



Algorithmic revision: The role of generative AI and digital media in exam preparation among journalism students

Mercy Munene

Kabarak University, Kenya

Article History

Received: 2026-02-20

Revised: 2026-08-17

Accepted: 2026-09-01

Published: 2026-09-04

Keywords

ChatGPT

Digital media

Generative AI

Journalism

How to cite:

Munene, M. (2026). Algorithmic revision: the role of generative AI and digital media in exam preparation among journalism students. *Journal of Visual and Performing Arts*, 4(1), 39-57.

Abstract

Generative artificial intelligence (GenAI) and digital media have rapidly become central to how students study, yet journalism students, whose profession is defined by verification, accuracy, and originality, remain under-researched, particularly in the high-stakes context of examination preparation and in settings outside the global North. This study describes the patterns, purposes, and perceived consequences of GenAI and digital-media use in exam revision among undergraduate journalism students, tests whether critical verification of AI output strengthens with academic seniority, and derives implications for curriculum design. A cross-sectional survey combining closed and open-ended items was administered to 65 undergraduate journalism students across four years of study during an examination period (April 2026). Closed items were analysed with descriptive statistics, cross-tabulation by year, and correlations; open-ended responses were analysed thematically using an inductive procedure. Adoption was widespread: 69.2% used these tools frequently or exclusively, led by ChatGPT (66.2%) and Google Gemini (35.4%). The dominant purpose was explaining complex concepts (mean = 3.92/5); 70.8% rated the tools very or highly effective and 61.5% reported reduced reliance on traditional materials. However, only 58.5% verified AI output consistently against authoritative sources, and, contrary to expectation, verification did not significantly increase with year of study. Qualitative themes were factual inaccuracy and "hallucination," divergence from lecturer-sanctioned content, and anxiety about eroded critical thinking; students' outlook was markedly ambivalent, framing AI as a "double-edged sword." A value-action gap characterises journalism students' AI use: uptake and perceived effectiveness are high while critical verification lags and does not mature with experience. Because students are already critically aware, programmes should teach AI verification and critical use as explicit, assessed professional competencies rather than relying on prohibition or on the assumption that judgement develops on its own.

Copyright © 2026



Introduction

The arrival of widely accessible generative artificial intelligence has altered the everyday study practices of university students faster than institutions have been able to respond. Within roughly two years of the public release of conversational large language models, tools such as ChatGPT, Google Gemini, and Claude moved from novelty to infrastructure for millions of learners, sitting alongside



an older layer of digital media, YouTube tutorials, podcasts, and search engines, that students had already woven into their revision routines. ChatGPT alone attracted an estimated one hundred million active users within two months of its launch (Tian et al., 2024), and surveys across widely differing contexts now report that a majority to a large majority of students use such tools for academic work, including assignment writing and exam preparation (Abbas et al., 2024; Alonso-Muñoz & Casero-Ripollés, 2026; Ajalo et al., 2025). The question facing educators is therefore no longer whether students use generative AI, but how, to what end, and with what consequences for the skills their degrees are meant to build.

For journalism students in particular, this shift carries a distinctive weight, because their discipline is simultaneously a site of study and the profession most directly disrupted by the same technologies. The newsroom they are training to enter is itself rapidly integrating generative AI into news production and distribution, even as practitioners voice concern that these tools threaten accuracy and credibility and demand new forms of algorithmic literacy (Cools et al., 2024; Pavlik, 2023). Journalism's foundational professional norm, that claims must be verified against authoritative sources before they are published, is precisely the norm that fluent, confident, and sometimes fabricated AI output most directly tests. A journalism student who accepts a chatbot's summary of media law without checking it is not merely risking a wrong exam answer; they are rehearsing the opposite of the habit their future work will require. This alignment between a study behaviour and a professional value makes journalism students an unusually revealing population in which to examine the consequences of uncritical AI use.

This study investigates a specific and consequential moment in that transition: how journalism students actually use generative AI and digital media when the stakes are highest, during examination preparation. Exam revision is a revealing lens because it compresses students' real learning priorities into a short, high-pressure window. What tools they reach for, what they ask those tools to do, and, critically, whether they check the answers reveals more about their genuine study epistemology than any statement of intention could. It is one thing to endorse critical thinking in the abstract; it is another to fact-check a model's output at two in the morning before an examination. The present study is further situated in an African higher-education context, a setting largely absent from the early literature, which has been dominated by samples from Europe, East Asia, and North America; evidence from comparable sub-Saharan settings suggests both very high adoption and an acute, explicitly voiced need for AI literacy (Ajalo et al., 2025).

Background and related work

Generative AI refers to a class of machine-learning systems, principally large language models, capable of producing fluent human-like text, summaries, explanations, and practice material on demand. Scholars have argued that such systems hold genuine promise as personalised tutors that adapt explanations to a learner's level, provide instant feedback, and lower the barrier to engaging with difficult material (Kasneci et al., 2023). Empirical work is beginning to substantiate specific benefits: ChatGPT-supported retrieval practice and self-regulated study designs have been associated with improved engagement, motivation, and, in some settings, higher examination performance (Ng et al., 2024). At the same time, the same body of work is unambiguous that these benefits are shadowed by well-documented risks, factual "hallucination," the reproduction of bias present in training data, and the encouragement of shallow, passive learning in which students substitute the model's output for their own cognitive effort (Kasneci et al., 2023; Rudolph et al., 2023).

