CALMA — The Center for Advances in Libraries, Museums, and Archives

AI and The Future of Holocaust Research & Memory

Benjamin Charles Germain Lee
Robert M. Ehrenreich
Todd Presner
Michael Haley Goldman
Jana Keck
Victoria Grace Richardson-Walden
Abigail Potter
Alexis Lerner
Andrew Dean
Anna Bonazzi
Ulysses Pascal
Aileen Tang
Sophia Toubian
Alex Wasdahl

March 2026

DOI: https://doi.org/10.6069/cx72-s697

Table of Contents

Introduction

Benjamin Charles Germain Lee

Artificial Intelligence (AI) is impacting all areas of academia. Scholars are utilizing AI in every component of research, including the collection, synthesis, analysis, and visualization of data; the drafting of articles; and the "peer" review of final submissions. The field of Holocaust studies is not immune to these trends. The question is how will the advent of AI impact the future of Holocaust studies? Will it provide new methods for analyzing data and displaying information for research and education that will benefit the field, or will the reduction of victim data to datasets and the problems of accuracy, distortion, and the stochasticity yet again strip people of their humanity?

In May, 2025, Robert M. Ehrenreich (Director of Academic Research and Dissemination at the United States Holocaust Memorial Museum), Todd Presner (Professor of European Languages and Transcultural Studies at the University of California, Los Angeles), and I (Assistant Professor in the Information School at the University of Washington) convened a workshop and public symposium at the University of Washington. The workshop evolved out of conversations that Robert, Todd, and I had been having periodically over the past decade. We recognized the intrinsically multidisciplinary nature of the issues and challenges before the field, including the wealth of digitized collections in need of new modes of discoverability and analysis. Accordingly, we convened experts in the fields of computer science, information science, history, sociology, anthropology, Jewish studies, museology, material culture, media and communication studies, literary studies, and art history:

- **Andrew Dean**, Lecturer in English, University College London
- **Michael Haley Goldman**, Executive Director, New Hampshire Humanities
- **Jana Keck**, Research Fellow, Digital History, German Historical Institute Washington
- **Alexis Lerner**, Assistant Professor of Political Science, United States Naval Academy

- **Miriam Posner**, Associate Professor, Information Studies, University of California, Los Angeles
- **Abbey Potter**, Senior Innovation Specialist, LC Labs, Library of Congress
- **Victoria Grace Richardson-Walden**, Professor of Digital Memory, Heritage, and Culture, University of Sussex
- **Noah Shenker**, Associate Professor of Jewish Studies and Film and Media Studies, Colgate University
- **Aileen Tang**, Data Scientist, AI and Cultural Heritage Lab, University of California, Los Angeles
- **Sophia Toubian**, PhD candidate, Information Studies, University of California, Los Angeles
- **Mike Trizna**, Data Scientist, Smithsonian Data Science Lab

Over the course of two full days, we investigated the following themes:

1. ***AI and Holocaust Studies research***: How can the advent of AI technologies advance research, and what are the perils of applying AI technologies to Holocaust studies in terms of accuracy, tone, reliability, and reproducibility?
2. ***AI and libraries, archives, and museums***: Can AI increase accessibility, discoverability, metadata generation, cataloguing, authentication, and preservation; what are the responsibilities of libraries, archives, and museums in using AI technology and presenting AI output; and what resources and expertise are required?
3. ***Limits of representation and reception***: Can AI be used to accurately and authentically perpetuate and promote Holocaust memory in education and public memory, and what are its limits and taboos?
4. ***AI and computational sciences***: What is the current state of AI technology, and what does the future hold? How can humanities scholars and computer scientists work together to forge responsible developments and applications of AI?

The contributions included in this paper have been written by participants and organizers in the months following the event. They include reflections, provocations, and refusals.[1]

In the first contribution, **Robert M. Ehrenreich** and **Michael Haley Goldman** examine the decline of traditional web search in favor of search queries using large language models and the consequences for Holocaust Memory.

The following three contributions ask how we might understand the role of AI in studying the Holocaust. **Todd Presner** and members of his AI and Cultural Heritage Lab, including **Anna Bonazzi**, **Ulysses Pascal**, **Aileen Tang**, **Sophia Toubian**, and **Alex Wasdahl**, investigate the question of what LLMs seem to understand about the Holocaust, arguing that what they primarily capture is the notion of the Holocaust *as vibe*. **Jana Keck** investigates how generative-AI models *fill in the gaps* of Holocaust records and the dangers posed by these interpolations. **Victoria Richardson-Walden** studies the outputs of AI-generated images pertaining to Holocaust representation and argues that they are *post-representational* in nature, "both beyond representation and yet deeply connected to it."

The subsequent three contributions explore practical uses of AI in Holocaust research and memory. **Alexis Lerner** explores AI as an analytical tool and its role within Holocaust pedagogy. **Abbey Potter** advocates for responsible stewardship within libraries, archives, and museums when using AI. **I** reflect on the concept of AI refusal and what forms of AI might benefit the field.

Lastly, **Andrew Dean** reflects on the role of forgetting in Holocaust memory and how it relates to technologies of remembrance.

Acknowledgments

We are grateful to the Jack, Joseph and Morton Mandel Center for Advanced Holocaust Studies at the United States Holocaust Memorial Museum; the University of California, Los Angeles; and the Information School, the Stroum Center for Jewish Studies, the Simpson Center for the Humanities, and the Center for the Advances in

[1] As defined by Violet Fox in "AI Refusal in Libraries: A Starter Guide" (2025): "AI refusal can refer to a spectrum of approaches to AI, whether that's refusing to use AI tools entirely, refusing the use of AI as much as possible, refusing to prioritize the use of AI, refusing to accept either boosterism or doomer narratives, refusing to accept the idea that AI is inevitable, or some other refusal."

Libraries, Museums, and Archives (CALMA) at the University of Washington for co-sponsoring the workshop and public symposium that inspired this white paper. We are also grateful to CALMA for providing a pathway for making this document available to the public.

We welcome feedback and questions from readers.

Holocaust Memory and the Death of Web Search

Robert M. Ehrenreich and Michael Haley Goldman[2]

Holocaust scholars, educators, and museum professionals have proclaimed the dire negative impacts of new genres and technologies to Holocaust history, representation, and memory for over 40 years. Literary scholars first began discussing the appropriateness, inappropriateness, and risk of Holocaust literature and representation in the 1980s.[3] Then came the boom in Holocaust movies, attracting ever larger audiences and making such discussions about literature seem almost quaint. A similar debate occurred with the introduction of the Shoah Foundation's *New Dimensions in Testimony* (currently appearing in a number of Holocaust museums and memorials worldwide simply as *Dimensions in Testimony*), which give the illusion of in-depth discussions with survivors. Now there are Generative-AI applications that allow users to "converse" with avatar-representations of victims, such as the Anne Frank chatbot.

