Is Artificial Intelligence the Future of Collective Memory?

Sarah Gensburger, Frédéric Clavert

BRILL
NIJHOFF

MEMORY STUDIES REVIEW 1 (2024) 195–208

memory studies review
brill.com/mesr

Introduction

Is Artificial Intelligence the Future of Collective Memory?

Sarah Gensburger | ORCID: 0000-0001-5329-9542
Sociology, CNRS / Sciences Po, Paris, France
sarah.gensburger@sciencespo.fr

Frédéric Clavert | ORCID: 0000-0002-0237-2532
History, C2DH, University of Luxembourg, Esch-sur-Alzette, Luxembourg
frederic.clavert@uni.lu

Received 1 December 2023 | Accepted 1 April 2024 |
Published online 10 December 2024

Abstract

This *Memory Studies Review* special issue explores the intricate relationship between artificial intelligence (AI) and collective memory. In the one hand, the emergence of generative AI, exemplified by ChatGPT's 2022 release, appears to herald a new infrastructure for collective memory. On the other, the memory studies work highlights the limits and the backlashes of this new form of memory in its social dimension. This leads to raise a provocative, open-ended question: Is artificial intelligence the future of collective memory? Our issue brings together diverse perspectives from memory studies scholars of different backgrounds and machine learning practitioners, fostering critical engagement with AI in memory practices. This multidisciplinary approach offers an initial exploration of the interactions between AI-powered software, platforms, and collective memory. The articles herein present a multifaceted analysis of AI's role in shaping collective memory's future. We advocate for increased

Published with license by Koninklijke Brill BV | DOI:10.1163/29498902-202400019
 | ISSN: 2949-8902 (online)

interdisciplinary collaboration and ethical reflection in this rapidly evolving domain, providing memory studies scholars with a foundation for understanding and engaging with these technological transformations.

Keywords

artificial intelligence – artificial memory – digital memory studies – machine learning

The definition of "collective memory" has been a major topic of discussion and dispute among scholars coming from humanities and social sciences interested in understanding the dynamics of remembrance (Gensburger, 2016). However, and besides their different views, most of them agreed on the fact that "collective memory" was nothing like a "computer memory": it is different from a stock of souvenirs and images of the past. In 2000, still, Liam Bannon suggested that "what has blinded us to the richness of human activities related to the theme of collective remembering has been the dominance of the computer model of mind in more recent cognitive psychology and artificial intelligence research, where models of human memory have been imported from computing" (Bannon, 2000, p. 278).

Twenty years later, however, when artificial intelligence (AI), and more specifically generative AI, seems to have taken a technological leap forward as a new kind of infrastructure of memory. This computer-generated memory of a new nature could appear as the future of collective memory. This is the question that lies at the heart of this special issue: is artificial intelligence (really) the future of collective memory? By raising this (slightly) provocative question, we aim to foster a dialogue between specialists in memory studies and experts in engineering and machine learning. It is now clear that these two fields could respectively benefit from an informed dialogue (Kvasnička and Pospíchal, 2015). And with this issue, we hence offer memory scholars a set of introductory articles by authors from diverse backgrounds.

In the aftermath of Covid, for example, the need for mourning and for an immediate visual memorialisation of the dead relied on artificial intelligence to produce a memorial of portraits of all the victims of the Covid-19 pandemic in the US, based not on their actual pictures, which would be impossible to gather in a non-biased, practical way, but through representative portraits generated by an artificial intelligence using data available about the dead (age, ethnicity, gender, etc.) (Korostoff n.d.; Fridman & Gensburger, 2024). What if generative AI could enable us to remember people forever, transforming existing data in actual embodied memories?

In this issue, Matthias Meitzler, Jessica Heesen, Martin Hennig and Regina Ammicht Quinn have started to answer those questions by pushing this topic a bit further, analysing what AI is doing to the afterlife in the digital era. Investigating the ethical (and commercial) dimension of the use of AI, they question the risk of deep fakes about the deceased, but also less obvious ethical problems, such as the respect of the deceased will. But they also track the role of AI in the afterlife: they investigate how digital media are changing grief practices, with new ways to stay connected to the dead, through a second "body" of memories and narratives about the deceased. It allows a "programming permanence" of the deceased that implies a deep ethical reflection on issues linked to data privacy, integrity and the social impact of digital afterlife. The authors also point out that AI-based digital afterlife has applications beyond private grief, such as in politics, entertainment, and business but also in education, including teaching of the history of the Holocaust.