This double-edged character recurs across the recent literature. Cotton et al. (2024) foregrounds the academic-integrity challenge posed by fluent, hard-to-detect output, while large-sample studies of



student behaviour have begun to trace the downstream consequences of heavy use. In a two-study investigation, Abbas et al. (2024) found that academic workload and time pressure predicted greater ChatGPT use, and that heavier use was in turn associated with increased procrastination, memory loss, and, counter to students' own expectations, lower academic performance. Analyses of assignment behaviour reach compatible conclusion. Das et al. (2025) report that frequent reliance on ChatGPT to learn new concepts correlate with over-dependence and warn that unchecked reliance may erode critical thinking and originality, calling for discipline-specific guidelines and "AI-aware" assessment design. Descriptive surveys converge on the same tension, finding that students perceive clear efficiency benefits, rapid task completion, personalised tutoring, round-the-clock access, while simultaneously acknowledging that over-reliance narrows exploration, critical thinking, and academic honesty (e.g., Chan & Hu, 2023).

Students' own perceptions have been the focus of a widely cited strand of research. Chan and Hu (2023), surveying undergraduates and postgraduates, found generally positive attitudes towards GenAI, valued for personalised support, writing assistance, and research capability, coupled with real anxieties about accuracy, privacy, ethics, and the impact on personal and professional development, a pattern the present study also observes. The journalism-specific literature is smaller but growing quickly and is especially pertinent here. Alonso-Muñoz & Casero-Ripollés (2026), surveying 281 journalism students in Spain, found that 93% used GenAI, that 77% used it for learning or studying, and, notably, that use was significantly higher among third- and fourth-year students, while students themselves recognised that they needed more AI literacy. Notably, Alonso-Muñoz & Casero-Ripollés (2026) reported no significant difference in evaluation behaviours by educational level, which pre-empted our exploration of whether critical verification matures with seniority. Chen et al. (2026), modelling data from 500 journalism and communication students, showed that AI integration improves professional skill development not directly but through a developmental sequence, first raising learning engagement, then media literacy, and only then professional competence, implying that critical-evaluation skills are a necessary intermediate stage rather than an automatic by-product of tool use. Complementing these survey studies, experiential and interview-based work has identified distinct student "archetypes" of AI engagement, ranging from those who treat the tool as a controllable instrument to those who experience it as an unreliable shortcut, and has argued for a "critical AI literacy" that builds metacognitive awareness of the relationship students form with the technology (Seven et al., 2025).

Digital media, especially video, had already been incorporated into informal learning long before conversational AI, and is particularly valued for the procedural and practical skills, camera operation, video editing, graphic design, that are central to journalism training. The popular assumption that today's students are innately capable "digital natives" (Prensky, 2001) has, however, been substantially challenged: Kirschner and De Bruyckere (2017) argue that fluency in using digital tools does not confer the critical or metacognitive skills needed to learn effectively with them, a distinction between operational skill and critical judgement that is central to interpreting the present findings. Institutional guidance has meanwhile shifted from prohibition towards managed, human-centred integration (UNESCO, 2023), even as reviews note that pedagogical practice has consistently lagged technical capability (Zawacki-Richter et al., 2019). A promising reframing comes from assessment research: observing engineering students permitted to use ChatGPT in open-book examinations, Mahmoud (2026) found that the strongest evidence of understanding appeared not when students retrieved answers but when they critically verified, debugged, and justified AI output, suggesting that verification and judgement, rather than answer production, are becoming the competencies assessment should target.



Three gaps motivate the present study. First, journalism students are still rarely examined as a distinct population despite the acute alignment between AI's risks and their profession's verification norms, and the few existing studies are concentrated in European settings (Alonso-Muñoz & Casero-Ripollés, 2026) and East Asian settings (Chen et al., 2026). Second, much existing work reports general attitudes or aggregate usage rather than behaviour in a specific, high-stakes context such as exam revision, where verification either happens or does not. Third, few studies pair quantitative usage data with students' own qualitative reasoning about the trade-offs. Furthermore, while Alonso-Muñoz & Casero-Ripollés (2026) tested the intuitive assumption that critical verification matures as students advance through a degree finding no statistically significant differences by educational level in their European sample there remains a critical need to replicate this inquiry in new, non-Western settings.

Aim and research questions

This study describes and interprets the patterns, purposes, and perceived consequences of GenAI and digital media use in exam preparation among undergraduate journalism students, and locates a critical-literacy gap with direct implications for curriculum design. It is guided by five research questions: (RQ1) How frequently, and with which tools, do journalism students use GenAI and digital media during exam revision? (RQ2) For what primary learning purposes are these tools used? (RQ3) How effective do students perceive them to be, and how has their reliance on traditional study materials changed? (RQ4) How consistently do students verify AI output against authoritative sources, and does this vary by year of study? (RQ5) What do students perceive as the main risks of AI reliance, and how do they anticipate it affecting their future journalism skills? RQ4 is of particular theoretical interest: exploring whether verification is a maturational skill acquired through disciplinary training that should rise with year of study; if it is not, the case for teaching it explicitly is strengthened.

The study makes three contributions. Empirically, it provides one of the few discipline-specific accounts of AI-supported exam preparation among journalism students, and does so from an African higher-education setting that is largely missing from a literature dominated by European, East Asian, and North American samples. Theoretically, it reframes verification, usually treated as a matter of academic honesty or study quality, as the specific competence at which a study behaviour and a professional norm coincide, and it tests, rather than assumes, whether that competence grows with experience. Practically, it converts a widely reported but diffuse worry about student "over-reliance" into a concrete, addressable target: not use of AI as such, which is widespread and often constructive, but the uneven and non-maturing habit of checking what AI produces. Framing the problem this precisely matters, because it points to an intervention, assessed verification, that is low-cost, professionally authentic, and within the reach of any programme, rather than to a prohibition that the evidence suggests would simply be ignored.

