

Yale MACMILLAN CENTER
Genocide Studies Program

Mass Atrocities in the Digital Era Initiative (MADE) Working Paper No. 6

March 2024

White Paper on Archiving Holocaust Digital Memorialization: The Lublin District Camps

Eve M. Zucker, Roberto Ulloa, David J. Simon, Mykola Makhortykh, and Daniel Bultmann

Abstract:

This white paper considers the potential for a digital archive for the memorialization of mass atrocities (abbreviated herein as DAMMA) that would integrate artificial intelligence (AI) both in the archive's creation and its use. Based on the proceedings of a workshop held in July 2023 in Bochum, Germany, at the Center for Advanced Internet Studies (CAIS), the workshop addressed conceptual and practical questions regarding the scope, form, benefits and limitations, usage, development, and ethics of DAMMA. The workshop also considered the parameters of a pilot version of DAMMA, focusing on the memorialization practices regarding Holocaust sites in the Lublin district (e.g., Majdanek or Sobibor).

To cite:

Zucker, E. M., Ulloa, R., Simon, D.J., Makhortykh, M., & Bultmann, B. (2024). *Archiving Holocaust Digital Memorialization: The Lublin District Camps. Digital Archive of Memorialization of Mass Atrocities (DAMMA) Workshop Whitepaper*. Yale University: Genocide Studies Program: Mass Atrocities in the Digital Era Working Paper #6.

Table of Contents

Introduction	1
I. Archives and Memorialization	2
II. The Role of AI in Memorialization and Archiving	4
III. Towards an Atrocity-Sensitive AI in Atrocity Memorialization Archive	5
IV. Offline and Online Memorialization of Lublin District Camps: Introducing Our Case Study	6
V. Evaluation: Opportunities and Risks	8
VI. Ethical and Legal Challenges	11
VII. Lessons Learned and Future Considerations	13
References	15

Introduction

The digital turn brought significant changes and challenges to memorializing historical and present mass atrocities. The variety of digital media made memorialization practices more accessible for affected individuals and communities and, potentially, more visible to the broader public. However, the digital turn also amplified concerns regarding privacy, misuse of sensitive material and data, and possible distortions of memorialization practices as part of online disinformation and propaganda campaigns. The new opportunities and concerns are relevant to how digital forms of memorialization shape how societies and communities remember or forget mass atrocities and inform the stories of who we are and how we exist in the world on individual and collective levels.

Understanding the implications of the digital turn for atrocity memorialization requires having the means to study its articulation and praxis in a systematic manner. Such an effort includes collecting and preserving memorialization-related data via a publicly accessible digital archive. For this aim, we decided to investigate the possibilities and risks of archiving digital forms of atrocity memorialization. In 2021, we hosted an international workshop that collected input from specialists in archiving the Web, including digital content dealing with mass atrocities. The main findings of the workshop that addressed the parameters, opportunities, and challenges of designing such an archive are documented in a whitepaper (see Bultmann et al., 2022). The workshop identified the next step to be the development of a pilot archive containing information about the memorialization of a specific instance of mass atrocity, including both online sources (e.g., memorial websites) and user engagement with them (e.g., comments on websites or social media) as well as attempts to deny the past (e.g., via distorted facts and false claims spread online).

As we prepared to design the pilot archive project, the medium of digital memorialization started to shift profoundly. In late 2022, text-generative forms of artificial intelligence (AI) became available to a broad online audience with the rise of chatGPT, followed by the growing integration of AI-powered chatbots and web search engines (e.g., Bing Copilot and Google's Bard). Since then, generative AI's role in shaping how individuals and societies perceive their past, present, and future has accelerated.¹ We recognized that the impact of generative AI on many dimensions of atrocity memorialization (Makhortykh et al., 2023a), including but not limited to ethical, legal, social, psychological, and political aspects, would directly influence the design and implementation of the digital archive project. While there is still no comprehensive overview of the use of generative AI in the context of memorialization, the growing volume of evidence of its uses (but also abuses) for representing the past² prompted the necessity to understand better how AI works or might work in the context of mass atrocity memorialization.

With these ideas in mind, the DAMMA (Digital Archive for Mass Atrocity Memorialization) group convened with a group of academics, heritage practitioners, archivists, and AI experts for four days in July of 2023 in Bochum (Germany) at the Center for the Advancement of Internet Studies (CAIS), our host and sponsor for the workshop.³ The purpose was to explore the feasibility of a pilot archive to prepare for a larger project for

¹ And continues to advance; for some studies discussing its implication for interaction with the past, see Kansteiner, 2022; Makhortykh et al., 2023b; Walden et al., 2023; Zucker et al., 2023.

² For examples, see, Makhortykh et al., 2023b.

³ Participants included: Daniel Bultmann, Daniel Gomes (Arquivo), Christian Groh (Municipal Archives of Mannheim), Marc Green (SimSpace), Łukasz Kukawski (Museum and Memorial in Sobibór), Mykola

developing an archive of online memorialization of mass atrocities (Bultmann et al., 2022) and conceptualize the implications of the growing rise of AI for such an archive.

We selected the Lublin District camps as a case study for such an endeavor, following a brainstorming session with representatives from the related memorial centers in 2022. The Lublin District camps were a network of concentration and extermination camps in Nazi-occupied Poland, including Sobibor, Majdanek, and Belzec, which played a primary role in *Aktion Reinhard*, the Nazi Germany plan to exterminate the entire Jewish population of Poland in 1942-1943. *Aktion Reinhard* resulted in the murder of “at least 1.8 million Jews in the extermination camps Bełżec, Sobibór, and Treblinka in 1942/43 – with less than 150 survivors” (Lehnstaedt, 2021). The complexity of the history of the Lublin camps, which are often targeted by Holocaust deniers (see, for instance, Hobbs, 2015) and instrumentalized by political actors (e.g., the Russian government; Sawkins, 2020), makes this a challenging and meaningful case that can provide a valuable precedent for future projects. Section III explores the merits of the specific case study in more detail.

