



Research Article

Deepfakes and Information Literacy: Implications for the Emerging Role of Academic Librarians

Susan Nnadozie Umeozor* 

Department of Library and Information Science, University of Port Harcourt, Port Harcourt, Nigeria

Abstract

The emergence of deepfakes challenges the traditional assumptions underlying information literacy because authenticity can no longer be established through conventional evaluation of source, appearance and audiovisual evidence alone. Consequently, academic librarians must expand information-literacy instruction towards synthetic-media literacy, verification, provenance assessment and critical evaluation of AI-generated content. The paper addressed this issue by examining the challenges deepfakes pose to conventional information-literacy practice and the implications for the emerging role of academic librarians. It discussed how deepfakes undermine traditional authority cues, complicate source corroboration, make authentic information difficult to distinguish from synthetic content, and expose limitations in existing information-evaluation practices of users. The paper further considered the evolution of information-literacy practice from source-centered approaches of CRAAP test towards SIFT and lateral reading, and established that these approaches, although valuable, may not sufficiently address the authenticity problem created by synthetic media. It proposed an expanded information-literacy practice incorporating synthetic-media awareness, provenance tracing, verification, contextual analysis, triangulation, AI literacy and critical judgement. The paper concluded that academic librarians should be repositioned as mediators of trust, verification and critical information judgement within contemporary academic information environments through curriculum integration, institutional partnerships, verification-tool training and continuous professional development in AI and synthetic-media literacy.

Keywords

Librarianship, Academic Library, Artificial Intelligence, Deepfake, Information Literacy

1. Introduction

Information has traditionally occupied a central position in academia, serving as the raw material through which teaching, learning and research are conducted. Academic libraries have long shouldered the responsibility of mediating this relationship, helping users discover, evaluate and use information in ways that support scholarly inquiry and informed decision-making. However, successive waves of technological devel-

opment have progressively reshaped the information environment within which this mediating role is performed, moving from print and conventional digital repositories toward content that is algorithmically generated, aggregated and increasingly manipulated. Indeed, the internet has made it quicker and easier to access, create and manipulate information, making the skills of evaluating information more important than ever [1]. This shift indicates that the challenge facing library

*Correspondence: Susan Nnadozie Umeozor (Susan.umeozor@uniport.edu.ng)

Received: 14 September 2026; Accepted: 27 September 2026; Published: 20 October 2026



Copyright: © The Author(s), 2026. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

users is no longer simply one of information scarcity or restricted access, but rather the far more complex task of determining what information can be trusted.

This transformation has been accelerated by the emergence of generative artificial intelligence, a subset of machine-learning systems capable of producing increasingly realistic text, images, audio and video. Generative AI algorithms learn the patterns and characteristics of a dataset and use this knowledge to generate synthetic content resembling the original data [14], a capability that has given rise to deepfakes as a particularly consequential manifestation of synthetic media. The term deepfake commonly refers to synthetic or heavily manipulated media, most notably images, videos and speech, whose realism is enabled by modern machine learning, particularly deep generative models [9]. Fabricated representations of academics, public figures or institutions, produced through such techniques, demonstrate, with unsettling clarity, how easily apparently authentic audiovisual resources can now be manufactured.

Traditional information literacy has outlined a defined set of competencies including identifying information needs, searching effectively, evaluating sources, determining authority, corroborating information across sources and using information ethically. Central to this tradition is the understanding that information resources reflect the expertise and credibility of the originator and are evaluated based on the information need and the context in which the information will be used, with authority itself being constructed and contextual [1]. Deepfakes complicate these established practices precisely because the apparent source, voice, image or audiovisual resource being evaluated can be artificially constructed. A user may therefore apply every conventional evaluative skill to a seemingly credible source and still encounter entirely fabricated content. This reality establishes why existing information-literacy approaches, built for an environment in which authenticity could largely be assumed, now require fundamental reconsideration.

