

ARTICLE

## AI Governance as Regulation of Collective Memory

Michal Shur-Ofry

The Network Law Review is pleased to present a special issue entitled “The Law & Technology & Economics of AI.” This issue brings together multiple disciplines around a central question: What kind of governance does AI demand? A workshop with all the contributors took place on May 22–23, 2025, in Hong Kong, hosted by Adrian Kuenzler (HKU Law School), Thibault Schrepel (Vrije Universiteit Amsterdam), and Volker Stocker (Weizenbaum Institute). They also serve as the editors.

\*\*\*

**Abstract.** The emerging regulatory landscape in the field of AI will substantially influence the construction of collective memory and play a critical role in shaping our “future’s past”. This contribution takes a close look at the traits of collective memory, maps its potential frictions with the emerging AI field, and demonstrates how multiple AI governance schemes—from direct AI regulation, to competition policy in AI markets, to copyright law applicable to training datasets—can either enhance or mitigate these tensions. The analysis does not offer an exhaustive, one-size-fits-all solution. Rather, it maintains that any regulatory initiative in the field of AI should consider its potential implications for our shared past, and the need to preserve a thriving sphere for collective memory.

## 1. Introduction

Collective memory is increasingly mediated through AI tools. Prominent examples include initiatives that focus specifically on mediating past events—such as “virtual witnesses” of the Holocaust<sup>1</sup>, the “library of the future”, developed as part of a collaboration between academia and Open AI<sup>2</sup>, or the use of AI for restoring cultural heritage artifacts<sup>3</sup>. Additional examples are general-purpose large language models (LLMs), which are gradually becoming a major source through which people obtain information about the past.<sup>4</sup>

Yet, the heated debates on AI governance in recent years are largely forward-looking, trying to predict how rapid developments in this field will shape our future economic, health, education, and other systems, and what the appropriate regulatory response should be. This short essay suggests that AI’s regulatory landscape will inevitably influence our past, or more accurately, our collective memory—the ways societies and groups will remember their joint past. It begins with a brief introduction to the concept of collective memory and its societal significance. It then maps potential tensions and frictions at the intersection of collective memory and AI, and explores how AI governance schemes across different fields may effectively—if inadvertently—regulate collective memory by exacerbating or mitigating those tensions.

I focus on three areas: competition and diversity in the market of LLMs; copyright law and its interface with AI training datasets; and direct regulation applicable to AI that mediates collective memory. These cases are merely illustrative. My purpose is neither to exhaust the exploration of the multiple interfaces between AI governance schemes and collective memory, nor to offer complete solutions that will align AI regulation with the societal need to preserve a robust sphere for collective memory. Rather, I aim to provide a framework for illuminating and discussing the role of AI governance in crafting a space that would allow collective memory to thrive.

---

<sup>1</sup> Michal Shur-Ofry & Guy Pessach, *Robotic Collective Memory*, 97 Wash. U. L. Rev. 975 (2020).

<sup>2</sup> Oxford and OpenAI Launch Collaboration to Advance Research and Education March 4, 2025), [https://www.ox.ac.uk/news/2025-03-04-oxford-and-openai-launch-collaboration-advance-research-and-education?utm\\_source=chatgpt.com](https://www.ox.ac.uk/news/2025-03-04-oxford-and-openai-launch-collaboration-advance-research-and-education?utm_source=chatgpt.com).

<sup>3</sup> Münster S et al., *Artificial Intelligence for Digital Heritage Innovation: Setting up a R&D Agenda for Europe*. 7(2) Heritage 794 (2024).

<sup>4</sup> Michal Shur-Ofry, *Multiplicity as an AI Governance Principle*, 100 Ind. L. J. 1 (2025).