A key question we addressed at the 2023 workshop in Bochum was whether an online archive of online memorialization related to the case study of our choice is feasible. We also sought to understand whether such an archive would help researchers and the public understand how individuals and societies contend with the difficult past and how the rise of digital technologies, including different forms of AI, changes these processes. In addition to these conceptual questions, we wanted to collect feedback on the practical ways of mapping the universe of memorialization and what type of interface and tools should be employed to implement a robust and sustainable resource. Finally, we were interested in how we could integrate AI-made content into our archive and what impact AI has as a source of information about atrocities and as a tool for online memorialization initiatives.

This whitepaper summarizes valuable insights from the workshop contributing to the ongoing discussion about the relationship between archives, memorials, and AI in the context of mass atrocities. It is structured according to the core themes discussed during the workshop. The first section examines the role of archives in memorialization and considers whether an archive can be a memorial in itself. The next section examines how AI is influencing memorialization and archive practice. We propose to develop a model of an “atrocity-sensitive” AI-driven archive in the ensuing section, the third. The fourth section turns to the subject of our pilot project and the workshop’s focus: archiving online memorialization practices regarding Holocaust sites in the Lublin district and the importance of atrocity-sensitive AI in this context. The fifth and sixth sections provide an overview of the risks and opportunities of the digital archive of atrocity memorialization identified during the workshop and then a discussion of related ethical and legal challenges. In a concluding section, we summarize the lessons learned and future work.

I. Archives and Memorialization

The starting point of the discussion at the workshop was to define the terms “archive” and “atrocity memorialization,” at least in an operational sense. Both concepts are complex constructs that can take on

Makhortykh, Łukasz Mroziak (The State Museum at Majdanek), Sara Day Thomson (University of Edinburgh), David Simon, Roberto Ulloa, Eve Zucker.

different meanings in different contexts, as evidenced by the extensive scholarship dealing with each subject ⁴. Following the definition that we adopted in our earlier whitepaper (Bultmann et al., 2022), based on Featherstone (2006, p. 591), we specify the archive as a “place for the storage of documents and records.” As defined by Moncur and Kirk (2014, p. 956), memorialization is “the act of marking a physical or conceptual space for posterity in remembrance of a person or an event.” Thus, memorialization involves a diverse set of practices that, in the case of atrocities, usually focus on preserving the memory of the victims and the experiences they endured.

Workshop participants debated what practices count as memorialization and whether an archive could serve as a memorial⁵. Is creating an archive an act of memorialization, or can it be used to study the processes of memorializing past atrocities without being treated *as* a form of memorialization itself? This question is pertinent to our project because it shapes the purpose, scope, and impact of preserving digital content concerning mass atrocities. It touches upon the balance between objective criteria of documenting the past and the subjective nature of memorialization, exploring whether archives can serve as spaces for critical examination and learning about how societies remember traumatic historical events. Within this context, the structure of archives (which includes sensitive decisions such as document labeling) and their accessibility (both in terms of what information is accessible and how it can be accessed) have direct implications for memorialization.

The boundary between memorialization and archiving is difficult to draw. Archives themselves preserve both fact-based evidence and subjective recollections about atrocity events. Many different actors, such as journalists, researchers, survivors (and their descendants), and educators, may use archives for purposes viewed as memorialization. Some examples of such uses involve seeking information about relatives and their fate, learning about a particular atrocity-related event, and gathering material for a book. The Arolsen Archive, formerly the International Tracing Service, is particularly illustrative in this context.⁶ Initially, the Arolsen Archive served as a collection point for documents to locate missing persons and determine the plight of individuals. Over time, it metamorphosed into a public resource for those seeking information for restitution or learning about their relatives. Later, recognizing that in some cases, the last traces of someone who perished or disappeared during the Holocaust were its records, the Arolsen Archive added “memorial” to its identity. This evolution demonstrates that the line between the memorial and an archive is shifting and can depend on how the archive is used at a specific point in time.

The case of Arolsen demonstrates how the preservation of information or relics of the past at risk of disappearance becomes part of the process of historical narrative-making by shaping what is collectively considered important or not important to preserve. Consequently, archiving is hard to separate from the act of memorialization because archives can influence memorialization-related processes and choices. One could argue that preservation always involves selection, which co-constitutes past narratives. That way, institutional logic (for instance, rules about what to preserve and how), individual choices of people working as archivists implementing and hence interpreting these rules, and the artifacts themselves, as well as the digital or non-digital platforms these artifacts enter, create sedimentations of human as well as non-human memorial narrative practice (Latour, 2005).

⁴ For some examples regarding the notion of an archive and its different operationalizations in the digital realm, see Moss (2016). For the examples in relation to atrocity memorialization, see Ketelaar (2008) and Recuber (2012).

⁵ See Haskins, 2007; Arvanitis, 2019; Bultmann et al., 2022 for prior discussions along these lines.

⁶ A former archivist from the Arolsen Archives participated in the July 2023 workshop that serves as the genesis of this whitepaper.

This argument that an archive is a memorial, however, was challenged more than a few times in the course of the workshop, with participants noting additional differences between the processes of archiving and memorialization, such as the question of intention (remembering vs. preserving), degree of information selectivity (selective representation of memorialization vs. comprehensive coverage of archiving), and target populations (general public vs. academic community). While the overlap between the two concepts remained debated, the participants agreed that memorialization is one of the goals of archiving. In cases of mass atrocities, this goal is more prominent.

Another similarity noted during the workshop was that both processes of archiving and memorialization are profoundly affected by the advancement of digital platforms and AI. The digital turn transforms how archival materials are produced and stored today and how they are employed in memorialization practices. Specifically, technological tools can be utilized at multiple stages of building the archive, from mining the documents on the Web to labeling and structuring them to facilitating user access to the archived content. The debate had implications for how we should assemble our project, and the points raised will be considered in the following sections, particularly when we evaluate our project and directly list the ethical challenges.