This problem is particularly challenging in academic librarianship. Academic librarians professionally play a central role in information-literacy instruction, research support and helping users evaluate the credibility of information. One thing is certain: the professional judgement of academic librarians cannot be relegated in this new environment. If the information ecosystem now includes increasingly sophisticated synthetic media, librarians cannot continue to restrict information-literacy instruction to conventional source evaluation alone. The everyday encounters with AI-generated media already show that navigating deepfakes draws on media and information literacy practices that extend well beyond standard source-evaluation skills [13], pointing to the need for competencies involving synthetic-media awareness, verification, provenance tracing, triangulation, contextual analysis, AI literacy and critical judgement. This paper argues that librarians may not be professionals in technical detection of deepfake, rather, their role will expand from teaching users how to find

and evaluate information toward helping them understand how information is produced, manipulated, authenticated and contextualized within an AI-mediated environment. Hence, this paper examines the challenges that deepfakes pose to information literacy and argues for the repositioning of academic librarians as mediators of trust, verification and critical information judgement.

2. Literature Review

Information literacy (IL) has long been a foundational concern of academic librarianship, evolving from identified competencies on the understanding of how individuals engage with information. The Information Literacy Competency Standards outlined six interrelated frameworks for understanding information literacy, each pairing a core concept with associated knowledge practices and dispositions [13]. Two of these frames are especially important in the present discussion. The first, *Authority Is Constructed and Contextual*, holds that information resources reflect the expertise and credibility of their creator and are evaluated based on the information need and the context in which the information will be used, with different communities recognizing different types of authority [1]. The second, *Information Creation as a Process*, reframes information not as a static, verifiable object but as the product of an iterative process of creation, revision and dissemination. Both frames implicitly assume that authorship and provenance of a resource can, in principle, be traced or verified.

Complementing these conceptual frames, Mullins and Boyd-Byrnes [15] affirmed that the distinct pedagogical value librarians bring to IL instruction. The authors found that when librarians taught IL modules embedded in first-year seminars, students met learning objectives more consistently than when the same modules were taught by campus faculty, substantiating librarians as uniquely positioned to deliver this instruction. This finding underscores that IL instruction is not merely a curricular add-on but a specialized pedagogical function. This provides justification that librarians, rather than being displaced by an increasingly synthetic information environment, are positioned to expand their instructional remit within it.

2.1. Generative AI and Synthetic Media: State of the Art

The rapid advancement of generative artificial intelligence has enabled the creation of highly realistic synthetic media, commonly referred to as deepfakes, which are increasingly multimodal and difficult to detect [9]. Deepfake generation techniques have evolved considerably since their emergence, moving from early Generative Adversarial Network (GAN) architectures toward diffusion models and neural rendering systems capable of producing near-photorealistic images, video and speech [9]. Generative AI algorithms, broadly defined, learn the patterns and characteristics of a dataset and use this knowledge to generate synthetic content resembling the

original data [14]. Deepfake outputs specifically resemble humans and their actions, spanning image, video and speech modalities, and may be either targeted (resembling a known identity) or untargeted (a wholly fabricated identity) [14]. The accessibility of these tools has expanded considerably as commercial and open-source generative systems have proliferated, a trend that has direct implications for academic contexts. Higher-education stakeholders have begun to reckon with this shift. Eberl *et al.*, [8] found substantial ambivalence about the use of deepfake and synthetic-media technologies in teaching and learning settings, with participants identifying both instructional opportunities and significant risks to trust and integrity. This ambivalence reflects the broader tension envisioned in this paper, that synthetic media is simultaneously a pedagogical resource and an epistemic threat yet to be fully grasped by academic institutions and libraries.