Methods

The study used a cross-sectional descriptive survey design combining quantitative closed-ended items with illustrative qualitative commentary. Quantitative metrics established the scale, frequency, and patterns of tool use, while open-ended qualitative responses were analysed using an inductive thematic procedure (Braun & Clarke, 2006) to provide context and nuance to the quantitative findings. A single questionnaire combined closed items, capturing frequency of use, tools relied upon, purposes, perceived effectiveness, change in reliance on traditional materials, and verification behaviour, with three open-ended questions inviting students to give a concrete example of helpful use, to name the biggest risk of AI reliance, and to reflect on how their current habits might affect their future journalism skills. This design allowed the quantitative items to establish the scale and shape of the phenomenon while the qualitative items surfaced the reasoning, nuance, and ambivalence behind



the numbers, following the logic that neither strand alone could adequately characterise a behaviour that is at once widespread and internally conflicted.

Participants and sampling

Sixty-five undergraduate journalism students completed the survey during the semester's peak continuous assessment and examination window, with responses collected across active revision periods in April 2026. Participation was voluntary, and every respondent recorded informed consent before proceeding. Because the survey was administered via an open online link across the department, a formal denominator of approached students and an exact response rate could not be tracked, limiting the sample to self-selected respondents. Before data collection, the study received institutional ethical approval and a research license from the National Commission for Science, Technology and Innovation (NACOSTI)

Table 1: Distribution of respondents by year of study (N = 65)

Year of study	Respondents (n)	Percentage
First year	19	29.2%
Second year	22	33.8%
Third year	16	24.6%
Fourth year	8	12.3%
Total	65	100%

The sample spanned all four years of undergraduate study (Table 1): second-year students formed the largest group (33.8%), followed by first-years (29.2%) and third-years (24.6%), with fourth-years the smallest cohort (12.3%). The sample was self-selected and drawn from a single programme context, and is best understood as a purposive, non-probability sample suited to analytical rather than strictly statistically generalisable inference; this is an appropriate and common design for an emerging phenomenon in a specific disciplinary and regional setting.

Instrument and measures

Frequency of use was captured on a four-point ordinal scale (Rarely, Sometimes, Frequently, Exclusively; noting that the instrument did not include a 'Never' option, which restricted measurement to active users). Tool reliance was a multiple-response item allowing students to select every tool they leaned on most, so tool percentages sum to more than 100%. The four purpose items summarising notes, generating practice questions, explaining concepts, and structuring essay outlines were each rated on a five-point agreement scale, designed to capture specific study and comprehension support roles rather than unmeasured answer-generation or submission behaviours. An attention and consistency check item ("I did not use AI tools") was included; 5 respondents produced inconsistent ratings by affirming non-use while reporting frequent tool interaction elsewhere in the survey, and a sensitivity analysis excluding these $n = 5$ cases confirmed that overall descriptive distributions and correlations remained virtually identical, justifying the retention of the full sample. Perceived effectiveness used a five-level ordinal scale, change in reliance on traditional materials a four-level scale, and verification frequency a five-level scale from Never to Always. For analysis, ordinal scales were assigned integer values (e.g., Never = 1 to Always = 5) to permit the calculation of means and Spearman's rank correlations; these are reported alongside full frequency distributions to preserve transparency and to avoid over-stating the precision of ordinal data.



Data analysis

Closed responses were analysed using descriptive statistics (frequencies, percentages, and means – with standard deviations omitted to align with non-parametric ordinal presentation), cross-tabulation by year of study, and Spearman's rank correlations (ρ) between ordinally-scored variables (reported with exact sample size $n = 65$ and p -values; the reliance-change item was coded on a scale from 1 = significantly decreased to 4 = increased, ensuring the positive sign of $\rho = .36$ correctly reflects direct co-variation with frequency of use); given the modest sample and ordinal measurement, these correlations are interpreted as indicative of association rather than as tests of formal hypotheses. The three open-ended questions were analysed thematically following an inductive, reflexive thematic analysis procedure (Braun & Clarke, 2006): all 65 respondents contributed open-ended data across the survey instruments, responses were read repeatedly for familiarisation, initial codes were generated close to participants' own language by the author, codes were grouped into candidate themes, and themes were reviewed against the full dataset with active searches conducted for negative or disconfirming cases to ensure analytical coherence and distinctiveness. Illustrative quotations are reproduced verbatim, including original spelling and grammar, to preserve authenticity of voice and to allow readers to judge the interpretation against the evidence.

Ethical considerations

Prior to data collection, formal ethical clearance was obtained from the Institutional Research Ethics Committee (IREC) of Mount Kenya University. The lead author is a member of staff at Kabarak University and is concurrently a PhD candidate at Mount Kenya University, where this study forms part of the doctoral requirements; ethical clearance for doctoral research is issued by the candidate's degree-awarding institution, which accounts for the review being conducted by Mount Kenya University's IREC. On the strength of that institutional approval, a single research licence was secured from the National Commission for Science, Technology and Innovation (License No. NACOSTI/P/26/4188904), in full compliance with the Science, Technology and Innovation Act 2013 of Kenya. Participants were recruited via an open departmental invitation and online distribution of the survey link across student cohorts. To mitigate any potential conflict of interest or perceived academic coercion arising from institutional affiliation, the author confirms that they held no direct teaching, grading, or assessment relationship with the participating students during the data collection window. The survey opened with an explicit consent question, and only respondents who confirmed that they understood the information provided and voluntarily agreed to participate were included; all 65 analysed records met this condition. No personally identifying information was collected, responses were anonymous, and participation carried no bearing on students' academic standing. Because the open-ended responses occasionally touched on students' study anxieties, quotations are attributed only by year of study to preserve anonymity.