## 2. Collective Memory Meets AI

The concept of collective memory, introduced by sociologists, refers to the joint recollection of the past by societies, nations, or other communities sharing a common identity.<sup>5</sup> Underlying this notion is an understanding that our memories (in the sense of ‘recollection of the past’) are not confined to our own individual experiences but are also a product of social construction—they are formed, in part, by the groups to which we belong, be they nations, minority groups, religious groups, or other communities. The emphasis in collective memory studies, then, is not on the cognitive dimensions of memory, but rather on the social dimensions and the factors that shape group memory. To quickly illustrate, we may say, in everyday parlance, that we remember the first landing of man on the moon, or the Holocaust, even though we did not personally experience these events, and may not even have been alive when they took place.<sup>6</sup>

Ample interdisciplinary literature from the past decades explored the significance of collective memory for the formation and flourishing of group identity. Collective memory is crucial for narrating the life stories of nations and communities and provides vital building blocks for their shared identities.<sup>7</sup> It can be particularly important for minority groups, since it allows such groups to “seek a voice and a shared identity that is distinct from and may conflict with that of the nation-state”.<sup>8</sup> Finally, collective memory does not merely benefit communities. Because our self-identity draws from and is entangled with the identities of the social and cultural groups to which we belong, it is also crucial for the formation of individual identity.<sup>9</sup>

Come AI. The rapid developments in the field of AI and its integration into almost every aspect of human life present new and significant opportunities for collective memory. For example, AI-generated memorials mediate historical

---

<sup>5</sup> For some of the burgeoning interdisciplinary literature on collective memory, its traits and social significance, see Maurice Halbwachs, *On Collective Memory* (Lewis A. Coser ed. and trans., 1992); Jeffrey K. Olick & Joyce Robbins, *Social Memory Studies: From “Collective Memory” to the Historical Sociology of Mnemonic Practices*, 24 *Ann. Rev. Soc.* 105, 106 (1998); Jeffrey K. Olick, Vered Vinitzky-Seroussi & Daniel Levy, Introduction to *THE COLLECTIVE MEMORY READER 3* (Jeffrey K. Olick, Vered Vinitzky-Seroussi & Daniel Levy eds., 2011); Jeffrey K. Olick, *Collective Memory: The Two Cultures*, 17 *Soc. Theory* 333 (1999).

<sup>6</sup> For this distinction, see Eviatar Zerubavel, *Social Memories: Steps to a Sociology of the Past*, 19 *Qualitative Soc.* 283 (1996).

<sup>7</sup> E.g., Olick, Vinitzky-Seroussi & Levy, *supra* note 5, 8-29.

<sup>8</sup> Guy Pessach & Michal Shur-Ofry, *Intangibles and Collective Memory: The Role (and Rule) of Law*, 25 *Jer. Rev. Legal Stud.* 227 (2022).

<sup>9</sup> E.g., Susan A. Crane, *Writing the Individual Back into Collective Memory*, 102 *The American Historical Rev.* 1372 (1997).

narratives to the public in interactive and user-friendly ways.<sup>10</sup> AI-based historical holograms can interact with users in a human-like way, years after the passing of the actual individuals they “represent”,<sup>11</sup> and general-purpose LLMs easily summarize vast amounts of historical information, presenting it in a friendly and accessible way.<sup>12</sup> The increasing integration of AI in mediating collective memory has led researchers to predict a future where AI driven systems “will effectively decide what information sources and what interpretations of the collective past gain more visibility, and thus, shape how this past is remembered.”<sup>13</sup>

As the latter quote indicates, alongside the significant advantages, the intersection of collective memory and artificial intelligence is also a site fraught with tensions and frictions. In fact, several prominent traits of collective memory make it particularly vulnerable to the disruptive effects of AI.