Indeed, while, with the last witnesses of the Holocaust dying, a crisis of witnessing is underway (Lothe et al., 2012), artificial intelligence has been increasingly used to give birth (again) to Holocaust survivors and enable interaction between their holograms, school students or museum visitors (Makhortykh, 2021b; Shur-Ofry & Pessach, 2020; Walden, 2021). Of course, this "new perspective on testimony", following the name of the program which created these holograms, has no reason to be limited to the Holocaust in the future. "Imagine, for example, an AI that integrates a large collection of testimonies from the Vietnam War, thus creating the 'ultimate witness' – one that delivers an integrated testimony about the War – or a 'virtual Abraham Lincoln' – one that relies on the 40,550 'Lincoln papers' stored at the Library of Congress to answer people's questions" (Shur-Ofry & Pessach, 2020, p. 988). Meanwhile, other works have highlighted what artificial intelligence and machine learning tools can bring to heritage and archives institutions (Barlindhaug, 2022; Bultmann et al., 2022; Jaillant, 2022; Pessanha & Salah, 2021; Rolan et al. 2019), assessing the extent to which they are "the future of memory institutions" (Horsley, 2020). Furthermore, the release of ChatGPT, the most well-known generative AI platform today (that is, towards the end of 2024) has been the starting point for the opening of numerous AI services of varying quality promising to chat with historical figures, as well as for the creation of business companies, rarely politically neutral, promising to bring the "past back to life".

Of course, some researchers from the field of memory studies have already investigated the relationship between (collective/cultural) memory and the digital: introducing the collective book he edited, Andrew Hoskins (2018) describes a past that has become "restless" after the connective turn that we

have been experiencing; a collective memory scattered by the illusions of the age of access into a "memory of the multitude". Hoskins notes that the Internet, as "the first medium that's actually bigger than us", is "the technology that makes visible our inability to encompass everything" (Hoskins, 2018, pp. 1–3).

Hoskins' (2018) pessimistic view claims that a new form of memory is being born: the memory of the multitude. However, such an analytical perspective is currently being challenged in several ways. For instance, the practice of digital memory can be seen as an implementation of theoretical claims which have been made by memory studies for several decades: collective memory is a process, constantly "mediated and remediated by multiple media, with the participation of dynamic communities that perform rather than represent the past" (Mandolessi, 2023, p. 1514). This view implies to take into account a series of transformations: first the emergence of digital archive, making databases the privileged form of collective memory, instead of narrative; second, an agency that is distributed between human and non-human agents and the importance of mnemonic assemblages instead of objects (Mandolessi, 2023).

In this issue, the question of non-human agency is indeed investigated from different perspectives. For instance, Pierre Depaz looks at how interfaces, chatbots and developers' memory (Github, Stack Overflow) are interacting. He considers the rise of the use of Large Language Models[1] (LLM) to write code as susceptible to change the collective memory of developers. The influence of non-human agents is visible here in the shift from code collaborative writing to the use of coding assistants. This shift will change the way memory is "recorded, retained and accessed technologically". In this scenario, LLM-powered chatbots become "memory interfaces", justified by writing code productivity. Developers' values are shifting from collaboration to immediate productivity, though still relying on some sort of collective memory, in the sense that those code-specialised LLMs need to be trained, usually on GitHub code repositories or Stack Overflow conversations: there is here a paradox of a new productivity ethos that is nevertheless based on data generated by the previous, community-based, ethos. A switch of interface is changing the representation of the knowledge that is accessed. In the end, this modifies the way collective memory is built. This raises the question of the sustainability of LLMs in building collective memory, all the more that, as time flies, LLMs will be increasingly trained on synthetic data, which could lead to a decline in linguistic diversity

1 Large Language Models are artificial intelligence systems based on artificial neural networks. They deduce statistical relationships between "tokens" (more or less words) from training on a large amount of text. They can achieve a general-purpose text generation. They are the core engines of chatbots such as OpenAI's ChatGPT, Anthropic's Claude or Mistral's Le Chat.

