Preprint
Forthcoming in *Cambridge Forum on AI: Culture and Society,* special issue 'AI and Archives'

AI assistants in the archive and the lure of 'instant history'

Finola Finn, University of Luxembourg, Centre for Contemporary and Digital History
finola.finn@uni.lu

Donal Khosrowi, Leibniz University Hannover, Institute of Philosophy
donal.khosrowi@philos.uni-hannover.de

Abstract

AI assistants are increasingly used for navigating and analysing the contents of major archives. Applying Retrieval Augmented Generation to existing large language models, these tools draw on indexes of the relevant archives to answer, in natural language, users' questions. In addition to being powerful finding aids, archival AI assistants are also presented as being capable of providing useful, automated answers to questions about the past. This article argues that such tools and how they are marketed result in major conceptual disruptions and uncertainties, placing pressure on our understanding of a range of roles, forms of information, and outputs involved in the production of historical knowledge. In particular, we argue that these tools may obscure well-established beliefs that 'sources' and 'archives' are not unmediated, clearly navigable, or necessarily comprehensive, and that the processes by which these are used to write 'history' are by no means straightforward or instantaneous. With the aim of mitigating these misunderstandings, the article makes suggestions for how deployers could more carefully frame and describe the intended use of archival AI assistants (especially for public users), to ensure that their benefits for accessibility are exploited while also avoiding misconceptions and safeguarding rigorous historical practice.

1. Introduction

AI is increasingly used for navigating and analysing the contents of major archives. Ambitious pilot projects have developed impressive AI-powered dashboards that offer new ways to search and analyse vast collections, such as through categorisations and visualisations. Offered standalone or built into the functions of such dashboards, LLM-based AI assistants (or 'chatbots') are designed to help users navigate the archives through natural-language queries (Archives of the European Parliament 2024; eLuxemburgensia 2023; MAPE 2025; see also NARA 2025). While the way in which the functionality of these systems is framed varies across projects, users are generally invited to locate relevant materials with more ease and speed – but

also to ask the AI assistant historical questions, such as "summarise what happened on 1 January 1944" (eLuxemburgensia 2025). In return, users receive blocks of brief, informative text, along with the references to the archival sources used to formulate the answer.

These developments raise important questions about what role AI assistants should play in navigating historical sources and archives; what kinds of tasks we may expect such systems to perform well or poorly on; and how the roles and limitations of these technologies should be communicated to users to avoid misplaced expectations. This article focuses on these archival AI assistants and their promise to offer accounts of the past, arguing that certain usages and framings of these tools result in major *conceptual disruptions* – placing pressure on our understanding of and expectations for a range of roles ('archivist', 'historian', 'expert'), forms of information ('document', 'collection', 'archive'), and outputs ('history') (Löhr 2023; Hopster and Löhr 2023; Colavizza 2021). In particular, we highlight ways in which certain uses and framings of the assistants may lure users to 1) lose sight of the inherently mediated nature of documents and the collections in which they reside; 2) erroneously think of digital sources as providing a comprehensive view onto the past; 3) mistake archives as being coherently navigable, tidy entities; and 4) come to view history in a more simplistic, one-dimensional way (cf. Jordanova 2019; Edenheim 2013; Agostinho et al. 2019). To address these concerns, the article makes suggestions for how providers could more carefully frame and describe the intended use of these AI assistants, in order to ensure that their benefits for research and accessibility are exploited while also safeguarding rigorous historical practice (see also Jaillant and Rees 2023).

Section 2 outlines the design, functionality, and current framings of archival AI assistants. Section 3 explains how a 'conceptual-disruptions' lens draws out the ways in which these novel technologies generate conceptual uncertainty and promote conceptual misapplication, before discussing the disruptions and concerns we identify here in relation to archival AI assistants. Section 4 moves onto practical suggestions, considering both usage by the public and more specialised researchers. Section 5 concludes.

2. Archival AI assistants

2.1 What they offer and how they work

As outlined earlier, archival AI assistants can now be used by the public to explore major archives online, including the European Parliament Archives and the Luxembourg National Library (Archives of the European Parliament 2024; eLuxemburgensia 2023). Similar systems are also being planned or piloted at the National Archives and Records Administration (NARA) in the United States and the Arquivo Histórico Ultramarino de Lisboa of Portugal (Koebler 2024; NARA 2025; MAPE 2025). Notwithstanding some differences in design, their basic architecture is highly similar: applying Retrieval Augmented Generation (RAG) to existing large language models (LLMs; such as Anthropic's Claude or Google's Vertex AI), the AI assistants draw on indexes of the relevant archives' documents to answer, in natural language, users' questions

pertaining to these archives. In particular, addressing well-known problems that off-the-shelf LLMs face in regard to producing outputs lacking factual accuracy and making 'reasoning' errors, RAG techniques seek to anchor LLMs' outputs more closely in specific sources of knowledge, such as indexes of digitised documents from an archive, to ensure that generated outputs stay closer to original source material (although not always successfully), and enable AI assistants to pinpoint specific sources on which answers are based (Lewis et al. 2021; Magesh et al. 2024). In the context of archives, then, RAG-based AI assistants can only access those sources which have been prioritised for digitisation (as opposed to the entire physical holdings), and expertise in LLM engineering, whether sourced internally or externally, is necessary for the development of these new tools. Some of the archival AI assistants following these design principles have been given names: 'Archibot' at the European Parliament Archives, 'Archie AI' at NARA, and simply 'Chat' at the National Library of Luxembourg (Archives of the European Parliament 2025; Koebler 2024; eLuxemburgensia 2025). While this article engages with these existing, pioneering archival AI assistants, it is not intended as a critique of these specific systems. Rather, the discussion draws on aspects of these existing systems as jumping-off points to engage the larger, ongoing project of developing similar (including more advanced) systems likely to be introduced at further archives, and to initiate epistemological and practical exchange on their risks, benefits, and appropriate design.

