

Hallucinated References in Student Academic Writing: Quantify AI Cheating via Post-Hoc Analysis

Markus O. Goldbach*, Johannes F. Knabe*

* R&D, Mentafy GmbH

Paper Publication Date: 26th January 2026

Abstract - The integration of generative AI tools in academic settings is reshaping student writing practices - and with it, the landscape of academic integrity. This study investigates the occurrence and implications of hallucinated references (i.e., fabricated or untraceable citations) in student submissions across German-speaking universities. Analyzing 368 undergraduate and graduate-level texts, we find that 36.4% contain at least one invalid reference, with an average false reference rate of 12.7% among affected texts. Drawing on interdisciplinary literature from 2023–2025, we interpret this pattern as a proxy for unchecked reliance on AI-generated outputs. As recent research shows, large language models like ChatGPT often produce plausible but fictitious citations, posing risks that extend from coursework to scholarly publishing. We conclude that hallucinated references are a measurable and actionable indicator of either academic misconduct or insufficient research literacy. To address this, institutions must combine reference auditing tools, updated academic integrity policies, and curriculum-integrated AI literacy training. The findings offer a practical, evidence-based roadmap for schools and universities seeking to maintain academic standards while navigating the challenges and opportunities of generative AI in education.

Index Terms - Academic Integrity, Generative AI in Education, Hallucinated Citations, Information Literacy

I. INTRODUCTION

Generative AI tools like ChatGPT are increasingly used in academic writing, raising new concerns for educational integrity. Surveys in 2024 found that over 86% of students globally use AI as part of their studies (with 66% using ChatGPT specifically) (Rong & Chun, 2024). While these tools can assist with drafting and editing, they also frequently “hallucinate” academic references - i.e. produce citations that look plausible but are incorrect or even entirely fabricated (Welborn, 2023) (Cole, Maher, & Rice, 2025). ChatGPT’s underlying language model is not connected to a real bibliographic database, so when asked to provide sources, it often “fills in the blanks” by generating fake article titles, authors, journals, and DOIs that do not exist (Welborn, 2023). If a student includes such AI-generated

references without verification, these errors can persist into final submissions, undermining the work’s credibility. Academic references are meant to anchor claims in verifiable sources; a false citation is thus an objective error that is straightforward to detect and indicative of insufficient research or misuse of AI. In contrast to an AI lectorate, subtle forms of plagiarism or ghostwriting, a hallucinated reference is a clear red flag: either the student did not actually obtain and read the source, or the source itself is fictitious. As one educational resource notes, “*hallucinated references*” have quickly become one of the most common issues in AI-assisted academic writing, highlighting the need for students to verify AI-proposed citations for accuracy (Cole, Maher, & Rice, 2025). This paper examines the prevalence of such hallucinated references in student work from German-speaking universities and situates the findings within the emerging body of literature on AI misuse in academia.

II. BACKGROUND AND RELATED WORK

Misuse of AI in academic writing has gained global attention since the public release of ChatGPT in late 2022. Researchers have documented that large language models often produce convincing but incorrect information when asked for scholarly support (Cole, Maher, & Rice, 2025). In particular, multiple studies in 2023 analyzed ChatGPT-generated academic text and found a high incidence of fabricated citations. A study from Manhattan University showed that when ChatGPT-3.5 was prompted to write literature reviews with references, 55% of the citations it provided were *fabricated* (i.e. the cited works could not be found), and even the more advanced ChatGPT-4 still made up 18% of its references (Walters & Wilder, 2023). Across six studies spanning fields from medicine to geography, typically 47–69% of ChatGPT’s cited sources were non-existent, averaging 51% fake citations overall. For example, in one trial on clinical radiology questions, ChatGPT-3.5 produced 343 references of which 64% were not traceable to any real publication. These startling figures demonstrate that ChatGPT (especially earlier versions) frequently outputs bibliographies filled with what one article calls “*ghost bibliographic references*”.

Importantly, even when the AI cites an *actual* publication, it often perpetuates errors in the reference details. Up to 87% of ChatGPT's "real" citations contained at least one substantive error - such as incorrect author names, article titles, journal name, volume, or page numbers (Gravel, D'Amours-Gravel, & Osmanliu, 2023). Another analysis similarly reported pervasive inaccuracies in the metadata of ChatGPT-generated references (Walters & Wilder, 2023). These errors can make it difficult or impossible for a reader to locate the source, effectively defeating the purpose of citation. The pattern arises because large language models prioritize plausible language over factual verification: the AI will confidently assemble a reference that looks legitimate in format, even if the content is fabricated. As a Duke University library blog cautioned early on, "*If you try to find these 'AI-provided' sources through Google or the library - you will turn up NOTHING*" (Welborn, 2023). Students and researchers who are unaware of this limitation might unwittingly accept such outputs at face value.