II. The Role of AI in Memorialization and Archiving

AI is broadly defined as the ability of human-made artifacts to engage in intellectual behavior (Nilsson, 1998). In this whitepaper, we refer to AI as an umbrella term denoting a set of autonomous tools that can provide analysis, selection, categorization, and interpretation purposes⁷. Frequently, AI is employed by humans to identify relevant content from a potentially vast universe of data through a human-machine interface and to explain choices made in this context. More recently, AI can generate content based on a set of user directives, which might often be vague (e.g., “tell me about the Holocaust” or “draw a depiction of the Holocaust”), allowing for a wide range of possible responses. The selection process within the AI algorithm is poorly understood and essential for the specific narratives that are disseminated.

In the workshop, we focused specifically on how AI can serve atrocity memorialization archive purposes. Throughout the discussion, several key functionalities of AI in this context emerged. The first of these functions is related to **mapping and navigating** collections of atrocity-related archives. Such archives can include diverse materials about institutions and activities related to public commemoration as well as private mementos and recollections. These materials may feature historical evidence (e.g., orders to perpetrators or testimonies of eyewitnesses) and derivative memory-related content (e.g., educational materials produced by memory institutions or scholarly analyses of specific atrocity episodes).

Locating materials for inclusion in the archival collection or for fulfilling the user information needs can be facilitated by AI. AI can accelerate information retrieval in a customizable fashion (e.g., by personalizing information outputs). It can also better cope with information retrieval of formats other than text, such as visual and audio. Considering the AI’s potential to learn from previous user interactions, it can overcome some limitations of the traditional archives by identifying new relationships between the stored material and user interests. Finally, AI may also be able to police its own use – or, more to the point, misuse – by identifying cases

⁷ At the same time, it is important to note that human input still plays an essential role in the functionality of many AI systems, including the ones used in the context of archiving. Humans define the outer parameters of the universe of possible data by providing a directive (such as an inquiry or task) that the AI endeavors to fulfill, and by giving feedback regarding the utility of the output.

where users prompt the AI to generate distortion or denial-related content by prohibiting access to such information or accompanying it with fact-checked evidence.

The latter aspect of the navigational functionality relates to another function of AI in archiving atrocity memorialization: **analytical functionality**. AI permits identifying connections between archived content, which can help recognize emergent forms of atrocity or engagements with its memory. Such recognition can reshape autobiographical memory and various forms of collective memory (e.g., communicative, social, or prosthetic memories). Through initial guidance and continuous feedback, AI can recognize instances of memorialization and then analyze patterns and trends to help users understand the scale and interconnectivity of memory practices.

The third functional form of AI that emerged from the workshop discussion was **interactional functionality**. AI enables new possibilities for interfacing with the past, not just through query-based inputs, which result in a selection of system outputs, but potentially through interactive forms of archive-based memory. An example of such AI-powered interactions is a chatbot that can reduce the problem of subjective assessment of archival source relevance, a common problem with search engines, for example, when it comes to ranking results related to a search query. A chatbot can alleviate this problem because it (1) generates text based on different sources and (2) allows the user to enter into a conversation to ask further questions.

Nonetheless, chatbots powered by modern forms of generative AI, specifically large language models (LLMs), face many critics. Two of the most problematic (and often mentioned) are (1) they do not attribute the content to a source, and (2) they are susceptible to biases and factual errors, often undetectable by individuals seeking knowledge. Although it can be argued that source attribution is not necessarily the case for memorialization and learning as long as the narrative is correct, the absence of source information limits the transparency of the chatbot performance⁸, crucial in closely related contexts (academia, journalism, judicial proceedings). The presence of biases and plainly incorrect information is extremely concerning and must be addressed. In the next section, we present the ideas regarding possible mitigations of these risks in the form of an atrocity-sensitive AI.

III. Towards an Atrocity-Sensitive AI in the Atrocity Memorialization Archive

Many workshop participants acknowledged the multi-faceted potential of AI in the context of atrocity memorialization archiving. However, they also raised several questions regarding what forms of AI most suit this purpose. Particular attention was devoted to LLMs, the technology behind conversational agents such as ChatGPT or Bing Co-Pilot. These forms of AI feed from data, which can originate from online memorialization resources such as Wikipedia and institutional websites, to generate content about the past, ranging from essays about the Holocaust to conversation-style responses about historical facts to imagined diaries of survivors. By making such content and sharing it with the users in a conversational format, generative forms of AI can potentially transform mass atrocity memorialization and what the public learns from it.⁹ To do it constructively, however, AI must recognize historical facts and ethical complexities involved in remembering mass atrocities;

⁸ Matching the generated text to the sources can also be seen as a follow-up step after generating a narrative based on the documents, similar to how Microsoft integrated a chatbot to their Bing search.

⁹ Users also make memorial content by using AI, including LLMs and multimodal models that support image, video, and voice, as an assistant. Users may passively read what AI tells them when they also steer conversations thereby shaping the stories told. Mediums like Google Search labs, on the other hand, are designed less to be a chat agent and instead provide a snapshot summary, making it less of a space that encourages human agency.

otherwise, it may create confusion and doubt or can be misused to distort the past and manipulate the present. Other risks of applying AI to atrocity memorialization include the possibility of it showing systematic under- or over-representations of different aspects of the past, hallucinating non-existing events, and undermining the trust of users due to the lack of transparency in its functionality and sources (e.g., Makhortykh et al. 2023a; 2023b).

To mitigate these risks, platform designers must develop mechanisms for countering possible AI malfunction in the context of archiving atrocity memorialization. Similar to the other sectors affected by the deployment of (generative) AI, such mechanisms within atrocity memorialization could include creating regulations to protect the rights of different stakeholders, ensuring accountability of companies developing AI models and platforms utilizing them, and nurturing AI literacies among the AI users to help them critically evaluate AI performance and minimize the risks associated with its misuse. However, we argue that these mechanisms must be supplemented by a basic awareness of the ethical and moral aspects of atrocity memorialization from the side of AI (which we discuss at greater length below).