2.2 Framing and marketing

2.2.1 Navigation or analysis?

How digital tools are designed, framed, and marketed shapes how they are used (Ehrmann et al. 2019; Ciravegna et al. 2008; Shah and Bender 2024), and therefore elements such as user interface design, instructions, disclaimers, marketing language, and media releases warrant careful attention. Current framings of archival AI assistants suggest two key uses: *navigation* and *analysis* of archival contents. The potential of the navigational usage is significant, giving researchers a new entry point to the archives that goes beyond traditional catalogues and keyword searches (MAPE 2025). Being able to query an archive's contents in natural language could facilitate flexible and in-depth research, if used in thoughtful, self-reflective ways (as is the case with any search tool). For example, it provides the ability to identify documents that share common themes, but might not be drawn together in traditional searches due to a lack of shared keywords or catalogue tags. While being highly useful in its own right, this navigational capacity of the AI assistants is not clearly separated from their other, analytical affordances. Instead, search and navigation are blurred with analysis and interpretation. This sometimes occurs through the placement of the AI assistants within wider dashboards that facilitate new forms of AI-driven categorisation and search, as well as in wider communications about the tools. For example, in LinkedIn posts regarding the launch of the European Parliament Archive's Archibot (based on Anthropic's Claude), Daniela Amodei (president of Anthropic) announced that the system "helps researchers and staff analyze data and create reports, transforming what once took

weeks of archival research into seconds of discovery" (Amodei 2024). The Anthropic website, too, blurs document retrieval with document interpretation, stating: "What once required deep searches into the EU's physical archives is now available instantly – at everyone's fingertips – bringing the story of European democracy to life" (Anthropic, n.d.). Perhaps most pertinently for users, the blurring also occurs on the webpages of the tools themselves, through conflicting instructional text, e.g. "Type your search here" followed by "Enter your question ensuring it is specific and relating years from 1952 to 1994 and debates from 1999-2004" (Archives of the European Parliament 2025). Suggestions for questions placed near the prompt bar or in instructional materials similarly make the intended role and capabilities of the tool unclear, going well beyond archival navigation: e.g. "What was the Maastricht Treaty?" (Archives of the European Parliament 2024) or "how did JFK's assassination affect the Cold War?" at NARA (see internal demo presentation NARA 2024; Koebler 2024).

The answers provided in response to such queries are in the form and tone we have now come to expect from popular LLMs: for example, in many cases, an opening sentence or two provide the main thrust of the response; lists of supporting points are provided (in this case, with references to specific archival documents); and a summarising sentence closes. The content of the answers of course varies from question to question, but some identified strengths include (sometimes) flagging potential bias in the selected documents and making it clear that the response is based only on the contents of those documents. The answers are also provided with disclaimers. The European Parliament Archibot allows users to either "show", or then "hide", a disclaimer which summarises how the tool works (Archives of the European Parliament 2025). In a more direct way, small print on the eLuxemburgensia chat page makes clear that the AI assistant is an "experimental feature" and that it "may give you inaccurate information" (eLuxemburgensia 2025). It is worth noting that, despite such warnings, users may fail to recognise when errors occur (given they may lack independent information pertaining to the queries they prompt), so it remains unclear how effective disclaimers are in helping users self-police into cautious and critical uses of archival AI assistants.

2.2.2 The archives "speak"?

The framing and marketing of the AI assistants do not only encourage certain uses and expectations, however, but also help shape an understanding of the process by which their answers are formulated. Here, a degree of potentially misleading anthropomorphism has been fostered in some cases (see also Cohn et al. 2024; Reinecke et al. 2025). For example, in a live demonstration for employees, NARA's pilot Archie AI referred to itself as an "Expert Archivist" in chats – much to the dismay of some of the archival staff (Koebler 2024). The European Parliament's ArchiBot, moreover, makes similar allusions, with the title: "Ask the EP Archives" (Archives of the European Parliament 2025). The implicit suggestion that the archives can speak for themselves is made even more strongly in communications on the website, where users are encouraged to think of their interaction with the AI assistant as an opportunity "to chat directly

with the Archive's documents." "The historical archives of the European Parliament," we are told, "now speak directly to citizens" (Archives of the European Parliament 2024).

3. Disruptions caused by archival AI assistants

3.1 Conceptual disruptions

What effects do such framings have, in the minds of users? When interfaces suggest that AI assistants can produce answers to historical questions based on the contents of archives, pressure is placed on established understandings on how, if at all, we can extract reliable information about the past from a collection of documents. In particular, pressure is placed on our understanding of and expectations for a range of roles ('archivist', 'historian', 'expert'); forms of information ('document', 'archive'); and outputs ('history') involved in that process. We suggest that one way this pressure can be productively understood is through the lens of what has been termed *conceptual disruption* (Hopster and Löhr 2023; Löhr 2023). As explained by Löhr, a technology can be considered conceptually disruptive if "it challenges or prompts an overturning of entrenched conceptual or classificatory norms and practices, i.e., practices pertaining to language and thought" (Löhr 2023, pp.2-3). In this case, the disruptions caused by archival AI assistants may generate *conceptual uncertainty*, as the largely unknown scope of these systems' abilities induce uncertainty around the application of entrenched concepts traditionally used in relation to archives.

While *archival concepts* are not as uniformly understood or entrenched in popular thinking as those discussed in the conceptual disruptions literature thus far (for example, disruptions to the notion of 'author' by the widespread uptake of LLMs such as ChatGPT; see Khosrowi et al. 2024), we believe this framework to also be useful here, especially in a heuristic way. First, a conceptual-disruptions lens allows us to focus on the important role of entrenched concepts within the historical and archival sciences, even if these do not always have well-defined, shared meanings across broader society. Second, it facilitates clear, organised discussion of the conceptual, as well as methodological and practical, uncertainties and questions that are raised by these new archival tools for both lay and expert users. For instance, the lens invites us to consider: are archival AI assistants 'assistants', comparable in competence to a human archival assistant or archivist? Can they competently draw on archival sources to synthesise 'histories' that attend to user queries, like a 'historian' or 'expert' could? How should users interpret what kind of outputs archival AI assistants produce (i.e. do they constitute 'history'? (cf. Khosrowi and Finn 2025)) and in what sense, exactly, do they provide 'access' to 'sources' and 'archives'? These issues are conceptually and practically unclear, despite deployers' more definitive framings. As we can see, archival AI assistants hence induce a range of conceptual and practical disruptions and uncertainties that affect crucial concepts we use to understand and organise epistemic practices revolving around archives. Let us further unpack them, highlighting concerns that arise around specific disruptions.

3.2 'Sources,' 'archives'

Using archival AI assistants to answer historical questions and framing them as tools that allow one to "speak directly" with the archives could obscure well-established expert beliefs that 'documents', 'sources' and 'archives' are not unmediated, clearly navigable, or necessarily comprehensive. Historians and archival theorists have long shown that records find their way to us via many intermediaries and can contain multiple points of view. For instance, Ketelaar has argued that "archival fonds, archival documents, archival institutions, and archival systems contain tacit narratives which must be deconstructed in order to understand the meanings of archives" (Ketelaar 2001, p.131). Indeed, by suggesting that an AI tool can immediately extract an answer to a historical question based on the most relevant documents it identifies in an archive, we risk encouraging or facilitating a strange sort of reversal on the very established belief, most famously asserted by Derrida and Foucault (1995; 2002), that archives are not to be approached as mere repositories of records that offer a clear window onto the past, but as complex systems that warrant careful examination themselves (Callahan 2024).