Mai *et al.*, [14] demonstrated that human ability to detect synthetic media is unreliable, even under conditions favorable for detection. The authors carried out a large-scale experimental study, listeners correctly identified speech deepfakes only 73% of the time, and providing participants with prior examples of deepfakes to increase their awareness improved detection only marginally. Notably, the same study found that participants were more likely to misclassify authentic audio as fake than to miss an actual deepfake, suggesting that increased vigilance produces general skepticism rather than precise discrimination. Automated detection systems achieved near-perfect accuracy when tested on data like their training conditions, but their performance deteriorated sharply when confronted with unfamiliar speakers or recording conditions, limiting their reliability as a stand-alone safeguard [14]. Similar findings have emerged in visual-deepfake detection research more broadly, where human participants across multiple studies performed only marginally better than chance and frequently expressed unwarranted confidence in incorrect judgments [13]. Hence, the roles of librarians should center on provenance and verification practices rather than perceptual detection skill.

Qualitative research on adolescents' everyday encounters with AI-generated media found that young people's engagement with deepfakes occurs largely through serendipitous, informal exposure across everyday platforms, and that their sense-making practices draw on media and information literacy skills that extend beyond formally taught evaluative heuristics [13]. This finding suggests that existing literacy instruction has not caught up with the informal, everyday ways that users actually encounter synthetic content.

Within library and information science (LIS) specifically, a systematic review examining deepfake technology and the role of LIS professionals concluded that deepfakes pose significant challenges for the profession, requiring practitioners to develop advanced skills for evaluating digital content and to take an active role in promoting digital literacy among their user communities [6]. This review represents one of the few pieces of LIS scholarship to engage directly with deepfakes as a professional and pedagogical concern, and its conclusions

align closely with the position developed in this paper. However, its treatment remains largely generic.

The reviewed literature establishes three things. First, IL frameworks such as ACRL [1] remain built on assumptions of traceable authority and verifiable provenance — assumptions that generative AI increasingly unsettles. Second, empirical evidence on human and automated detection capability shows that neither users nor current technical tools can reliably distinguish synthetic from authentic media, undermining detection-centered responses to the problem [14]. Third, while LIS scholarship has begun to acknowledge deepfakes as a professional concern [6], this engagement remains general in scope, and has not been developed into a sustained argument for how academic libraries specifically should reposition their information-literacy instruction and their librarians' professional role in response. The emergence of deepfakes represents not merely a new form of misinformation but a challenge to the epistemic assumptions underlying information literacy, because content that appears authentic can now be artificially produced or manipulated. Consequently, academic librarians must move beyond conventional information-literacy instruction towards a broader role encompassing synthetic-media literacy, verification, provenance assessment and critical evaluation of AI-mediated information.

2.1.1. The Deepfake Challenge to Traditional Library Information Literacy Practice

Central to the position of this paper is the meddlesomeness of artificial intelligence in the traditional role of librarians. Deepfake has compounded a critical aspect of information use which is authenticity of information resources. The under-listed are the challenges deepfake poses to traditional library information literacy practice.

2.1.2. Breakdown of Traditional Authority Cues

The *Authority Is Constructed and Contextual* frame of the ACRL framework trains users to read authority from surface cues such as a face on camera, a recognizable voice, an institutional logo, a byline [1]. This heuristic has always required some contextual judgment, but it rested on one basic assumption, that the cue itself was produced by the entity it appeared to represent. Generative systems now dissolve that assumption. These systems learn the patterns of a dataset adequately to generate content that convincingly resembles it [14], meaning that the same visual and auditory signals users are trained to trust. The face of a professor, the voice of a news anchor, the letterhead of a university, among others, can be produced by anyone with access to commodity tools, independent of any actual affiliation with the person or institution being represented. The cue and the credential have come apart. Hence, Authority markers are no longer reliable signals, a user applying every ACRL-sanctioned evaluative skill can still arrive at a confident, well-reasoned judgment about a source that does not exist. What is being evaluated, in that case, is not authority

at all, but a simulation of its outward form.