Results

Adoption of generative AI and digital media during the exam period was notably intense. Almost seven in ten respondents (69.2%) reported using these tools frequently ("almost every day," 44.6%) or exclusively ("it was my main method of studying," 24.6%). A further 24.6% used them sometimes, and only 6.2% reported rare use (noting that the survey instrument did not include a 'Never' option, meaning non-use could not be empirically captured). As Figure 1 shows, the modal student was a daily user, and a substantial quarter of the sample had made AI and digital media the central pillar of their revision. This level of frequent engagement (69.2%) compares closely to, and slightly exceeds, the 61.2% of journalism students reported by Alonso-Muñoz & Casero-Ripollés (2026) who use GenAI



daily or several times weekly, aligning also with findings from comparable regional settings (Ajalo et al., 2025).




Figure 1: Frequency of generative AI and digital media use during the exam period (N = 65)

Tools relied upon

Tool use was concentrated but not monolithic (Figure 2). ChatGPT was dominant, relied upon by 66.2% of respondents, followed by Google Gemini at 35.4%. Beyond these two conversational models, a distinct second tier of digital media persisted: YouTube and other video platforms (23.1%), standard search engines (20.0%), and dedicated AI note-summarising tools such as ChatPDF and Quillbot (16.9%). Newer or niche tools, Grok, Claude, Gamma, and Google Scholar, appeared only occasionally. On average each student drew on 1.69 tools, and 41.5% combined two or more, indicating that many students assemble a small personal "stack" rather than depending on a single application. The persistence of video is significant for journalism specifically: as the qualitative data below show, students turned to tutorials for exactly the practical, procedural skills, editing, camera work, graphic design, that text-based models handle poorly, a division of labour between tool types that any account of "AI use" must accommodate.




Figure 2: Tools relied on most for exam preparation; multiple selections permitted, so values exceed 100% (N = 65)

Purposes of use

Across the four purpose items, mean agreement scores spanned a narrow range from 3.37 to 3.92 on the five-point scale: explaining complex journalism theories and concepts in simpler terms received the highest mean (3.92; 63.1% agreed or strongly agreed), followed by summarising long lecture notes and PDFs (mean = 3.58), generating practice questions for self-testing (mean = 3.51; 53.8% agreement), and structuring thoughts or outlines for essay questions (mean = 3.37). The pattern (Figure 3) indicates that students utilize the tools across these measured study support functions, though these descriptive means reflect a relatively compressed range without formal tests of difference between items.




Figure 3: Primary purposes for using AI and digital tools in revision: mean agreement (1-5) and percentage agreeing or strongly agreeing



The qualitative examples strongly corroborate this ordering. The most common account of helpful use was condensing dense material into revisable form, closely followed by simplifying difficult concepts and terminology, self-testing with generated questions, and learning practical production skills through video:

"I use Gemini to summarize note and get possible questions from the pdf." (First-year student)

"I used ChatGPT to simplify and summarize difficult topics in media studies, such as media theories and research methods. It helped me understand concepts like Agenda Setting Theory and Gatekeeping Theory... I also used digital videos to understand practical skills such as video editing techniques in Adobe Premiere Pro." (Third-year student)

"Watching of tutorials ... enhance my understanding of a particular concept like photography, camera handling and video editing." (Second-year student)

Perceived effectiveness and shift from traditional materials

Students' evaluation of these tools was strongly positive (Figure 4). A combined 70.8% rated them either "highly effective" (21.5%) or "very effective" (49.2%) for improving their understanding of course material measured on a 5-point ordinal scale ranging from 1 (not effective at all) to 5 (highly effective) with a further 21.5% judging them moderately effective. Negative assessments were rare: only 6.2% found them slightly effective and a single respondent (1.5%) found them not effective at all, yielding a mean effectiveness score of 3.83. This near-consensus that the tools "work" is important context for the more cautious findings that follow: the students in this sample are not sceptics who tolerate AI but enthusiasts who nonetheless harbour serious and specific reservations.




Figure 4: Perceived effectiveness of AI and digital tools for improving understanding of course material (N = 65)

Consistent with intensive adoption, a clear majority reported reduced reliance on traditional study materials since starting university (Figure 5). Combining the two "decreased" categories, 61.5% reported a shift away from handwritten notes and printed handouts, comprising 35.4% whose reliance had "significantly decreased" (relying mostly on AI and digital sources) and 26.2% whose reliance had



"slightly decreased." Roughly a third (32.3%) reported no change, using both equally, and only 6.2% reported increased reliance on traditional notes, a small but theoretically interesting group who said they now trust traditional materials more precisely because of experience with AI's fallibility. Frequency of use was positively associated with this retrospective shift (Spearman's $\rho = .36, p < .05, n = 65$), showing that heavier users tended to report greater movement away from traditional materials. However, because these data rely on cross-sectional retrospective self-reports, causal direction cannot be determined (i.e., whether heavier AI use caused the reduction or resulted from it). Furthermore, the framing of change "since starting university" encompasses vastly different time horizons for first-year respondents (29.2% of the sample) compared to senior cohorts, a temporal artifact pooled within the aggregate distribution.




Figure 5: Reported change in reliance on traditional study materials since starting university ($N = 65$).