With the growth of Large Language Models (LLMs), Holocaust scholars and educators worry that people are increasingly obtaining answers to their Holocaust queries via generative AI versus traditional search engines that provide links to reliable sources. The unease is warranted. AI is changing the way people find material on the Internet. Almost 80 percent of arts and cultural organizations have witnessed decreases in visitation via traditional searches (e.g., organic Google searches) during the first six months of 2025,[4] and sessions for the USHMM's online *Holocaust Encyclopedia* originating from ChatGPT referrals, for example, have increased dramatically for January-September 2025 when compared with the same period the previous year. The actual impact of this emerging landscape of visitation to Holocaust online reference resources is difficult to assess, however, since we cannot determine how many people saw the source links attached to the LLM

[2] The views expressed in this article are those of the authors and do not necessarily reflect those of the United States Holocaust Memorial Museum or New Hampshire Humanities.
[3] Berel Lang, *Writing and the Holocaust* (New York: Holmes & Meier, 1989).
[4] Beth Downs and Chris Unitt, "Report: The Impact of AI Overviews in the Cultural Sector," One Further (blog), September 1, 2025. p. 8: https://www.onefurther.com/blog/ai-overviews-cultural-sector-report.

responses but were sufficiently satisfied by the summaries to preclude visiting the referred sites.

The root of the concern is that these increasingly popular LLM summaries do not meet the same standards of accuracy, rigorousness, and respect for Holocaust history and memory previously grounded in public and classroom presentations by survivors and liberators—who sadly are becoming increasingly rare in this post-Holocaust era—and championed in the sources provided by institutions like the USHMM and Yad Vashem, which expend sizeable resources to ensure their work meets the highest academic standards and the greatest respect for this history. Although some level of artistic license is permitted in fiction and film, traditional web searches are expected to offer a range of credible sources that can be assessed and interrogated by the reader. Could the allure of the probabilistically conceived, conversationally toned, academically unvetted, yet seemingly authoritative answers provided by LLMs deluge searchers with inaccurate, misleading, distorted, or even antisemitic information that they would be unable to discern and unwittingly accept as true and rigorous? Put simply, have LLMs usurped expert authority? This concern is compounded by the ubiquitous nature of generative AI. Never before has a medium been able to reach so many people so easily. Unlike movies or books, in which people have to actively decide to participate, the growing reach of LLMs into daily life increases the possibility of the spread of disinformation at rates never before seen.

Although Holocaust scholars, educators, and museum professionals could not slow the development and deployment of previous Holocaust-related genres and media, they were able to mitigate many of the more egregious issues by developing and propagating effective and powerful educational strategies, reinforcing these strategies with reliable resources, optimizing websites for digital searches, and working with internet companies to rectify potential problems with search results. For example, the USHMM conducted a major initiative in roughly 2000 to combat the rise of antisemitic and Holocaust-denial sites in early search-engine results. Searches initially tended to return malicious sites in answer to such questions as “Did the Holocaust really happen?” because search-engine developers and trusted Holocaust-related sites never thought to account for such queries. Once aware of the situation, the Museum expanded its online resources, including the precursor

to the current version of its online *Holocaust Encyclopedia* (i.e., the Wexner Learning Center online). It also worked with Google, through the Google Grant Program, to strengthen the presence of USHMM's website in search results. Could similar methods be employed in response to the rise of LLMs?

A coordinated effort is required to explore the problem-space of learning about the Holocaust with and around generative AI by creating test sites worldwide for creative experimentation in learning. Each site should design approaches to balancing the issues raised by LLMs in Holocaust education with a thorough grounding in Holocaust pedagogy and LLM prompting in conjunction with current efforts by the digital-humanities community to wrestle with the implications of AI more broadly. The research should be conducted in secondary school and university classroom settings, incorporating students of Holocaust history as well as in affiliated and unrelated fields. The results should then be evaluated, analyzed, and compared across sites, concluding with a report detailing which methodologies were most effective and recommending potential ways in which the more egregious issues can be mitigated through content-focused partnerships with companies developing LLMs and user-focused, pedagogical methods aimed at LLM prompting.

The Dangers of Holocaust Vibe

UCLA AI and Cultural Heritage Lab
(Todd Presner, Director; Anna Bonazzi, Ulysses Pascal, Aileen Tang, Sophia Toubian, and Alex Wasdahl)

Large Language Models (LLMs) like ChatGPT, Gemini, Claude, and Grok are increasingly shaping public understandings of history. Recent studies have even predicted the 'death of the web' because users are relying on AI summaries, AI-generated answers, and AI chatbots rather than visiting authenticated websites.[5] But what do LLMs *know* about the Holocaust and, more significantly, how do they *narrativize* the Holocaust? Answers to these questions are particularly urgent for a field shaped by debates over responsible forms of narrative emplotment, questions over the limits of representation, the risks of aestheticization, the interconnection of traumatic memory and history, and the politics of comparison.

Unlike human historians who are accountable to disciplinary conventions, professional ethics, communities of memory, and responsibilities to the dead, LLMs generate historical narratives and images through myriad calculations and probabilistic predictions. These not only defy scrutiny but are ultimately accountable to no one but themselves, as their ethical guardrails are not public, and they are primarily meant to protect their parent companies from legal liability. In our previous research, we found that while LLMs often (but certainly not always) provide factually accurate accounts of well-documented events, they also exhibit concerning behaviors: stochastic omissions that leave out or downplay critical information and nuance, an optimism bias that tends to reframe atrocities through redemptive narratives and moralizing lessons, sanitized language that tends to avoid politically charged terms, and a propensity to reproduce stereotypes and flatten complex histories through forms of "uncritical fabulation."[6] LLMs commonly

[5] "AI Is Killing the Web: Can Anything Save It?," Economist, July 14, 2025, https://www.economist.com/business/2025/07/14/ai-is-killing-the-web-can-anything-save-it.

[6] Cf. Todd Presner, Anna Bonazzi, Ulysses Pascal, Aileen Tang, Sophia Toubian, and Alex Wasdahl, "Uncritical Fabulation and Holocaust Vibe: How AI Constructs the Past," *Digital Memory Dialogues* 2, no. 4 (2026). https://doi.org/10.20919/SAIM6630.

produce what we call “history as vibe”—coherent-sounding (or looking) responses that simulate authority but lack epistemic grounding, aesthetic intentionality, and ethical reflexivity. In other words, the “feeling” may strike the reader as right, but, ultimately, the AI-generated response reifies stereotypes, reinforcing biases or even historical inaccuracies because its adherence to accuracy is coincidental, not foundational.

The term “vibe” is widely used today as a shorthand for mood, style, feeling, or atmosphere. Its lineage stretches from the physics of sound vibrations and engineering applications to the percussive music of the vibraphone in the early-20th century and its later uptake in popular music (“Good Vibrations”). In the context of Holocaust representations, one might argue that certain films, such as *Schindler’s List*, also produced “vibes” -- feelings of what atrocity and trauma might have looked and felt like. Or one might further argue that a project like *Dimensions in Testimony* (interactive, holographic-like “conversations” with real survivors produced by the USC Shoah Foundation) gives rise to the vibe or feeling of talking to the dead. For the latter, this effect is achieved through the careful choreography of questions, answers, and gestures during the interview process as well as the extensive human training of the machine learning algorithm that selects answers from a large database of prerecorded answers when users ask questions on the platform.[7] Here, a vibe is an affective impression that emerges from certain epistemic commitments, ethical decisions, aesthetic choices, and intentional hermeneutical labor.