Beyond experimental settings, concerns have arisen that fabricated citations are leaking into real academic outputs. A 2023 letter in *Scientometrics* warned that some research papers indexed online contained references to non-existent articles, presumably due to authors relying on ChatGPT without fact-checking (Orduña-Malea & Cabezas-Clavijo, 2023). The letter coined the term "*ghost references*" for these fictitious citations and urged journals and publishers to be vigilant, lest fake references start "flooding their pages." At the time, the true scope of the problem was unknown - it was speculated to be mostly confined to preprints, predatory journals, or student theses - but the potential threat to scholarly integrity was deemed serious enough to sound an early alarm. By 2024–2025, evidence of this phenomenon had concretized. Camp et al. (2025) uncovered *eight published papers* (in lower-tier journals) that all cited a common pool of counterfeit references. In each case, the reference list included articles with official-sounding titles in legitimate journals, but none of the cited articles actually existed when editors attempted to retrieve them. Tellingly, *none of the DOI links worked*, and journal issues purportedly containing those articles had no such content. This suggests that the authors of those papers may have used the same ChatGPT-generated bibliography or copied from one another's AI-fabricated references. The incident prompted the journal editors to investigate and consider retractions. Camp et al. dubbed this situation a "*citation catastrophe*" and cautioned that propagation of fake citations could have "*seismic impacts on the literature that underpins modern knowledge creation*" if left unaddressed. In other words, uncorrected AI hallucinations in references risk polluting the scholarly record, leading to false leads and a breakdown of trust in citation networks.

Educators and academic integrity specialists likewise report anecdotal cases linking hallucinated references to student

misconduct. In one notable case in Germany, a business student was expelled ("exmatrikuliert") in 2023 after submitting a bachelor's thesis suspected to be written with AI assistance. Examiners found that the work not only plagiarized passages but also contained numerous invented sources in its footnotes (EF1, 2023). For example, the thesis cited a specific article in a well-known journal - the journal itself exists, but the cited issue and article were entirely fictitious, as confirmed by library checks. Several references listed fake DOIs and misspelled author names, immediately raising red flags. When questioned, the student insisted he had used an AI "as a helper" but claimed to have reviewed everything; he was unable to produce or explain the phantom sources he had cited, which the university panel took as evidence of intentional deception. University officials noted that "*whoever invents sources or copies whole passages undermines the foundation of our scholarly work*". Indeed, fabricating sources strikes at the heart of academic integrity - it violates the fundamental requirement that claims be supported by *actual, traceable evidence*. As one science ethics expert bluntly stated, providing made-up references is a clear case of academic misconduct: "*Automatic text generators are no substitute for genuine scholarly work. If you use them to generate content and even invent sources, you are acting negligently.*" In the aftermath, student representatives did not dispute the severity of the offense but called for clearer guidance: they demanded explicit AI usage policies and training so that students understand how to use AI tools "*sensibly and in line with the rules*". This reflects a broader trend in academia to address the AI influx not only with sanctions, but also with education on responsible use.

Why do hallucinated references slip past students (and even some authors)? A key factor is overtrust in AI outputs combined with poor information literacy. ChatGPT's fluent and authoritative tone can create a false sense of security - students see a neatly formatted citation and assume it must be real (Welborn, 2023). If they are under time pressure or lack research experience, they might include those citations without cross-checking in a library database. This is exacerbated by the fact that verifying references requires effort: one must search for the article or book to confirm its existence. Some students may not realize that AI can fabricate sources from whole cloth, and thus they don't recognize the need to verify each reference. Others might knowingly take the shortcut, hoping the instructor won't verify every citation. In either scenario, the result is that fabricated references make it into submitted work, where they become an objective indicator of AI-generated content. Unlike subtle plagiarism, a non-existent reference is readily falsifiable by any reviewer with internet access. This makes hallucinated citations a valuable "tell" for AI involvement, as noted by Walters & Wilder: instructors can often spot ChatGPT usage by the presence of bizarre or untraceable citations in a bibliography (Walters & Wilder, 2023). Some

educators have already incorporated source-verification steps into their grading or have alerted students that they will be checking the legitimacy of references. There is also a push for technological solutions, such as automated reference checkers, but these are still nascent. Ultimately, the consensus in the literature is that students

