ISSN 2324-7878 (online)
palrap.org

PaLA Pennsylvania Library Association | College and Research Division (CRD)

Pennsylvania Libraries: Research & Practice

Commentary

Librarian Responsibility in the Face of Generative AI

George J. Aulisio

Ph.D., Dean of the Library, The University of Scranton

Generative AI represents a fundamentally different kind of technological disruption than those libraries have previously navigated. This commentary argues that librarians are uniquely positioned to be GenAI leaders: not as uncritical champions, nor as reflexive opponents, but as conscientious information professionals trained to ask hard questions about how knowledge is produced, who controls it, and what interests shape it. Drawing on the consistent pattern of society outpacing its capacity to govern new technologies, the article makes an ethical and professional case for library leadership in AI education and engagement. I also warn that inaction carries institutional risk.

Introduction

Libraries have never been static institutions, and librarians have continually met the challenges of a shifting information landscape. I will skip the first two millennia and ignore clay tablets and scrolls. In the last forty years, we have embraced electronic resources that have reshaped collections, launched new librarian roles, and wrestled with the dilemma of ownership versus license-based access. The rise of the internet demanded an entirely new vocabulary of information literacy, one centered not on *finding* information, but on *critically evaluating* it.[1] We all know these transitions were significant, and each time I assume there was some debate about whether to quickly adopt the next big thing or to wait and see.

Generative AI (i.e., GenAI) is the next big thing, and I think it may be unavoidable. It is in the news, it has its own television commercials, and it is being integrated continuously into the platforms we use. More to the point, it has infiltrated the research process, how people write, how they make decisions, how they learn, how they create, and so on.[2]

I want to back up for a moment to provide more context on GenAI. The fundamental difference between GenAI and previous technological disruptions (e.g., Google web crawling, social media platforms, etc.) is that those technologies changed *how* people accessed information, whereas GenAI ostensibly takes a more active role in information creation. Truth be told, *creation* is not the right word; the better word is *generation*. What GenAI is doing is synthesizing ideas from its training data and presenting those ideas to the user as something new.[3]

Another important clarification about GenAI is that it is not retrieving knowledge nor information from a database. The danger here is that it repackages the ideas of others and presents them as original and authoritative.[4] But information professionals know that GenAI is not reliably authoritative. Many of you are thinking I am referring to hallucinations, which I am in part, but not entirely. There is another authority problem. When AI generates something completely accurate, which it often does, it should still not be viewed as authoritative, because it is not an authority. It is not sharing its own ideas; it is not pulling information from a database; it is generating results drawn from its training data, training data composed of others' ideas. That makes its apparent accuracy as dangerous as its inaccuracy for individuals who do not realize these details. Consider, for example, the student who uses AI to generate all or parts of a required paper. The generated idea might be entirely accurate, but it is someone else's idea, and the student might be accused of plagiarism if they don't understand information literacy and the point of the whole endeavor.

There are many broader concerns about AI worth naming, including its environmental impact,[5] the concentration of wealth and power it enables,[6] its infringements on intellectual property,[7] and the potential decimation of entire sectors of the workforce.[8] The sad truth is that individual librarians cannot meaningfully affect those things directly. What we can do is educate people about those issues and help them make informed decisions. That is no small thing; that important work matters, but real regulatory change requires legislation. Short of regulatory changes, I believe librarians have a professional ethical responsibility to serve as AI leaders. Not necessarily AI champions, but that is fine, too, for those of us who champion it.

A Familiar Pattern

What we have witnessed over the last forty years is a consistent pattern where society changes faster than policy and ethical awareness. The tools we build outpace our capacity to govern them, and the consequences affect us all.

GenAI carries similar if not greater and more pervasive potential. What is even more concerning is that this technology has had almost no runway. Sure, the internet moved quickly in its early days, but it was not immediately ubiquitous. It took time to go from dial-up to DSL to broadband to 5G WiFi in every cafe. Social media has also evolved over the decades. AI seems to be evolving every few months, and it moves more like a train without brakes than like a technology seeking early adopters. The best evidence of this is probably that each software update seems to include an integrated GenAI component. If it is not already in your email, web browser, or favorite library database, it will be soon. The question is not whether these tools will reshape society. The question is whether anyone will be positioned to step back and reflect on the potential consequences in time to still make a difference.