(Guo et al., 2024) and hence damage the very dynamics of generative memory, which is always a combination of repetition and recreation.

The deep entanglement between memory and technology is also well illustrated and demonstrated by Alina Volynskaya's article. In her article "Collective Memory Through Computer Memories: Retracing and Interpreting the Archive of the Stanford Artificial Intelligence Laboratory", Volynskaya investigates the SAILDART archive, a (rescued) repository of the Stanford Artificial Intelligence Laboratory (SAIL) files from 1972 to 1990. SAILDART, Volynskaya argues, represents a new form of machine-recorded collective memory that reflects the day-to-day, personal and collective history of the pioneering AI lab. She also demonstrates – which brings us back to the notion of "dark archive" (Jaillant, 2022) – that AI itself (in this case, ChatGPT) may allow us to understand and interpret this archive that is made of past and obsolete technology futures. Generative AI here, is a means to make sense of machine memories, that will be progressively more useful to human and social scientists.

For instance, practical research to study online echoes of commemorations show that large web platforms, which can be partially AI-driven, allow memory practices that can be better studied with tools based on machine learning (Clavert, 2021). This illustrates a second strand that diverges from *digital memory studies* as Hoskins grounded them, that might come precisely from what Mandolessi (2023) sees as a redistribution of agency: among non-human agents, AI-based systems are probably, with generative AI platforms, the most numerous today. AI here can be seen as a way to answer the age of access: as generative artificial intelligence systems are fed with large amounts of data generated during our digital era, they could be a (albeit imperfect) way to explore those large amounts of data.

Those generative AI systems have led to the emergence of new memory studies literature on the issue of the relationship between collective memory and artificial intelligence. However, it is striking to see that so far most of these articles have neglected the infrastructural shift implied by this interaction. They have focused on its ethical dimensions and on the policy needs implied by the development of artificial intelligence in the field of memory institutions, history and witnessing (Presner, 2016).

Memory studies scholars Victoria Grace Walden-Richardson, Mykola Makhortykh, and Kate Marrison, have coordinated an initiative to establish *Recommendations for using Artificial Intelligence and Machine Learning for Holocaust Memory and Education* (Walden et al., 2023). Additionally, law specialists Shur-Ofry and Pessach have called for "memory fiduciaries" to be imposed to artificial intelligence (2020), and historian Wulf Kansteiner (2022) has been dreaming of a GPThistory, a potential adaptation of an LLM enabling

it to become an actual and reliable auxiliary to the production of history and, from there, of collective memory: "If we think that the stories and images we consume influence our memories, identities, and future behavior, we should be very wary about letting AI craft our future entertainment on the basis of our morally and politically deeply flawed cultural heritage" (Kansteiner, 2022, p. 124). These reflections on the regulation of the use of machine learning and artificial intelligence in the field of historical narration and collective memory are crucial and are also one of the ambitions of this special issue in close relation to the consideration of infrastructural implications. In their article "Imagining Human-AI Memory Symbiosis: How Re-remembering the History of Artificial Intelligence Can Inform the Future of Collective Memory", Makhortykh and Walden-Richardson emphasise the lack of critical appraisal of AI. Going back to one of the computing pioneers, Alan Turing, they assess the scarcity of critical analysis regarding AI's functionality in the cultural memory sector (museums, archives, etc.) – most notably the lack of understanding of AI's mathematical functioning as well as the fact that their infrastructure is often overlooked. They underline, instead, the human tendency to attribute human-like features to AI and its uses, such as anthropomorphism – a phenomenon that is known since Weizenbaum's chatbot ELIZA in the 1960s and 1970s (Weizenbaum, 1976). As such, anthropomorphism tends to erase the fundamental mathematical logic of AI; but going back to AI as media memory rather than as a simulation of human memory is necessary to find out which productive (and, we would say, ethical) uses can be made of it in cultural memory institutions. Like Depaz in this issue, and following Mandolessi (2023) elsewhere, Makhortykh and Walden-Richardson also remind us that AI brings non-human agency to the making of collective memory, hence introducing different logics and different temporalities. They state: "Instead of perceiving the past as a narrative, AI views it as a sequence of values which have to be aligned according to 'the line of best fit' [...] Often, this alignment is operationalised not based on ethics or the public interest but on maximising user engagement with (memory-related) content". Based on this logic, they posit further analysis regarding AI where AI is not a simulation of human agency: "The scale of AI is intrinsically connected to *modularity* that results not only in AI's tendency to retrieve data fragments instead of coherent narratives but also allows AI to connect a vast range of data points and potentially re-contextualize them through constant rearrangement. It allows AI to re-interpret memories, thus allowing humans to look at the past from a diverse range of perspectives that may be culturally anchored or computationally transformed to highlight newly revealed data patterns." Those perspectives can emerge if developers, heritage practitioners, survivors, historians and memory activists work together, in order "to embed