Verification behaviour

Asked how often they fact-check AI output against lecturer notes or textbooks, responses spanned the full frequency spectrum: 35.4% answered "always," 23.1% "often," 24.6% "sometimes," 15.4% "rarely," and 1.5% "never" (Figure 6). While 58.5% verify consistently when combining "always" and "often," shifting the threshold alters the descriptive framing: 83.1% verify at least sometimes, whereas 16.9% verify rarely or never. In a discipline whose defining professional norm is verification, this full



distribution highlights that while strict routine checking is practiced by a majority, a substantial minority engages in sparse or absent verification.




Figure 6: Frequency of fact-checking AI output against official lecturer notes or textbooks (N = 65)

Directly addressing RQ4, cross-tabulation by year of study (Figure 7) shows that verification rates do not significantly vary across academic seniority $\chi^2(3) = 2.456, p = .483$; Cochran-Armitage $Z = 0.354, p = .723$ and $\rho = .36, \rho = .32$. Consistent verification was reported across cohorts (ranging from 45% among second-years to 69% among third-years), but given the modest sample size and statistical power, these cross-sectional findings remain inconclusive regarding whether verification frequency changes with academic progression. Because no data were collected on specific course structures or assessment designs, alternative explanations—such as sampling noise or cohort effects, given that fourth-year respondents began their degrees before ChatGPT was publicly released—cannot be ruled out, and attributions to pedagogical practices or individual dispositions cannot be supported.




Figure 7: Verification behaviour by year of study, as a percentage within each year group



A modest but suggestive relationship emerged between verification and perceived effectiveness. Students who verified consistently reported higher mean effectiveness (4.05) than those who verified rarely or never (3.64). However, reporting only these extreme groups conceals a non-monotonicity in the data: the middle group of students who verified only "sometimes" reported the lowest mean effectiveness (3.44). Therefore, while the two variables were positively correlated overall (Spearman's $\rho = 0.32$), this non-linear ordering breaks the assumption of a monotonic "virtuous circle," further indicating that causation (e.g., checking directly producing understanding) cannot be inferred from this cross-sectional association.

Perceived risks

The open-ended question on the biggest challenge or downside of relying on generative AI produced the richest and most consistent qualitative data, and three dominant themes emerged. The most frequently cited was factual inaccuracy, "hallucination" and misinformation, often expressed with real concern about its consequences for exam performance:

"Misinformation because it tends to hallucinate generating non-existent or wrong information."

(Fourth-year student)

"The biggest downside ... is that it can sound confident even when it's wrong or misses key details. Since AI generates responses quickly, it's easy to accept the information without double-checking sources." (First-year student)

The second theme, distinct from general inaccuracy, was divergence from lecturer-sanctioned content, a worry that AI answers would not match what the lecturer taught or expected, or would be too shallow relative to classroom material:

"Some answers or points will not align with what the lecture might have taught and may not give me the chance to expound on my thinking." (Second-year student)

"Sometimes the information provided is not detailed a lot like the ones provided in the classroom." (First-year student)

The third, and arguably most striking, theme was a self-aware anxiety that over-reliance erodes cognition, critical thinking, creativity, and independent reasoning. Many students named, in their own words, precisely the risk that educators most fear:

"It prevents your brain from wandering, thus less creativity." (Second-year student)

"Overlying on generative AI lowers the thinking capacity." (First-year student)

"The biggest risk is ... the danger of passive learning, where relying on quick summaries prevents me from actually reading and deeply understanding the source material." (Second-year student)

A related cluster of responses raised academic-integrity concerns, plagiarism and the danger of homogenised answers when many students prompt the same tool in the same way, while a small minority reported experiencing no challenges at all. Taken together, the risk responses show a student body that is far from naive: it can articulate the hazards of the tools it uses with considerable precision, even while continuing to use them intensively.

Outlook on future journalism skills

Asked how their current AI study habits might affect their actual journalism skills, writing, critical thinking, research, in the newsroom, respondents divided into three broad orientations. An optimistic group emphasised efficiency gains: faster background research, improved writing structure and grammar, idea generation, and quicker synthesis of complex topics. A pessimistic group feared the



atrophy of core competencies, worrying that dependence would leave them unable to generate ideas or conduct research independently. The largest and most interesting group, however, was ambivalent, framing AI explicitly as a "double-edged sword" whose ultimate effect would depend on how it is used:

"If i use it purely to automate tasks, this habit might create a "lazy" researcher; if i use it to augment, it will produce a highly efficient, data-driven journalist in me." (Fourth-year student)

"Over reliance will definitely reduce my capabilities to come up with material organically. However, it can speed up processes like editing copy and gathering research material." (Fourth-year student)

"As a future media professional, using AI ... can improve my research speed, writing structure, and idea generation ... However, overreliance on it may weaken my critical thinking, originality, and ability to conduct independent research. Therefore, I should use AI as a supportive tool while still practicing my own writing, verifying information, and engaging in real reporting." (First-year student)

This conditional, "augment-not-automate" framing recurred across all years of study and represents, in effect, a folk theory of responsible AI use that students have largely arrived at on their own.

Discussion

This study set out to describe how journalism students use generative AI and digital media during exam preparation, to test whether critical verification strengthens with seniority, and to identify implications for their education. The findings support a single organising claim, a value-action gap, which the following subsections develop by interpreting the pattern of use, situating it against the emerging literature, and translating it into concrete proposals for journalism curricula. The section closes with the study's limitations and its conclusions.