Today, algorithms operationalize the vibe: Google Maps offers “vibe checks” of neighborhoods. “Vibe coding” allows developers to create applications through conversational interactions with LLMs rather than typing code. What is new about AI is that it produces vibes about historical events that can pass as authentic representations but without any experiential, intellectual, or emotional understanding of or investment in those events. An AI has no epistemological, aesthetic, or ethical commitments; instead, it has probabilities based on the vectorization of a massive set of cultural artifacts. When an LLM produces an

[7] See Todd Presner, *Ethics of the Algorithm: Digital Humanities and Holocaust Memory* (Princeton: Princeton University Press, 2024), ch. 7.

account of the Holocaust, the vibe is the product of a hyper-rationalized, black-box logic involving hundreds of millions of calculations.

AI-produced "Holocaust vibe" is not the result of representation; rather, it is the byproduct of prediction, optimization, and statistical generalization. The conventional work of historical representation has been transformed into predictive pattern-matching. The AI randomly cannibalizes available narratives, styles, and data of the past to produce accounts that feel coherent and emotionally resonant but are completely unmoored from experience, intention, responsibility, and media. As our team has argued previously, through their training data, attention mechanisms, and learnt weights, LLMs abstract and assemble statistically likely patterns from what has already been 'said'.[8] The result is stereotyped vibes about the past marked by a new depthlessness and lack of historicity, not unlike Fredric Jameson's sense of pastiche: stylistically imitative while stripped of critical reflexivity, aesthetic intentionality, and ethical accountability.[9]

This is because the language of AI skews toward template-like answers, with standardized narrative arcs and consumable, on-demand content filled with overused metaphors ("humanity's darkest chapter"), tidy moral lessons ("a powerful symbol," "an important lesson"), and a marked optimism bias ("resilience," "beacon of hope"). Historical specificity and particularity are flattened and homogenized. The LLM ultimately relies on the law of averages to locate the central tendency, creating linguistic coherence devoid of understanding. The resulting "vibe" functions like generic stock footage in a film—providing atmosphere and recognizability while smoothing over dissent, contradictions, and less commonly heard voices while training the user to equate historical knowledge with reassuring familiarity. In the worst case, when it lacks appropriate training data and source material, the AI simply makes up plausible-sounding narratives, names, accounts, and explanations—all wrapped in the appropriate affective cadence.

[8] For a fuller discussion, see: Presner et al. "Uncritical Fabulation and Holocaust Vibe: How AI Constructs the Past"

[9] This is not unlike how Fredric Jameson describes the aesthetic styles of "postmodernism." See Fredric Jameson, "Postmodernism, or the Cultural Logic of Late Capitalism," in *Postmodernism: A Reader*, ed. Patricia Waugh (London: Routledge, 2016), 62–92. Here, pp. 17-18.

When “vibe” becomes the primary measure of success, we risk replacing the nuanced, dialogical, and open-ended work of historical representation with plausible-sounding, emotional simulations. We need to critically read, measure, and assess these AI-generated vibes to understand both their appeal and their dangers, and to ask what is gained or lost when the past comes to us as an algorithmically generated approximation or a probabilistic feeling.

Generative AI and Authenticity

Jana Keck

Generative AI is increasingly used to imagine the unimaginable. In the context of Holocaust research and memory, this poses both technical challenges and profound ethical questions. *Diffusion-model-based systems*, such as MidJourney and Stable Diffusion, have enabled the creation of hyper-realistic portraits and fictional scenes that never existed. These tools draw from vast and largely uncurated training datasets scraped from across the web, including copyrighted material, manipulated imagery, and culturally biased aesthetics. Their use raises the urgent question: What does *authenticity* mean when images are being generated—which bear different names ranging from synthetic fabrications and deepfakes to AI slop—that simulate source-like representations and are reshaped to meet the aesthetic and media expectations of our present?

Cultural heritage institutions are already experimenting with AI to "fill in the gaps" of historical records—particularly where no visual material exists.[10] Some projects have used AI to colorize black-and-white photographs of victims. Others have gone further, generating synthetic images or videos of Holocaust survivors as imagined children or teens, using adult photographs as input. These synthetic reconstructions are often presented as “emotional bridges” between past and present, meant to foster empathy, particularly with younger generations. But this empathy comes at a cost. The generated images are rarely labeled clearly or contextualized critically in the physical or digital exhibition space. The result is a distortion of historical realities, aesthetic flattening, and a profound breach of representational responsibility. The effects that may be identifiable to experts but are often imperceptible to children and general audiences. This risks reinforcing what Presner et al. in the previous contribution describe as a “Holocaust vibe”: even when images *seem* right and feel plausible to these audiences, that does not make them authentic. Moreover, as the technology advances, it will become increasingly difficult even for trained eyes to distinguish fact from fabrication.

[10] To respect the anonymity of specific projects and avoid assigning public blame, I do not name them here; instead, I draw on a range of applications to formulate general observations about best practices—and what to avoid—when using Generative AI in Holocaust-related context.

Even when many of these images *seem* right, much about them does not feel right at all. This is particularly troubling in the case of women: diffusion models repeatedly render them as implausibly young girls aligned with Western beauty standards: white, slim, conventionally attractive, and often subtly sexualized. This is not merely an aesthetic issue but a structural effect of training data and model design, which reproduce internet biases—social as well as technical—rather than historical specificity.[11] The output does not only privilege a narrow cultural type; it also helps circulate unrealistic body ideals. Such biases risk re-inscribing harmful stereotypes under the guise of education or engagement.

Holocaust history is deeply tied to processes of dehumanizing *datafication* (lists, photographs, records, numbers). There is also a deeper discomfort: the very idea of using Holocaust victims as training material.[12] Does feeding these traces into generative systems that merge them with millions of unrelated images (often without consent or contextual control) risk stripping them of their historical specificity once again? Even well-intentioned attempts to make victims look "more Jewish" through prompt engineering risk essentializing identity and aesthetics. In addition, they raise critical issues about the impulse to "fix" people: an impulse that evokes not only twentieth-century eugenic ideologies but also twenty-first century obsessions with beauty-driven correction and enhancement.

AI will not disappear. On the contrary, social media executives increasingly expect AI-generated content to dominate platforms in the near future. On January 13, 2026, the "Network Digital History and Memory" published an open letter raising their concern about the surge of "AI-slop" about National Socialism: they distort history, trivialize the past, and erode the trust in "authentic" historical documents.[13] They have the power to fuel doubt: *did this really happen*? These activities are not

[11] Cf. Joy Buolamwini and Timnit Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," Proceedings of Machine Learning Research 81 (2018): 77–91, https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf

[12] Todd Presner, *Ethics of the Algorithm: Digital Humanities and Holocaust Memory* (Princeton: Princeton University Press, 2024).

[13] Netzwerk Digital History and Memory, "Offener Brief: Konsequentes Vorgehen gegen KI-generierte Holocaust-Verfälschungen auf Social-Media-Plattformen," Gedenkstättenforum, January 13, 2026, https://www.gedenkstaettenforum.de/gedenkstaetten/netzwerke/netzwerk-digital-history-and-memory/offener-brief-konsequentes-vorgehen-gegen-ki-generierte-holocaust-verfaelschungen-auf-socia-media-plattformen.

only driven by economic interests (sensational posts generate more engagement on platforms—or “clicks”—that are monetized by the platform via advertising, etc., and paid to the poster) but also by motives that fuel Holocaust denial and distortion.

