the same thread. In that scenario, given the CPRA’s preference for transparency, it is generally recommended to produce both the responsive and nonresponsive information unless the nonresponsive information is itself exempt from disclosure because redactions are only justified where a valid exemption applies.

IV. Managing the Benefits of Technology Against the Drawbacks

A. Potential Uses for AI in Processing Public Records Requests

In a recent survey by the National Archives and Records Administration, almost one-fifth (18.6 percent) of respondent agencies report using AI and/or machine learning in processing Freedom of Information Act (“FOIA”) requests, which is the federal public records law.³⁷ The survey report sums up the AI issue for public records by noting that it is important for agencies to consider AI technologies that will reduce processing times while reiterating the importance of human input when making exemption decisions.³⁸ In recognition of the potential efficiency savings, some federal agencies have tested machine-learning models and algorithms to search extensive document repositories which are only growing with the proliferation of electronic documents.³⁹ Other tools have also emerged, which not only help with searching for responsive documents, but can also identify private information and redact it or even flag documents which appear to contain “deliberations.”⁴⁰ Additionally, AI integrations on public websites are being

³⁶ A public entity has a duty to provide such a record to the requester in redacted form if the nonexempt information is “reasonably segregable” from what is exempt, unless the burden of redacting the record becomes too great. (Gov. Code, § 7922.525, subd. (b); *American Civil Liberties Union Foundation v. Deukmejian* (1982) 32 Cal.3d 440, 452-454, 458; *Becerra v. Superior Court* (2020) 44 Cal.App.5th 897, 929–934.

³⁷ National Archives and Records Administration, *Assessing Freedom of Information Act Compliance through the National Archives and Records Administration’s 2024 Records Management Self-Assessment*, <https://www.archives.gov/files/ogis/documents/ogis-2024-rmsa-final-9.25.2025.pdf>.

³⁸ *Id.*

³⁹ NBC News, “Some U.S. government agencies are testing out AI to help fulfill public records requests” (Aug. 1, 2023), <https://www.nbcnews.com/news/us-news/federal-agencies-testing-ai-foia-concerns-rcna97313>.

⁴⁰ See, e.g., [FOIA Assistant](#) (“Key to the FOIA Assistant’s capabilities is highlighting *deliberative language*—that is, language containing multiple linguistic features such as suggestion, opinion, emotion, supposition, and conjecture—in PDF documents. Since deliberative language frequently points toward potential exemptions, this instant identification helps government analysts choose whether to redact the suggestion.”); [Everlaw](#) (“Use

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

developed to help requesters draft requests or quickly locate documents, which are already available online so that the requester does not need to submit a request at all.⁴¹ While many of these tools appear to be in development, they are constantly emerging and more tools are being introduced to the market every day. These technologies will benefit agencies and the public if they effectively create efficiencies in the era of electronic public records.

However, of course, AI itself is new and developing, which means that there are legitimate concerns to consider when deciding whether to purchase AI tools. As government watchdogs have pointed out, AI tools run the risk of creating complacency, inconsistency in decision-making, and inadvertent redactions or disclosures when used without consistent standards and enforcement of those standards.⁴² Thus, while public agencies should consider implementing AI technologies to assist with processing CPRA requests, those agencies must be prepared to work internally and with vendors to create standards and assurances that AI work product will be reviewed by a human, particularly when it comes to invoking exemptions.

B. Managing the Creation and Production of Public Records

While AI technologies have great potential for improving the process of responding to CPRA requests, AI and chat applications also contribute to the number of public records that may need to be searched, reviewed, and disclosed. Thus, internal policies are critical for managing the creation of public records, the volume of data to be retained, and the efficient processing of requests for records housed in personal devices of public officials.

1. Internal Policies

To manage both the number of public records and content of those records, it is crucial that public entities create policies to guide the appropriate use of AI and chat application technologies. There are several key issues which should be addressed in these policies:

- **First**, internal policies should provide approved uses for such technologies, including the appropriate use of AI (e.g., drafting, translating, etc.) and messaging applications (e.g., transitory matters only).
- **Second**, policies should address what types of information should not be inserted when using technology, including a clear prohibition on uploading private information to an AI tool that is not a part of the entity's closed system, an

Everlaw^{AI}'s Coding Suggestions to quickly organize documents based on relevance, privilege, or custom coding instructions.”)

⁴¹ NBC News, “Some U.S. government agencies are testing out AI to help fulfill public records requests” (Aug. 1, 2023), <https://www.nbcnews.com/news/us-news/federal-agencies-testing-ai-foia-concerns-rcna97313>.

⁴² *Id.*

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

explanation of what information may be entered into a closed system AI tool, and guidance on avoiding the reference of private information in chats, particularly those with external participants.

- **Third**, guidance on when and how to save important (i.e., non-transitory) AI Interactions and chats should be included in policies to ensure compliance with retention schedules and implement consistent record-keeping procedures.
- **Fourth**, policies should recommend or require that public entity business only be discussed on entity-specific accounts to avoid commingling of personal and public business, and reduce the need to search personal accounts when responding to records requests. Additionally, to the extent that public officials may inherently receive messages related to public business on their personal accounts, policies should give guidance on how to save and search for public records on personal devices or accounts. (*See* Section IV(B)(2) below.)
- **Fifth**, policies should address when and how to indicate that AI was used to assist with creating a record to ensure internal accountability over content and transparency with the use of assistive technologies in public operations.
- **Sixth**, policies should provide specific guidance on opting out of data collection on approved applications to reduce the risk of inadvertent exposure of sensitive information.