Principal findings

The central tension in these data is what may be called the value-action gap (or attitude-behaviour gap). Journalism students have adopted generative AI widely (69.2% frequent or exclusive users), rate it highly effective (70.8% positive), and are shifting away from traditional materials (61.5%). Yet more than four in ten do not verify its output consistently, and, contrary to the intuitive expectation tested in RQ4, verification does not improve with year of study. Adoption has plainly outpaced the critical literacy needed to use these tools safely. This mirrors the broader critique of the "digital native" myth (Kirschner & De Bruyckere, 2017; Prensky, 2001): operational fluency with AI is now widespread, but the evaluative judgement required to interrogate its output is not, and the two must not be assumed to travel together. The paradox is sharpened, not softened, by the discipline in question. For most students, failing to verify a chatbot risks a poor grade; for journalism students it also means rehearsing a habit that is the precise inverse of professional practice, in which unverified claims are, by definition, not yet publishable.

The non-monotonic relationship between verification and year of study deserves emphasis because it complicates a comfortable assumption, yet it converges directly with prior work. While, Alonso-Muñoz & Casero-Ripollés (2026) found that more senior journalism students used GenAI more, and Ajalo et al. (2025) found that younger students adopted AI tools more readily than older ones; the present data suggest that whatever changes with seniority, disciplined verification is not reliably part of it. If critical checking were an emergent property of disciplinary training, something students absorb simply by spending more years in a journalism programme, one would expect a clear upward gradient. Its absence implies that verification behaviour is contingent on local and modifiable factors, most plausibly the design of specific assessments and the expectations of individual lecturers. That is



a hopeful reading: contingent behaviours can be engineered, whereas maturational ones can only be waited for.

Finally, the modest positive association between verification and perceived effectiveness (Spearman's $\rho = 0.32$) indicates that consistent verification and higher perceived utility are linked. One might hypothesize that students who check AI output against authoritative sources are engaging with the material more deeply, and that this active engagement is associated with higher perceived understanding. However, given the cross-sectional design and the non-monotonic subgroup ordering noted earlier, it cannot be concluded that active engagement through verification directly linked to the understanding students attribute to the tool. The association remains strictly correlational and may reflect other underlying student characteristics rather than a straightforward causal mechanism.

Patterns and purposes of use

The purpose data suggest a constructive dominant use: students reach for AI above all to understand, to have complex theories explained simply and dense notes compressed into revisable form. This aligns with the promise identified by Kasneci et al. (2023) of AI as a personalised explanatory tutor and with evidence that well-structured, AI-supported retrieval and self-regulated study can raise engagement and performance (Ng et al., 2024). Whether that offloading supports or supplants germane learning is, however, exactly the question the verification gap raises. Summarisation that is checked and engaged with builds understanding; summarisation that is passively consumed may only simulate it, and the large-sample findings that heavy, unchecked use is associated with procrastination, memory loss, and weaker performance (Abbas et al., 2024; Das et al., 2025) mark the downside of the same behaviour. The dominant purpose being comprehension is therefore genuinely encouraging, but only conditionally so: it is a good use that depends on a habit, verification, that four in ten students do not reliably practise.

The continued prominence of video platforms (23.1%) is not a residual habit but a rational division of labour, and it cautions against treating "AI use" as a single, homogeneous behaviour. Journalism is a practical discipline, and students turned to video specifically for the procedural skills, editing, camera operation, graphic design, that conversational text models cannot demonstrate. The generative AI for conceptual compression, video for embodied skill, and search engines for breadth, an intervention that treats all digital study as interchangeable would miss this structure entirely.

Student awareness and critical literacy

Perhaps the most actionable finding is that students already possess the conceptual vocabulary of responsible AI use. They name hallucination, over-reliance, homogenised answers, and the erosion of originality; many independently articulate an "augment, not automate" ethic. The three orientations identified here, optimistic, pessimistic, and, most commonly, ambivalent, echo the engagement archetypes documented in experiential studies of journalism students, in which learners position themselves along a spectrum from confident director to wary antagonist of the tool (Seven et al., 2025). Educators are therefore not starting from ignorance or denial but from a substantial base of genuine, if unstructured, critical awareness. The task is less to convince students that risks exist than to convert their diffuse anxiety into concrete, teachable practices. That students voice these concerns while continuing to use the tools intensively is not hypocrisy but the ordinary condition of anyone using a powerful technology whose costs are diffuse and whose benefits are immediate; the pedagogical opportunity lies in closing the gap between what students already know and what they routinely do.

This interpretation also reframes the widely reported student demand for more AI literacy (Alonso-Muñoz & Casero-Ripollés, 2026; Chan & Hu, 2023) as a request that educators are well placed to meet. The developmental model proposed by Chen et al. (2026), in which AI use raises engagement, which



builds media literacy, which in turn supports professional competence, implies that critical-evaluation skill is the load-bearing intermediate stage. The verification gap documented here is, in those terms, a bottleneck at exactly that stage: students have cleared the engagement threshold almost completely but have not uniformly built the evaluative literacy that would let engagement translate into professional competence rather than dependence.

Verification and professional practice

A distinctive contribution of studying this behaviour in journalism students, rather than students in general, is that it exposes verification as the point where study habit and professional identity converge. In most disciplines, checking a chatbot's claim is a matter of academic accuracy; in journalism it is a scaled-down enactment of the profession's central act, in which a claim has no standing until it has been corroborated against an authoritative source. Read this way, the 41.5% who verify only sometimes, rarely, or never are not simply cutting a study corner: they are practising, night after night, the disposition their future editors will most need them to have unlearned. The concern is amplified by the newsroom context these students are entering, where practitioners are simultaneously adopting generative AI for efficiency and warning that it threatens accuracy and credibility unless matched by robust algorithmic literacy (Cools et al., 2024; Pavlik, 2023). The habits formed in exam revision are, in effect, a rehearsal for the newsroom, and an unverified rehearsal rehearses the wrong thing.

