The University of Texas at Austin
Center for Media Engagement
Moody College of Communication

AI-TAILORED NEWS FOR GEN Z AND BEYOND:

What We Learned About Journalistic AI Use, Detection, and Public Reaction

Ava Motes, Natalie (Talia) Jomini Stroud, Christian DeBrady, Naomi Levine, Yuting He, Josephine Lukito, Gina M. Masullo, and Ashwin Rajadesingan

SUMMARY

As newsrooms become increasingly interested in both artificial intelligence (AI) and innovative solutions to capture Gen Z audiences, the Center for Media Engagement explored how generative AI could be used to tailor news articles for younger readers and whether this strategy improves how people evaluate and learn from the news.

Although our findings show that using AI to rewrite news articles for young adult audiences has little impact on how both Gen Z and older readers perceive and learn from articles, a surprising pattern emerged. Many assumed that AI had been used to write the articles — even when we shared original, non-AI reporting. People across age demographics often missed the AI disclosure labels on tailored news articles. Members of Gen Z estimated greater AI use depending on the way AI was prompted and, when they did, they also rated the articles more negatively and learned less. This suggests generational differences in AI-detection abilities and serves as a cautionary tale for newsrooms seeking to disclose their use of AI through conventional news labels.

SUGGESTED CITATION:

Motes, A., Stroud, N. J., DeBrady, C., Levine, N., He, Y., Lukito, J., Masullo, G. M., & Rajadesingan, A. (April, 2026). AI-Tailored news for gen Z and beyond: What we learned about journalistic AI use, detection, and public reaction. Center for Media Engagement. https://mediaengagement.org/research/ai-tailored-news-gen-z-beyond

THE PROBLEM

News consumption is declining in the United States, particularly among members of Gen Z (the generation born between 1997 and 2012).[1] Although a majority of Gen Z and Millennials (the generation born between 1981 and 1996) are getting some amount of news from traditional outlets, they express little enjoyment of news,[2] as well as boredom and disconnection with news content.[3] This raises concerns about the future of the journalism industry and its ability to reach younger generations.

In the past decade, newsrooms and media companies have tried to attract young audiences through strategies such as investing in social media and video content.[4] Newsrooms are also thinking about how to provide Gen Z with news that is seen as relevant and interesting.

Advancements in generative AI are igniting interest in tailoring the news to different audiences, and there's good reason for this optimism. In the field of health communication, for example, research has shown that AI can be used to tailor messages for specific demographics,[5] sometimes outperforming human-created content.[6] And previous research has suggested that readers are drawn to news content about similar-aged people.[7] Generative AI can now be prompted to write in a youthful style, quickly revising traditional news content with the goal of making stories more engaging, comprehensible, and relatable for younger news audiences.

To test this approach, we prompted ChatGPT to rewrite published news articles about immigration and AI in two ways: (1) for Gen Z as a whole and (2) for 18-to-21-year-olds, the youngest subset of Gen Z readers.[8] Both prompts resulted in shortened, digestible articles, often with statements emphasizing the importance of the issue for young people. The Gen Z versions also included distinctly informal language and internet slang, such as "keeping it real," while the versions for 18-to-21-year-olds used more standard journalistic language in a streamlined format.

This research is part of our connective democracy initiative funded by the John S. and James L. Knight Foundation.

KEY FINDINGS

- **Articles rewritten by AI for young adult audiences are not evaluated more positively than traditional news articles.** Both Gen Z and older respondents had similar impressions of the articles, intentions to engage, and learning outcomes, regardless of which version of the article they saw.
- **AI disclosure labels on news stories are easily missed.** When asked what percentage of the article people thought was written by AI, viewing an article with an AI disclosure label made no difference in how much of the article people thought was written by AI.
- **Audiences perceive that journalists use AI, even when they do not.** Study participants perceived that a portion of the articles — including the original versions — were written with AI. Overall, only 14% of participants thought that no AI was used.
- **Older audiences estimate higher use of AI than Gen Z.** Older participants' impressions of what percentage of the article was written by AI were higher on average than those of Gen Z readers. On average, older participants estimated that just over half of the article that they read was written by AI across all versions, including the original version, whereas Gen Z readers estimated that just over a third of the article, on average, was written by AI.
- **There are generational differences in AI perception — Gen Z forms impressions about AI based on the article tone.** Older participants' perceptions of AI use did not vary significantly across the articles. Gen Z participants, however, detected significantly more AI when they read the Gen Z-targeted article compared to the 18-to-21-year-old-targeted article and the original article. This was the case even though both the Gen Z and 18-to-21-year-old-targeted articles explicitly noted that AI had been used.
- **Those detecting greater AI use are also less likely to favor, engage with, and learn from the article.** Across all study participants, there were negative correlations: The more AI people detected, the less favorably they felt about the article, the lower their intentions to engage with it, and the lower their score on a quiz of what they learned from the article.
- **Attitudes towards the use of AI in journalism without human oversight are relatively unfavorable**. However, there is more support for AI use with human oversight, especially among older participants.