There may be further issues that should be addressed in internal policies, depending on the technologies used by a public entity, the entity's goals and preferences, and any limitations imposed by local laws or regulations.

2. Personal Device Searches

AI assistants and chat applications have the added benefit of being widely available to the public. This means that public officials or employees may use private, non-official accounts for services like Signal or ChatGPT. As noted above, if public officials utilize such applications for public business, *even if on a private device or account*, the resulting records remain public and must be disclosed unless they are exempt.⁴³ The ease and efficiency of online tools, including social media, means that officials are increasingly using private accounts, even if just to receive unsolicited messages from constituents or colleagues. Accordingly, the need to search personal devices and accounts in order to comply with the CPRA is ever-growing. To facilitate efficient personal device searches and ensure demonstrable compliance with the CPRA, the following procedures are recommended:

⁴³ *City of San Jose v. Superior Court*, *supra*, 2 Cal.5th at 625-29.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

- Inform all public officials early and often that any record relating to public business is subject to disclosure, even if contained in a personal device or account;
- Recommend that public officials only send correspondence regarding public business on public-facing accounts, reducing the potential responsive documents on personal accounts;
- Provide step-by-step instructions for searching for and extracting social media messages, chat messages, AI Interactions, and personal emails that constitute public records;
- Require that public officials conduct a search of personal devices/accounts whenever implicated by a request and require that the officials sign an affidavit attesting to the search and whether responsive documents were located⁴⁴; and
- Create enforceability mechanisms if possible to hold officials accountable for completing the searches. Such mechanisms could include a Council adopted policy that: (1) requires public officials to conduct thorough searches and to complete an affidavit confirming that was done; and (2) permits the public entity to inform the requester of the failure of any public official to comply with that requirement after a reasonable effort to obtain compliance.

V. Key Takeaways

AI technology, chat platforms, and other technological tools are becoming increasingly important for public entities to aid in efficient operation. At the same time, these technologies pose clear consequences and concerns in light of their novelty, impact on data storage, and potential security risks. Overall, public entities should keep an open mind to technologies which can assist them with processing records requests or otherwise, but should concurrently work to develop internal policies and protocols which reign in the creation of public records, educate employees on the benefits and pitfalls of such tools, and prevent ill-advised use of these technologies. The primary considerations for the different record types discussed in this paper are summarized in the chart below:

Record Type	Public record?	Retention necessary?	Possible Exemptions?	Policy Considerations
AI Interactions	Presumed public if it relates to entity business. Transcriptions of	Standard 2 year retention applies if record relates substantively to public business and	<u>Deliberative process</u> : may be difficult to argue given lack of	Comprehensive policies with explanation of appropriate use-cases, data collection

⁴⁴ The affidavit procedure has been expressly approved of by the California Supreme Court. (*City of San Jose v. Superior Court*, *supra*, 2 Cal.5th at 628.)

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

	some meetings are likely public. Interactions for the purpose of obtaining simple information may not constitute a public record (i.e., asking AI to solve a math problem or give computer instructions)	has some long-term value. If record relates to a specific record category, retain for corresponding period. Simple exchanges that do not have long-term value or relevance to public entity business may be eligible for destruction prior to 2 years.*	human deliberation. <u>Drafts</u>: may apply when AI exchange relates to or is part of a preliminary draft; citations to AI in final draft are recommended. <u>Balancing test</u>: case-specific and only applies if circumstances demonstrate a clear interest in nondisclosure.	settings, saving important records, and citing to AI are recommended. Policies should define when AI notetaking is permissible and require human review. Policies should specify types of information that should never be input into open system AI tools versus closed system AI tools, including personal info, security info, and privileged content.
Chats	Presumed public if it relates to entity's business. Informal messages that do not substantively relate to public business may not constitute public records (e.g., city employee wishing another a happy birthday)	Standard 2-year retention applies if record relates substantively to public business and has long-term value. If record relates to a specific record category, retain for corresponding period. Transitory exchanges may be eligible for destruction prior to 2 years.*	Many exemptions could potentially apply based on purpose, content,) audience, and public interest considerations. Should be reviewed similar to emails; redactions may often be necessary to segregate exempt from non-exempt and responsive from non-responsive material.	Policies should encourage minimal and informal-only use of chats, using public accounts only for entity business, and minimizing exchange of sensitive information. Specific policies and guidance for chat records contained on private accounts of officials are recommended.

From Emails to Algorithms: Emerging Public Records Challenges in the AI Era

J. Leah Castella

Alexandra Mulhall

Burke, Williams & Sorensen LLP

Metadata	Presumed public if the primary document associated with the metadata relates to entity business. If the primary document is not a public record, neither is its metadata.	Subject to the same retention period as the primary document.	Subject to the same exemption as the primary document. May be specifically exempt even if primary document is not where disclosure would expose private information or create a security risk, which is not outweighed by interest in disclosure.	CPRA policies should include guidance on when metadata should be withheld to protect privacy and/or security. Internal policies should also address the technological processes for ensuring adequate retention of vital metadata.
*Retention periods of less than 2 years may be subject to legal challenge. (See Section II(B).)				