“Atrocity-sensitive AI” is thus a mechanism for atrocity-related archiving and memorialization that harnesses AI benefits while mitigating its threats. The concept draws inspiration from the scholarship on value-sensitive design (e.g., Friedman et al., 2002; Umbrello, 2019; Van Der Hoven & Manders-Huits, 2020). The design approach, which accounts for human value throughout the system design process, has been applied to systematically address the risks posed by non-generative forms of AI powering, for example, recommender systems (Chen et al., 2022) and conversational agents (Wambsganss et al., 2021). Specifically, the value-sensitive design was applied to embed principles such as privacy and diversity into AI system design to minimize the likelihood that these systems would disclose private information or keep individuals in their information bubbles by limiting their access to alternative viewpoints.

Similar to privacy- and diversity-sensitive design principles, AI can be designed with ethical principles related to atrocity memorialization. While there is currently no exhaustive list of such principles, some examples include respect for the victims, beneficence, fairness (Makhortykh, 2023), and historical accuracy and acknowledgment of the complex nature of individual and collective suffering associated with mass atrocities. The value-sensitive design includes a series of recommendations to operationalize the production, from identifying value, technology, or context to mapping benefits and harms onto those Values (Friedman et al., 2002). In the context of atrocity-sensitive AI design, the practical implementation of these recommendations will vary depending on the actual purpose of the AI model or the system powered by it. For instance, non-generative AI systems powering archival search may require a different implementation of an atrocity-sensitive design than a conversational agent powered by an LLM. However, some tools for implementing such design can include, for instance, the careful selection of training data for AI models that strongly focus on memorialization-related principles, the implementation of the safeguards preventing the potential abuse of AI models for undermining the ethical values, and the development of AI logic with the certain memorialization-related aims in mind. The activities and discussions during the workshop covered many of these aspects, but the principles will help us identify gaps in the future and guide our decisions.

IV. Offline and Online Memorialization of the Lublin District Camps: Introducing Our Case Study

We selected the digital memorialization of the Lublin District camps for the pilot study for several reasons. Central among them was the historical complexity of this specific episode of the Holocaust and its contemporary commemoration (e.g., Wilson, 2019; Lehnstaedt, 2021; Tingler, 2023). As previously noted, the Lublin camps played a key role in the course of Aktion Reinhard, which was characterized by the “unparalleled” (Diepenbroek

et al., 2021, p. 2) secrecy from the side of perpetrators both in the course of the atrocity and its aftermath. Toward the end of World War II, the Nazis hid the evidence of the atrocity by burning and burying the bodies of the victims and covering and disassembling the camps. Some areas, such as Sobibor, were “partially leveled, asphalted, or planted with trees after 1945 – thus destroying further evidence” (Lehnstaedt, 2021).

Unlike Auschwitz, there were few witnesses to tell what had happened at the Lublin camps and few visible traces of extermination, such as crematoriums. Most of the Lublin camps’ inmates were immediately gassed or shot on arrival to the camps, leaving fewer than 150 survivors out of the nearly 2 million taken to Sobibor, Belzic, and Treblinka (Lehnstaedt, 2021). The local population in the area certainly knew about the camps but generally did not speak up for many reasons (including a few instances of looting of the mass graves by local residents; Wilson, 2019; Tingler, 2023). The Soviet control over Poland after the end of the war further contributed to the collective silence, with the Soviet authorities tending to subjugate memories about the Holocaust and treating them as part of the overarching narrative of the Soviet civilians’ suffering. Despite the massive number of victims murdered during Aktion Reinhard (Diepenbroek et al., 2021), after the war, “Aktion Reinhardt was out of sight and out of mind” (Lehnstaedt, 2021).

It was not until the 1960s that the first studies on the Lublin area camps appeared, with more interest developing in the 1970s and 1980s (Wilson, 2019). In 1987, the film “Escape from Sobibor” was released based on the novelesque book by Richard Rashke that was published under the same name in 1982 and gathered substantial attention (Wilson, 2019). The film, a British production, aired on CBS and is available today for streaming on platforms like Amazon Prime and Netflix. As one of the Lublin camps, Sobibor, became part of popular culture, it prompted the growing interest in the Holocaust sites in the region, with a particular emphasis on the Sobibor uprising led by Alexander Peckersky. The memory of the uprising remains active in contemporary discourses that various groups and individuals deem analogous today -- such as posters on the platform X who have sought to draw (or debunk) comparisons between the Sobibor uprising and the October 7, 2023 incursion and massacres by Hamas in southern Israel.

Despite the importance of the film and the book about the Sobibor uprising, both were based on relatively few accounts of survivors, with a particular focus on only one Lublin District camp and the uprising there. Little was known about other Lublin camps, such as Majdanek or Treblinka, until a series of archeological excavations in the late 1990s, including the ones in Belzec in 1997-99 (Kola, 2000) and Sobibor in the early 2000s and more extensively in 2011-2016. Particularly, the discovery of the remains of a gas chamber at Sobibor in 2014 brought the site and its memory into the public discourse, with the news about the find being shared globally (Wilson, 2019).

The excavations also prompted the memorialization process, with the Sobibor Museum opening in 2021 and featuring many of the archeological remains unearthed in the investigations. It also increased attention towards other Lublin camps, where the memorials and monuments existed already. For instance, a museum was established in Majdanek as early as November 1944. A memorial for Sobibor, by contrast, was not established until 1965. Moreover, this memorial did not initially indicate that those who perished were largely Jews. In Belzec, the museum was founded in 2004 and is under the auspices of the Majdanek State Museum.

Today, the memory of the Lublin District camps and the events that unfolded there during the Holocaust is still fraught with contestation, particularly as an ideological vehicle in Russia's war against Ukraine and also in the current Middle East conflict, where references to the Lublin camps are used to justify anti-Ukrainian and anti-Israeli violence. However, the contestation of the past is not only harnessed to these recent conflicts but also circulates online in other forms, for example, through Netflix’s “The Devil Next Door” and the 2018 Russian

movie “Sobibor.” Moreover, the relatively recent 2015 revelation of the Niemann photographs, known now as the Sobibor Perpetrator Collection, has stirred legal controversy relating to the role of John Demjanjuk at the camps, whose story “The Devil Next Door” is based on (Katz, 2020). This collection was digitized and made available online. The politicization of memory, including the control over and use of materials from archives, is part of the motivations for this project, and we see the Lublin camps as an ideal subject for the pilot project.