With this in mind, it becomes clear that the 'documents' and 'collections' that these AI tools are drawing on to provide their answers to users are not unmediated sources that can instantly "unlock decades of democratic history," as Amodei of Anthropic phrased it in her announcement, nor consistently provide one singular, 'true' account of past events in the form of an 'answer'. Rather, these digitised documents have been selected and shaped by the actions and decisions of various human actors, along with the biases, institutional norms, archival practices, and sometimes sheer luck that governs the survival of historical material. As various scholars have demonstrated, even what seem to be the driest administrative documents can be unreliable and contain biases, being the products of what an institution wanted to show and preserve (Prescott 2008; discussing Hunnisett 1971; Galbraith 1964). What is more, the selective and political processes of digitisation have often only perpetuated this issues, further narrowing the scope of materials most examined by users, and entrenching the absences of marginalised groups in archives (Thylstrup 2019; Milligan 2022; Ortolja-Baird and Nyhan 2022).

Moreover, significant epistemic gaps exist between historical phenomena/events, physical sources recording these phenomena, and digital representations of those sources. Acts of encoding, translation, and re-presentation underpin the transformation from one form to the other, and these require careful scrutiny and acknowledgement (Jordanova 2019). In particular, awareness of what information is often lost when sources are digitised and/or transcribed (e.g. a telling change in ink colour or hand; differences between versions of a text; or redactions and edits) and that errors may be introduced (e.g. through OCR) is vital (Prescott 2008; Smith and Vine 2024). These necessities are ever more pressing in the case of archival AI assistants, as these tools insert yet another intermediary layer and shift in medium, providing users with a new re-presentation and indeed *interpretation* of the underlying sources.

Making the distinction that the AI assistants are likely providing an interpretation (even if mindlessly and unintentionally), and not merely raw, unmediated information, is important – and

stems from the understanding that sources are not "quarries of raw factual material" (Prescott 2008, p.13).[1] Rather, as Ludmilla Jordanova has highlighted, sources "stem from the messiness of lives as they are lived – including institutions and arms of government" and "it is worth reminding ourselves of the anachronisms of taking sources and making them work for the discipline of history" (2017, p.43). As has been argued in other disciplines in relation to the misleading use of the term "data literacy" to describe the transformation of data into comprehensible narratives (as if the insights were waiting, fully-formed, within the data all along) (Veel 2018, section 3), current framings of AI assistants producing accounts of the past based on the digitised contents of archives obscure the complexity of the underlying process, as well as the role of the human user (e.g. through their query or 'prompt') (see also Khosrowi et al. 2024). With their promise of fast answers, archival AI assistants could encourage lay users to fall into the trap of assuming that 'chatting' with 'documents' and 'collections' can quickly and easily produce a window onto past events, and of imagining a misleadingly neat, straight line between historical accounts and their evidential base. Experts with historical training, too, might be susceptible to these misconceptions, especially in terms of failing to appreciate the extent to which AI assistants affect their engagement with a collection. Indeed, historians are unlikely to be immune, as supported by various uncritical approaches taken by scholars to new digital affordances (see, for example, calls for greater methodological reflection made by Putnam 2016; Hill 2016; Milligan 2013, 2022; Hitchcock 2013). While concerns about both the public's ability to exhibit relevant skills in interpreting historical documents and historians' uses of new digital tools are neither novel nor distinctive of interactions with AI systems, our point here is that framings around archival AI assistants may further exacerbate such concerns in significant ways.

3.3 'History,' 'historian,' 'archivist'

Oversimplifications of how historians come to formulate accounts of the past – which, in reality, is a multi-step and complex process – brings us to our next set of disruptions, as our understanding of 'history' and the roles played by those involved in its production come under pressure. Given users are encouraged to ask these chatbots historical questions about the past, the implicit suggestion is that the output is some form of 'history', but is it? By jumping directly from a limited number of documents, predicted to be most relevant to the query, to automated reports, archival AI assistants skip and/or skim over fundamental steps of the historical process – and any distinction between information/data/sources/evidence and, then, interpretation/analysis/argument/historical knowledge becomes blurred. In order to produce the latter, the process includes, but is not limited to, wide-ranging research (including, perhaps, engagement with serendipitous or less obvious finds which have fallen under the radar of e.g. keyword searches); thoughtful selection of sources according to articulated criteria; decisions of

[1] This is of course not to say that sources cannot contain raw, factual material that is more robust against transformation and distortion, e.g. dates or basic information, such as who was in attendance at a certain event. However, the robustness of such information only goes so far and does not necessarily extend to the larger context it is embedded in.

foregrounding and backgrounding; source criticism; intensive contextualisation; the disentangling of the many layers of information present in sources; and reflection on one's findings in light of others' research (Jordanova 2019; Currie and Walsh 2019; Prescott 2008).

What is more, the process by which historical accounts are produced also arguably involves applying human empathy to historical topics, and showing an awareness of a topic's relation to the present, both of which an archival AI assistant or any LLM cannot currently do (Götter 2024; discussing Droysen 1875). As this brief description of at least some of a historian's process suggests, archival AI assistants could encourage (especially nonexpert) users to overlook the complexity of producing accounts of the past, and to view history in a more simplistic, one-dimensional way. Indeed, these tools could also obscure well-established understandings that the practice and outputs of history are often disputable, malleable, opinionated, and so on.

At the same time, it is generally considered important that those engaging with historical knowledge are able to assess the possible biases or influences which may have shaped the work of its author (e.g. personal beliefs, intellectual heritage, etc.), and while this is at least somewhat possible with a human author, it would be a far greater challenge with an AI assistant (see also Clavert and Muller 2024, p.15). Despite being an overused term at this point, these AI systems are indeed 'black boxes' and it is therefore difficult for users to understand why they might, for example, privilege or hone in on some archival documents over others (see, for example, De Ninno and Lacriola 2025; Makhortykh et al. 2023a).

Further pressure is placed on how we should conceptualise archival AI assistants' workings and outputs by the fact that LLMs are not carrying out 'analysis' or 'reasoning' in the way humans do: they are, at bottom, statistical models of the distribution of word forms, designed to produce likely and plausible sequences of words (Shah and Bender 2024; Stone et al. 2024).[2] At this point, more far-reaching and fundamental questions also arise, such as: what are we losing when we hand over the process of formulating historical knowledge to AI systems? Is the 'point' of the practice of history simply the end result, therefore arguably justifying faster routes to that – or does its value also lie in the process, and the fact that humans undertake the bulk of the labour? Although not within the scope of this article, these are important and relevant questions which warrant future discussion.