TAKEAWAYS FOR NEWSROOMS

- **Use caution with AI personalization.** Although AI can tailor language for younger audiences, this approach is not necessarily effective, and could even backfire, as readers tended to rate articles they thought had higher AI involvement less favorably than articles thought to have lower AI involvement.
- **Prioritize substance over style.** Audiences' impressions of news content may be more affected by their perceptions of the media and the specific topics covered than by the way articles are written. Featuring Gen Z people and perspectives in stories may be more effective at engaging young adult readers than using youthful language.
- **Consider new avenues for disclosure.** Readers of all ages frequently missed conventional AI disclosure labels, meaning more attention-grabbing alternatives may be necessary to communicate information about AI's involvement in news production.
- **Clearly explain where AI is being used and where it isn't.** While younger generations intuit the use of AI based on text features, older generations may default to assuming some AI is being used, regardless of how news content looks or sounds. Be clear about your newsrooms' specific AI policies and boundaries. Here, work by Poynter and Trusting News could be useful as starting points.

FULL FINDINGS

To understand how people react to articles revised by AI (and then checked by human editors), we conducted an online experiment with 1,007 participants. We purposefully quota sampled 504 Gen Z respondents (those born in or after 1997) and 503 non-Gen Z respondents (those born before 1997).

Roughly one-third of respondents were assigned to each article style: the original article, the article revised by AI for 18-to-21-year-olds, and the article revised by AI for Gen Z. We also tested four different articles — two about AI and two about immigration. We did this to look for general trends, as opposed to looking at how people react to one specific article or topic. The AI-revised articles for 18-to-21-year-olds and Gen Z were accompanied by a disclosure statement that clearly indicated AI had been used.

Example Article: Gen Z Prompt with AI Disclosure Statement

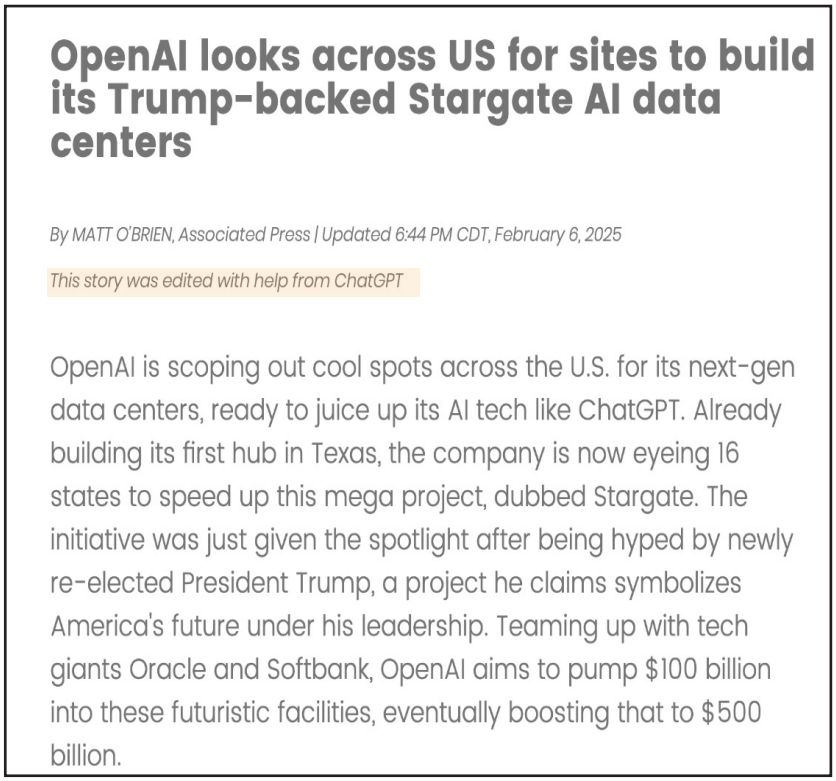

OpenAI looks across US for sites to build its Trump-backed Stargate AI data centers