III. METHODOLOGY

To gauge the prevalence of AI-related reference errors in student writing, we conducted an empirical analysis of bibliographies from student papers in the German-speaking academia (the DACH region: Germany, Austria, and Switzerland). Our sample comprised 368 academic submissions from undergraduate (217) and graduate (151) students across various universities in 2025. These included a mix of term papers, bachelor's theses, and master's theses in fields ranging from social sciences and humanities to natural sciences. All samples were *final submitted versions* (either as PDF or Word documents) obtained with permission of the authors or via public institutional repositories and corpora free for scientific investigation, ensuring we analyzed the actual text that instructors would evaluate. We focused exclusively on the reference lists (bibliographies) of these documents, thereby anonymizing the data by default.

Data Collection: Each bibliography was extracted with a script - where necessary we manually post-processed the format to make the references machine-readable - and parsed using a custom reference analysis tool. The tool attempted to identify each cited source's key metadata (author, title, journal or book, volume, year, DOI, etc.). We then cross-checked each reference for validity. This verification involved querying multiple scholarly databases and search engines (e.g. Google Scholar, library catalogs, DOI registries, and subject-specific databases) to determine if the referenced work *exists* and matches the details provided. For instance, if a reference was given as an article in *Journal X*, volume 12, 2019 by Smith et al., we searched for that exact combination. A reference was flagged as *potentially false* if we could not find the work or if the details were grossly incorrect (e.g. right journal but no such volume/issue, or title/author mismatches with the claimed DOI). Because remaining formatting quirks or minor typos can cause false negatives in automated checks, we manually reviewed all flags to eliminate false positives. In many cases, what looked like a nonexistent source at first was actually due to a spelling mistake or a formatting issue (for example, an author's name garbled by a PDF converter, or an incorrect punctuation in the title). We corrected such errors and re-searched to confirm the source. Only if a reference still could not be corroborated (or clearly did not match any real publication) was it deemed a "wrong reference" for the purposes of our study. We use the term *wrong reference* to encompass both fabricated references (entirely made-up works) and fundamentally erroneous references (where an attempt to cite a real work went so wrong that the

must learn to treat AI outputs with skepticism and verify any reference through independent research (Cole, Maher, & Rice, 2025). The issue is not just one of catching cheaters, but of teaching proper research practice: reliance on fabricated evidence undermines the learning process and the scholarly conversation.

reference as given is effectively invalid). A trivial typo or stylistic deviation was not counted as wrong if the source was identifiable; we focused on substantive errors that affect verifiability (such as wrong year, volume, or DOI that leads to no real record).

By manually vetting each flagged item, we aimed to ensure that our estimates of false references are conservative and not inflated by minor citation style mistakes. For additional reliability, two researchers independently reviewed each questionable reference and reached consensus on its status. We also recorded the total number of references in each paper's bibliography to compute the proportion of incorrect references per paper.

Analysis: We calculated two main measures: (1) the percentage of student papers containing any wrong references, and (2) among those papers, the proportion of references that were wrong in each, with overall averages. We also noted cases of extreme misuse (e.g. papers with very high fractions of fake citations). These metrics were then compared to contextual benchmarks from the literature (e.g. known rates of ChatGPT citation fabrication) to infer the likelihood of AI involvement. While we did not directly interrogate the students about AI use (as our analysis was retrospective on submitted texts), the patterns in the references provide an indirect indicator of how frequently AI might have been used without proper oversight.

IV. RESULTS

Our analysis revealed a significant presence of incorrect or non-existent references in the student submissions. Key findings include:

- **Prevalence of Wrong References:** Out of 368 student papers, 134 (36.4%) contained at least one wrong reference in their bibliography. In other words, more than one in three students had submitted work with an objectively false or untraceable source citation. This aligns with widespread anecdotal reports of instructors discovering bizarre citations in student work and provides quantitative confirmation that the issue is not isolated.
- **Proportion of False Citations:** Focusing on the affected papers (those 36.4%), we found that on average 12.7% of the references in each of these bibliographies were false. In practical terms, if a paper's reference list had 20 citations, about 2–3 of them, on average, were bogus. This is an alarming rate of error for something as critical

as references. It suggests that many students who included any AI-fabricated sources did not just have one slip-up, but multiple problematic references in their work.