First, collective memory has a “multifaceted nature”. At the most fundamental layer, it relies on historical events, and the accuracy and veracity of historical information constitutes a crucial building block in collective memory formation. Yet, collective memory is not synonymous with history. Rather, the same historical event can have a certain shared meaning in the collective memory of one group and an entirely different meaning in the collective memory of another community.<sup>14</sup> Collective memory can thus encompass a multiplicity of interpretations and voices, allowing different communities to narrate their own narratives and seek a shared identity that at times may be different from, and even conflict with, that of majority groups.<sup>15</sup>

Generative AI, particularly general purpose LLMs, seriously challenges this intricate nature. The first and obvious challenge is, of course, the ease with which representations of past events and historical figures can be misrepresented or

<sup>10</sup> E.g., Ștefania Matei, *Generative Artificial Intelligence and Collective Remembering: The Technological Mediation of Mnemotechnic Values*, 2 J. Hum.-Tech. Rel., 4 (2024).

<sup>11</sup> Mykola Makhortykh, *No AI After Auschwitz? Bridging AI and Memory Ethics in the Context of Information Retrieval of Genocide-Related Information*, in A. Mukherjee, J. et al. (eds.), *Ethics in Artificial Intelligence: Bias, Fairness and Beyond* 71, 73 (2023); Cf. Pataranutaporn et al., *Living Memories: AI-Generated Characters as Digital Mementos*, ACM Conf. Hum. Factors Comput. Syst. (Apr. 2023), <https://dl.acm.org/doi/fullHtml/10.1145/3581641.3584065>.

<sup>12</sup> Mario Carretero & Elisa Gartner, *Artificial Intelligence and Historical Thinking: A Dialogic Exploration of ChatGPT / Inteligencia Artificial y Pensamiento Histórico: Una Exploración Dialógica del ChatGPT*, 45 Stud. Psychol. 80, 81 (2024).

<sup>13</sup> Mykola Makhortykh, *Shall the Robots Remember? Conceptualising the Role of Non-Human Agents in Digital Memory Communication*, 3 Memory, Mind & Media e6, 6 (2024).

<sup>14</sup> For example, the atomic bomb during WWII has a different role in the collective memory of the American and the Japanese people. See Stefanie Fishel, *Remembering Nukes: Collective Memories and Countering State History*, 1 Critical Mil. Stud. 131, 136–37, 141 (2015).

<sup>15</sup> Pessach & Shur-Ofry, *supra* note 8.

manipulated through generative AI, whether intentionally or as a result of unintentional inaccuracies and ‘hallucinations’. These, in turn, can penetrate and harm the fundamental information layer, on which collective memory is then built. Yet, AI can potentially impact collective memory in other, more subtle ways, by challenging the additional, important, layer, that allows for diverse narratives and social constructions. As I argued elsewhere, due to the technological paradigm underlying LLMs that relies on statistical frequency, their outputs are likely to be concentrated and geared toward the popular and mainstream—projecting a ‘narrow universe’ rather than a breadth of voices and narratives.<sup>16</sup> For example, in exploring the default outputs generated by LLMs when asked about prominent historical figures, we found that their answers converged around a small number of personae (such as Lincoln, or Darwin), and were substantially less diverse than human responses.<sup>17</sup> This, in turn, implies that when asked about the past, LLMs’ default outputs might oust niche and broader perspectives that allow smaller communities to narrate their collective memory.

A second, related, tension concerns raw materials. Collective memory is facilitated by, and often depends on, access to raw materials. These materials can include iconic documentations of major events but also piecemeal recordings of more mundane occurrences, embodied in photographs, diaries, videos, witness statements, letters, drawings, and other artifacts, that together enable to draw the multi-faceted narratives comprising collective memory.<sup>18</sup> Access to raw materials may be particularly important for the collective memory of extreme and radical events. As historian Saul Friedlander explained in the context of the Holocaust, the extremity, depravity, and scale of events, raise a “problem of representation.”<sup>19</sup> It is difficult for people under normal circumstances to grasp the event as a reality that actually happened. In such circumstances, providing broad access to raw materials—not through paraphrase or summary—could be crucial for bridging and alleviating the problem of representation.<sup>20</sup>

Here, too, the tension with AI becomes apparent. One of the great advantages of generative AI is its ability to process huge amounts of texts and data and summarize

---

<sup>16</sup> Shur-Ofry: Multiplicity, *supra* note 4. Cf. Rishi Bomassani et al., Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?, *NeurIPS*, 1 (2022), <https://arxiv.org/abs/2211.13972v1jll>.