2.1.3. Challenge of Source Corroboration

Cross-checking multiple sources is standard IL practice precisely because independent convergence is treated as evidence of reliability: if several unconnected sources say the same thing, the claim is assumed more trustworthy than any single source could establish alone. Synthetic media breaks the independent assumption this logic depends on. Research on coordinated inauthentic behavior shows that manipulative actors increasingly operate across, not just within, individual platforms, constructing networks of accounts that share and amplify the same fabricated content in ways designed to look like organic, independent agreement [4]. What looks to a user like several unrelated sources converging on the same story may in fact be a single fabricated origin, distributed through coordinated accounts precisely to manufacture the appearance of corroboration [16]. A student or researcher performing corroboration correctly, checking several sources, and looking for agreement can still be corroborating a single fabrication wearing several faces.

2.1.4. The Difficulty of Distinguishing Authentic from Synthetic Information

The difficulty of distinguishing authentic information from synthetic content represents a major challenge to information-literacy practice because deepfakes can closely reproduce the visual, auditory and contextual features of genuine information. This makes it increasingly difficult for users to determine authenticity through observation and conventional evaluation practices, even when they are confident in their judgement. iProov [12] found that only 0.1% of the participants were able to correctly identify all the real and fake items presented to them, although 60% of the participants expressed confidence in their ability to identify deepfakes. This indicates that confidence in identifying manipulated content may not necessarily correspond with actual ability to distinguish authentic content from deepfakes. Diel *et al.* [7], in a systematic review and meta-analysis of 56 studies involving 86,155 participants, found that the overall accuracy of humans in detecting deepfakes was 55.54%, with the 95% confidence interval crossing the 50% chance level. This demonstrates the difficulty individuals experience in reliably distinguishing authentic content from manipulated media. The challenge is also evident in the identification of audio deepfakes. Mai *et al.* [14] found that participants correctly identified speech deepfakes in only 73% of cases, while prior exposure to examples of deepfake speech produced only a marginal improvement in their ability to detect them. The study further found that participants were more likely to classify genuine audio as fake than to fail to identify an actual deepfake [14]. These findings suggest that increased awareness of deepfakes does not necessarily result in accurate discrimination between authentic and

manipulated information. The challenge is further complicated by the limitations of automated detection tools, as Chandra *et al.* [3] found that detection systems that perform strongly under controlled benchmark conditions may experience substantial reductions in accuracy when applied to deepfakes encountered in real-world settings. Taken together, these findings demonstrate that neither users' confidence and experience nor existing technological detection tools provide a sufficiently reliable basis for distinguishing authentic information from deepfakes, thereby reinforcing the need for stronger information-literacy competencies for evaluating synthetic and manipulated content.

2.1.5. The Problem of Prior Informal Exposure

Information literacy instruction typically assumes that users are being introduced to a structured process of identifying, evaluating and using information. However, this assumption becomes problematic in the context of synthetic media, because individuals are often exposed to deepfakes through routine and unplanned encounters on social media platforms. Their interpretation of such content is, therefore, likely to develop through informal experiences, platform practices and interactions with peers rather than through formally taught information-evaluation principles [13]. This has particular implications for academic libraries because users who participate in information-literacy instruction may already have developed personal ways of judging the authenticity and credibility of synthetic content before receiving formal instruction. Such practices may not necessarily be systematic, reliable or consistent with established principles of information evaluation. Consequently, information-literacy instruction on deepfakes should not assume that learners are approaching the problem without prior knowledge or experience. Rather, academic librarians need to recognize existing informal practices and build upon them by introducing structured approaches to verification, source evaluation, contextual analysis and the responsible assessment of synthetic media.