- High-Incidence Cases:** Looking at the whole sample, 18.2% of all submissions (nearly one in five) contained a *particularly high fraction of false references* - we define this threshold as over 20% of the bibliography being wrong. In these cases, at least one out of five citations was invalid, and in some extreme instances, nearly half of the listed references were unverifiable. Such high percentages strongly imply that entire bibliographies were generated or heavily supplemented by AI, with minimal fact-checking by the student. It is unlikely for a diligent student to accidentally include so many incorrect sources; this pattern is consistent with blindly trusting an AI-generated reference list.
- Nature of Errors:** The types of reference errors we encountered mirrored those documented in the literature for AI hallucinations. Many false references had a kernel of plausibility that could mislead a casual checker. For example, several bibliographies listed *real journal titles* or *known author surnames*, but when we attempted to retrieve the articles, we discovered the article titles did not exist, or the journal volumes were wrong. One student paper in our sample cited “*International Journal of Economics*, 47(3), 2021, pp. 112-126” for a topic on labor markets - the journal is real, but it had only 30 volumes by 2021, and the article title (and given page range) did not match any content in that journal. This resembles the German case described earlier, where a journal existed but the specific cited article was fabricated. We also saw instances of fake DOIs (identifiers that looked legitimate but led nowhere) and mismatched author–title pairings (e.g. an actual article’s title paired with different author names). These errors indicate a typical ChatGPT pattern: mixing and matching pieces of real references into a Frankenstein-like fake citation.
- Comparison with AI Output Benchmarks:** The observed average of 12.7% false references (among affected bibliographies) is somewhat lower than the ~50% fabrication rate reported for raw ChatGPT-3.5 outputs (Walters & Wilder, 2023), which suggests that students have been using more advanced models (GPT-4/5) - very plausible, given our data is ~2 years younger - or doing partial verification. GPT-4, according to Walters & Wilder, hallucinated citations at a lower rate (18%), closer to what we see as an average. However, the fact that 18.2% of papers had over 20% false citations shows

many students likely *did* encounter large numbers of AI-fabricated references and still submitted them. Notably, in those high-incidence cases, the share of wrong refs could reach 30–40%, approaching the magnitude expected if a student used ChatGPT-3.5 with minimal checks. Thus, our data implies a spectrum of AI misuse: as mentioned above some students have very likely used newer models or checked some references (limiting wrong citations to ~10%), whereas others copied AI output wholesale (resulting in very high wrong-reference rates).

Measure	Value
References per bibliography	Mean = 29.0 (SD = 12.0)
Papers with ≥ 1 invalid reference	134/368 = 36.4% (95% CI 31.7–41.4)
Invalid-reference rate among affected papers (≥ 1 invalid)	Mean = 12.7% (SD = 9.0%); 95% CI of mean 11.2–14.2
Automated screening (candidate flags)	Flagged items: 2,835 ; False positives removed (manual correction): 1,554 (54.8%) ; Remaining validated flags: 1,281 (45.2%)

Table 1. Descriptive results of the study sample and automated screening workflow. Reported are the mean number of references per bibliography, the prevalence of papers containing at least one invalid (non-verifiable) reference, the mean invalid-reference rate among affected papers, and the number of algorithm-flagged items with the proportion removed as false positives following manual correction. Invalid references are defined as bibliographic entries that remained non-verifiable after manual correction of parsing/format-related false positives.

Overall, these results confirm that hallucinated references are not a rarity in student work - they appear in over one-third of papers, and in a non-trivial fraction of cases they appear in bulk. The presence of such objectively verifiable errors provides a clear window into AI-related cheating or misuse, as students would have little reason to concoct false sources on their own. In the next section, we discuss what these findings imply for academic integrity and how they correlate with global observations.

Invalid-reference rate in a paper's bibliography	<i>n</i>	% of papers	95% CI (Wilson)
0% (no invalid references)	234	63.6	58.6–68.3
1–10% invalid references	55	14.9	11.7–18.9
11–20% invalid references	12	3.3	1.9–5.6
>20% invalid references	67	18.2	14.6–22.5
Total with ≥1 invalid reference	134	36.4	31.7–41.4

Table 2. Distribution of invalid-reference rates across papers ($N = 368$). Categories indicate the percentage of invalid references within each paper's bibliography (0%, 1–10%, 11–20%, >20%), with Wilson 95% confidence intervals for each proportion. *Notes:* Categories are based on (invalid references / total references) per paper. Confidence intervals are computed per row using the Wilson method for binomial proportions.