This framing also clarifies why the divergence-from-lecturer theme matters beyond the exam hall. Students who worried that AI answers "will not align with what the lecture might have taught" were, in their own idiom, describing a source-authority problem: which account governs when two sources disagree? Learning to adjudicate that question, by returning to the primary source, weighing provenance, and reconciling discrepancies, is precisely the judgement journalism demands, and it is a judgement that the mere act of prompting a model does nothing to build. The data thus suggest that verification should be taught not as defensive "study hygiene" against a flawed tool, but as the positive, transferable professional competence at the heart of the discipline; the same skill that catches a hallucinated case in a media-law summary is the one that will later catch an unsourced claim in a press release. Positioned this way, teaching verification is not a concession to the presence of AI but an opportunity to teach journalism itself through a medium that students are already using every day.

Implications for journalism education

Three implications follow, each aimed at converting the value-action gap into practice. First, and most directly, journalism programmes should treat AI verification as an explicit, assessed skill rather than an assumed one. Given the substantial proportion of students across all cohorts who do not consistently verify AI output against authoritative sources, this critical skill cannot be left to develop implicitly; it should be taught and graded. Concretely, coursework can require students to submit, alongside an AI-assisted piece, a short record of how each AI-derived claim was checked against a primary or lecturer-sanctioned source, turning verification into a visible, credit-bearing act. This directly rehearses professional method. While this recommendation draws on emerging, unrefereed preprint literature such as Mahmoud (2026), who studied engineering students completing open-book, take-home examinations, and Das et al. (2025) the underlying pedagogical principle translates directly to journalism, in AI-mediated assessment, verification and judgement are the competencies worth measuring.

Second, assessment design should be revisited in light of what AI now trivially produces. Where a conversational model can generate a passable answer to an essay prompt, the assessment no longer discriminates the skills the degree exists to build, and "AI-aware" designs that reward original



reporting, source triangulation, live interviewing, and demonstrable process are needed to develop and reveal those skills (Das et al., 2025). For journalism specifically, assessments anchored in original news-gathering, material a model cannot fabricate without being caught, both resist misuse and cultivate the discipline's core competence. Third, rather than banning tools that effectively the entire cohort already uses, programmes should teach a task-appropriate "AI stack," making explicit which tasks suit conversational models (conceptual explanation, first-draft structure), which suit video (procedural and production skills), and which demand unmediated human judgement (verification, source evaluation, ethical decision-making). This reframes the institutional posture from prohibition to guided integration, consistent with UNESCO's (2023) human-centred guidance and with the recognition that AI and algorithmic literacy are becoming baseline newsroom competencies rather than optional extras (Cools et al., 2024; Pavlik, 2023).

These implications carry particular force in the study's regional context. In settings where class sizes are large and individualised feedback is scarce, generative AI is attractive precisely because it supplies on-demand explanation and practice at no marginal cost, and evidence from comparable sub-Saharan contexts shows both very high uptake and an explicit institutional need for AI-literacy provision (Ajalo et al., 2025). The recommendation to teach verification as a credit-bearing skill is therefore not a luxury import from well-resourced systems but a low-cost, high-leverage intervention: it requires reframing existing assessment rather than new technology, and it turns a tool student already use into an instrument for building the professional judgement the region's newsrooms will demand.

Limitations and future research

The cross-sectional design captures a single moment in a fast-moving technological landscape and cannot establish how habits evolve within individuals over time. Crucially, the reverse-coded "I did not use AI tools" item produced internally inconsistent responses from 5 of the 65 participants (7.7%); some heavy users nonetheless rated it highly, indicating that it was misread or answered inattentively. Because a failed attention check can cast doubt on the reliability of those respondents' answers across the entire instrument, a sensitivity analysis was conducted with these $n = 5$ cases excluded. The analysis confirmed that the overall descriptive distributions, majorities, and correlation coefficients remained practically identical. Consequently, the full sample was retained for the primary analysis, though the specific reverse-coded item was excluded from the composite interpretation of purpose. To align with these documented methodological limitations, all claims of "universal" or "total" adoption were removed from the abstract and conclusion. Finally, the study measures perceived effectiveness rather than actual learning outcomes or examination performance, and perceptions and outcomes are known to diverge (Abbas et al., 2024).

These limitations map onto a clear research agenda. Longitudinal designs could track whether verification habits and reliance change as a cohort advances, testing the maturational question this study could only address cross-sectionally. Experimental or quasi-experimental studies could link AI-supported revision, and, specifically, verification-embedded assessment, to measured learning outcomes rather than perceptions. Comparative work across disciplines could test whether journalism students' acute awareness of verification norms translates into better practice than their peers, or whether, as the present data hint, awareness and behaviour remain stubbornly decoupled. Finally, observational or think-aloud methods during actual revision, capturing the moment a student chooses to verify or not, would illuminate the decisions that self-report cannot reach, extending to exam preparation the transcript-based approaches now being applied to AI-mediated assessment (Mahmoud, 2026; Seven et al., 2025).



Conclusion

Within this sample of 65 undergraduate students, adoption of generative AI and digital media was widespread and intensive. These students value the tools primarily for study support such as summarising notes and simplifying complex concepts and judge them to be highly effective, leading to a marked shift away from traditional revision materials. However, this fluency with the technology is not matched by equivalent critical literacy: a substantial proportion of the sample engages in sparse or inconsistent verification of AI outputs against authoritative sources, resulting in a pronounced value-action gap.