The central question, therefore, remains: should cultural-heritage institutions use generative AI to imagine what cannot be seen? Is a distorted image better than no image at all, and if it is not authentic, should it be used at all, even if it is labelled as such? But is it authentic when it awakens an emotion? Authenticity is not only a matter of factual correspondence: images can feel “authentic” because they elicit emotion, even when they are synthetic. If we defined authenticity strictly as “only what really happened,” we would also have to dismiss fiction—from novels and TV series to animated documentaries and graphic narratives—as illegitimate for historical understanding. These tensions (and questions) require further research and sustained conversation among scholars, cultural heritage professionals, and the public; however, the first question can, for now, be answered clearly: *not* with current tools and practices. Instead of relying on "generalized" AI platforms to (re-)imAIgine the lives of the lost, cultural heritage institutions might turn to human-created art rooted in testimony and research. For example, artist and cartoonist Amy Kurzweil has drawn evocative illustrations of her grandfather, a Holocaust survivor, to represent the undocumented fragments of his life. This human-mediated approach preserves interpretive nuance without the illusion of visual accuracy.

Nevertheless, if institutions decide to use generative AI, we recommend three principles:

1. Do not use diffusion models like *MidJourney* for generating Holocaust-related visuals.
2. Do not use generative tools for life stages or individuals with no historical record unless clearly marked as speculative.
3. And always, always prioritize transparency. AI-generated content must be labeled as such—clearly, accessibly, and unambiguously. Words like “with the help of AI” are not enough. Transparency also requires clear documentation of the models, methods, and data involved including above all any

partnerships, especially when external collaborators were commissioned to create the AI-generated *content*.

If we are to use generative AI in Holocaust research and memory work, we must do so with curatorial responsibility—but also with tolerance for making mistakes.

Post-Representation

Victoria Grace Richardson-Walden

AI poses a fundamental challenge to Holocaust research and memory, but arguably not in the way that has penetrated popular discourse. We should be wary of the 'AI hype,'[14] from adopting AI uncritically as the norm to blaming it for any perceived crisis in Holocaust representation, memory, and education. To really understand what AI might or might not be good for in the context of Holocaust research and memory, we first need to develop an appropriate discourse to discuss it.

For so long, media engaging with this history has been discussed through the lens of 'Holocaust representation'. Yet, if we are to understand media representation as systems that enable people to "build up a culture of shared understandings" so they can "interpret the world in roughly the same ways,"[15] then generative AI does not operate in this way.

We can contrast generative AI with the broadcast media through and with which Holocaust representation developed: television and film outputs all involved deliberate processes of selection, construction, and editing to present specific messages about the past. Thus, we have seen academic volumes pertaining to national Holocaust cinemas, suggesting that particular national groups have a shared understanding of this past.[16] Yet, generative AI platforms like *MidJourney* do not synthesise their images in the same way. They rely on mathematical modelling, translation of meaningful language-based prompts into tokens to be processed through probability algorithms and then amalgamated into visual presentations. *MidJourney* neither understands nor makes meaning through the semantic values or representational codes that it encounters. Rather, what it presents are semantic

[14] Alva Markelius et al., "The Mechanisms of AI Hype and Its Planetary and Social Costs," *AI and Ethics* 4, no. 3 (August 2024): 727–42, https://doi.org/10.1007/s43681-024-00461-2.
[15] Hall, Stuart. (1997) *Representation: Cultural Representations and Signifying Practices.* SAGE, 4.
[16] Examples include Judith E. Doneson, *The Holocaust in American Film*, 2nd ed. (Syracuse, NY: Syracuse University Press, 2002); Olga Gershenson, *The Phantom Holocaust: Soviet Cinema and the Jewish Catastrophe (New Brunswick*, NJ: Rutgers University Press, 2013); Marek Haltof, *Polish Film and the Holocaust: Politics and Memory* (New York: Berghahn Books, 2012); Giacomo Lichtner, *Film and the Shoah in France and Italy* (London: Vallentine Mitchell, 2008); Judith Petersen, "The Holocaust on British Television: Shaping Collective Memory Since 1945" (PhD diss., University of Southampton, 2006).

diffusions—distinct representation codes detached from their cultural context and rearranged through mathematical logics.

The representational approach of mass media (e.g., film, television, and radio), then, offers a particular way of organising the world, but it is not the only one. Recognising this would encourage us to deter from using language like 'understanding', 'meaning', 'producing', 'language', and 'image', when referring to generative AI outputs, and instead think about how 'semantic diffusions' and 'mathesis' offer specific ways of organising knowledge about and memory of the Holocaust. This is not the same 'mathesis' that Foucault identified in the Classical ordering of things: applying rationale, scientific thought to both the sciences and classification in the humanities. Here, 'mathesis' is detached from rationality and reason.[17]

That is not to say that AI outputs cannot be read through a representational lens. Indeed, we have been encultured to do this: we see something visual, and our almost unconscious response is to start deconstructing its cultural values and meaning. Nevertheless, we must consider how AI's specific mathesis logic might complement or conflict with our representational tradition. AI models at an operational level are distinctly *non-representational.* Nevertheless, it would be better to understand their outputs as *post-representational* – they are both beyond representation and yet deeply connected to it. They do not attempt to construct culturally rooted meaning, they are untethered from localised cultural contexts in terms of their operations, whilst relying on training data that is deeply embedded in a specific culture. Then, when interpreted by users, these users only have their own epistemological references through which to read the outputs (relying on their known representational systems).

Reconceptualising AI in these terms—on its own terms—encourages us to confront a particular question: what can mathesis do for Holocaust memory and research? How does it produce different ways of organising and conceptualising both the past and how we remember it? To what ends can these distinctions complement or challenge existing paradigms within our field(s)?

[17] Michel Foucault, *The Order of Things: An Archaeology of the Human Sciences* (London: Routledge, 1994).

We can begin to address these questions in two ways:

(1) **Practical:** If generative AI (and indeed AI and machine learning models more broadly) are grounded in mathesis logics, then we should apply them to contexts where probability and stochasticism would be useful. This is not in the reproduction of testimony (as we see demonstrated in ChatBots like Deep AI's "Anne Frank"[18]) or in the remediation of conversations with survivors (as with *Dimensions in Testimony*). Such approaches echo the illusionary simulation of human intelligence, which gives a false impression of what AI actually are (and I emphasise the plural here, not *is*).[19] AI models then would be most useful to Holocaust memory-work where big computational data sets need sorting or analysing (although it is important to recognise that it is not simply a case of putting the human-in-the-loop, but rather that these are human processes, into which AI can be 'in-the-loop'). Presner et al.'s (2024) call for an "integrated methodology" is crucial here.[20]

(2) **Literacies:** Both in the classrooms and public spaces where Holocaust education is delivered as well as behind-the-scenes in museums, memorial spaces, and other organisations who do the delivering of this education, digital literacies need to be embedded in practice. That is, anyone considering using AI in this context, or indeed wanting to be prepared to teach those engaging with it, needs to understand what they are—their logics—and be able to approach them critically and not just reactively. The recently revised IHRA Recommendations for Teaching and Learning about the Holocaust (2026) state that digital literacies are now a core competency in Holocaust education. However, we must not simply dictate what is 'appropriate' or 'inappropriate' use of such new technologies, rather we need to encourage critical reflection on the possibilities and challenges that their specific logics offer, and the variety of ways they can be used. This does mean serious work is needed in reconfiguring Holocaust education to be fit for the 21st Century—introducing mathesis as a logic of knowledge

[18] https://deepai.org/chat/anne-frank

[19] Such concerns can be traced back in the literature as far as Josef Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: W.H. Freeman, 1976).