Given this context, archiving online forms of memorialization using atrocity-sensitive AI can help educate present and future generations worldwide about the importance of Lublin District camps while preventing memory about them from being used for manipulation and distortion of the past. To achieve this aim, we argue that it is essential to synthesize existing expertise on the potential of using AI for genocide memorialization in the research community and ensure cooperation with relevant stakeholders, particularly representations of Lublin Region memorial sites.¹⁰

The first step of such cooperation should be to map existing forms of digital memorialization of the Holocaust in the Lublin area and to track its evolution over time. Following the mapping that can combine human expert- and AI-assisted approaches,¹¹ it will be possible to establish the archive of memorialization and distortion practices related to the camps. In the context of such an archive, AI can facilitate user interactions with the archived content (e.g., via a conversational LLM-powered interface) and enhance the analysis of the archived data (e.g., by automatically labeling archived content features).

Furthermore, archived data can be employed for developing an atrocity-sensitive LLM to counter limitations of existing forms of AI regarding the representation of information about Lublin District camps (e.g., hallucinations and the risks of jailbreaking that can circumvent AI’s ethical guardrails for denialism/manipulation) and to improve quality of AI-generated content on the Lublin District camps. The combination of theoretical insights about the transformations of digital memorialization and the publicly available atrocity-sensitive LLM model will expand our understanding of the impact of technology on Holocaust remembrance and create new possibilities for Holocaust education and combatting denial. It will also refine research strategies for studying how the rise of AI influences memory and education practices and finding solutions to prevent these practices from being disrupted by the digital turn.

V. Opportunities and Risks

During the workshop, we discussed the opportunities and risks of developing a pilot archive of online memorialization for the case study outlined above. In this section, we summarize the points that emerged from that discussion; while particularly relevant for the case study, many of these points can likely apply to other contexts and cases.

Apart from giving visibility to the particular case study of the Lublin area camps, a strategic motivation to work on a pilot is to test and refine a model that can later be scaled, either in the form of a larger archive or multiple archives of digital memorialization of different atrocities. Additionally, working on a pilot would give us the flexibility to adapt and integrate new technological solutions that continue to emerge with the integration of AI

¹⁰ “Lublin District” was the Nazi administration term during the occupation period. So the camps were in the Lublin District, but the memorials (as present entities) are considered in the “Lublin area”.

¹¹ One example of such a combination can be the identification of the key memorialization projects based on interviews with human experts together with the detection of projects treated as more relevant for the topic by AI systems (e.g., search engines).

in memorialization practices. At the same time, the pilot can increase visibility and facilitate the establishment of an international network focused on exploring challenges and opportunities of preserving digital atrocity memorialization practices, e.g., by connecting to existing researchers and practitioner groups and divulging information about similar projects and initiatives.

Additionally, the pilot will enable the long-term storage of non-institutional forms of atrocity memorialization. This resource can contribute to a more comprehensive understanding of how mass atrocities are remembered in the current digitally mediated age. Capturing the diversity of digital memorialization practices is integral for understanding the opportunities and challenges of technology in this context. For example, there could be extended possibilities for comparative analysis of atrocity memorialization, executed by archiving online memorialization projects developed and maintained by different national and cultural communities. Similarly, archiving can be used to create snapshots of memorialization-related outputs of AI-driven systems. Such snapshots are essential to contrast AI-identified priorities regarding interpretations and sources related to specific instances of the atrocity (e.g., top Google search results) with human expert evaluations. In such a way, the pilot can provide unique possibilities for tracing similarities and differences between how humans and AI prioritize information sources and make interpretations of the past.

Another advantage of the long-term storage of diverse forms of memorialization relates to its potential as a training and evaluation corpus for AI development and research in retrieving and generating information about mass atrocities, particularly considering the possibility of making such corpus multi-lingual. In the context of our case study, such a corpus may include content in Hebrew, Polish, German, Ukrainian, Czech, Yiddish, Russian, French, Dutch, and other languages. The multi-lingual nature of the corpus is particularly important considering much evidence of unequal performance of AI in high- and low-resource languages, which can be particularly detrimental to atrocity memorialization. The exact uses of the corpus can vary from its employment in developing new atrocity-sensitive LLM models to evaluating the performance of the existing models.

The possibility of using the pilot to track evolving mass atrocity narratives was also discussed during the workshop. Such tracking involves archiving references to memorial practices in the context of current events. The purpose here is to map and analyze changes in societal narratives about these practices and their changing roles. Tracking the evolution of atrocity narratives and their impacts is important to ensure the preservation of historical evidence and identification of possible instances of distortion. Part of it is the use of the pilot for understanding changes in atrocity discourses and their relationship with individual and collective trauma and the silencing of certain aspects of the past. Lastly, the pilot might trigger innovations within the realm of memorialization and archiving of mass atrocities. Several new opportunities are specifically related to integrating an AI component (e.g., in the form of LLMs helping users to navigate archival collections or analyze them) in the pilot design. Such integration can partially address the challenge of making subjective choices regarding operationalizing the relevance of archived content. Instead of relying on ranking search results, which always imply the existence of a certain hierarchy of retrieved data, AI can allow dialogic exchanges between the archival interface and the user, which can provide an additional benefit of embedding additional contextual information within answers. This approach would ensure a more nuanced understanding of information regarding archived materials, which is of paramount importance in the context of highly sensitive topics dealing with mass atrocities. Another important opportunity focused on applying AI for data analysis, grouping, and pattern recognition within research processes. Addressing the problem of misinformation and denial stood out as a key end objective, with several proposals from participants stressing the promise of AI in general and its integration in the archive, in particular in this context.

During the workshop, participants also highlighted several conditions for ensuring successful AI integration with the pilot archive of memorialization, summarized in Table 1.

Table 1. Conditions to integrate AI into the pilot.