3.4 Public responses: disruptions evident

Emerging public discussions around archival AI assistants on social media platforms illustrate how the conceptual disruptions and uncertainties we describe here are unfolding amongst lay

[2] To be sure, there is a recent surge in ML research to build LLM-based systems that exhibit reasoning-like behaviors, including chain-of-thought-based approaches, various fine-tuning approaches, as well as generator-evaluator and other multi-agent architectures (Ke et al. 2025). While we are aware of these developments, it is unclear whether archival AI systems are yet designed to take advantage of these approaches, and whether these approaches will be successful in enabling LLM-based AI systems to exhibit reasoning-like behaviors that are on par with those to be expected from expert human researchers.

audiences and users. Consider, for instance, the comments made under the announcement of the European Parliament's Archibot on LinkedIn. One comment notes:

> "This is truly amazing, finally this critical part of knowledge of our history is available in a way that is not intermediated, is not filtered, not opinionated, and is accessible to the hungry mind of a 5yo who is doing a school project, as well as to anyone who wants to know how history went down. Finally we don't need some out of reach person to tell how it went, it is there for us, just ask the best way you can and it will *unveils* [sic] to you" (LinkedIn 2024; emphasis added).

Such uncritically enthusiastic responses underwrite concerns about the potential for lay users setting misplaced confidence in archival AI assistant's abilities, being drawn towards epistemic short cuts that bypass the role of expert intermediaries and instil illusions of comprehensiveness. At the same time, some users also highlight reasonable concerns about censorship and disinformation (e.g. "I hope it's not going to 'rewrite history' then. 'For safety'" (LinkedIn 2024)), demonstrating how AI assistants could also cause *erosions* of trust and confidence in archival institutions (see also Makhortykh 2023b).

Other comments proceed to highlight how Archibot facilitates "unlocking 2.1 million documents, spanning seven decades, for the world? That's not just nice to have, that's a fundamental shift in transparency and accessibility" (LinkedIn 2024). While it seems plausible to think that Archibot indeed helps mitigate barriers to archive accessibility for citizens (though less comprehensively than the discussion suggests), such comments again raise concerns about misplaced trust and confidence.

Beyond social media commentary, other, journalistic responses similarly fall short of distinguishing the outputs of archival AI assistants from concepts like 'history' or 'historical knowledge'. For instance, an article about the European Parliament's Archibot equates the outputs of Archibot to "history" without distinguishing 'history' from 'historical sources' or considering or acknowledging the significant differences presented by an AI-mediated process of gathering, learning, and engaging with 'history.' "This initiative is not merely an archival project; it's a democratic endeavour aimed at enhancing the public's understanding of the European Union's legislative history," the article reads; "It enables citizens, researchers, and policymakers to dive into the rich legislative history of the European Union, offering insights into the decision-making processes and debates that have shaped current policies and regulations" (Kimaid 2024).

As these impressions from public reception suggest, it is important, especially for lay audiences, to conceptually situate archival AI assistants on more solid foundations, especially their capabilities, outputs, and the nature of our interactions with them. Crucially, we believe that doing so helps users avoid collapsing distinct steps in the research process and the production of historical knowledge into a misguided vision of streamlined, AI-facilitated conversational processes that directly extract historical knowledge from archival contents.

The concerns outlined here are particularly pressing in view of the similarities of archival AI assistants to popular conversational systems built on LLMs like Claude, ChatGPT, and others. Recent years have seen rapid growth in the user base of such systems, with users increasingly exploring use-cases, getting accustomed to prompt-based interactions with AI tools, and integrating these systems into everyday workflows to optimise and partly automate knowledge-related tasks. While a significant portion of users show scrutiny and caution in their use of these systems, public discussions on platforms like Reddit suggest that many people also use these systems in ways that 'skip steps' associated with manual research processes; for example by moving straight from prompt to answer; employing chatbots as teacher-like entities; cementing the belief that there is one 'correct' answer to complex questions (that one can reach with the right prompting); and reporting relief in speeding up tedious aspects of traditional research (e.g. Reddit 2024). A growing number of studies are making similar claims about lax LLM usage, and how it may atrophy various critical thinking and research skills (Lee at al 2025; Shah and Bender 2022). Given the strong similarities between archival AI assistants and familiar conversational AI systems, it seems reasonable to think that the public may engage with archival AI assistants in similar ways, thereby importing modes of interaction that (even more so than in more standard knowledge-retrieval settings) threaten to induce simplistic views on how to acquire knowledge of the past.

3.5 Conceptual disruptions: the right lens?

As we have argued so far, the lens of conceptual disruptions seems helpful to draw out how archival AI assistants put pressure on how we apply familiar concepts in new contexts where novel technologies may assume roles previously occupied by humans and generate outputs that may, at face value, resemble those produced by human experts. Even so, we may ask what, exactly, the lens of conceptual disruptions adds to understanding these emerging issues. Let us elaborate and respond to some potential reservations about it.

One way to push back on our suggestions is to insist that at least some of the problems we highlight here are simply new forms of familiar problems: unreliable historical accounts have always existed; some people already lacked strong skills in source criticism; all new digital methods change the practice of and/or engagement with history in some way; and general forms of misunderstanding around how much we can know for certain about the past have been pertinent well before the advent of archival AI assistants. With these points in mind, one might worry that such a lens is misleading because it presupposes a relatively fixed pre-existing conceptual fabric that is then disrupted by archival AI assistants. As the objection goes, it may seem that some of the terms discussed here (e.g. 'history,' 'archives') did not, and do not, have well-defined, shared conceptual meanings across society – and thus the putative disruptions we highlight do not cut as deep as, say, how generic LLM-based conversational agents like Claude and ChatGPT disrupt the public's notion of what it is to be an 'author' (Khosrowi et al. 2024), or how the invention of the mechanical ventilator interrupted our understanding of 'death' (Löhr

2023; de Boer and Hoek 2020). Unlike disruptions affecting such household concepts, we may think that, even before the advent of archival AI assistants, many members of the general public might not have considered archival documents to be heavily mediated remnants of the past, so it remains unclear in what sense archival AI assistants induce *novel* misunderstandings, as such misunderstandings have recognizably existed before.

Pushing back on these reservations, we point out that while some concerns we highlight here are familiar, they are aggravated by the fact that archival AI assistants are, for now, considerably more prone to error and opaque than human experts whose roles they may conceptually and practically encroach on. It is not just that laypeople may fail to appreciate the considerable amount of analysis, interpretation, and expertise necessary to bring archival materials to bear fruitfully on queries about the past, but that the novel roles played by archival AI assistants make these epistemological tenets even more difficult to recognise, in light of the lure of automation and straightforward-seeming interaction with sources they provide to users. Whereas laypeople may have falsely, but somewhat safely, disregarded the crucial role of expertise in interrogating archival materials in the past, they may now be drawn to rely similarly on systems that are likely to be significantly more brittle, opaque, and prone to error than human experts are (Borji 2023).

Moreover, at least within the specific domain of historical research itself, there are indeed more stable conceptual fabrics that archival AI assistants disrupt, e.g. concerning entrenched concepts such as 'source', 'evidence', 'history', and so on. Historical researchers, perhaps even more than the general public, are pressed to make conceptual choices about how to understand and frame the (appropriate) roles of archival AI assistants, and these issues are affected by considerable conceptual uncertainty that spurs further methodological uncertainties downstream around how to use these systems in epistemically responsible ways.