What concerns me most is not the obvious risks of misinformation, job displacement, and further environmental degradation, though those are all very real. My rationale was already stated: I can not personally affect those things directly. My real concern, the one I think I can meaningfully affect, is actually quieter and more insidious. It is the risk of ceding our capacity for critical thinking to a system that will do our thinking for us if we let it. The brain is not really a muscle, but like a muscle, it too can atrophy.[9] If GenAI absorbs not just the mundane tasks of our professional and intellectual lives but the engaged tasks where we wrestle with difficult ideas, sit with uncertainty, and take on the hard work of forming an original thought, then we are trading ease and speed for something that will be hard, if not impossible, to reclaim. The promise of AI has been that it will free us from drudgery to make more time for dedicated higher-order thinking. What I have seen so far seems to suggest otherwise.

The answer, I think, is not abstention but conscientiousness.

The Case for Librarians

Given those realities, I would rather see librarians take on the AI leadership role than corporate interests. I would also rather see us in that role than those in industries that, by and large, are composed of conscientious objectors (e.g., writers, graphic designers, etc.). The reason is simple: AI is here, and it seems increasingly unlikely to go away. Many of the people who come into our libraries are already using it and will use it throughout their careers and in their personal lives. We should take on a leadership role so that our thoughtful, tech-forward, and information-literate approach can be shared with others.

Consider the landscape of institutional actors responding to GenAI on a typical university campus. Business schools are, by disposition and mission, oriented toward fast adoption. They want to leverage this tool because it presents itself as a competitive advantage, and employers are explicit that they are integrating it into their businesses; therefore, business school graduates should know how to work with GenAI. In contrast, humanities programs often approach technological disruption with more skepticism. Sometimes that skepticism turns into outright opposition that leads to disengagement. I do not think either posture serves students, or really society at large, all that well.

All of which brings me to what is, at its core, an argument based on ethical and professional responsibility. I am calling on librarians to be GenAI leaders not because it is convenient or has strong potential to put us in the spotlight. I am making the call because what is needed right now is something that was absent when social media launched; namely, a groundswell of conscientious information professionals willing to ask hard questions and share honest perspectives.[10] This is markedly distinct from the uncritical enthusiasm of many business leaders, and from the reflexive opposition of those who see only threat. There is a balanced, thoughtful, and principled approach available to us, and I believe librarians are the ones for the job.

Consider the library's structural position within society. Libraries are the great equalizer of the information landscape. We are already embedded in information literacy work, reference and research consultations, curriculum development, student success initiatives, lifelong learning, sustainable practices, and civic engagement.[11] We are, in short, already doing the work that genuine fluency in AI demands of us.

But it is not just what we do, it is how we are trained. Librarians are perhaps the only profession formally educated in how knowledge is produced, who controls it, how it circulates, and what interests shape it. That is an incredibly important critical apparatus to have in place in a world where we are witnessing the runaway adoption of GenAI firsthand.

The Risk of Standing Still

This brings us to the risk of inaction, which is real and underappreciated. If libraries do not take an active, visible role in AI education and ethical engagement, that vacuum will not go unfilled. In higher education, Information Technology divisions, Centers for Teaching Excellence, and Student Success offices are exploring ways to expand their AI portfolios. If the library's voice is absent from that conversation, we face a compounding risk. We would be ignoring a tool that, at its core, is an information tool. Ceding ground on an information resource that is moving faster than any we have previously encountered is professionally dangerous. The clearest example is Google—the tool once deemed to be the library's biggest threat—which appears to be actively trying to replace itself with AI-generated summaries.[12] If we cede this moment, we risk losing the core of our literacy mission.

The moment we are in is unprecedented in its speed and stakes. It calls for the particular combination of innovation and ethical seriousness that defines the library profession at its best. We have always helped people navigate information they did not fully understand, in systems that were not always designed with their interests in mind. In this moment, inaction is a choice, and it is not consequence-free.

The tools and systems will always change; the mission will not.

Notes

[1] Association of College & Research Libraries, 2015

[2] I grant that it's not used by and affecting everyone, but there is evidence of increasing adoption rates. See, for example, Kemp, 2026; Makridis, 2025; and Marken, 2026.

[3] For a more in-depth introductory explanation, see Martineau, 2023.