V. DISCUSSION

Recent evidence underscores that hallucinated references are a global phenomenon undermining academic integrity. Our findings from German-speaking universities parallel reports of AI-generated citations slipping into work at all levels. Notably, an analysis of NeurIPS 2025 - a top AI conference - found over 50 accepted papers containing fabricated or altered citations that reviewers failed to catch (Goldman, 2026). Even scholarly publishers have retracted AI-assisted works after discovering that up to two-thirds of their references were nonexistent (Aksenfeld, 2025). Such cases illustrate how hallucinated references violate a core principle of scholarship: references are meant to ground claims in verifiable sources and demonstrate that authors engaged with relevant literature. A single fake citation can thus signal a severe lapse, casting doubt on an author's integrity or research literacy.

Crucially, a student submitting AI-fabricated citations may be exhibiting either misconduct or an educational gap. Some institutions explicitly classify the use of false AI-generated citations as academic falsification, on par with fabricating data (FSU, 2026). In many documented incidents, however, students appeared unaware their AI-provided references were fake (Walters & Wilder, 2023). ChatGPT often produces plausible-looking references blending real author names or journals with bogus details (Watson, 2024). This can mislead novices who lack the information literacy skills to vet sources. Thus, hallucinated references are a "red flag" indicating that the student either knowingly tried to deceive or did not know how to check the credibility of sources. In both scenarios the pedagogical message is clear: students must learn to check generative AI output. Detection and Prevention: Academic institutions are now

responding with both technical checks and educational measures. Unlike unreliable AI-writing detectors (FSU, 2026), verifying citations is comparatively straightforward. Educators and librarians are advised to cross-check all references in suspicious work, since generative AI often produces broken links or non-existent articles. New tools are emerging to automate this process: for example, conference editors used an AI "citation checker" to flag fake references by searching databases for each source (Goldman, 2026). Such tools reportedly achieve high accuracy (nearly 99%) by confirming whether a cited work truly exists. Implementing routine reference audits - manually or with software - can thus help detect integrity breaches that would otherwise go unnoticed. On the prevention side, institutions are strengthening training and policies. Integrating AI literacy into the curriculum has been recommended to ensure students understand the limitations of tools like ChatGPT (Tlili, et al., 2025). When students are made aware of how and why LLMs "hallucinate," they are more likely to approach AI outputs critically and avoid unverified content. Universities are updating honor codes to explicitly forbid invented sources and to require disclosure of AI assistance in research work (Kofinas, Tsay, & Pike, 2025). Clear guidelines about acceptable AI use, that are communicated via syllabi and academic integrity policies, help set expectations and deter misconduct. Additionally, assessment design is being rethought to reduce opportunities for AI misuse. Recent research suggests that purely take-home, essay-style tasks are vulnerable in the AI era. In response, instructors are incorporating more authentic, process-focused assessments (e.g. requiring annotated bibliographies, iterative drafts, or oral defenses) that make it harder to fake research engagement. By shifting emphasis from the final output to the research process and critical analysis, educators can both discourage reliance on AI for shortcuts and better evaluate students' true understanding.

VI. CONCLUSION

In sum, hallucinated references have emerged as a tangible indicator of academic writing falling short of scholarly standards. Their presence in an assignment or publication is an unambiguous warning sign, either of deliberate integrity violation or of insufficient scholarly preparation. The most recent literature makes it evident that addressing this issue requires a multifaceted response. Universities must reinforce a culture of verification: students (and even faculty) should be expected to confirm that every citation is real and relevant. At the same time, cultivating students' research and AI literacy is imperative so they do not inadvertently propagate AI falsehoods. The convergence of technology, policy, and pedagogy is key. By deploying effective reference-checking mechanisms, establishing clear ethical guidelines, and redesigning learning activities, academic institutions can mitigate the risks of generative AI misuse while harnessing its benefits. Ultimately, the challenge of AI hallucinated citations highlights an opportunity to *strengthen* research integrity frameworks. It calls on educators, librarians, and administrators to collaborate in fostering an environment where proper evidence-based writing is non-negotiable, and

where students are equipped to use AI tools responsibly as aids for learning rather than engines of fabrication. Such concerted efforts will help uphold the credibility of scholarly work in the AI-supported writing era.