[20] Todd Presner, *Ethics of the Algorithm: Digital Humanities and Holocaust Memory* (Princeton: Princeton University Press, 2024),

organisation and challenging existing paradigms not only about the potential limits of representation, but whether representation is in itself the only way to conceive of accessing data about the Holocaust.

Responsible Stewardship

Abigail Potter

Introduction: A Critical Juncture for Historical Memory

The stewardship of historical memory, particularly the legacy of the Holocaust, has reached a critical inflection point. As the last of the eyewitness generation passes away, artificial intelligence is reshaping how history is documented, shared, and understood. As Holocaust Studies comes to rely exclusively on primary sources, the very technologies that promise to "unlock" these collections also introduce fundamental questions about evidence, interpretation, and truth in a public sphere already rife with mistrust of experts and institutions.

While this moment is fraught with uncertainty, it is not without precedent. Libraries, archives, and museums (LAMs) have decades of shared experience navigating profound digital transformations, from machine-readable cataloging to mass digitization. This collective expertise provides a solid foundation for developing responsible, human-centric AI strategies. By acting thoughtfully and collaboratively, we can harness the power of this technology while upholding our basic mission: to ensure the integrity of the historical record for generations to come.

The Fraught Present: Why AI Challenges the Foundations of Research

AI presents an opportunity to transform workflows and deepen public engagement, but it also amplifies the realities of stewardship. LAM collections are always incomplete, representing only a fraction of the historical record. They are also unevenly digitized. It is typical for major international libraries to have less than 10% of their collections digitized and available online.[21] For archives and museums,

---

21

Maureen Pennock and Michael Day, "Managing and Preserving Digital Collections at the British Library," in *Managing Digital Cultural Objects: Analysis, Discovery and Retrieval*, eds. Allen Foster and Pauline Rafferty (London: Facet, 2016), 111–28, accessed February 16, 2026.
German National Library, Facts and Figures – Digitization: Annual Report 2024 (Frankfurt am Main: German National Library, 2024), accessed February 16, 2026, https://jahresbericht.dnb.de/Webs/jahresbericht/EN/2024/ZahlenUndFakten/zahlenFakten_node.html#doc2152336bodyText3.

the percentage is even smaller. This means that the vast majority of our holdings—most of which are under copyright—are unavailable to train the AI models that are increasingly shaping the public's understanding of the past.

This dynamic introduces a series of additional risks to the research process, on top of those already outlined above:

- Data Bias and Incompleteness: The limited and skewed nature of the data used to train large AI models creates a distorted knowledge-base. Because the vast majority of LAM holdings are not digitized or are under copyright, AI models are trained on an incomplete slice of human history, not to mention that LAMs do not contain a perfectly representative or complete historical picture. Scholars have repeatedly documented how LAMs fall short of evenly or fairly representing marginalized communities.[22] LAMs do have professional standards and ethics for systematically documenting, cataloging their information, and sharing the evolution of collection and digital policies, which inform how the historic record is understood and the biases that the collections reflect. AI tools that do not include the collections of institutions that have been the (imperfect) repository of vast and historic stores of knowledge for hundreds of years distort the foundation of research.
- Disruptive AI Marketplace: Commercial AI products are widely available, but they have not been tested for LAM-specific use cases and are not designed for workflows demanding error-free, accountable operations. In previous digital transformations, LAMs have had to develop and share their own technical standards and software to meet their needs.[23] Open source LAM or

[22] Hope Olson, "The Power to Name: Representation in Library Catalogs," *Signs: Journal of Women in Culture and Society* 26, no. 3 (Spring 2001): 639–68.
Lori Jahnke, Kyle Tanaka, and Chris Palazzolo, "Ideology, Policy, and Practice: Structural Barriers to Collections Diversity in Research and College Libraries," *College & Research Libraries* 83, no. 2 (March 2022): 166, https://doi.org/10.5860/crl.83.2.166.
[23] Federal Agencies Digital Guidelines Initiative, "Technical Guidelines for Digitizing Cultural Heritage Materials: Third Edition," by the Still Image Working Group (Washington, DC: FADGI, 2023), accessed February 16, 2026, https://www.digitizationguidelines.gov/guidelines/digitize-technical.html.
Jenny Mitcham, "Introducing the New NDSA Levels of Preservation," Digital Preservation Coalition (blog), November 14, 2019, accessed February 16, 2026, https://www.dpconline.org/blog/introducing-the-new-ndsa-levels-of-preservation.

research-focused AI tools are few, and the community to develop the tools and software is immature and small. Further complicating the development of technical expertise and understanding that would support shared quality standards for LAM AI tools and data are the constant releases of new AI tools, breakthroughs, or paradigms. Keeping abreast of the growing AI products, options, and functions is extremely challenging. There are also barriers to using some AI tools. Many AI tools are not approved for use on networked IT systems because they are not compliant with security rules and privacy laws. Yet these are exactly the AI tools that are widely used by the public and have indiscriminately ingested as much data as possible, including copyrighted information and data from "chronically underfunded" libraries.[24] These factors further disrupt the development of trustworthy AI services in LAMs.

A Framework for Responsible Stewardship: The Library of Congress AI Planning Framework for Experimentation

The Library of Congress's LC Labs AI Planning Framework for Experimentation offers a practical guide for how libraries, archives, and museums (LAMs) can use AI responsibly.[25] The framework is based on several years of hands-on experimentation and evaluation of AI tools and approaches to enhance a variety of cataloging, metadata generation, and summarization workflows. The results of these experiments were mixed and only after further tests and evaluations will use cases be implemented. Responsible AI is not about banning or avoiding new technology, but about actively understanding potential risks and benefits, testing assumptions and hypotheses through structured experiments and evaluations, and only implementing AI when it adds value.

---

[24] Aaron Schaffer, Will Oremus, and Nitasha Tiku, "Inside an AI Start-up's Plan to Scan and Dispose of Millions of Books," *Washington Post*, January 27, 2026, accessed February 16, 2026, https://www.washingtonpost.com/technology/2026/01/27/anthropic-ai-scan-destroy-books/.

[25] Abigail Potter, "Introducing the LC Lab Artificial Intelligence Planning Framework," *The Signal*, Library of Congress, November 15, 2023, https://blogs.loc.gov/thesignal/2023/11/introducing-the-lc-labs-artificial-intelligence-planning-framework/.

The core components of the LC Labs AI Planning Framework for Experimentation are an *Understand* phase and an *Experiment or Test* phase before the *Implement* phase for AI in order to provide important lessons-learned to institutions considering similar use cases and to signal to the research community a commitment to trustworthy collections in the age of AI.