<sup>17</sup> For details, see Michal Shur-Ofry, Bar Horowitz-Amsalem, Adir Rahamim, Yonatan Belinkov, Growing a Tail: Increasing Output Diversity in Large Language Models, *arXiv* (Nov. 5, 2024), <https://arxiv.org/abs/2411.02989> (the prompt asked to name three influential people from the 19th. century).

<sup>18</sup> Guy Pessach & Michal Shur-Ofry, Copyright and the Holocaust, 30 *Yale J.L. & Human.* 121, 133 (2018).

<sup>19</sup> See, generally, Saul Friedlander, Introduction, in Saul Friedlander, Ed., *Probing the Limits of Representation: Nazism and the “Final Solution”*, 1–21 (1992).

<sup>20</sup> *Id.*

them in a condensed and coherent way. But, in so doing, it also conceals large parts of the world, distancing us from raw materials that in some cases may be necessary for collective memory.

A third potential friction between AI and collective memory concerns trust. Traditionally, collective memory has been mediated through public or quasi-public institutions, such as museums, libraries, and archives that view the preservation and mediation of intergenerational memory as their primary mission.<sup>21</sup> Even when private parties, such as Google’s Cultural Institute, are engaged in the provision of collective memory, they are still perceived, and often describe themselves, as public-oriented.<sup>22</sup> As a result, people have a social expectation that entities active in the field of collective memory would adopt a public-oriented approach and are inclined to trust their outputs.<sup>23</sup>

Here too, AI complicates things. Even setting aside, for a moment, intentional manipulation of authentic historical materials through generative AI, it is far from clear that AI providers, which are typically private tech-giants would prioritize the public interest when it comes to potential externalities of their activities on collective memory. Users’ intuitive trust, in other words, may be jeopardized or compromised.

<sup>24</sup>

Altogether, aggregating these properties of collective memory against the realities of AI clarifies that the governance schemes applicable to AI can either exacerbate the disruption of collective memory or conversely mitigate those tensions.<sup>25</sup> The following sections elaborate and demonstrate.

---

<sup>21</sup> Shur-Ofry & Pessach, *supra* note 1, p. 992.

<sup>22</sup> See, e.g., Google’s description of its Cultural Institute, <https://about.artsandculture.google.com/> (“our mission is to preserve and bring the world’s art and culture online so it’s accessible to anyone, anywhere”).

<sup>23</sup> Cf. James Cuno, ed., *Whose Muse? Art Museums and the Public Trust* (2004); Guy Pessach, *The Role of Libraries in A2K: Taking Stock and Looking Ahead*, Mich. St. L. Rev. 257 (2007).

<sup>24</sup> As the following discussion clarifies, promoting a robust sphere for collective memory may be costly and involve, *inter alia*, diversifying training datasets, digitizing mass amounts of materials, or taking steps to ensure security and integrity of outputs, which, in the absence of legal incentives, private stakeholders may prefer to avoid.

<sup>25</sup> Cf. Pessach and Shur-Ofry: *The Role (and Rule) of Law*, *supra* note 6, at 228-29 (arguing that “the traits of collective memory imply that many ... legal fields, ... have impact on collective memory”).