Beyond capturing some of these downstream street-level issues (of, e.g. how exactly to divide labour between humans and machines in archival contexts, how much credibility to assign to archival AI assistants' outputs, etc.), it is important to appreciate that the lens of conceptual disruptions also helps us in two other crucial ways. First, it encourages us to focus on the bigger picture. Concepts are crucial for understanding and organizing the world: they are the currency that underlies and shapes our thinking and our norms. By taking a step back and focusing on concepts, we are able to consider the effects of AI in archives not just in terms of archival and historical theory but within a broader landscape of how AI technologies are disrupting "entrenched conceptual or classificatory norms and practices" across different industries and areas of society (Löhr 2023, p.3). This broader focus helps us better attend to seemingly mundane but cognitively significant choices, such as the rhetoric we use in framing archival AI assistants and the implicit messages regarding their roles, capabilities, and affordances that are embedded in that rhetoric.

Second, focusing on concepts helps us think about new solutions to managing the impacts of novel technologies, and how we might negotiate their proper roles and facilitate users'

understanding of their limitations. We are not confined to applying only familiar concepts to emerging AI systems, their roles, and their outputs; we can also come up with new, more suitable concepts. The philosophical literature on *conceptual engineering* (Cappelen 2018) offers a range of approaches for how we might engineer new, or better, concepts that respond to disruptions, minimise conceptual uncertainty, and shape our interactions with archival AI systems (Hopster and Löhr 2023). In what follows, we explore several options for how conceptual, communicative, and methodological work can help with these aims.

4. Suggestions

This section offers some suggestions for how the conceptual disruptions identified may be mitigated. As with the preceding discussion, these ideas are not aimed solely at the archival AI assistants currently available, but at this category of tool more generally, both existing and potential. It is important to bear in mind that the needs and requirements of any specific archival AI assistant will depend on the archive it is operating on, and these suggestions are intended to be productive contributions at a more general level. We will first consider how disclaimers and instructions can provide clarity, before turning to alternative approaches in the marketing and design of these tools, with a view especially to wider public engagement. We will then propose ways in which historians might ideally use these 'assistants.'

4.1 Clarity in disclaimers and instructions

Like search environments, AI assistant interfaces create an implicit model in the minds of users of what questions can be asked and what answers might appear (Ciravegna et al. 2008). A key guiding aspect of these interfaces is therefore their disclaimers and instructions. In order to mitigate the conceptual disruptions of archival AI assistants, detail and clarity should be provided to users on: the intended and appropriate uses of these tools; the extent of the documents they draw on; and the limitations of their outputs, especially in comparison to other historical outputs (see also Clavert and Muller, esp. p.21). This information should be prominent, and equally present in interfaces, information pages, announcements, marketing content, and training materials.

To elaborate on these points in turn, it would be beneficial to users to flag that, while useful in locating materials, the archival AI assistant's ability to provide historical accounts is more limited and will likely necessitate significant fact-checking work. Moreover, the distinction between these types of queries (locating materials versus seeking historical knowledge) should be made clear. The source base used to formulate responses should also be clearly described at the outset, in terms of extent and type (e.g. a maximum of 10 documents deemed most relevant to the query from a subset of 75,000 quality-checked files out of the 1.5 million available at the European Parliament Archives (Kimaid 2024)), in order to facilitate users forming relevant expectations and understanding the depth of the response. It would also be beneficial to clearly emphasise that the AI assistant's response is based centrally on the archive's contents, rather

than only a diffuse set of web-scale training data, as this might help to distinguish the assistants from other common and popular LLMs (ChatGPT, Claude, CoPilot etc.) and mark them more uniquely as a tool for conducting archival research.

In terms of limitations, users may benefit from clear communication in disclaimers and instructions of the risk of the AI assistant 'hallucinating'/'confabulating,' but also other less popularly known possibilities, such as the effects of underlying LLMs safety mechanisms (which might e.g. censor certain sources or affect the accuracy of historical text translations (Tekgürler 2025)); inaccurate referencing (Lopatto 2024); and the contemporary perspective of these models which, having been trained on modern texts from the internet, could imbue outputs with various anachronisms in terms of source criticism and translation (Neudecker 2024). In addition, clear and detailed documentation of how the archival AI assistant has been developed (especially in terms of any fine-tuning of the underlying off-the-shelf LLMs and any expert evaluations of model performance) should be made easily available from the main interface, so that users are encouraged to consider what limitations the AI assistant might have. Ideally, this would be made available in both accessible and advanced forms, so that users of differing technical abilities can engage with the documentation effectively.

4.2 Alternatives for marketing and design

Slogans such as "Ask the Archives" echo the language used in marketing for other similar technologies (such as "speak to the PDF" or "chat with the book") and, as discussed in Section 2, can lead to misunderstandings. While these sorts of phrases are perhaps a predictable or unavoidable evolution of language in response to these new tools, they nevertheless require reflective consideration. Perhaps it is more productive for archival institutions to avoid replicating this form of (private sector) communication about its tools, given that straightforward and transparent descriptions of systems' capabilities could better facilitate productive and critical engagement with their collections. In particular, consistency in describing the goals of the tool might be preferable, i.e. asserting that it either increases accessibility *or* provides analysis and interpretation; and if both, then a clearer distinction between these goals is needed.

Relatedly, the option of conceptual engineering could be considered, in which more appropriate terms and associated concepts that avoid uncertainty are created for new technologies (Cappelen 2018). In the context of archival AI assistants, terms that clearly categorise and speak more clearly to the tool's intended use and capabilities could be deployed (e.g. terming the assistant's output as a 'document synthesis' rather than the more vague 'answer,' and eschewing potentially misleading terms like 'historical knowledge').

Avoiding replicating the choices of private technology firms is perhaps also beneficial in the realm of tool design. While the format of LLM-based chatbots may now be well-known and popular, this may not be the most appropriate choice for archival research; perhaps other ML/AI methods are better suited to preserving and/or replicating authentic, transparent human engagement with archives. As the wider dashboards which house some of the current archival AI

assistants show, there are many more ways to use AI in the archives than LLM chatbots, and other exciting projects are underway (e.g. Widegren 2024; Jaillant et al. 2025a; de la Rosa 2025 and other contributions in Jaillant et al. 2025b; Szatucsek 2025). If LLM chatbots are to be built into archival information retrieval systems, then a rethinking of the question-answer format might aid in steering the tool away from participating in the wider, problematic trend of LLMs being used as the new Google Search, which arguably results in even greater commodification of knowledge and less critical engagement with the information generated (Lopatto 2024; Karunaratne 2023). The generation of articulated 'answers', which summarise a group of documents, rather than providing solely lists of sources may unintentionally give "the illusion of a full vantage point" (to borrow a phrase from the writer Ayşegül Savaş, in regards to the power of narrative), and it is worth considering what these outputs obstruct from a user's view or consideration (Savaş 2019). What is more, due to the sampling behaviours of LLM-based conversational systems like Claude or ChatGPT, outputs can display significant variability for one and the same prompt and may differ even more strongly between semantically similar (or identical) prompts that are phrased differently. As a result, outputs might change so that different users receive different 'answers' for otherwise identical questions. The extent and significance of this variability and what measures may be taken to improve output consistency has yet to be researched in the context of archival AI assistants. It seems that society is now moving on from the 'Googlization' of knowledge (Vaidhyanathan 2012; Prescott 2008; Hitchcock 2008) to the LLM-ization of knowledge, and archival institutions have the opportunity to take an autonomous role in shaping that shift in ways that take user behaviour into thoughtful account and actively facilitate critical thinking.