Understand

The *Understand* phase of the LC Labs AI Planning Framework for Experimentation provides time, space, and activities to center LAM principles and missions; map human capabilities, needs, and expertise; and maximize long-term benefits to the public. The primary goal of AI use cases in LAMs is to *help* staff, stakeholders, and users, not replace or placate them. The deep subject matter and collections knowledge of curators, catalogers, and other professionals are essential for identifying relevant AI use cases, and, most importantly, evaluating the quality and accuracy of AI outputs. The *Understand* phase is also when the readiness of collection data for AI tools is analyzed. For good results with AI tools, often significant data preparation and review steps are necessary. This is also the phase to clarify how successful AI use cases will be understood and measured. Ideally, staff, stakeholders and potentially impacted communities are included in the design of AI workflows and tests, especially those that pose significant risks or benefits. AI use-cases that involve sensitive or unique content, AI that could impact large groups of people—particularly vulnerable people or communities, or AI tools that are open to unmoderated public use—can pose significant reputational, trust, security, and cost risks. While other AI use-cases like transcribing audio visual materials to make them more discoverable and usable will pose some risk, mitigation strategies could lower risks and have real benefits to users and researchers. The LC Labs AI Planning Framework for Experimentation offers worksheets to help determine risks and benefits of specific use cases and offers a rating scale inspired by the FADGI 5 star scale, associating higher risks to higher costs and longer timelines for testing.[26]

[26]See: https://github.com/LibraryOfCongress/labs-ai-framework/blob/main/Understand/Data_Assessment_Worksheet_2023-11-14.docx, https://github.com/LibraryOfCongress/labs-ai-framework/blob/main/Understand/PhaseII_Risk_Assessement_-_2023-10-30-draft.docx

Experiment

Robust Experimentation and Evaluation. The core of Responsible AI adoption is experimentation and testing. Public leaderboards that share performance metrics for various AI tools are often not relevant to the kinds of tasks and workflows where LAMs are seeking to apply AI. Only by testing specific AI tools with actual institutional data in real-world use cases will it be clear if the AI tool or process works well enough. The determination of what is good enough comes from the experience and knowledge of staff, users, and communities. Often, the experiment process is repeated with different data, approaches, or tools to determine which combination of AI model, data, and use case is optimal. There are meaningful differences in the functions, performance, and usage terms of AI tools. Understanding those differences and choosing tools that align with the policies and values of LAMs is a key part of using AI responsibly. Building a shared set of performance benchmark data for key LAM data and tasks with shared evaluation factors will help establish consistent, high-quality LAM AI implementations.

During the experimenting and testing phase, it is also essential to document the data transformations with the predicted and actual AI model performance for specific tasks. The Framework offers a Data Processing Plan as a model for this documentation.[27] It combines elements from a model card, data coversheets, and documents curatorial provenance.[28] This type of documentation informs future researchers as well as stewarding and contributing to content provenance and authenticity, all of which supports trustworthy collections.

To help guide the evaluation process, the Framework offers a guide for assessing the outcomes of AI experiments. The factors of Responsible, Effective, and Practical could include a multitude of questions and measures customized to specific

---

[27] See: https://github.com/LibraryOfCongress/labs-ai-framework/blob/main/Experiment/Data-Processing-Plan-template-2021-12-01-draft.docx

[28] Margaret Mitchell et al., "Model Cards for Model Reporting," in *Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency* (New York: ACM, 2019), 220–29, https://doi.org/10.1145/3287560.3287596.
Timnit Gebru et al., "Datasheets for Datasets," *Communications of the ACM* 64, no. 12 (December 2021): 86–92, https://doi.org/10.1145/3458723.

organizations or use-cases to account for such key factors as utilization of natural resources for compute and the integration of review and feedback from key user groups or communities.

Evaluation Factor	Questions to Evaluate
Responsible	Do the potential benefits outweigh the risks, and is the process compliant with all LAM values and policies (copyright, privacy, security)?
Effective	Does the output meet quality standards after being tested with LAM data and reviewed by LAM staff?
Practical	Is the process cost-effective, can it be integrated into existing infrastructure, and can quality be ensured over time??

Implement

LAMs operate in constraints. Budget, expertise, and deep engagement are often lacking in comparison to the volume, complexity, and constant digital transformations impacting the research and stewardship lifecycle. Experiments, hands-on demonstrations, and tests of AI tools garner useful performance indicators, as well as informing the grouping of use-cases and mapping of them into categories or domains, which is a helpful first step in understanding how to prioritize AI use cases.

Prioritization is necessary because opportunities and available budgets and staff need to be aligned at an institutional level to accelerate strategic goals and make efficient use of resources. The real long-term costs of AI systems are not yet known and not every task is appropriate for AI.[29] In an environment when choices need to be made, AI must be deployed selectively and responsibly. The LC Labs AI Planning

[29] For example, in workflows that require consistent, highly accurate output, AI may not be the right solution.

Framework for Experiments offers the following rubric for evaluating AI's potential impact for LAM use cases to guide higher level prioritizations. These ratings are done after a single or series of experiments have gathered detailed staff and user feedback and speak to organizational principles, strategic goals, and resource availability.

a. *Harmful*: Outcomes prove detrimental to staff, users, or stakeholders by reducing access, impeding rights, perpetuating errors, or eroding quality or trust.

b. *Inferior*: Risks, costs, errors, or vulnerabilities outweigh the value to users, staff or stakeholders.

c. *Iterative*: Outputs are mixed but lessons-learned can be applied to future development and risk mitigation strategies.

d. *Successful*: Output meets expected quality level while being compliant with privacy, rights, and security requirements.

e. *Innovative*: Successful outputs plus adding value for users, staff, or stakeholders by improving quality, satisfaction, efficiencies, or capacities.

f. *Transformational*: Meets or exceeds quality while significantly improving user stakeholder or staff capabilities of a workflow, product, or service.

The goal of this ranking system is to hold space, resources and aspirations for 'transformational' AI use cases and not spend all our time, effort, and resources on software or tools that are generally 'successful' in terms of individual productivity but not as potentially impactful for users, staff, or stakeholders.

A successful AI use case can be a lot of different things based on the context of implementation. Some uses of AI pose more risks to communities or users while offering efficiencies or enhanced access to others. Especially as AI technologies are new, constantly changing, and untested, implementing AI according to higher level principles can help guide implementation decisions. Returning to the basic goal of a

LAM is helpful: to ensure the integrity of the historical record for generations to come.

Building a Shared Vision for a Trusted Future

The ultimate goal for Responsible AI in Libraries, Archives, and Museums is to ensure that our collections remain trusted, identifiable, and discoverable. We must advocate for and help build AI tools that are open, transparent, and customized for the unique needs of the cultural-heritage sector. Above all, our efforts must be aimed at empowering researchers and the public, not misleading them.

This is not a task for LAM professionals alone. Historians, researchers, technologists, and the general public all have a part to play in shaping this future. We must engage in open dialogue about the values to be embedded in these systems and the standards to be upheld to maintain a trustworthy historical record.