diverse and potentially dissenting perspectives about specific memory events into training data and computational logic". Otherwise, a risk of switching to a new hegemonic memory exists.

Another group of questions which arise from this special issue tackles the nature of the relationship between artificial intelligence and collective memory. To what extent AI and machine learning tools can help social sciences and humanities to overcome the methodological loopholes social sciences and humanities have been confronted with, when trying to grasp and circumscribe "collective memory" in its shared narrative dimensions? Can AI help to make apparent, often indivisible, social frameworks of collective memory, or will the emergence of AI, and the social uses it will lead to, change the very nature and functioning of "collective memory" (Makhortykh, 2021a)? Can we speak of "robotic collective memory" or "algorithmic memory" beyond simple metaphors? Should we consider these AI and Machine learning tools as a way to make visible "collective memory", or should we acknowledge the fact that they generate and promote a new form of "collective memory"?

In his article "Nostophiliac AI: Artificial Collective Memories, Large Datasets and AI Hallucinations", Phivos-Angelos Kollias prefers to explore the concept of "artificial collective memories" as an interplay between human collective memory and the large datasets used by AI algorithms, such as LLMs or text-to-image models. Using his audio-visual project "nostophiliac AI", he examines the relationship between collective and individual memory as reflected and manipulated through AI. He hence argues that a portion of our collective memory is codified in the large datasets used by AI, which represent a "digital twin" of collective memory and terms this "artificial collective memory". Kollias's creative approach involves a continuous transformation of digital "found objects" (images or audio snippets) through AI, defining keyframe-like states while allowing interpolation. This creates a flux of meaningful transformations and a dialogue between the algorithm and the artist. Kollias distinguishes between curated datasets, representing a refined collective memory, and larger uncurated datasets, mirroring a more generalised collective memory. He also discusses the concept of AI-generated "pseudo-memories" that may interlace with human collective memory. Throughout his article, he emphasises AI's potential to mirror and alter representations of human memory and perception. The article urges deeper exploration into how these digital reflections may impact our understanding and evolution of collective memory itself.

Kollias' article can help us formulate another question: can scholars' works on collective memory use these new AI products as innovative research tools to critically and experimentally engage with the often taken for granted "historical

analogies" mechanism, or the relationship between moral values and collective memory? Could we imagine that the very biased nature of AI – which is the reason why comments have so far been focusing on ethical issues – make it a promising tool for exploring "collective memory" dynamics when defined as a socially structured and spatially located point of view of the past? The debate about ethics, collective memory, and AI invites us to concretely and explicitly list what is left of the concept of collective memory if it is separated from all its moral, and often implicit, implications. In other words, this special issue considers that thinking critically about the relationship between collective memory and artificial intelligence can help the field of memory studies move forward.

We also want to stress on that memory studies can bring something to the field of artificial intelligence. Makhortykh and Walden-Richardson insist on the necessity to integrate humanities perspectives in AI, in how AI is taught and conceived in computing, in industry, and in heritage sectors. That, according to them, implies engagement between humanities scholars, heritage practitioners and computing researchers and industry. They also acknowledge the necessity of defining the future imaginaries of AI in the context of collective remembrance: acknowledging the potential of AI to support different forms of engagement with the past. AI tools are first and foremost memory products, as, at least for connectionist AI, they rely on the notion of training which implies the use of datasets – such as CommonCrawl, a sort of archive of the web – which are a product of human activities and memories. However, most of the current remarks on the implications of this simple fact are very rarely informed by the knowledge memory studies works have established over these past twenty years. Crossing these two literatures could help research on AI to move beyond the very notion of bias to pay attention to the constantly organisational, spatial, and socially framed nature of all memory dynamics. Previously, different AI systems relied on different conceptions of memory and internal memorialisation dynamics (Romero, 2021). Therefore, memory studies have a lot to teach people developing "multi-agent systems" meant to deal with memory, but so far very few works have crossed the two literatures. It is necessary for these cross-fertilisations to be ordered, because of, for example, the necessity to build certain typologies regarding this matter. Likewise, the articulation between accuracy, authenticity, and exemplarity, as between the fictional and non-fictional, has been at the core of public history and memory studies work for some time. However, so far, this has failed to shed light on the future of artificial intelligence and its social uses.