[4] For instance, Appel, Neelbauer, and Schweidel (2023) write, "[…] it may seem like these new AI tools can conjure new material from the ether, that's not quite the case. Generative AI platforms are trained on data lakes and question snippets—billions of parameters that are constructed by software processing huge archives of images and text. The AI platforms recover patterns and relationships, which they then use to create rules, and then make judgments and predictions, when responding to a prompt."

[5] See, for example, Zewe, 2025.

[6] See Frank, 2025.

[7] Appel, Neelbauer, and Schweidel, 2023.

[8] Richardson, 2025.

[9] See, Kosmyna et al., 2025. In their preprint that has not yet undergone peer review, they note, "[a]s the educational impact of LLM use only begins to settle with the general population, in this study we demonstrate the pressing matter of a likely decrease in learning skills based on the results of our study."

[10] My colleagues Donna Witek and Teresa Grettano were early to the challenge of Facebook. See Witek and Grettano, 2012 and Witek and Grettano, 2014. Another one of my colleague's Sheli Pratt-McHugh is working on the challenge of TikTok. See Pratt-McHugh, 2025. I was ignorant of the problem of Facebook at the time but see it now, and I am aware of the TikTok challenge but feel too disconnected from it.

[11] Another relevant example is the PA Forward Initiative. PA Forward is organized around five key literacies: Basic, Information, Civic and Social, Health, and Financial. In my opinion, all of them are relevant to the point I'm trying to make about wide scale societal adoption of AI and librarian responsibility to address it. See Pennsylvania Library Association, 2025.

[12] See Reid, 2024.

References

Appel, G., Neelbauer, J., & Schweidel, D. A. (2023, April 7). Generative AI Has an Intellectual Property Problem. *Harvard Business Review*. https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem.

Association of College & Research Libraries. (2015). *Appendix 2: Background of the Framework Development* (pp. 15–16). https://www.ala.org/sites/default/files/acrl/content/issues/infolit/Framework_ILHE.pdf

Frank, R. (2025, August 10). *AI is creating new billionaires at a record pace*. CNBC. https://www.cnbc.com/2025/08/10/ai-artificial-intelligence-billionaires-wealth.html

Kemp, A. (2026, January 25). *Frequent Use of AI in the Workplace Continued to Rise in Q4*. Gallup.com; Gallup. https://www.gallup.com/workplace/701195/frequent-workplace-continued-rise.aspx

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025, June 10). *Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task*. ArXiv.org. https://arxiv.org/abs/2506.08872

Makridis, C. (2025, May 15). *Play the Long Game With Human-AI Collaboration*. Gallup.com; Gallup. https://www.gallup.com/workplace/660572/play-long-game-human-ai-collaboration.aspx

Marken, S. (2026, April 2). *AI Is Routine for College Students, Despite Campus Limits*. Gallup.com; Gallup. https://news.gallup.com/poll/704090/routine-college-students-despite-campus-limits.aspx

Martineau, K. (2023, April 20). *What is generative AI?* IBM. https://research.ibm.com/blog/what-is-generative-AI

Pennsylvania Library Association. (2025, November 3). *PA Forward - Pennsylvania Libraries*. Paforward.org. https://paforward.org/

Pratt-McHugh, S. (2025, October 10). Information Literacy for Real Life: Debunking TikTok Videos with Library Resources. *Pennsylvania Library Association Annual Conference*.

Reid, E. (2024, May 14). *Generative AI in Search: Let Google do the searching for you*. Google. https://blog.google/products-and-platforms/products/search/generative-ai-google-search-may-2024/

Richardson, N. (2025, August 26). *Yes, AI is affecting employment. Here's the data.* ADP Research. https://www.adpresearch.com/yes-ai-is-affecting-employment-heres-the-data/

Witek, D., & Grettano, T. (2012). Information literacy on Facebook: an analysis. *Reference Services Review*, *40*(2), 242–257. https://doi.org/10.1108/00907321211228309

Witek, D., & Grettano, T. (2014). Teaching metaliteracy: a new paradigm in action. *Reference Services Review*, *42*(2), 188–208. https://doi.org/10.1108/rsr-07-2013-0035

Zewe, A. (2025, January 17). *Explained: Generative AI's environmental impact*. MIT News; Massachusetts Institute of Technology. https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117