AI in the Loop

Alexis Lerner[30]

Undergraduates know that generative AI hallucinates. This understanding drives a game that Naval Academy midshipmen play in my class: They pit the Large Language Models (LLMs) like ChatGPT, CoPilot, and Gemini against each other to see which embellishes first and with greater dramatic effect. The goal of the exercise is to make clear the dangers of permitting artificial intelligence to take over human critical thinking. Beyond the classroom, when generative AI concocts false details about historical events, the stakes become even higher. In May 2025, for example, the Auschwitz-Birkenau State Museum sounded alarms over Meta's allowance of Instagram accounts to produce AI-generated images of the camp and its prisoners.[31] Scholars like Benjamin Charles Germain Lee have also expressed concern that AI-driven ChatBots were making up details about Holocaust victim trajectories in the name of user engagement.[32]

This spread of misinformation is dangerous for two reasons: first, false details about the Holocaust could be accepted and shared as fact, leading to distortions of historical events; and, second, these hallucinations could discredit belief in the historical record altogether. After all, if ChatBots trained on real testimonies were producing falsehoods, users might begin to question something deeper: the integrity of the original testimonies themselves. In the case of the latter, our rich archival data on the Holocaust may be passed over in favor of photos and videos designed to grab sustained attention and generate click-throughs.

This is all further complicated by the fact that generative AI is human-built. Indeed, algorithms reflect our sentient biases and limitations, and are far from neutral.

---

[30] The views expressed here are those of the author and do not reflect the official policy or position of the U.S. Naval Academy, the Department of the Navy, the Department of War, or the U.S. Government.

[31] "Auschwitz Museum Sounds Alarm over 'Harmful' AI Images of Holocaust Victims," *The Times of Israel*, May 24, 2025, accessed November 3, 2025, https://www.timesofisrael.com/auschwitz-museum-sounds-alarm-over-harmful-ai-images-of-holocaust-victims/.

[32] Benjamin Charles Germain Lee, "Uncanny Testimony: AI, Holocaust Survivors & Memory," *Longreads*, September 25 2025, accessed November 3 2025, https://www.longreads.com/2025/09/25/ai-holocaust-survivors-memory/.

Moreover, the training data they rely on depends on what information is available. This means that Holocaust narratives that died with their source—those that were never recorded—may be completely excluded from these databases, reinforcing survivorship biases in a cyclical manner.

AI as Analytical Tool

I once heard Hank Greenspan talk about the power of teenagers hugging survivors in Holocaust museum conference rooms. Survivor narratives, shared firsthand, were of course powerful. But the electric connection of human touch was what cemented widespread commitment among listeners to preserving the memory of the Holocaust and protecting humanity from its repeat. As time outlasts all of us, our opportunities to hear these stories firsthand today have become painfully few. Almost paradoxically, artificial intelligence has entered the realm of Holocaust memory to fill that gap.

In some respects, AI can be a useful analytical tool. Scholars like Todd Presner write of using AI as a method for zooming both in and out of existing databases, allowing for close readings in the former or computational trend analysis of larger numbers of observations in the case of the latter.[33] Much has also been written on the USC Shoah Foundation's *Dimensions in Testimony* project—a holographic reenactment of a Holocaust survivor, based on footage of them giving testimony solely for this mode of delivery.[34] And when it comes to exceptionally large datasets like the Arolsen Archives, which holds over 30 million individual pages of primary source information on the Holocaust, AI allows for the analysis of documents that would take decades for a team to complete by hand.

But we must not allow machines to operate without human oversight, especially when it comes to maintaining intercoder reliability, recognizing cultural and linguistic intertextuality, and applying intuitive and substantive gut checks for accuracy. AI cannot have free rein; humans must remain central to scholarly inquiry, while keeping 'AI in the Loop.'

---

[33] Todd Presner, *Ethics of the Algorithm: Digital Humanities and Holocaust Memory* (Princeton, NJ: Princeton University Press, 2024).

[34] USC Shoah Foundation, *Dimensions in Testimony*, interactive-biography project, developed in association with Illinois Holocaust Museum and Education Center, USC Institute for Creative Technologies, and Conscience Display (Los Angeles: USC Shoah Foundation, 2014–).

AI in Holocaust Education Pedagogy

As educators, it is imperative that we teach students to navigate this ever-changing technological world with cautious optimism, a collaborative methodological mindset, and careful attention to ethics. With this objective in mind, I designed a multi-disciplinary course that I taught at the University of Toronto and the United States Naval Academy. In this course, students engage with Holocaust testimony from the USC Shoah Foundation's Visual History Archive.[35] After learning the historical context of the Holocaust and exploring the archive's structure, ethical limitations, and research potential, students apply a combination of analytical and *in vivo* codes to an assigned testimony. Their codes are compiled into a greater dataset—now at over 100 observations—that students then analyze computationally to identify patterns and outliers.

At this stage, students work in tandem with machines to systematically code—and carefully spotcheck—a much larger set of observations, expanding the scope of the project substantially. By keeping humans at the core and 'AI in the Loop', scholars can expand their analyses to cumbersome datasets while keeping generative AI, its predilection for hallucination, and the dangerous spread of misinformation in check.

[35] Alexis Lerner and Andrew Gelman, "In Pursuit of Campus-Wide Data Literacy: A Guide to Developing a Statistics Course for Students in Nonquantitative Fields," *Journal of Statistics and Data Science Education* 32, no. 3 (2024): 241-252, https://doi.org/10.1080/26939169.2023.2276844.

Sorting AI Out: Refusals & Futures

Benjamin Charles Germain Lee

My interest in Holocaust research and memory is a consequence of my family history: my grandmother survived Auschwitz-Birkenau concentration camp. It is from this vantage point, rather than an academic one, that I first became interested in and developed concerns with efforts to re-animate Holocaust survivors using AI.[36] It is also through my pursuit of familial history that I discovered an entirely different dimension of AI in relation to Holocaust memory: within the archive, a setting in which applications of AI have a long and productive history.[37] For decades, digitization pipelines have employed optical character recognition (OCR) algorithms for transcriptions; audio transcription tools make recorded materials more accessible; machine learning algorithms undergird the search and discovery platforms that enable patrons to browse digital collections. Archives pertaining to survivors and victims of the Holocaust are not exempt from such algorithmically mediated affordances. In fact, such archives are accessible precisely *because* AI has been adopted. To speak of the future of AI and Holocaust research and memory requires understanding this enmeshment of past and present.

Within libraries, a growing number of voices have begun advocating for the idea of AI refusal, which, broadly construed, encourages librarians, researchers, and patrons to recognize and consider their own agency in decision-making surrounding the application of AI in libraries.[38] As defined by Violet Fox, whose work has been at the forefront of this movement:

> "AI refusal can refer to a spectrum of approaches to AI, whether that's refusing to use AI tools entirely, refusing the use of AI as much as possible,

[36] Benjamin Charles Germain Lee, "Uncanny Testimony," *Longreads*, September 5, 2025, https://longreads.com/2025/09/25/ai-holocaust-survivors-memory/; Benjamin Charles Germain Lee, "Speaking for the Past," *Real Life Magazine*, August 12, 2021, https://reallifemag.com/speaking-for-the-past/.

[37] Benjamin Charles Germain Lee, "Machine Learning, Template Matching, and the International Tracing Service Digital Archive: Automating the Retrieval of Death Certificate Reference Cards from 40 Million Document Scans," *Digital Scholarship in the Humanities* 34, no. 3 (September 2019): 513–35, https://doi.org/10.1093/llc/fqy063.