Condition	Purpose
Controlling the quality of data used for training the atrocity-sensitive AI.	Ensure accurate and ethical treatment of the contested aspects regarding the history of the Lublin camps.
Remaining aware of new advancements in the rapidly developing field of AI (including archive-related tools).	Potentially integrate AI advancements into the pilot to streamline its implementation and enhance the archive's efficiency and potential impact. It is important to critically question the necessity of integrating such advancements - to avoid doing it just for novelty - and their potential impact on the different stakeholder groups.
Developing a user-friendly interface.	Ensure accessibility of the knowledge hosted within the archive.
Ensuring the preservation of diverse content.	Enable visibility of diverse perspectives on the past without promoting distortion and denialism and account for a broad range of views on how the atrocity can be memorialized.
Engaging different groups of stakeholders when developing and maintaining AI solutions integrated with the archive.	Identify how these solutions can accommodate the specific interests and needs of the stakeholder groups.

Participants also identified a series of risks and challenges that need to be carefully considered and mitigated. Under **sustainability risks**, the uncertainties surrounding the need to secure financial resources for the development and maintenance of the archive were often noted. The limited availability of funds may curtail the project scope, hinder planned activities, and disrupt the long-term preservation of archived data. Additionally, managing and securely storing sensitive data presents multiple ethical risks (e.g., compromising individual privacy; see the next section for more information). The ethical risks can be amplified by insufficient human resources and expertise, emphasizing the importance of adequate staffing and expertise allocation.

Management risks constitute another set of challenges, including a need for clarity regarding project objectives and the potential for its scope to cause confusion and misalign project outcomes. Uncertainties regarding what content exactly to collect and preserve might result in delays and miscommunications. Ineffective project leadership could lead to directionlessness, coordination issues, and decision-making delays, hindering progress. Moreover, the risk of the limited use of the resulting archive raises questions about its utility, making it crucial to assess and maximize the pilot's impact.

Selection bias challenges present risks arising from biases introduced during content collection and archiving, potentially leading to the misrepresentation of relevant information regarding the memorialization of Lublin area camps. The input from human experts may also introduce additional biases, distorting project insights or conclusions. The limited scope of the archive, for instance, due to the incomplete list of projects that are to be archived or the shortage of resources for long-term data preservation, has the potential to skew the composition of the archive. Additionally, reliance on AI-driven systems (e.g., search engines) to identify projects to be archived might perpetuate existing biases within those search engines.

While there are multiple risks associated with implementing the pilot, it is worth noting that *not* implementing it *also* bears risks. The lack of consistent and sustainable archiving of online memorialization practices increases the likelihood of their disappearance or distortion. While alternatives such as the Internet Wayback Machine, Google indexing, or arquivo.pt, including commercial web data repositories exist, dedicated to the systematic and atrocity-sensitive preservation of online memorialization practices. Consequently, the scope and accuracy of preserved information about memorialization are limited. Moreover, in an age of hyperspeed development and adoption of AI systems, the absence of a dedicated archive of digital memorialization that can preserve and protect sensitive historical and memorial material on the web risks the loss or distortion of Holocaust memory. By collecting snapshots of memorial activity and resources on the web, ideally before AI enhancement or corruption, we can ensure these resources are available for future research, education, and memorial purposes.

In conclusion, focusing our efforts on the Lublin camps offers distinctive advantages, utilizing web-based archived resources and tapping into extensive network sources. The envisaged concrete output of a dedicated web archive tailored for museum use at the Lublin area camps promises to be a significant achievement. Moreover, the project's outcomes are poised to inspire diverse use cases, demonstrating scalability for broader applications beyond its immediate scope.

VI. Ethical and Legal Considerations

In line with the value-design methodology proposed for the pilot, the workshop included a dedicated discussion addressing ethical, political, and legal challenges that might arise when implementing and maintaining a digital archive, such as the one proposed here. Many issues surfaced, ranging from ethical dilemmas of selecting content for archiving to legal intricacies of accounting for the privacy norms and regulations and practical hurdles, for instance, regarding data management. Table 2 groups these issues into several broad categories to facilitate the examination of various challenges. Each of these categories of challenges will need to be assessed and addressed during the implementation of the pilot outlined in Section “Towards an atrocity-sensitive AI.”

Table 2. Challenges of implementing a digital archive. Each set of challenges (second column) is associated with a specific aspect of the archive (first column).

Topic	Challenges
Case selection and preservation	<ul style="list-style-type: none"> - Navigating political and ethical considerations that can influence case selection and unequal visibility of specific instances or aspects of the atrocity that is memorialized - Addressing concerns about the potential ephemerality of online memorialization practices due to platform data retention or discontinuation of specific platforms or memorial projects
Resource allocation	<ul style="list-style-type: none"> - Determining how much content and of what types can be archived; managing the gradual expansion of the archive and its long-term maintenance - Addressing questions regarding the potential influence of funders on what can and cannot be done (e.g., in terms of acceptable uses of an archive or the ways information about it is dissemination and to whom access if granted)
Database and international interest	<ul style="list-style-type: none"> - Leveraging the dataset to generate international interest in memorialization efforts - Managing the adverse impact of national and cross-national policies and practices on interpretation of archived materials or access to collected data
Mapping and labeling	<ul style="list-style-type: none"> - Developing criteria to decide what can and what needs to be collected and archived - Deciding on metadata schemata and related constructs (e.g. how to categorize individuals mentioned in the archives - e.g. as victims, perpetrators, or rescuers - or how to organize information about specific aspects of the atrocity) - Considering how justified is the preservation of the denialist and distortive materials and designing strategies for managing access to such materials to prevent their undesired uses
Data management	<ul style="list-style-type: none"> - Deciding to which instances of online memorialization the existing legal norms (e.g., the the right to be forgotten or the right to erasure) are applicable -Dealing with the potentially incomplete data due to limitations of the scraping (e.g., caused by firewalls, paywalls, or dynamically loading content) - Developing usability and accessibility strategy to avoid creating just another huge data collection without it making an impact and being used - Operationalizing constructs required for automated annotations, e.g., for detecting misinformation, content classification, and filtering of sensitive content (see Mapping and labeling) -Mitigating risks of collecting and making accessible sensitive content, such as visual evidence of suicides, torture, murder, or child abuse
Community Involvement	<ul style="list-style-type: none"> - Effectively involving relevant groups of stakeholders (e.g., heritage institutions and individual practitioners, survivors and descendants, researchers) in archiving efforts - Ensuring the archive is representative of diverse spectrum of groups involved in atrocity memorialization and their experiences to ensure accuracy of representation - Expanding mapping and archiving efforts beyond mainstream platforms and projects - Ensuring that archiving efforts support and acknowledge the work of stakeholders (e.g., archivists) and not undermine it.
Regulations	<ul style="list-style-type: none"> - Addressing legal considerations regarding including and making accessible embargoed or confidential material in archives - Addressing privacy concerns related to capturing private experiences regarding atrocity memorialization from different range of platforms -Contending with the legal exceptions and commitments associated with the archive (e.g., the accreditation process for an archive in the UK that provides GDPR exemptions and a commitment to custodianship) - Handling social media and other platforms' content that may break the law (e.g., showing child abuse or propagating hate speech) and dealing with the presence of personal information - Addressing other concerns, for instance, opt-in and opt-out options and copyright