Moreover, artificial intelligence memory is peculiar given the fact that the data used during the machine learning process have their own temporality and

spatiality (Clavert et al., 2022). The more we go into the past or the more we go into so-called "data poor" regions of the world – as already argued by Patrick Manning when writing about big data of the past (2013) – the less data is available to "nourish" the AI system. The danger is, as Rik Smit, Thomas Smits, and Samuel Merrill, remind us, the construction of a hegemonic collective memory, all the more hegemonic that we are prone to anthropomorphism, and, hence, to not use our critical abilities towards the collective memory resulting from non-human agency. AI systems, here large language models, are, they argue, based on the work of Bender et al., a stochastic rendering of memory. Instead, LLM's influence on collective memory should be seen as the result of a continuous human/non-human agencies interaction. The fact that LLM s like ChatGPT are returning an attractive and mainstream narrative of the past is not so much linked to biases per se, but rather to an unequally distributed agency between human and non-human actors on the one hand, yet also between different parts of humanity on the other. They push forward the hypothesis that regular users have their own imaginaries of AI, that will shape their perception and their interaction with the system. ChatGPT is then seen as a "friendly and helpful" assistant, though it will forget quickly and is not context-aware. Expert users might better understand how AI systems are working but are not able to change how they function. Designers have access to and can change the system for specific uses. They themselves are dependent on investors and shareholders. ChatGPT aims to serve the most agreeable answer to the most users possible, which leads to a hegemonic view on the past.

In fact, the very first and main "bias" of any AI system is one of temporality and spatiality, topics which have been at the core of memory studies since the very birth of the field. Julien Schuh, qualifying AI as artificial memory, reminds us that tech giants are ambitiously positioning these generative AI models as primary mediators for information access, integrating them into popular platforms. If those AI systems can present an innovative way to explore collective memory, weaving together fragmented pieces of the past, their widespread use raises profound ethical concerns. Challenges arise regarding historical truth, representativeness, and the vast data's origin, leading to debates on data ownership and the governance of digital collective memory. Schuh depicts AI as artificial memory in the sense that it allows first and foremost the automation of tasks linked to the making of memory (including information retrieval). Taking the example of the images of France Gall in several AI-based systems, including the data analysis that led Google to create a 75-year anniversary doodle of the French singer, or the retouching of photographs of the singer, or generated image of France Gall and singer and

songwriter Serge Gainsbourg, Schuh depicts embedded biases. In this matter, we could say that the knowledge built by memory studies these past decades is more than valuable to understand and explicate some of the biases that are at the core of any AI system, which in so many ways can be defined as the product of collective memory processes in the Halbwachsian definition of the term.

Last but not least, it can be said, regarding the collective memory of artificial intelligence, as both an economic product and a symbolic object/actor, that its mythologies and its narratives are an important part of the investigation this special issue wants to foster. It is fair assumption to suggest that the way that the "history" of AI has been written throughout time has only recently started to be studied (Gefen, 2022). This special issue intends to engage with these thought-provoking questions, and yet, at the same time, to engage with the extant literature regarding the effects of AI and Machine learning tools on the future of collective memory.