By MATT O'BRIEN, Associated Press | Updated 6:44 PM CDT, February 6, 2025

This story was edited with help from ChatGPT

OpenAI is scoping out cool spots across the U.S. for its next-gen data centers, ready to juice up its AI tech like ChatGPT. Already building its first hub in Texas, the company is now eyeing 16 states to speed up this mega project, dubbed Stargate. The initiative was just given the spotlight after being hyped by newly re-elected President Trump, a project he claims symbolizes America's future under his leadership. Teaming up with tech giants Oracle and Softbank, OpenAI aims to pump $100 billion into these futuristic facilities, eventually boosting that to $500 billion.

Center for Media Engagement

Note: AI disclosure statement was not highlighted in the experiment.

AI Detection

We asked participants to tell us what percentage of the words in the article they thought were written by an AI editor, rather than by a human journalist.

Older participants estimated that roughly half (51.5%) of the words were written by AI, regardless of which version they saw. Gen Z, across all the conditions, estimated that just over a third (36.2%) of the article, on average, was written by AI.

Estimates of how much of the article was written by AI also differed depending on which version of the article they saw. Gen Z participants reading the Gen Z version of the AI-rewritten article detected significantly more AI involvement than Gen Z participants reading the version of the AI-rewritten article for 18-to-21-year-olds or the original article.[9]

There is no evidence that participants were affected by the presence of an AI label, which was present for the Gen Z and 18-to-21-year-old versions (and not for the original version, where no AI was used). Older participants reported a similar level of AI involvement for all the articles. Gen Z participants detected more AI in the Gen Z version — but not for the similarly labeled 18-to-21-year-old version.