4.3 Historians' engagement

The training that historians have received will of course aid them in avoiding a number of the potential disruptions that these tools present, encouraging them to make appropriate sorts of queries; not take the archival AI assistants' answers at face value; check and evaluate the sources provided; merely use the outputs as a launching pad for further thinking and research; and so on. That said, they may not be entirely immune to the risks posed (as indicated earlier) and it is still beneficial to consider how exactly historians might best use these tools, in ways that maintain rigorous historical practice. Standard transparency in describing methodological choices would of course extend to openness regarding the use of archival AI assistants (whether for locating sources, generating interpretations, or anything else), and historians should be conscious of how their use of an archival AI assistant might affect their research's route and contents. Just as catalogues, keyword searches, OCR searchable text, recommendation systems (e.g. content-based filtering), and specialised digital methods (such as topic modelling) can currently impact the pathways of researchers' findings and thought processes, archival AI assistants will now also play a guiding role, pointing the researcher towards some records and topics, and implicitly obscuring others. While this is not a problem in and of itself, it requires awareness and intentionality on behalf of the researcher. In similar ways to how researchers came to make

productive use of the fresh perspectives topic modelling provides while being clear about the limitations of the method – being aware, intentional, and transparent will also allow historians to make the most of the advantages of archival AI assistants (e.g. navigation, categorisation, distillation) while not losing sight of the complexity of the sources they find.[3] Moreover, facilitating researchers' ability to accurately document their use of AI assistants, such as by making chat histories downloadable as txt files (as included in the MAPE Engine's AI assistant; MAPE 2025) will be crucial to executing this transparent approach.

5. Conclusion

This article has argued that archival AI assistants cause conceptual disruptions and uncertainties, placing pressure on a number of key concepts we use to understand and organise epistemic practices around archives and the production of historical knowledge. Using the design, framing, and marketing of the first archival AI assistants as a launching pad, we have argued that these tools may lure users into soliciting 'instant history' and to severely underestimate the complexity of the process that lies between authentic historical accounts and their evidential base. With these risks in mind, the article has also presented a number of practical suggestions for mitigating or averting the disruptions caused by these new tools.

17 years ago, Tim Hitchcock asserted that:

> "Digitisation, new search facilities, new ways of representing and connecting information, fundamentally changes the nature of the archive - what it means and how it is used, and how we as historians experience it. If our claims to cultural authority are built on our relationship with that archive and the sources they contain, then we need to rethink how the social authority of history can be reconstituted to reflect the changing nature of those holdings" (2008).

The advent of archival AI assistants, and LLMs more generally, presses historians to again rethink how their value and authority in drawing insightful, reliable meaning out of archives can be effectively communicated and, in turn, appreciated by wider society. Working with archives and archival experts to develop, shape, and frame tools like the AI assistants discussed here will be one important step in this process, as well as responsibly mastering any advantages of these tools. It seems that AI-driven systems increasingly exert "pro-active engagement in memory communication," shaping individual and collective perspectives on the past (Makhortykh 2024), and historians now have the opportunity to try influence both the extent and form of that engagement. We are confident that the lure of 'instant history' can be overcome.

[3] We thank Marten Düring for this helpful point.

Acknowledgements

We are grateful to the anonymous reviewers, Frédéric Clavert, Eliane Schmid, and members of the Digital History and Historiography research group at C^{2}DH for their helpful comments on this work.

Funding statement: Finola Finn's work on this article was supported by the Fonds National de la Recherche Luxembourg (FNR), grant number 13307816.

References

Agostinho, Daniela, Catherine D'Ignazio, Annie Ring, Nanna Bonde Thylstrup, and Kristin Veel. 2019. "Uncertain Archives: Approaching the Unknowns, Errors, and Vulnerabilities of Big Data through Cultural Theories of the Archive." *Surveillance & Society* 17 (3/4): 422-41. https://doi.org/10.24908/ss.v17i3/4.12330.

Amodei, Daniela. 2024. "The European Parliament is using Anthropic's Claude." LinkedIn post, October 24. https://www.linkedin.com/feed/update/urn:li:activity:7254172155595169793/.

Anthropic. n.d. "European Parliament expands access to their archives with Claude." Accessed February 27, 2025. https://www.anthropic.com/customers/european-parliament.

Archives of the European Parliament. 2024. "Ask the Archives: The EP Archives Unit Launches its first Generative AI tool." Accessed February 27, 2025. https://historicalarchives.europarl.europa.eu/en/sites/historicalarchive/home/cultural-heritage-collections/news/ai-dashboard.html.

Archives of the European Parliament. 2025. "Ask the EP Archives." Accessed February 27 2025. https://archidash.europarl.europa.eu/ep-archives-anonymous-dashboard.

Borji, Ali. 2023. "A Categorical Archive of ChatGPT Failures." arXiv preprint. https://doi.org/10.48550/arXiv.2302.03494.

Callahan, Sara. 2024. "When the Dust Has Settled: What Was the Archival Turn, and Is It Still Turning?" *Art Journal* 83 (1): 74–88. https://doi.org/10.1080/00043249.2024.2317690.

Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford University Press. https://doi.org/10.1093/oso/9780198814719.001.0001.

Ciravegna, Fabio, Mark Greengrass, Tim Hitchcock, Sam Chapman, Jamie McLaughlin, and Ravish Bhagdev. 2008. "Finding Needles in Haystacks: Data-mining in distributed historical datasets." In *The Virtual Representation of the Past*, edited by Mark Greengrass and Lorna Hughes. Ashgate.

Clavert, Frédéric, and Caroline Muller. 2024. "L'histoire au temps des algorithmes: Une réflexion prospective sur l'introduction de l'intelligence artificielle en histoire au 21e siècle." *20 & 21. Revue d'histoire* 162 (2): 13-26. https://doi.org/10.3917/vin.162.0013.