Traditional information literacy practice was developed largely to evaluate credibility after the identity and authenticity of the information source have been reasonably established. Deepfakes challenge this assumption by making the authenticity of the content itself uncertain before credibility can be meaningfully assessed. This challenge is intensified by the scale of synthetic media, the gap between users' confidence and actual detection ability, and their prior informal exposure to deepfakes through digital platforms. Consequently, an information-literacy framework that remains focused only on source-level credibility, however rigorously applied, may be inadequate for addressing the distinctive problem posed by deepfakes. Emerging technologies often introduce changes whose implications are not immediately foreseeable, and their eventual impact may extend beyond what was initially anticipated. In such circumstances, the effects of technological advancement cannot necessarily be stopped; rather, they become

developments that must be encountered and carefully managed within the existing information environment. Consequently, technological changes may present both opportunities for improving existing library and information services and challenges whose negative consequences may need to be carefully avoided.

A similar situation was experienced during the COVID-19 pandemic. The emergence and rapid spread of the virus caught the world largely unprepared, as the nature and extent of the threat could not have been adequately anticipated. Although the existence and spread of the virus could not be immediately stopped, measures were introduced to reduce the possibility of contracting and transmitting it. The use of face masks, hand gloves, hand hygiene and social distancing became measures adopted to reduce exposure to the virus and limit its consequences. The response to the pandemic therefore illustrates how, when an emerging threat cannot be completely stopped, measures can be adopted to carefully avoid or reduce its harmful effects. Such circumstance can be observed in the emergence of deepfake technology and the increasing intricacies of synthetic media. The development and sophistication of synthetic media may not be completely curtailed, particularly as technological advancement continues to produce increasingly complex forms of artificially generated content. This, therefore, raises the need for information-literacy practice to move beyond credibility assessment towards provenance, authenticity and verification as essential starting points for evaluating synthetic media. The role of academic librarians, therefore, is not to out-detect deepfakes, but to help users develop the competencies required to establish where information originated, how it has been produced or modified, and whether it can reasonably be accepted as authentic.

2.2. The New Information Literacy: Implications for Academic Libraries and Librarians

The library as a growing organism, has continually witnessed phases of transformation in response to changes in the information environment and the needs of its users. This transformation is evident in the movement from a philosophy centered on content to user-oriented philosophy, and from the librarian as an intermediary in information retrieval and use to the emerging role of the librarian as an apo-mediator who guides users while allowing them greater autonomy in engaging with information. The emergence of deepfake information presents yet another situation that calls upon the library to adapt its information-literacy practice, because the challenge is no longer only how users retrieve and evaluate information, but how they establish the authenticity and provenance of information in an environment where content can be convincingly fabricated.

Library information-literacy instruction has continually adapted to changes in the information environment, moving

from source-centered approaches such as the CRAAP test towards approaches such as SIFT and lateral reading, which encourage users to leave the source under review and examine what other independent sources say about it. Brodsky *et al.*, [2] opined that the value of lateral reading is improving users' ability to evaluate the credibility of online information. However, the emergence of deepfake information presents a further challenge because lateral reading is primarily directed towards evaluating claims and sources, whereas synthetic media introduces the additional question of whether the content itself is authentic. A user may, therefore, verify a claim across several sources and still fail to establish whether the image, audio or video presented as evidence has been manipulated. This suggests that the evolution from CRAAP to SIFT and lateral reading, although important, may not be sufficient for the deepfake information environment and that information-literacy practice requires further adaptation towards media authenticity, provenance and verification. Hence, the librarians need to improvise expanded competencies. Several competencies for deepfake-IL may include:

- 1) *Synthetic-media awareness* – synthetic media awareness is the baseline competency on which every other one depends. It refers to an explicit understanding, taught rather than absorbed incidentally, that video, audio and images are no longer reliable evidence of authenticity by default. This is not an intuitive lesson, because it runs against decades of accumulated user habit in which seeing or hearing something functioned as its own form of proof. That habit is precisely what makes user overconfidence so persistent: people continue to trust their own perceptual judgment even where the evidence shows that judgment cannot reliably distinguish authentic from synthetic content. A large-scale controlled study on deepfake detection found that the vast majority of participants reported confidence in their ability to identify fabricated content, regardless of whether their actual judgments were correct [12]. Synthetic-media awareness, as a discrete instructional objective, exists to close that gap not by teaching users to detect fakes more accurately, but by teaching them to distrust their own confidence in the first place, which is a more modest and more achievable goal.
- 2) *Verification and provenance tracing* – this shifts the operative question in source evaluation from *who published this* to *can the source of the information be independently established*. This is a meaningfully different task from conventional fact-checking, because it does not ask whether a claim is plausible or corroborated but whether the media object itself can be traced back to a verifiable point of creation. Emerging technical infrastructure is beginning to make this possible. The Coalition for Content Provenance and Authenticity (C2PA) standard attaches cryptographically signed metadata to a media file recording who created it, when, and whether it has since been altered, giving users a mechanism to