More broadly, the findings invite a shift in how educators frame the arrival of generative AI. The instinctive institutional response (to police, restrict, or detect) treats students as adversaries and the technology as contraband, a posture this study's data render both unrealistic and unnecessary. Unrealistic, because adoption is already widespread and perceived as effective; unnecessary, because students are not credulous consumers but ambivalent, self-critical ones who can already name the technology's dangers. What they lack is not awareness but structure: an assessed, teachable practice that channels their existing scepticism into consistent action. The most consequential move available to journalism programmes, therefore, is also among the least expensive: to make the verification of AI-assisted work a visible, credited, and repeatedly practised part of the curriculum. If they do, the generation of students described here need not choose between the efficiency they value and the professional integrity their field requires; they can be taught, through the very tools they already trust, to hold both at once.

References

- Abbas, M., Jam, F. A., & Khan, T. I. (2024). Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students. *International Journal of Educational Technology in Higher Education*, 21(1), 10. <https://doi.org/10.1186/s41239-024-00444-7>
- Ajalo, E., Mukunya, D., Nantale, R., Kayemba, F., Pangholi, K., Babuya, J., Akuu, S. L., Namiiro, A. M., Nsubuga, Y. B., Mpagi, J. L., Musaba, M. W., Oguttu, F., Kuteesa, J., Mubuke, A. G., & Munabi, I. G. (2025). Widespread use of ChatGPT and other artificial intelligence tools among medical students in Uganda: A cross-sectional study. *PLOS ONE*, 20(3), e0313776. <https://doi.org/10.1371/journal.pone.0313776>
- Alonso-Muñoz, L., & Casero-Ripollés, A. (2026). Adoption of generative AI in higher education: Perceptions of journalism students. *Information*, 17(2), 189. <https://doi.org/10.3390/info17020189>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Chen, L., & Chen, M. (2026). Technological integration and professional competency development: A structural equation analysis of generative AI, learning engagement, and media literacy in journalism education. *Frontiers in Education*, 11, 1812350.
- Cools, H., & Diakopoulos, N. (2024). Uses of generative AI in the newsroom: Mapping journalists' perceptions of perils and possibilities. *Journalism Practice*. Advance online publication. <https://doi.org/10.1080/17512786.2024.2394558>



- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228–239. <https://doi.org/10.1080/14703297.2023.2190148>
- Das, S., & Eliseev, A. (2025). Predicting ChatGPT use in assignments: Implications for AI-aware assessment design [Preprint]. *arXiv*. <https://arxiv.org/abs/2508.12013>
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kirschner, P. A., & De Bruyckere, P. (2017). The myths of the digital native and the multitasker. *Teaching and Teacher Education*, 67, 135–142. <https://doi.org/10.1016/j.tate.2017.06.001>
- Mahmoud, Q. H. (2026). Reimagining assessment in the age of generative AI: Lessons from open-book exams with ChatGPT [Preprint]. *arXiv*. <https://arxiv.org/abs/2605.12363>
- Ng, D. T. K., Tan, C. W., & Leung, J. K. L. (2024). Empowering student self-regulated learning and science education through ChatGPT: A pioneering pilot study. *British Journal of Educational Technology*, 55(4), 1328–1353. <https://doi.org/10.1111/bjet.13454>
- Pavlik, J. V. (2023). Collaborating with ChatGPT: Considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator*, 78(1), 84–93. <https://doi.org/10.1177/10776958221149577>
- Prensky, M. (2001). Digital natives, digital immigrants. *On the Horizon*, 9(5), 1–6. <https://doi.org/10.1108/10748120110424816>
- Rudolph, J., Tan, S., & Tan, S. (2023). ChatGPT: Bullshit spewer or the end of traditional assessments in higher education? *Journal of Applied Learning and Teaching*, 6(1), 342–363. <https://doi.org/10.37074/jalt.2023.6.1.9>
- Seven, M., Ozan, Ö., & Emirtekin, E. (2025). Prompting future journalists to prompt: An experiential study on GenAI, critical literacy, and reflective practice in data news. In *Proceedings of the 6th International Conference on AI Research (ICAIR 2025)* (Vol. 6, No. 1, pp. 681–685). Academic Conferences and Publishing International. <https://doi.org/10.34190/icaire.5.1.4358>
- Tian, W., Ge, J., Zhao, Y., & Zheng, X. (2024). AI Chatbots in Chinese higher education: adoption, perception, and influence among graduate students—an integrated analysis utilizing UTAUT and ECM models. *Frontiers in Psychology*, 15, 1268549. <https://doi.org/10.3389/fpsyg.2024.1268549>
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), 39. <https://doi.org/10.1186/s41239-019-0171-0>

***Appendix A: Supplementary summary of quantitative findings***

Measure	Result
Total valid responses	65
Frequent or exclusive users	69.2%
Most-used tool (ChatGPT)	66.2%
Second tool (Google Gemini)	35.4%
Video platforms (YouTube etc.)	23.1%
Mean tools used per student	1.69
Top purpose, explain concepts (mean)	3.92 / 5
Rated tools very/highly effective	70.8%
Mean perceived effectiveness	3.83 / 5
Reliance on traditional materials decreased	61.5%
Verify AI output always/often	58.5%
Verify AI output rarely/never	16.9%
Mean verification score	3.75 / 5
Effectiveness: consistent vs infrequent verifiers	4.05 vs 3.64
Correlation: frequency × reliance shift	$\rho = .36$
Correlation: effectiveness × verification	$\rho = .32$

Table A1. Consolidated quantitative results referenced throughout the paper.