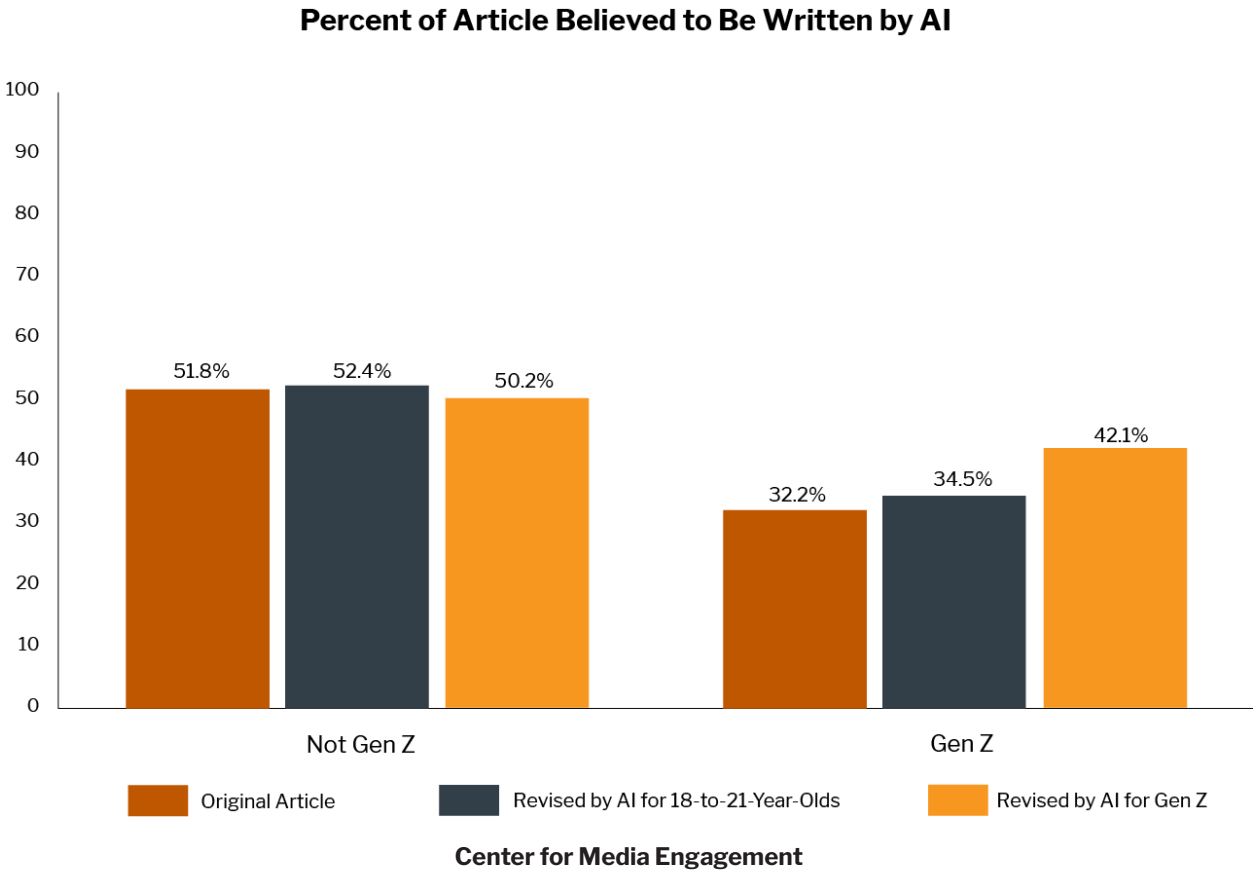

Notes: Participants were asked to report what percentage (0 to 100) of the words in the article they thought were written by an AI editor, rather than a human journalist.

Readers estimating higher levels of AI involvement in the articles also tended to view those articles less favorably and reported less interest in engaging with them. They also tended to know less about the contents of those articles.[10] While these were correlations, rather than causal relationships, we did detect effects in the experimental results. Younger audiences seeing the Gen Z AI-rewritten article (vs. the original article) detected more AI and, in turn, had less favorable impressions of the article, expressed fewer intentions to share it, and learned less from it.[11]

Attitudes Toward AI in Journalism

At the end of the study, we asked participants to rate how acceptable they found several newsroom uses of AI, from completely unacceptable to completely acceptable. These perceptions were unrelated to the article that people read.[12]

Participants had strong reactions about how newsrooms should use AI. As this was not a random, nor nationally representative, sample, exact results may differ when applied more broadly, but we expect the following patterns to hold:

- **Participants had negative reactions to newsroom use of AI without a human journalist reviewing the work**. Eighty-four percent of participants thought that it was "unacceptable" or "completely unacceptable" for a news organization to use AI to write news articles without a human journalist reviewing them. People also had a negative reaction to newsrooms using AI to edit news articles without human review, with 80% finding this "unacceptable" or "completely unacceptable."
- **Participants had more positive reactions toward the use of AI when a human journalist reviewed the work**. Nearly 50% thought it was "acceptable" or "completely acceptable" to have AI write news articles with a human reviewing afterward. AI editing with human review yielded the most positive response, with more than 65% finding this "acceptable" or "completely acceptable."

There were also interesting differences in attitudes by age.

- **Although both Gen Z and older participants had negative reactions to newsroom use of AI without a human journalist reviewing the work, Gen Z had more negative reactions.**
    - Whereas 54% of older participants thought that it was "acceptable" or "completely acceptable" for news organizations to use AI to write news articles with a human journalist reviewing them afterward, only 43% of Gen Z participants shared this view.
    - Whereas 72% of older participants thought it was "acceptable" or "completely acceptable" for news organizations to use AI to edit news articles with a human journalist reviewing them afterward, only 61% of Gen Z participants felt the same.