check provenance that does not depend on visual or auditory judgment at all [5]. This is important because even trained instinct performs close to chance when distinguishing real from synthetic content. At the same time, this competency must be taught with caution because provenance standards such as C2PA are not a guarantee of authenticity in themselves, since their specifications only claim to detect tampering with the credential and the protected bits of the file, not to certify truth [11]. Teaching provenance tracing, then, means teaching a genuinely more reliable class of tool, while being equally clear about what that tool does and does not.

- 3) *Triangulation resistant to coordination* – this should update the familiar corroboration habit of checking whether multiple sources agree, for an information environment in which apparent agreement can itself be manufactured. Generative tools have made it considerably cheaper to produce and distribute messages designed to look as though they originate from many independent individuals, creating the appearance of grassroots consensus around a claim that in fact traces back to a single coordinated source [10]. This is the core mechanism of coordinated information manipulation, in which multiple connected sources are made to appear independent, creating an artificial impression of agreement that conventional triangulation is not designed to detect. A user who dutifully checks three, five or ten sources and finds them in agreement has, under these conditions, no guarantee of having found anything more than one well-distributed fabrication. Teaching triangulation resistant to coordination means teaching users to ask not only whether sources agree, but whether their apparent independence can itself be verified.
- 4) *Contextual and structural analysis* – this moves the evaluative task beyond the artefact itself and toward the conditions that produced and circulated it. This means asking not just whether a claim is corroborated but why a given piece of synthetic or manipulated content exists in the first place, who stands to benefit from its circulation, and what platform dynamics and incentive structures shaped its spread. This kind of analysis has long been central to more advanced information-literacy scholarship, which has argued that students need contextual and structural awareness not just source-checking mechanics, to reach ethical, higher-order judgments about information [7]. Bringing this competency explicitly into library IL instruction reframes the librarian's teaching task rather than asking *is this true* only, students are taught to also ask *why does this exist*, and *why is it reaching me now*.
- 5) *AI literacy and critical judgement* – It refers to a working understanding of how generative systems produce convincing synthetic content not the technical detail of model architectures, but enough grounding to reason about a piece of plausibility of a content, likely origin

and likely purpose. LIS scholars have already begun to formalize what this looks like specifically for librarians' instruction, proposing an AI-literacy framework built explicitly from the perspective of librarians already teaching source evaluation, rather than importing a generic AI-literacy model from computer science or data science [4]. AI literacy, in a library IL context, is not a separate subject bolted onto existing instruction, but a natural extension of the critical-evaluation habits being taught by librarians, applied to a new category of content that did not previously need to be accounted for.

Academic librarians have long been positioned as evaluators-of-evaluation, professionals that teach others how to arrive at that information judgment responsibly. Das and Das [6] concluded that professionals in the field must actively develop the skills to evaluate synthetic digital content and to promote digital literacy around it. This posits that librarians already occupy the structural and pedagogical position needed to take this on, and the professional judgement built into that position cannot be delegated to a browser extension or a detection algorithm, particularly given how unreliable current detection tools remain outside their training conditions [3].