[38] Violet Fox, "AI Refusal in Libraries: A Starter Guide," 2025, https://acrlog.org/2025/06/11/ai-refusal-in-libraries-a-starter-guide/.

> refusing to prioritize the use of AI, refusing to accept either boosterism or doomer narratives, refusing to accept the idea that AI is inevitable, or some other refusal."[39]

Most, if not all, of the contributions in this paper engage implicitly with AI refusal, probing where the limits of appropriate and responsible stewardship and scholarship lie in relation to AI, computational technologies, and datafication more broadly.

AI refusal, then, offers us a dialectical mechanism by which we can understand what value AI might hold. I suspect that the more "invisible" forms of AI (e.g., OCR, audio transcription, and search algorithms) that have been integrated into the core infrastructure of Holocaust research and memory withstand interrogation: despite these algorithms' well-documented limitations from the vantage point of algorithmic bias, the intrinsic value of making archives discoverable changes these questions from *if* the technologies should be used to *how* we can use them more responsibly.

In a world of rapidly evolving generative AI, it is clear that many of the contributors to this paper resist the encroachment of other forms of AI ingratiation: Ehrenreich and Haley Goldman question the integrity of Holocaust history in the face of AI-generated summaries replacing well-sourced search results; Keck points out the dangers of relying on generative AI systems to fill in the gaps of the historical record; Presner *et al*. document the limits of large language models and their shallow summarization best understood as "Holocaust as vibe." But amidst these refusals, what might other possibilities be for enriching our understanding of Holocaust history and memory with AI, in between the extremes of spell check and Anne Frank chatbots?

In speaking for myself, I remain optimistic about the capacity for AI to reshape how scholars and the public navigate the vast archives that cultural heritage institutions have so thoughtfully preserved—from the hundreds of thousands of hours of audiovisual testimony holdings to the hundreds of millions of pages of documents. Recognizing the expansiveness of these holdings and how it exceeds the capacity of any one human to navigate them is an opportunity to re-imagine how we draw

[39] Violet Fox, "A Librarian Against AI, or I Think AI Should Leave," https://violetbfox.info/against-ai/.

connections, chart associations, and understand affective experiences within the archive. That is not to say that we should adopt AI due to extrinsic pressures: as enumerated by Potter, there are clear and compelling reasons to resist this temptation. However, a dedication to AI refusal and its dialectical intent might enable us to conjure AI futures for Holocaust research and memory that will serve the survivors, victims, scholars, and patrons much in the same way that OCR and audio transcription have.

Forgetting

Andrew Dean

Hayim Josef Yerushalmi concludes his book *Zakhor* on Jewish history and memory with what he styles as 'Reflections on Forgetting.' Quoting Nietzsche, he wonders if "the unhistorical and the historical" might be "equally necessary to the health of an individual, a community, and a system of culture."[40] Expanding on these thoughts, Yerushalmi asks: "How much history do we require? What kind of history? What should we remember, what can we afford to forget, what must we forget?"[41]

Yerushalmi's provocative questions are only becoming more pressing in Holocaust remembrance as we confront what is known as the "post-survivor era." According to the Claims Conference, there are only around 221,000 living Holocaust survivors. Of these nearly all are classified as "child survivors" (those born between 1928 and 1946).[42] At the time of writing, the median age of survivors is 87.

While Yerushalmi imagined an inevitability to the process by which the past turns into history, contemporary technologies are being used to make the survivors newly and constantly present. The best-known and most comprehensive of these is the *Dimensions in Testimony* (DiT) project, run by the USC Shoah Foundation. As part of the project, several dozen survivors late in their lives gave interviews in a studio designed to record them in 360 degrees.

Noah Shenker and Dan Leopard describe in some detail their experiences of both using 'Pinchas DiT' (the interactive testimony of Holocaust survivor Pinchas Gutter) and observing others attempting to interact with it. In the event, some degree of user training has tended to be required, even if the system has advanced

[40] Yosef Hayim Yerushalmi, *Zakhor: Jewish History and Jewish Memory* (University of Washington Press, 1996), p. 107.
[41] Yerushalmi, *Zakhor*, p. 107.
[42] *Holocaust Survivors Worldwide: A Demographic Overview*, Claims Conference Center for the Study of Demography (Claims Conference: Conference on Jewish Material Claims Against Germany, 2025), pp. 1–38 (p. 4).

considerably since its first public release.[43] Nevertheless, they describe users in the space of the museum interacting with the projection as though the survivors were really in the room. There are senses of presence even true for users such as Leopard and Shenker, who have interacted with the system numerous times, and have tended to do so in the light of their scholarship on film and Holocaust testimony.[44]

In projects of this nature, there is a certain resistance to memory's transformation: it is this which troubles us as scholars of Holocaust memory. After all, one of the central appeals of *DiT* is that the post-survivor era, at least for certain museum-goers, will never arrive. We might wonder if we would be better not to make the past present but rather to find ways to live in and with history. This is what Yerushalmi refers to as a kind of *halakhah*[45] of historical thinking, whereby our relationship to the materials of the past is structured by interpretation of text, custom, principles, and precedent—with the life of the community in mind. The opposite of forgetting, he reflects, might not be "remembering" but rather "justice."[46]

The post-survivor era marks a transition as new generations seek to make sense of the Shoah in their lives. Technological solutionism will not undo the problem of human finitude, nor will it take us any closer to understanding what we are meant

---

[43] Dan Leopard and Noah Shenker, with image-texts by Otto M., "Pinchas-DiT: Simulation and the Imagined Future of Holocaust Survivor Memory," in: *Lessons and Legacies XV: The Holocaust; Global Perspectives, National Narratives, Local Contexts,* eds. Erin McGlothlin and Avinoam Patt (Evanston: Northwestern University Press, 2024), , pp. 256–59.

[44] Leopard, Shenker, and M., 'Pinchas-DiT', p. 259.

[45] *The Oxford Dictionary of the Jewish Religion* defines *halakhah* as:

> The legal, as opposed to the nonlegal or aggadic [...], aspect of Judaism. It is also used to indicate a definitive ruling in any particular area of Jewish law. The Hebrew root of *halakhah* means "to go," and the Bible refers to the fulfillment of the Torah as the way in which the people "are to go" (*Ex*. 18.20).

The processes of interpretation are central to *halakhah.* In fact, interpretation begins "in the Bible itself," as Ezra is described as interpreting Torah (*Ezr.* 7.10). Yerushalmi's plea for a *halakhah* of historical thinking in that sense also describes the norms of dispute over the interpretation of law.

M. Grossman, *Halakhah*, in: A. Berlin, ed., *The Oxford Dictionary of the Jewish Religion* (Oxford: Oxford University Press, 2011).

[46] Yerushalmi, *Zakhor*, p. 117.

to do with the legacies of both Holocaust survivors and victims. Even as our tools for recording, repeating, analyzing, and sharing their experiences become ever more powerful, the problem of how to do justice for the victims of genocide and their families remains. New technologies in Holocaust remembrance must be accompanied with deeper philosophical reflections on the historical, cultural, and psychological dynamics that are at the heart of our relationships with Holocaust memory—rather than substituting for it.