Cohn, Michelle, Mahima Pushkarna, Gbolahan O. Olanubi, Joseph M. Moran, Daniel Padgett, Zion Mengesha, and Courtney Heldreth. 2024. "Believing Anthropomorphism: Examining the Role of Anthropomorphic Cues on Trust in Large Language Models." In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*: Article 54. Association for Computing Machinery. https://doi.org/10.1145/3613905.3650818.

Colavizza, Giovanni, Tobias Blanke, Charles Jeurgens, and Julia Noordegraaf. 2021. "Archives and AI: An Overview of Current Debates and Future Perspectives." *Journal on Computing and Cultural Heritage* 15 (1): Article 4. https://doi.org/10.1145/3479010.

Currie, Adrian, and Kirsten Walsh. 2019. "Frameworks for Historians and Philosophers of Science." *HOPOS* 9 (1): 1–34.

de Boer, Bas, and Jonne Hoek. 2020. "The Advance of Technoscience and the Problem of Death Determination: A Promethean Puzzle". *Techné: Research in Philosophy and Technology* 24 (3): 306-31. https://doi.org/10.5840/techne2020128111.

De Ninno, Fabio, and Michele Lacriola. 2025. "Mussolini and ChatGPT. Examining the Risks of A.I. Writing Historical Narratives on Fascism." *Journal of Modern Italian Studies* (February): 1-23. https://doi.org/10.1080/1354571X.2025.2457250.

Derrida, Jacques. 1995. *Archive Fever: A Freudian Impression*. Translated by Eric Prenowitz. The University of Chicago Press.

Droysen, Johann Gustav. 1875. *Grundriss der Historik*. 2nd edition. Leipzig.

Edenheim, Sara. 2013. "Lost and Never Found: The Queer Archive of Feelings and Its Historical Propriety." *Differences: A Journal of Feminist Cultural Studies* 24 (3): 36-62. https://doi.org/10.1215/10407391-2391950.

Ehrmann, Maud, Estelle Bunout, and Marten Düring. 2019. "Historical Newspaper User Interfaces: A Review." *IFLA WLIC 2019*: Session 85. https://library.ifla.org/id/eprint/2578/.

eLuxemburgensia. 2023. "eLuxemburgensia.lu receives new chatbot: Artificial intelligence supporting Luxembourg's printed heritage." October 4. https://bnl.public.lu/en/a-la-une/actualites/communiques/2023/chatbot-eluxemburgensia.html.

eLuxemburgensia. 2025. "Let's Chat." Accessed February 27, 2025. https://chat.eluxemburgensia.lu/.

Foucault, Michel. 2002. "The Archaeology of Knowledge." 2nd edition. Routledge. https://doi-org.proxy.bnl.lu/10.4324/9780203604168.

Galbraith, V.H. 1964. *Introduction to the Study of History*. C.A. Watts.

Götter, Christian. 2024. "Künstliche Intelligenz und Geschichtswissenschaft. Mehr als ein neues Werkzeug für die Digital History?" *Historische Zeitschrift* 319 (2): 299-330. https://doi.org/10.1515/hzhz-2024-0026.

Hill, Mark J. 2016. "Invisible Interpretations: Reflections on the Digital Humanities and Intellectual History." *Global Intellectual History* 1 (2): 130–50. https://doi.org/10.1080/23801883.2017.1304162.

Hitchcock, Tim. 2008. "Digital searching and the re-formulation of historical knowledge." In *The Virtual Representation of the Past*, edited by Mark Greengrass and Lorna Hughes. Ashgate.

Hitchcock, Tim. 2013. "Confronting the Digital: Or How Academic History Writing Lost the Plot". *Cultural and Social History* 10 (1): 9–23. https://doi.org/10.2752/147800413X13515292098070.

Hopster, Jeroen, and Guido Löhr. 2023. "Conceptual Engineering and Philosophy of Technology: Amelioration or Adaptation?" *Philosophy & Technology* 36 (70). https://doi.org/10.1007/s13347-023-00670-3.

Hunnisett, R.F. 1971. "The Reliability of Inquisitions as Historical Evidence." In *The Study of Medieval Records: Essays in Honour of Kathleen Major*, edited by D.A. Bullough and R.L. Storey. Oxford University Press.

Jaillant, Lise, and Arran Rees. 2023. "Applying AI to digital archives: trust, collaboration and shared professional ethics." *Digital Scholarship in the Humanities* 38 (2): 571-85. https://doi.org/10.1093/llc/fqac073.

Jaillant, Lise, Olivia Mitchell, Eric Ewoh-Opu and Maribel Hidalgo Urbaneja. 2025a. "How can we improve the diversity of archival collections with AI? Opportunities, risks, and solutions." *AI & Society*. https://doi.org/10.1007/s00146-025-02222-z.

Jaillant, Lise, Claire Warwick, Paul Gooding, Katherine Aske, Glen Layne-Worthey, and J. Stephen Downie. 2025b. *Navigating Artificial Intelligence for Cultural Heritage Organisations*. UCL Press. https://doi.org/10.14324/111.9781800088375.

Jordanova, Ludmilla. 2019. *History in Practice*. 3rd edition. Bloomsbury.

Karunaratne, Thashmee, and Adenike Adesina. 2023. "Is it the new Google: Impact of ChatGPT on Students' Information Search Habits." In *22nd Conference on e-Learning (ECEL 2023)*, edited by Shawren Singh and Sarah Jane Johnston, 147-55. https://papers.academic-conferences.org/index.php/ecel/issue/view/27/29.

Ke, Zixuan, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, Peifeng Wang, Silvio Savarese, Caiming Xiong, and Shafiq Joty. 2025. "A Survey of Frontiers in LLM Reasoning: Inference Scaling, Learning to Reason, and Agentic Systems". arXiv preprint. https://doi.org/10.48550/arXiv.2504.09037.

Ketelaar, Eric. 2001. "Tacit Narratives: The meanings of archives." *Archival Science* 1: 131-41. https://doi.org/10.1007/BF02435644.

Khosrowi, Donal, Finola Finn, and Elinor Clark. 2024. "Engaging the many-hands problem of generative-AI outputs: a framework for attributing credit". *AI and Ethics*. https://doi.org/10.1007/s43681-024-00440-7.

Khosrowi, Donal, and Finola Finn. 2025. "Can Generative AI Produce Novel Evidence?" *Philosophy of Science*, 1-28. https://doi.org/10.1017/psa.2025.10123.

Kimaid, Luís. 2024. "Artificial Intelligence-Driven Archibot: Transforming Access to European Union Parliament Archives." *LegisTech Library*, April 24. https://library.bussola-tech.co/p/artificial-intelligence-archibot-eu-parliament.

Koebler, Jason. 2024. "'AI-Mazing Tech-Venture': National Archives Pushes Google Gemini AI on Employees." *404 Media*, October 15. https://www.404media.co/ai-mazing-tech-venture-national-archives-pushes-google-gemini-ai-on-employees/.