What changes is the scope of the teaching, not the nature of the role. The librarian's task has never been to personally verify every claim a user encounters, rather, it has been to help users develop the capacity to do so. If the information environment now includes increasingly sophisticated synthetic media, that capacity-building simply has to extend one layer deeper from *how do I evaluate this source* to *how do I establish that this is a source at all*. This paper does not argue that librarians should become technical deepfake detectors, a role that may not be reliably achievable by anyone, expert or otherwise. It argues, instead, that the librarian's role expands from teaching users how to find and evaluate information toward helping them understand how information is produced, manipulated, authenticated and contextualized in an AI-mediated environment. That is a repositioning of emphasis within an existing professional identity, not a demand for a new one.

3. Conclusion

This paper has established that deepfakes present information literacy with a challenge that goes beyond the conventional evaluation of sources. Academic library information-literacy practice has evolved with changes in the information environment, from source-centred approaches to methods such as SIFT and lateral reading that encourage users to verify information across sources. However, deepfakes introduce a further problem because the information being evaluated may itself be artificially produced or manipulated. A user may identify an apparently authoritative source, corroborate its claims across several platforms and still fail to establish whether the image, audio or video being presented is authentic. The difficulty of human detection, the limitations of technological detection tools, coordinated circulation of fabricated

content and users' prior informal exposure to synthetic media further demonstrate that conventional information-literacy competencies cannot sufficiently address the problem. This paper, therefore, concluded that librarians and libraries should not abandon existing information-literacy practice, but to extend it towards synthetic-media awareness, provenance tracing, verification, contextual and structural analysis, triangulation that considers coordination, AI literacy and critical judgement. This development is consistent with the nature of the academic library as a growing organism that continually adapts to changes in information, technology and user behaviour. The emerging role of the academic librarian is, consequently, not that of a technical deepfake detector, but that of a professional who helps users understand how information is produced, manipulated, authenticated and contextualized within an AI-mediated information environment.

For this repositioning to become meaningful in practice, academic libraries must translate these emerging competencies into their instructional programmes, institutional practices and professional development. At the curriculum level, synthetic-media literacy should be embedded within existing information-literacy sessions and courses so that users learn to question not only the credibility of information but also its authenticity, provenance and mode of production. At the institutional level, libraries should develop partnerships with journalism and media-literacy programmes and provide users with practical exposure to provenance and verification approaches, including reverse image and video searching and C2PA Content Credentials, while making clear that such tools support rather than replace human judgement. At the professional level, librarians require continuous development in generative AI, synthetic media, verification practices and responsible AI use to remain competent guides within the changing information environment. It is when librarians and academic libraries achieve these curriculum-level, institutional-level and professional-development transformations that the profession can be repositioned to respond effectively to the challenges of artificial intelligence and its exigencies, while preserving its enduring responsibility for trust, verification, critical judgement and scholarly integrity.

Abbreviations

AI	Artificial Intelligence
CRAAP	Currency, Relevance, Authority, Accuracy, and Purpose
SIFT	Stop, Investigate the Source, Find Better Coverage, and Trace Claims, Quotations, and Media to the Original Context
ACRL	Association of College and Research Libraries
IL	Information Literacy
GAN	Generative Adversarial Network
LIS	Library and Information Science
C2PA	Content Provenance and Authenticity

Author Contributions

Susan Nnadozie Umeozor: Conceptualization, Data curation

Conflicts of Interest

The author declared no funding or conflict of interest.