### 3. AI Governance as Collective Memory Regulation

#### 3.1 Competition

##### 3.1.1. Competition and Diversity in the Market for Models

Competition is not intuitively associated with collective memory. However, collective memory's multifaceted nature clarifies the potential influence of a diverse AI arena, especially when it comes to general-purpose LLMs. The previously discussed tendency of LLMs to generate mainstream, concentrated outputs highlights the significance of ensuring users' access to several LLMs. While each LLM is still likely to project a narrow and concentrated view of the past, our study indicates that aggregating outputs from several LLMs can increase diversity, especially when combined with other diversity-inducing methods, such as raising the models' temperature or prompting for diversity.<sup>26</sup> To return to the previous example, when asked about personae from the 19th century, an aggregation of models generated a longer "tail" of outputs, adding additional figures to the initial condensed list, including Albert Einstein, Florence Nightingale, Friedrich Nietzsche, Marie Curie, alongside others.<sup>27</sup>

Even if each individual user were to consult only one or two LLMs when seeking information about the past, the availability of multiple models on the market means that different users will encounter somewhat different narratives.<sup>28</sup> Taken together, this diversity creates more opportunity for collective memory to flourish. Moreover, a diverse market of models would also increase users' ability to identify errors, manipulations, and hallucinations in historical outputs generated by a single LLM.

Altogether, a competition policy that ensures the availability of, and users' access to, a range of models will disperse the power to control collective memory, rather than leave it in the hands of a few.

---

<sup>26</sup> See Shur-Ofry, Horowitz-Amsalem, Rahamim and Belinkov, *supra* note 18 (further reporting, interestingly, that applying such steps when there are multiple possible outputs did not significantly harm models' accuracy).

<sup>27</sup> *Id.*, at p. 6-7.

<sup>28</sup> As the foregoing discussion indicates, one should distinguish between historical facts (e.g.: leaders who lived during the 19th Century), where the desirable standard should be accuracy, and collective memory narratives (e.g.: who was the most prominent leader in the 19th. Century?), which vary across cultural, national, and other communities.

### 3.1.2. Public LLMs?

According to reports, several countries, including Japan, Germany, France, and South Korea, have been engaged in the LLM market as active stakeholders by building, funding or otherwise supporting different “public LLMs”.<sup>29</sup> Such LLMs would presumably be trained on datasets in the local languages, which means that their outputs would give more prominence to the cultural and historical narratives of local societies.<sup>30</sup>

The lens of collective memory helps to evaluate such policies. On the one hand, direct state involvement in media markets can raise free speech and other concerns. Indeed, public LLMs could allow states to influence collective memory by delineating the boundaries of their training datasets, which in turn would impact their outputs. On the other hand, the foregoing analysis indicates that collective memory is a public good, characterized by non-rivalry and non-excludability.<sup>31</sup> This, according to standard public goods theory, means that leaving collective memory entirely to the market may result in a market failure: private AI companies will have no incentive to prioritize robust access to the diverse building blocks of collective memory. Cultural and historical heritage materials of local and niche communities, especially in languages other than English, may not penetrate training datasets on which AI is based. The omission of these materials from the models’ outputs will push them (and the narratives built from them) to the far margins of collective memory. “Public LLMs” may alleviate this concern and ensure the collective memory of specific societies and communities does not “drown” in an ocean of global or mainstream outputs.<sup>32</sup> Their introduction—alongside diversity in the private LLM market—may thus be desirable.

<sup>29</sup> See, Tim Hornyak, *Why Japan is Building Its Own Version of ChatGPT?* Nature (Sept. 14, 2023) (Japan); OpenGPT-X, *About the Project*, <https://opengpt-x.de/en/about/> (Germany); Release of largest trained open-science multilingual language model ever, CNRS (Jul. 12 2022), <https://www.cnrs.fr/en/press/release-largest-trained-open-science-multilingual-language-model-ever> (France); Shin Ha-Nee, *Korea Aims to Develop World-Class AI Model Through New Initiative* Korea JoongAng Daily (Feb. 20, 2025), <https://koreajoongangdaily.joins.com/news/2025-02-20/business/industry/Korea-aims-to-develop-worldclass-AI-model-through-new-initiative/2246580> (South Korea).

<sup>30</sup> Cf. Hornyak, *id.* (describing the Japanese attempt to build a version that will “grasp the intricacies of Japanese language and culture”).