VII. REFERENCES

- Aksentfeld, R. (2025, 06 30). *retractionwatch.com*. Retrieved from Springer Nature book on machine learning is full of made-up citations: <https://retractionwatch.com/2025/06/30/springer-nature-book-on-machine-learning-is-full-of-made-up-citations>
- Camp, N. T., Bengston, J. A., & Sandstrom, J. C. (2025). The citation catastrophe: Propagation of AI-generated counterfeit citations in scholarship. *The Journal of Academic Librarianship*, 51(4), 103065. doi:10.1016/j.acalib.2025.103065
- Cole, R., Maher, L., & Rice, R. (2025). *Understanding and Avoiding Hallucinated References*. Retrieved 1 15, 2026, from <https://wacclearinghouse.org/repository/collections/continuing-experiments/august-2025/ai-literacy/understanding-avoiding-hallucinated-references/>
- EF1. (2023). *Gefälschte Quellen, kopierte Passagen, Exmatrikulation: Wie eine KI-Bachelorarbeit zum Aus für einen Studierenden wurde*. Retrieved from efactory1.de: <https://efactory1.de/gefaelschte-quellen-kopierte-passagen-exmatrikulation-wie-eine-ki-bachelorarbeit-zum-aus-fuer-einen-studierenden-wurde>
- FSU. (2026, 01 26). *Florida State University*. Retrieved from Artificial Intelligence - Academic Integrity: <https://ai.fsu.edu/academic-integrity>
- Goldman, S. (2026, 01 21). *Fortune.com*. Retrieved from NeurIPS, one of the world's top academic AI conferences, accepted research papers with 100+ AI-hallucinated citations, new report claims: <https://fortune.com/2026/01/21/neurips-ai-conferences-research-papers-hallucinations/>
- Gravel, J., D'Amours-Gravel, M., & Osmanliu, E. (2023). Learning to Fake It: Limited Responses and Fabricated References Provided by ChatGPT for Medical Questions. *Mayo Clinic Proceedings: Digital Health*, 1(3), 226-234.
- Kofinas, A. K., Tsay, C. H.-H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*, 2522-2549. doi:10.1111/bjet.13585
- Orduña-Malea, E., & Cabezas-Clavijo, Á. (2023). ChatGPT and the potential growing of ghost bibliographic references. *Scientometrics*, 128(9), 5351-5355. doi:10.1007/s11192-023-04804-4
- Rong, H., & Chun, C. (2024). *Digital Education Council Global AI Student Survey 2024: AI or Not AI: What Students Want*. Singapore: Digital Education Council.
- Scarfe, P., Watcham, K., Clarke, A., & Roesch, E. (2024). A real-world test of artificial intelligence infiltration of a university examinations system: A "Turing Test" case study. *PLoS One*, 19(6), e0305354. doi:10.1371/journal.pone.0305354
- Shao, A. (2025). New sources of inaccuracy? A conceptual framework for assessing different types of misinformation produced by generative AI. *Harvard Kennedy School Misinformation Review*. doi:10.37016/mr-2020-182
- Tlili, A., Bond, M., Bozkurt, A., Arar, K., Chiu, T. K., & Rospigliosi, P. (2025). Academic integrity in the generative AI (GenAI) era: a collective editorial response. *Interactive Learning Environments*, 1819-1822. doi:10.1080/10494820.2025.2471198
- Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific Reports*, 13(1), 14045. doi:10.1038/s41598-023-41032-5
- Watson, A. (2024). Hallucinated Sources: An Analysis of Student-Submitted AI-Generated. *Brick & Click Libraries Conference Proceedings* (pp. 38-50). Maryville, MO: Northwest Missouri State University Press.
- Welborn, A. (2023, 3 9). *ChatGPT and Fake Citations*. Retrieved 1 15, 2026, from Duke University Libraries: <https://blogs.library.duke.edu/blog/2023/03/09/chatgpt-and-fake-citations/>

Disclaimer: This paper was prepared with the support of generative AI tools. AI was used to assist with literature discovery (e.g., identifying potentially relevant publications and keywords) and to improve readability through language editing and stylistic optimization. All core research activities, including study design, data collection, data cleaning (including manual correction of false positives), analysis, interpretation of results, and the substantive arguments and conclusions - were developed by the authors. All cited sources were selected and checked by the authors, and responsibility for the content, accuracy, and final wording of this manuscript rests entirely with the authors.