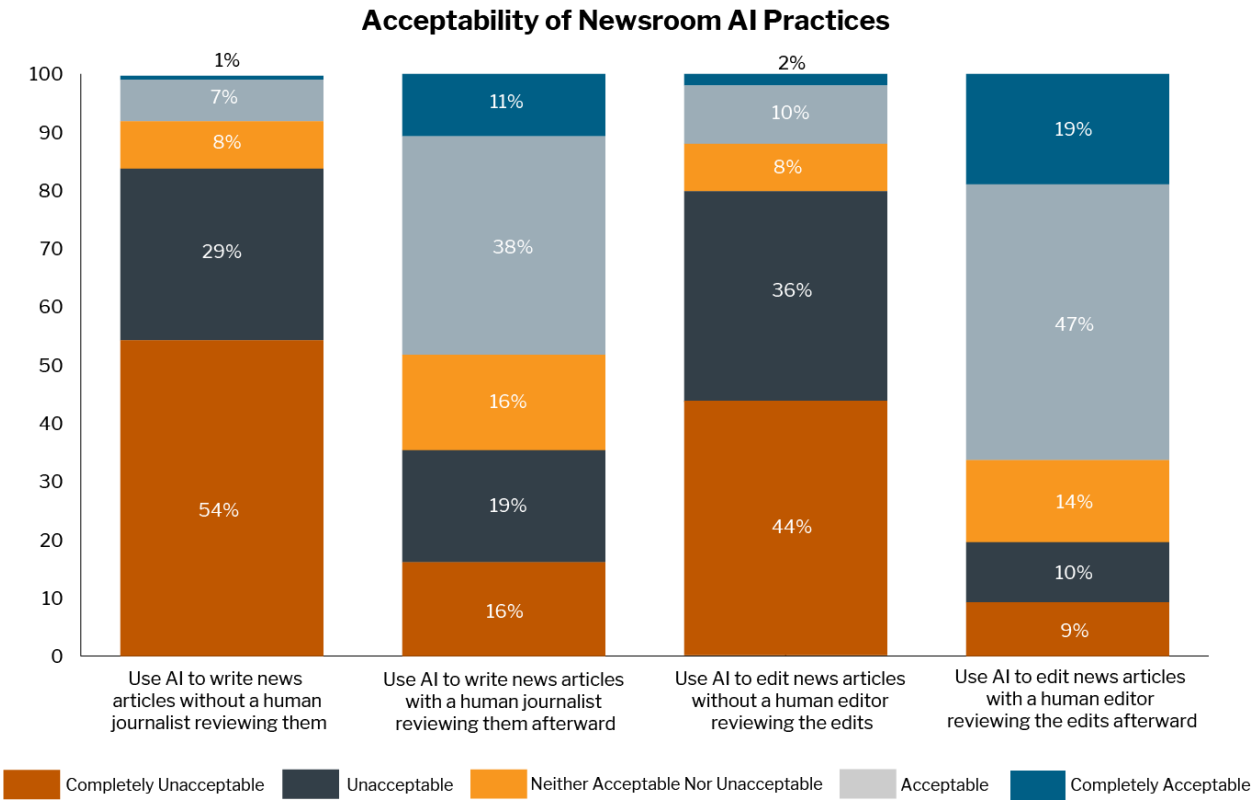

Center for Media Engagement

Notes: Participant attitudes towards newsrooms engaging in various AI practices w assessed using Likert-type options ranging from completely unacceptable (1) to completely acceptable (5).

Article Perceptions and Engagement Intentions

We asked people what they thought of the article after they read it. Specifically, we asked participants to tell us whether they thought the article was accurate, trustworthy, and enjoyable to read, among other descriptors. Overall, participants had moderately positive attitudes about the articles, rating them a 3.5 on average on a measure from 1 to 5, with higher values indicating more favorable attitudes.

Participants did not evaluate the article differently depending on whether they saw the original article, the 18-to-21-year-old AI version, or the Gen Z AI version. This was true of both Gen Z participants and older participants.[13]

Efforts to entice younger audiences through AI-generated article revisions did not change Gen Z's perceptions of, or learning from, the articles. These revisions also had no impact on how older audiences evaluated and what they learned from the articles.

When we asked people whether they would be interested in reading another article, recommending the article, or sharing the article, the results were similar: study participants reported the same level of interest regardless of whether they saw the original article, the AI-rewritten version for 18-to-21-year-olds, or the AI-rewritten version for Gen Z (3.1 on a measure from 1 to 5, with 5 indicating more agreement). This was again true for both Gen Z and older participants.[14]

Recall and Recognition

We analyzed what people knew after reading the articles. We wanted to know whether people came away with different levels of knowledge depending on the article that they read.