References

- [1] Association of College and Research Libraries (ACRL). (2015). *Framework for information literacy for higher education*. American Library Association. <https://www.ala.org/acrl/standards/ilframework>
- [2] Brodsky, J. E., Brooks, P. J., Scimeca, D., Todorova, R., Galati, P., Batson, M., Grosso, R., Matthews, M., Miller, V., & Caulfield, M. (2021). Improving college students' fact-checking strategies through lateral reading instruction in a general education civics course. *Cognitive Research: Principles and Implications*, 6(1), 23. <https://doi.org/10.1186/s41235-021-00291-4>
- [3] Chandra, N. A., Murtfeldt, R., Qiu, L., Karmakar, A., Lee, H., Tanumihardja, E., Farhat, K., Caffee, B., Paik, S., Lee, C., Choi, J., Kim, A., & Etzioni, O. (2025). *Deepfake-Eval-2024: A multi-modal in-the-wild benchmark of deepfakes circulated in 2024* (arXiv:2503.02857). arXiv. <https://arxiv.org/abs/2503.02857>
- [4] Cinus, F., Minici, M., Luceri, L., & Ferrara, E. (2025). Exposing cross-platform coordinated inauthentic activity in the run-up to the 2024 U. S. election. In *Proceedings of the ACM Web Conference 2025* (pp. 541–559). Association for Computing Machinery. <https://doi.org/10.1145/3696410.3714698>
- [5] Coalition for Content Provenance and Authenticity. (2024). *C2PA and content credentials explainer* (Version 2.4). <https://spec.c2pa.org/specifications/specifications/2.4/explainer/Explainer.html>
- [6] Das, S., & Das, M. (2025). Use of deepfake technology and the role of library and information science professionals. *College Libraries*, 40(3), 82–88. <https://collegelibraries.in/index.php/CL/article/view/217>
- [7] Diel, A., Lalgı, T., Schröter, I. C., MacDorman, K. F., Teufel, M., & Bäuerle, A. (2024). Human performance in detecting deepfakes: A systematic review and meta-analysis of 56 papers. *Computers in Human Behavior Reports*, 16, 100538. <https://doi.org/10.1016/j.chbr.2024.100538>
- [8] Eberl, A., Kühn, J., & Wolbring, T. (2022). Using deepfakes for experiments in the social sciences: A pilot study. *Frontiers in Sociology*, 7, Article 907199. <https://doi.org/10.3389/fsoc.2022.907199>
- [9] Gazis, A., Karypidis, E., Santamouri, K., Vavouras, T., Mastorakis, N. E., & Pappas, S. (2026). Deepfakes and synthetic media: Generation, detection, and governance. *Encyclopedia*, 6(8), 165. <https://doi.org/10.3390/encyclopedia6080165>

- [10] Hervieux, S., & Wheatley, A. (2024). *Building an AI literacy framework: Perspectives from instruction librarians and current information literacy tools* [White paper]. Choice / American Library Association.
- [11] Hervieux, S., & Wheatley, A. (2024). Comparing generative artificial intelligence tools to voice assistants using reference interactions. *The Journal of Academic Librarianship*, 50(5), 102942. <https://doi.org/10.1016/j.acalib.2024.102942>
- [12] iProov. (2025, February 11). *iProov study reveals deepfake blindspot: Only 0.1% of people can accurately detect AI-generated deepfakes* [Press release]. <https://www.iproov.com/press/study-reveals-deepfake-blind-spot-detect-ai-generated-content>
- [13] Lao, Y., Hirvonen, N., & Larsson, S. (2025). Everyday encounters with deepfakes: Young people's media and information literacy practices with AI-generated media. *Journal of Documentation*, 81(7), 216–235. <https://doi.org/10.1108/JD-01-2025-0007>
- [14] Mai, K. T., Bray, S., Davies, T., & Griffin, L. D. (2023). Warning: Humans cannot reliably detect speech deepfakes. *PLOS ONE*, 18(8), Article e0285333. <https://doi.org/10.1371/journal.pone.0285333>
- [15] Mullins, K., & Boyd-Byrnes, M. K. (2024). Academic librarians' contribution to information literacy instruction and learning. *College & Research Libraries*, 85(3). <https://crl.acrl.org/index.php/crl/article/view/26292>
- [16] Murero, M. (2023). Coordinated inauthentic behavior: An innovative manipulation tactic to amplify COVID-19 anti-vaccine communication outreach via social media. *Frontiers in Sociology*, 8, Article 1141416. <https://doi.org/10.3389/fsoc.2023.1141416>