Lee, Hao-Ping (Hank), Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers." In *CHI '25: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121. Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." arXiv preprint. https://doi.org/10.48550/arXiv.2005.11401.

LinkedIn. 2024. Comments on Daniela Amodei's post, October 24. https://www.linkedin.com/feed/update/urn:li:activity:7254172155595169793/.

Löhr, Guido. 2023. "Conceptual disruption and 21st century technologies: A framework." *Technology in Society* 74: 102327. https://doi.org/10.1016/j.techsoc.2023.102327.

Lopatto, Elizabeth. 2024. "Stop using generative AI as a search engine." *The Verge*, December 5. https://www.theverge.com/2024/12/5/24313222/chatgpt-pardon-biden-bush-esquire.

Magesh, Varun, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, and Daniel E. Ho. 2024. "Hallucination-free? Assessing the reliability of leading AI legal research tools." arXiv preprint. https://doi.org/10.48550/arXiv.2405.20362.

Makhortykh, Mykola. "Shall the Robots Remember? Conceptualising the Role of Non-Human Agents in Digital Memory Communication." *Memory, Mind & Media* 3 (2024): e6. https://doi.org/10.1017/mem.2024.2.

Makhortykh, M., V. Vziatysheva, and M. Sydorova. 2023a. "Generative AI and contestation and instrumentalization of memory about the Holocaust in Ukraine." *Eastern European Holocaust Studies* 1 (2): 349-55. https://doi.org/10.1515/eehs-2023-0054.

Makhortykh, M., E.M. Zucker, D.J. Simon, et al. 2023b. "Shall androids dream of genocides? How generative AI can change the future of memorialization of mass atrocities." *Discov Artif Intell* 3 (28). https://doi.org/10.1007/s44163-023-00072-6.

Mapping the Atlantic Portuguese Empire (MAPE). 2025. "AI-Powered Archival Research: MAPE Engine for the Arquivo Histórico Ultramarino de Lisboa." *Youtube*, February 13. https://www.youtube.com/watch?v=is-PqW6Cn2Q.

Milligan, Ian. 2013. "Illusionary Order: Online Databases, Optical Character Recognition, and Canadian History, 1997-2010." *Canadian Historical Review* 94 (4): 540–69. https://doi.org/10.3138/chr.694.

Milligan, Ian. 2022. *The Transformation of Historical Research in the Digital Age.* In Elements in Historical Theory and Practice. Cambridge University Press. https://doi.org/10.1017/9781009026055.

National Archives and Records Administration (NARA). 2024. "AI-mazing Tech-venture," internal presentation slides from NARA Information Services, Virtual Town Hall, June 5. https://cdn.muckrock.com/foia_files/2024/09/19/2024_Information_Services_Town_Hall_Presentation_Redacted.pdf?ref=404media.co

National Archives and Records Administration (NARA). 2025. "Inventory of NARA Artificial Intelligence (AI) Use Cases." Last revised January 7. https://www.archives.gov/ai.

Neudecker, Clemens. 2024. "Context matters. Chancen und Herausforderung bei der Arbeit mit Künstlicher Intelligenz und Kulturellem Erbe". Digital History Berlin Forschungskolloquium, November 27. https://dhistory.hypotheses.org/9017.

Ortolja-Baird, Alexandra and Julianne Nyhan. 2022. "Encoding the haunting of an object catalogue: on the potential of digital technologies to perpetuate or subvert the silence and bias of the early-modern archive". *Digital Scholarship in the Humanities* 37 (3): 844–67. https://doi.org/10.1093/llc/fqab065.

Prescott, Andrew. 2008. "The Imaging of Historical Documents". In *The Virtual Representation of the Past*, edited by Mark Greengrass and Lorna Hughes. Ashgate.

Putnam, Lara. 2016. "The Transnational and the Text-Searchable: Digitized Sources and the Shadows They Cast". *The American Historical Review* 121 (2): 377–402. https://doi.org/10.1093/ahr/121.2.377.

Reddit. 2024. "R/singularity: What are you guys using Claude/Chat GPT the most for right now?" [Discussion]. https://www.reddit.com/r/singularity/comments/1e7p1gt/deleted_by_user/.

Reinecke, Madeline G., Fransisca Ting, Julian Savulescu, and Ilina Singh. 2025. "The Double-Edged Sword of Anthropomorphism in LLMs." *Proceedings* 114 (1): 4. https://doi.org/10.3390/proceedings2025114004.

Savaş, Ayşegül. 2019. "The Cost of Reading." *Longreads*, July 12. https://longreads.com/2019/07/12/the-cost-of-reading/.

Shah, Chirag, and Emily M. Bender. 2024. "Envisioning Information Access Systems: What Makes for Good Tools and a Healthy Web?" *ACM Transactions on the Web* 18 (3): Article 33. https://doi.org/10.1145/3649468.

Shah, Chirag, and Emily M. Bender. 2022. "Situating Search." In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval (CHIIR '22),* 221-32. Association for Computing Machinery. https://doi.org/10.1145/3498366.3505816.

Smith, Harry, and Emily Vine. 2024. "Material and Digital Archives: The Case of Wills." *Transactions of the Royal Historical Society* 2: 441–70. https://doi.org/10.1017/S0080440124000094.

Stone, Matthew, Lauren M. E. Goodlad, and Mark Sammons. 2024. "The Origins of Generative AI in Transcription and Machine Translation, and Why That Matters." *Critical AI* 2 (1). https://doi.org/10.1215/2834703X-11256853.

Szatucsek, Zoltán. 2025. "Experience with AI based tools in the National Archives of Hungary". Presentation at *KI und Archive - Eine Zeitreise in die Zukunft: Zweitägiger Workshop des Bundesarchivs und der Landesarchivverwaltung Rheinland-Pfalz*, June 24-5. https://lav.rlp.de/wir-ueber-uns/pressemitteilungen/detail/ki-und-archive-eine-zeitreise-in-die-zukunft.

Tekgürler, Merve. 2025. "LLMs for Translation?: Historical Texts and Contemporary AI Models." South African Centre for Digital Language Resources lecture series. *Youtube*, February 19. https://www.youtube.com/watch?v=59JMt-u-1os.

Vaidhyanathan, Siva. 2012. *The Googlization of Everything: (And Why We Should Worry).* University of California Press.

Veel, Kristin. 2018. "Make data sing: The automation of storytelling." *Big Data & Society*, 5(1). https://doi.org/10.1177/2053951718756686.

Widegren, Johannes. 2024. "AI for improving access to archives pertaining to the Sámi: An overview of current approaches and future possibilities." Presentation at *DHNB 2024: Digital Humanities in the Nordic and Baltic Countries 8th Conference*, May 27-31. https://urn.kb.se/resolve?urn=urn:nbn:se:lnu:diva-129846.