<sup>31</sup> For the analysis of collective memory as a public good, see Pessach & Shur-Ofry: *Role (and Rule) of Law*, *supra* note 8, p. 229 (2022), <https://doi.org/10.1093/jrls/jlac011>; Cf. Richard S. Whitt, “Through a Glass, Darkly”: Technical, Policy, and Financial Actions to Avert the Coming Digital Dark Ages, 33 *Santa Clara Computer & High Tech. L.J.* 117, 178–79 (2016).

<sup>32</sup> Shur-Ofry: *Multiplicity*, *supra* note 4, at 35. Cf. Inna Kizhner et al., *Digital Cultural Colonialism: Measuring Bias in Aggregated Digitized Content held in Google Arts and Culture*, 36(3) *DHS* 607 (2020) (finding, *inter alia*, that art from provinces is underrepresented in Google’s A & C project, relative to art from capital cities).

### 3.2 Copyright Law and AI Training Datasets

Another indirect regulatory interface concerns copyright law and its impact on the datasets that ‘feed’ algorithmic memory.

One intense debate centers around the question of whether the training process itself constitutes permissible use or whether it requires the consent of the owners of copyright in the training materials (this question exceeds our scope and I do not explore it here).<sup>33</sup> Yet, from the perspective of collective memory, the influence of copyright starts even earlier, at the stage of digitization. Digitizing historical materials and archival collections—typically held by remembrance institutions, such as museums and archives—not only makes vast amounts of items accessible to broader audiences online, but also plays a critical enabling role in AI. Because a digital format is a precondition for inclusion in the training data on which AI is based, such digitization can also allow these heritage materials to enter the universe of algorithmic collective memory.

Copyright may hinder such digitization. Due to its typically long duration, it can pose an obstacle to the digitization of historical materials even if they were created a few generations ago.<sup>34</sup> This, in turn, may result in a ‘file drawer’ effect: what is not included in the digital space will only be seen in the (limited) space of the physical museum, or worse, will end up in a file-drawer. In the long term, AI’s disregard for such materials could lead to their marginalization—and the marginalization of the narratives built upon them—to the outermost periphery of collective memory.

Additionally, when it comes to materials documenting historical events—especially extreme events—there is an inherent tension between copyright law’s “fair use” exception and digitization efforts. The doctrine of fair use, adopted by various jurisdictions, allows the use of copyrighted materials without the consent of the copyright owners, under certain circumstances. However, the exception prioritizes partial and transformative uses, that alter the original work and do not copy it in its entirety.<sup>35</sup> This may discourage AI developers from providing users with access to complete, verbatim, copies of cultural and historical heritage materials. Yet, as the previous discussion suggests, providing full access to raw materials—rather than

---

<sup>33</sup> For a review, see, e.g., Matthew Sag and Peter Yu, *The Globalization of Copyright Exceptions for AI Training*, 74 *Emory L. J.*, (2025, forthcoming) <https://ssrn.com/abstract=4976393>.

<sup>34</sup> For concrete examples concerning materials created by Holocaust victims and perpetrators, see Pessach & Shur-Ofry: *Copyright and the Holocaust*, *supra* note 19.

<sup>35</sup> See, for example, the Fair Use provision in section 107 of the U.S. Copyright Act, which specifies among the fair use factors “the character of the use” and “the amount and substantiality of the portion used in relation to the copyrighted work as a whole”.

paraphrasing or summarizing—is crucial to mitigate the problem of representation.

<sup>36</sup>

These misalignments can be mitigated by calibrating copyright law in different ways: for example, by crafting specific exceptions that would enable digitization of entire collections held by remembrance institutions, as recently adopted by some jurisdictions,<sup>37</sup> or by a purposeful interpretation of the fair use doctrine that would facilitate the digitization and digital-display of historical materials in their entirety.<sup>38</sup> A detailed review of these schemes exceeds our scope. My point here is more general: When determining the different ways in which copyright may or may not apply to AI, legal policy makers should take into consideration the possible impact of any copyright governance scheme not only on current and future stakeholders, but also on our collective past.