⁘

Some points are unfortunately missing in this special issue. One of the most important is probably the question of digital labour. Following Antonio Casilli's *En attendant les robots*(Casilli, 2019),[2] ethical considerations on using AI that was trained by workers, often from the Global South, should be considered when integrating AI in memory studies and practices. To avoid the most toxic generated texts, ChatGPT's outputs, for instance, were corrected by Kenyan workers, which was a painstaking amount of work (Perrigo, 2023). But a second consideration should be added here: this human work to polish non-human agents' content production is often non-advertised and the instructions received by those workers and the firms that employ them are not known. Here AI functioning relies on a dynamic of oblivion of its internal production. In other terms, large parts of our collective memories can be undermined at this stage because they are considered as "non-aligned" with the conditions needed for the commercialisation of large language models. All AI-based systems are released within a precise context and a specific society. If tracking censorship implemented in an image generating system like the Russian ruDALL·E is relatively easy (all prompts with a word identified as related to Ukraine will give birth to images of flowers, for instance), it is more complex to track with

2 English version to be published in Autumn 2024: *Waiting for robots* (The University of Chicago Press).

other systems, such as ChatGPT (Ermoshina, 2023). Those "alignments" of AI systems are influencing how we can access and retrieve collective memory-related information. Furthermore, the invisibilised memory of digital labour workers should be investigated too.

In November 2023, a month after the bloody Hamas attack on Israeli territory, a famous stock photographs website added to its catalogue AI generated images of the Israel-Hamas conflict (Tangermann, 2023). To our mind, this fact is one more illustration of the emergency we are in: memory studies should engage with AI in general and generative AI in particular. This special issue is a first step in this endeavour. To move forward and to promote the study of AI by memory scholars, we have thus set up a collaborative bibliography on Zotero,[3] created a discussion list, and launched a website.[4]

We would like to thank the *Memory Studies Review* team for their support, especially Hanna Teichler, as well as the reviewers whose comments have obviously led to better articles and a really innovative issue. We also warmly thank our authors and all the participants of the seminar we organised in September 2023 at the University of Luxembourg's Centre for Contemporary and Digital History.

Frédéric Clavert is Assistant Professor at the Centre for Contemporary and Digital History (C²DH, University of Luxembourg). A member of the digital humanities community, he turned his attention to the study of the relationship between historians and their primary sources in the digital age on the one hand, and the use of massive data from web platforms in memory studies on the other. He led the *#ww1* project around the Centenary of the Great War on Twitter. With Caroline Muller (Université Rennes 2), he is co-editor of the online book *Le Goût de l'archive à l'ère numérique*. He is managing editor of the *Journal of Digital History*.

Sarah Gensburger is Full Professor of Sociology and Political Science at the French National Centre for Scientific Research and Sciences Po-Paris. She is the author of *Beyond Memory. Can We Really Learn from the Past?* (with Sandrine

3 https://www.zotero.org/groups/4874770.
4 https://www.artificial-memory.net/.

Lefranc, Palgrave, 2020, also in Arabic, French and Spanish), *Memory on My Doorstep. Chronicles of the Bataclan Neighborhood* (Paris, 2015–2016) (Leuven University Press, 2019) and the co-editor of *Administrations of Memory* (with Sara Dybris McQuaid, Springer, 2022). Her current research interests concern the relationship between neoliberalism, the crisis of the welfare state and the contemporary memory boom as well as the impact assessment of memory policies and the unequal access to participatory archiving. In 2021, she was elected President of the international Memory Studies Association.

Works cited

Bannon, L. (2000). Towards artificial memories? *Le Travail Humain*, *63*(3), 277–285.

Barlindhaug, G. (2022). Artificial intelligence and the preservation of historic documents. *Proceedings from the Document Academy*, *9*(2).

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. https://doi.org/10.1145/3442188.3445922.

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: Mass Atrocities in the Digital Era Working Paper*, 3.

Clavert, F. (2021). History in the era of massive data. *Geschichte und Gesellschaft*, *46*(1).

Clavert, F., Mahroug, S., & Schafer, V. (2022). Préservation et distorsion: l'espace-temps des réseaux socio-numériques et du web archivé. *Revue d'histoire culturelle*, *2022*(5).

Ermoshina, K. (2023). Censure, auto-censure et les ruses du prompt: coconstruction des filtres dans les IA génératives. Unpublished. https://observatoire-ia.pantheonsorbonne.fr/evenements/penser-et-creer-ia-generatives.

Fridman, O., & Gensburger, S. (2024). *The Covid-19 pandemic and memory: Remembrance, commemoration, and archiving in crisis*. Palgrave.

Gefen, A. (2022, November 7). CulturIA. A cultural history of artificial intelligence. [Blog post].