Despite these challenges, participants recognized the need to continue this initiative, particularly in light of the rapid transformations occurring in both digital and political landscapes. Considering all these points, the consensus was that a focused effort would represent a significant advancement compared to the currently available generic Web archives. Moreover, the pilot should emphasize the need to develop an ethical code that addresses privacy concerns, navigates the complexities of labeling, and that human-centered values are kept at the forefront of an archiving process (Zucker et al., 2023).

VII. Lessons Learned and Future Considerations

The CAIS workshop in Bochum brought together an exceptional group of researchers and practitioners in archive studies, social sciences, and computer science to discuss how we can preserve and study atrocity-related digital memorialization practices. The preservation of these practices in our digital and increasingly AI-shaped age is integral for maintaining historical truth regarding mass suffering and ensuring the nuanced understanding of the past in the social, political, and technological tsunami that is our present. Such understanding is particularly important at a time when memories of atrocities are instrumentalized to justify new acts of mass violence and promote hate toward vulnerable groups (for some examples in the context of the Russian aggression against Ukraine, see Gaufman, 2023; Makhortykh & Aguilar, 2023).

The insights achieved in the course of the workshop can be applied to the diverse cases of online atrocity memorialization, even if, for now, we focus on a single case, that is, the Lublin area camps and atrocities committed there during the Holocaust. Our discussions regarding the perspectives of archiving online memorialization practices related to this iconic instance of atrocity highlighted the need for careful consideration of how to handle archived materials. This task is particularly challenging given that the documents, photographs, and films that comprise the memorial material are intimately connected to the memory of victims of genocide. Moreover, we need to ensure that our proposed project is carried out with sufficient forethought to prevent its misuse in the future. By focusing on the Lublin area camps and soliciting the experiences and ideas of experts (including representatives of the Majdanek and Sobibor memorial sites), we aim to contour ways to bring together online archives and AI to account for the present and future needs of practitioners and the broader public in a manner that is informed, effective, ethical, and sensitive to the memory of the victims. Through this process, we hope to learn how the memories of the horrific events that took place in the Lublin area under the Nazis are affected by the digital turn and for what purposes these memories are engaged (or even appropriated) in online environments. Unlike the long-term preservation of historical evidence of the Holocaust that existing heritage institutions provide, online memorialization practices that engage with this evidence are at higher risk of disappearance due to their often individualized and not institutionalized nature and ephemerality of online platforms and their data. We argue that enabling possibilities for archiving and studying these memorialization practices is essential to understanding how atrocities are remembered and to identify solutions for avoiding the repetition of such atrocities.

The present whitepaper summarizes the key themes discussed during the workshop, encompassing debates on whether an archive is a form of memorialization and if such a distinction is necessary, what constitutes memorialization, the role of AI in the memorial and archival contexts, how to navigate the ethical issues inherent to sensitive materials, and how AI might handle these materials (e.g., filtering for and then containing particularly sensitive items in a separate space), the challenges and benefits of building such an AI-archive, technology questions, and the involvement of a dedicated network of practitioners and related communities. Following the workshop, the DAMMA core members reviewed the participants' input (drawn from several activities) and our own notes to assemble a concept for a pilot project described in this paper.

At the end of the Bochum workshop, the overall consensus among the participants was that an AI-enhanced archive of online activities related to the Holocaust in the camps of the Lublin area could effectively capture and preserve digitized memorial artifacts and practices involving them. The benefits beyond preservation would include opportunities for research and analysis, as well as for facilitating personal and collective memorial projects and endeavors online and offline. As such, the workshop was concluded with the intention to write the present whitepaper summarizing the workshop proceedings, develop and submit a proposal for the pilot project, and publish a research paper on the topics raised. We will continue to adapt the project to keep pace with rapidly developing technologies and their potential to impact the project. As delineated in our risk section, the project's sustainability related to securing financial resources is fundamental to our ability to continue our effort to implement the pilot project. With it, we hope to have developed the expertise and the tools to expand beyond the camps of the Lublin area and the Holocaust to other mass atrocities worldwide.