### 3.3. Direct Regulation: AI as a Memory Mediator

Finally, regulation applicable directly to the AI field could clearly ease the frictions between AI and collective memory. To illustrate, obligations imposed on AI developers to abide by data security standards could reduce the vulnerability of those systems to hacking and technical failures that could harm the authenticity and integrity of the historical contents they mediate, and mitigate the risks of collective memory manipulation.<sup>39</sup> Duties of explainability and transparency<sup>40</sup> can alert users to the shortcomings and limitations of AI as a memory agent—for example by detailing the datasets on which ‘virtual witnesses’ or ‘generative archives’ were trained. Additionally, incorporating “multiplicity” in AI regulation as a high-level governance principle, as I suggested elsewhere, could benefit collective memory by obligating-or incentivizing-AI providers to expose users or alert them to the existence of multiple contents and narratives beyond the generated outputs, and encourage them to seek additional information.<sup>41</sup>

More specifically, Guy Pessach and I previously proposed conceptualizing entities that mediate collective memory through AI as “memory fiduciaries”, who are subject to fiduciary obligations toward their users. Such conceptualization, we argued, would provide a flexible governance scheme, rooted in the common law, that

<sup>36</sup> Cf. Copyright and the Holocaust, *supra* note 35, at pp. 151-154.

<sup>37</sup> For a WIPO study describing these exceptions and their specific conditions, see Kenneth D. Crews, Study On Copyright Limitations And Exceptions For Libraries And Archives: Updated And Revised (2017),

<sup>38</sup> Pessach & Shur-Ofry, *supra* note 19, at 170.

<sup>39</sup> For example, Art. 15 of the EU AI Act maintains that “High-Risk” AI systems must achieve an appropriate level of “accuracy, robustness, and cybersecurity”—see Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June, 2024 O.J. (L 1689)1 (“AI Act”).

<sup>40</sup> For example, AI Act, *Id.*

<sup>41</sup> Shur-Ofry: Multiplicity, *supra* note 4 (further elaborating how the principle can be implemented).

could adapt to new developments in the interface between AI and collective memory and respond to novel tensions as they arise.<sup>42</sup>

As part of those fiduciary obligations, we proposed recognizing a duty of integrity. Integrity means that entities that mediate collective memory through AI need to conform with “standards of accuracy and aspire [...] to be consistent with historical facts”.<sup>43</sup> The analysis above suggests that this duty is needed for mitigating concerns of distortion and manipulation of collective memory. Notably, the reminiscent “accuracy” obligations, recently introduced in the EU AI Act, could similarly promote this objective, and the foregoing analysis suggests that there are compelling reasons and for applying them to the mediation of collective memory.<sup>44</sup>

#### **4. Conclusion**

As Derrida wrote three decades ago, “[t]here is no political power without control of the archive, if not of memory”.<sup>45</sup> AI will increasingly function as our modern “Archive”, the major source for obtaining information about our shared past. Our analysis suggests that regulation affecting AI, both directly and indirectly, could substantially influence this Archive, and with it, collective memory. To maintain a robust space in which collective memory can flourish, regulators in the AI field must adopt a dual vantage point—directing their gaze simultaneously both toward the future and toward our future’s past. Collective memory deserves a seat at the table of AI regulation.

**Michal Shur-Ofry**

---

<sup>42</sup> Shur-Ofry & Pessach: *Robotic*, supra note 1. See also Jack Balkin’s “information fiduciary” framework, proposed with respect to social media providers—Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. Davis L. Rev. 1183 (2016).

<sup>43</sup> *Id.*, at 1003

<sup>44</sup> See AI Act, supra note 40. However, the question whether general-purpose LLMs are subject to the obligations applicable to “High-Risk” AI exceeds the scope of this analysis.

<sup>45</sup> Jacques Derrida, *Archive Fever: A Freudian Impression*, 25(2) *Diacritics*, 9 (1995), at pp. 10-11, fn. 10.