Gensburger, S. (2016). Halbwachs' studies in collective memory: A founding text for contemporary 'memory studies'? *Journal of Classical Sociology*, *16*(4), 396–413.

Guo, Y., Hu, J., Gong, Y., Bai, G., Kochkina, E., Ye, H., Shen, S., Zhang, Y., Sønderby, C. K., Kadav, A., Comeau, L., Molloy, M., Krueger, D., Zhao, T., & Yogatama, D. (2024). The curious decline of linguistic diversity: Training language models on synthetic text. *arXiv*.

Horsley, N. (2020). Intellectual autonomy after artificial intelligence: The future of memory institutions and historical research. In N. Lushetich (Ed.), *Big Data – A New Medium?* Routledge.

Hoskins, A. (Ed.). (2018). *Digital memory studies: Media pasts in transition*. Routledge.

Jaillant, L. (Ed.). (2022). *Archives, access and artificial intelligence: Working with born-digital and digitized archival collections*. Bielefeld University Press/transcript Verlag.

Kansteiner, W. (2022). Digital doping for historians: Can history, memory, and historical theory be rendered artificially intelligent? *History and Theory*, *61*(4), 119–133.

Korostoff, M. (n.d.). 535,000 faces: Comprehending the death toll of Covid-19. Retrieved February 22, 2023, from https://mkorostoff.github.io/hundred-thousand-faces/.

Kvasnička, V., & Pospíchal, J. (2015). Artificial intelligence and collective memory. In P. Sinčák, P. Hartono, M. Virčíková, J. Vaščák, & R. Jakša (Eds.), *Emergent trends in robotics and intelligent systems* (pp. 283–291). Springer International Publishing.

Lothe, J., Suleiman, S. R., & Phelan, J. (2012). *After testimony: The ethics and aesthetics of Holocaust narrative for the future*. The Ohio State University Press.

Makhortykh, M. (2021a). Memoriae ex machina: How algorithms make us remember and forget. *Georgetown Journal of International Affairs*, *22*(2), 180–185.

Makhortykh, M. (2021b, February 4). Algorithmic auditing, the Holocaust, and search engine bias. Digital Holocaust Memory.

Mandolessi, S. (2023). The digital turn in memory studies. *Memory Studies*, *16*(6), 1513–1528.

Manning, P. (2013). *Big data in history*. Palgrave Pivot.

Perrigo, B. (2023, January 18). Exclusive: The $2 per hour workers who made ChatGPT safer. *Time*. https://time.com/6247678/openai-chatgpt-kenya-workers/.

Pessanha, F., & Salah, A. A. (2021). A computational look at oral history archives. *Journal on Computing and Cultural Heritage*, *15*(1).

Presner, T. (2016). The ethics of the algorithm: Close and distant listening to the Shoah Foundation Visual History Archive. In C. Fogu, W. Kansteiner, & T. Presner (Eds.), *Probing the ethics of Holocaust culture* (pp. 167–202). Harvard University Press.

Rolan, G., Humphries, G., Jeffrey, L., Samaras, E., Antsoupova, T., & Stuart, K. (2019). More human than human? Artificial intelligence in the archive. *Archives and Manuscripts*, *47*(2), 179–203.

Romero, A. (2021, June 1). GPT-3 scared you? Meet Wu Dao 2.0: A monster of 1.75 trillion parameters. *Towards Data Science*. https://towardsdatascience.com/gpt-3-scared-you-meet-wu-dao-2-0-a-monster-of-1-75-trillion-parameters-832cd83db484.

Shur-Ofry, M., & Pessach, G. (2020). Robotic collective memory. *Washington University Law Review*, *97*(3), 975–1005.

Tangermann, V. (2023). Adobe caught selling AI-generated images of Israel-Palestine violence. *Futurism*. https://futurism.com/adobe-caught-selling-ai-generated-images-israel-palestine-violence.

Walden, V. G. (2021). Afterword: Digital Holocaust memory futures: Through paradigms of immersion and interactivity and beyond. In V. G. Walden (Ed.), *Digital Holocaust memory, education and research* (pp. 267–296). Palgrave.

Walden, V. G., Makhortykh, M., Jasper, C., Bade, D., & Balazs, K. (2023). Recommendations for using artificial intelligence and machine learning for Holocaust memory and education. *REFRAME*.

Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. W. H. Freeman.