We asked participants to list any details from the article that they could remember reading. The number of accurate details that participants recalled was not affected by which version of the article they read.[15]

We then gave participants a six-question, multiple-choice quiz about facts found in the articles. Again, the number of correct responses was not affected by which version of the article people read.[16]

METHODOLOGY

Stimuli

We selected four different articles: one about building Stargate AI data centers from The Associated Press; one comparing DeepSeek and ChatGPT from Salon; one about immigration authorities being permitted to arrest migrants at churches, schools, and hospitals from The Guardian; and one about creating a registry for undocumented immigrants from The Associated Press. We then prompted ChatGPT to rewrite each original news article for both a Gen Z and an 18-to-21-year old audience and fact-checked the AI rewrites to ensure consistency with the original text. To protect the copyright of the news articles used as stimuli in this study, we disabled the "improve the model for everyone" data-sharing option in ChatGPT. As a result, neither the original news content nor the generated news versions were used for model training. The result was 12 potential article conditions for participants to evaluate.

Procedure

We recruited participants using Prolific, an online platform that helps connect researchers with people who are willing to participate in research for a small fee. Participants were limited to U.S. residents aged 18 or older, and we used a quota sample to ensure that half of the participants were members of Gen Z (born in or after 1997) and half were non-Gen Z (born before 1997). The resulting sample (n = 1,007) was neither random nor representative of the population at large.

After providing informed consent, participants were asked which types of media they used to stay up-to-date on current events. They were then randomly assigned to see one of the 12 article conditions. After reading their assigned article, participants were asked several questions to assess their perception of the article, intention to engage, learning outcomes, perceived use of AI in the text, and attitudes towards AI in journalism.

Following completion of the study, participants were paid $3 for their participation. Data was collected between July and September, 2025. The study participants' demographics are included below.

Demographics

Age

Range	18 - 81 years old
Mean	35.3 years old
Gen Z	50.05%
Non-Gen Z	49.95%

Race

White	70.9%
Black	15.0%
Asian	9.9%
Other	4.2%

Ethnicity

Hispanic/Latino	8.3%
Not Hispanic/Latino	91.7%

Education

High School Diploma or Less	15.6%
Some College or Associate's Degree	31.2%
Bachelor's Degree	40.1%
Graduate or Professional Degree	12.4%

Partisanship

Democrat	54.1%
Republican	32.5%
Independent	13.4%

Gender

Male	49.6%
Non-male	50.4%

Center for Media Engagement

ENDNOTES

[1] See Forman-Katz, N. (2025, October 24). *Americans are following the news less closely than they used to*. Pew Research Center. https://www.pewresearch.org/short-reads/2023/10/24/americans-are-following-the-news-less-closely-than-they-used-to/ and U.S. Census Bureau. (2025, May 6). Birth cohorts geographic mobility report. *U.S. Census Bureau.* https://www.census.gov/newsroom/press-releases/2025/birth-cohorts-geographic-mobility-report.html

[2] The Media Insight Project. (2022). *Fatigue, Traditionalism, and Engagement: The News Habits and Attitudes of the Gen Z and Millennial Generations*. American Press Institute. https://americanpressinstitute.org/wp-content/uploads/2022/08/MIP-Report-Release-1-Final-8.29.pdf

[3] Tamboer, S. L., Kleemans, M., & Daalmans, S. (2022). 'We are a neeeew generation': Early adolescents' views on news and news literacy. *Journalism, 23*(4), 806-822. https://doi.org/10.1177/1464884920924527

[4] Poynter. (2021). *VidSpark*. https://www.poynter.org/vidspark

[5] Chu, H., & Liu, S. (2025). Generating targeted and tailored health communication narratives with AI. *Risk Analysis, 45*(11), 3505-3518. https://doi.org/10.1111/risa.70076

[6] Leung, J., Sun, T., Stjepanović, D., Vu, G., Yimer, T., Connor, J. P., Hall, W., & Chan, G. C. K. (2025). Generative artificial intelligence with youth codesign to create vaping awareness advertisements. *JAMA Network Open, 8*(7), e2514040. https://doi.org/10.1001/jamanetworkopen.2025.14040

[7] Knobloch-Westerwick, S., & Hastall, M. R. (2010). Please your self: Social identity effects on selective exposure to news about in-and out-groups. *Journal of Communication, 60*(3), 515-535. https://doi.org/10.1111/j.1460-2466.2010.01495.x

[8] All interactions with ChatGPT (GPT-4o) in this study were conducted with the "improve the model for everyone" data-sharing option disabled; therefore, none of the conversations and news articles were used for model training.