References

- Arvanitis, K. (2019). The 'Manchester Together Archive': researching and developing a museum practice of spontaneous memorials. *Museum and Society*, 17(3), 510-532.
- Bultmann, D., Makhortykh, M., Simon, D., Ulloa, R., & Zucker, E. M. (2022). *Digital Archive of Memorialization of Mass Atrocities (DAMMA) Workshop Whitepaper*. Yale University Genocide Studies Program. <https://gsp.yale.edu/digital-archive-memorialization-mass-atrocities-damma-workshop-whitepaper>.
- Chen, Z., Piao, J., Lan, X., Cao, H., Gao, C., Lu, Z., & Li, Y. (2022). Practitioners Versus Users: A Value-Sensitive Evaluation of Current Industrial Recommender System Design. *Proceedings of the ACM on Human-Computer Interaction*, 6 (CSCW2), 1-32.
- Diepenbroek, M., Amory, C., Niederstätter, H., Zimmermann, B., Szargut, M., Zielińska, G., Dür, A., Teul, I., Mazurek, W., Persak, K., Ossowski, A., & Parson, W. (2021). Genetic and phylogeographic evidence for Jewish Holocaust victims at the Sobibór death camp. *Genome Biology*, 22(1), 1-16. <https://doi.org/10.1186/s13059-021-02420-0>
- Featherstone, M. (2006). Archive. *Theory, Culture & Society*, 23(2-3), 591-596.
- Friedman, B., Kahn, P., & Borning, A. (2002). Value sensitive design: Theory and methods. *University of Washington Technical Report*, 2(8), 1-8.
- Gaufman, E. (2023). Damsels in distress: Fragile masculinity in digital war. *Media, War & Conflict*, 16(4), 516-533.
- Haskins, E. (2007). Between archive and participation: Public memory in a digital age. *Rhetoric Society Quarterly*, 37(4), 401-422.
- Hobbs, M. (2015). Writing 'history' for Hitler: Holocaust denial since 1945. In C. Carmichael & R. C. Maguire (Eds.), *The Routledge History of Genocide* (pp. 196-206). Routledge.
- Kansteiner, W. (2022). Digital doping for historians: can history, memory, and historical theory be rendered artificially intelligent?. *History and Theory*, 61(4), 119-133. <https://doi.org/10.1111/hith.12282>
- Katz, B. (2020, January 29). Newly Released Photos May Place the 'Devil Next Door' at Sobibor Death Camp. *Smithsonian*. <https://www.smithsonianmag.com/smart-news/newly-released-photos-may-place-devil-next-door-sobibor-death-camp-180974082/>
- Ketelaar, E. (2008). Archives as spaces of memory. *Journal of the Society of Archivists*, 29(1), 9-27.
- Kola, A. (2000). *Bełżec: the Nazi camp for Jews in the light of Archaeological Sources: Excavations 1997-1999*. Council for the Protection of Memory of Combat and Martyrdom.
- Latour, B. (2005): *Reassembling the Social. An Introduction to Actor-Network-Theory*. Oxford University Press.

- Lehnstaedt, S. (2021). Aktion Reinhardt—Sources, Research and Commemoration in the last 30 years. *Témoigner. Entre histoire et mémoire. Revue pluridisciplinaire de la Fondation Auschwitz*, 132, 62-70.
- Makhortykh, M. (2023). No AI After Auschwitz? Bridging AI and Memory Ethics in the Context of Information Retrieval of Genocide-Related Information. In *Ethics in Artificial Intelligence: Bias, Fairness and Beyond* (pp. 71-83). Springer Nature Singapore.
- Makhortykh, M. & González-Aguilar, J. (2023). Is it fine? Internet memes and hate speech on Telegram in relation to Russia's war in Ukraine. In B. Di Fátima (Ed.), *Hate Speech in Social Media: A Global Approach* (pp. 75-95). LabCom.
- Makhortykh, M., Zucker, E. M., Simon, D. J., Bultmann, D., & Ulloa, R. (2023a). Shall androids dream of genocides? How generative AI can change the future of memorialization of mass atrocities. *Discover Artificial Intelligence*, 3(1), 28. <https://doi.org/10.1007/s44163-023-00072-6>.
- Makhortykh, M., Vziatysheva, V., & Sydorova, M. (2023b). Generative AI and Contestation and Instrumentalization of Memory about the Holocaust in Ukraine. *Eastern European Holocaust Studies*, 1(2), 349-355.
- Moncur, W., & Kirk, D. (2014). An emergent framework for digital memorials. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (pp. 965-974). ACM Press.
- Moss, M. (2016). Opening Pandora's Box: What is an archive in the digital environment?. In L. Craven (Ed.), *What are Archives?* (pp. 71-87). Routledge.
- Nilsson, N. J. (1998). *Artificial intelligence: A new synthesis*. Morgan Kaufmann.
- Recuber, T. (2012). The presumption of commemoration: Disasters, digital memory banks, and online collective memory. *American Behavioral Scientist*, 56(4), 531-549.
- Sawkins, I. (2020). Russia's State Mobilization of the Holocaust Onscreen—Konstantin Khabensky's Film Sobibor (2018). *Modern Languages Open*, 1(23), 1-8. <https://doi.org/10.3828/mlo.v0i0.314>
- Tingler, J. (2023). A local history of the Sobibór death camp and Nazi occupation. In C. Gerlach (Ed.), *On the Social History of Persecution* (pp. 223-242). De Gruyter Oldenbourg.
- Umbrello, S. (2019). Beneficial artificial intelligence coordination by means of a value sensitive design approach. *Big Data and Cognitive Computing*, 3(1), 1-13. <https://doi.org/10.3390/bdcc3010005>
- Van Der Hoven, J., & Manders-Huits, N. (2020). Value-sensitive design. In K. W. Miller & M. Taddeo (Eds.), *The Ethics of Information Technologies* (pp. 329-332). Routledge.
- Wambsganss, T., Höch, A., Zierau, N., Söllner, M. (2021). Ethical Design of Conversational Agents: Towards principles for a value-sensitive design. In: Ahlemann, F., Schütte, R., Stieglitz, S. (Eds.) *Innovation Through Information Systems. Volume I: A Collection of Latest Research on Domain Issues* (pp. 539–557). Springer.

Walden, V. G., Makhortykh, M., Marrison, K., Arnold-de Simine, S., Balis, A., Clavert, F., ... & Wierenga, S. (2023). Recommendations for using Artificial Intelligence and Machine Learning for Holocaust Memory and Education. REFRAME.

Wilson, H. (2019). Sobibór death camp: awareness, memorialisation and re-conceptualization. *Holocaust Studies*, 25(3), 400-421. <https://doi.org/10.1080/17504902.2019.1567675>

Zucker, E., Makhortykh, M., Ulloa, R., Bultmann, D. & Simon, D. (2023). AI and Archives: How can Technology Help Preserve Holocaust Heritage Under the Risk of Disappearance? *Eastern European Holocaust Studies*, 2(1), 357-363. <https://doi.org/10.1515/eehs-2023-0052>