[9] Participants were asked to report what percentage of the words in the article they thought were written by an AI editor, rather than a human journalist. They responded by entering a number between 0 and 100 in a textbox. On average, participants estimated that 44% of the words in the article they read were written by AI (SD = 31). Although the main effect of the article version did not significantly predict AI detection (F(2,1000) = 1.84, p = .16), the interaction between article version and Gen Z/older did (F(2, 998) = 3.83, p = .02).

[10] The higher the percentage of AI detected in the article, the less favorable the article perceptions (r = -.25), the lower the engagement intentions (r = -.20), and the lower the recognition outcome (r = -.16).

[11] A mediation analysis showed a significant indirect effect of the Gen Z condition versus the original article condition on article perceptions (95% CI [-0.11, -0.02]), engagement (95% CI [-0.13, -0.02]) and recognition (95% CI [-0.14, -0.02]) through AI detection for Gen Z participants.

[12] Responses were recorded using Likert-type options ranging from completely unacceptable (1) to completely acceptable (5). The four measures were averaged to create an AI in journalism attitude index (M = 2.56, SD = 0.89, Cronbach's alpha = 0.77). Attitudes did not vary by article version (F(2,1000) = 1.97, p = .14) or by the interaction of article version and Gen Z/older (F(2, 998) = 0.08, p = .92).

[13] To measure how positively or negatively participants perceived the articles, they were asked to indicate how much a list of terms, including: accurate, informative, trustworthy, confusing, biased, interesting, enjoyable to read, and engaging, applied to the article they read using Likert-type response options ranging from not at all (1) to a great deal (5). The items "confusing" and "biased" were reverse-coded so that higher values consistently reflected more positive perceptions. These items showed high reliability (Cronbach's alpha = 0.83) and were averaged to create an article perception index (M = 3.48, SD = 0.70). In an ANOVA predicting article perceptions, article version had no effect (F(2, 1000) = 0.99, p = .37) nor did the interaction between article version and Gen Z/older (F(2, 998) = 0.39, p = .68).

[14] To assess participants' intention to engage with the articles, they were asked to indicate how much they agree or disagree with a list of statements, including: I would be interested in reading another article from this news outlet, I would recommend this article to others, and I would like to share this article with others. Responses were recorded using Likert-type options ranging from strongly disagree (1) to strongly agree (5). These items were then averaged to create an engagement intention index (M = 3.10, SD = 1.10, Cronbach's alpha = 0.90). Engagement intentions did not vary based on the article version ($F(2, 1000) = 0.18$, $p = .84$) nor based on the interaction between article version and Gen Z/older ($F(2, 998) = 0.66$, $p = .52$).

[15] Participants were asked to list any details from the article that they could remember up to 10 (only 115 participants completed all 10 fields, on average, participants left 5.26 details, (SD = 2.33). A codebook was iteratively built from the participants' responses to capture each detail mentioned. We conducted a reliability analysis with 13% of 5,294 open-ended responses (up to 10 fields per participant) and a second coder who was not involved in the codebook development (Krippendorff's alpha = 0.87). The details were then reviewed for factual accuracy based on the article, a reliability analysis of 665 coded responses yielded a Krippendorff's alpha of 0.80. To compute the recall measure, we summed the number of unique details per participant and subtracted any detail that was invalid, irrelevant or inaccurate (M = 3.81, SD = 1.74, Range = 0 to 9). Recall did not vary based on the article version ($F(2, 1000) = 1.24$, $p = .29$) nor based on the interaction between article version and Gen Z/older ($F(2, 998) = 1.80$, $p = .17$).

[16] To gauge recognition of factual details, respondents were asked six multiple-choice questions specific to the article they were assigned, each with one correct answer. Responses were scored as a 1 for a correct answer and a 0 for an incorrect answer or non-response. Individual scores were then summed across all questions to create a total recognition score for each participant (M = 5.03, SD = 1.14). Recognition did not vary based on the article version ($F(2, 1000) = 1.81$, $p = .16$) nor based on the interaction between article version and Gen Z/older ($F(2, 1000) = 0.01$, $p = .99$).