

DOI: https://doi.org/10.48009/1_iis_2026_143

From Curiosity to Integration: A Repeated Cross-Sectional Study of Student Perceptions and Usage of Artificial Intelligence in Higher Education (2024–2026)

Carl M. Rebman Jr., *University of San Diego*, carlr@sandiego.edu

Steven Levkoff, *University of California San Diego and University of San Diego*, slevkoff@ucsd.edu, slevkoff@sandiego.edu

Julie Allen, *Middle Georgia State University*, julie.allen@mga.edu

Hayden Wimmer, *Georgia Southern University*, hwimmer@georgiasouthern.edu

Queen E. Booker, *Metropolitan State University*, queen.booker@metrostate.edu

Abstract

Artificial intelligence (AI) has progressed rapidly from curiosity-driven, experimental technology into an increasingly integrated component of student academic, professional, and personal workflows. While early discussions surrounding generative AI focused primarily on novelty, awareness, and concerns regarding academic disruption, more recent conversations have shifted toward routine adoption, workflow integration, institutional governance, verification behavior, and dependency concerns. This study re-administered a 2024 survey instrument examining student perceptions of artificial intelligence to a new cohort of students at the same institution in 2026, producing a repeated cross-sectional comparison spanning a period of rapid AI normalization (2024 $n = 233$; 2026 $n = 217$). Between-wave differences were tested with Welch t-tests, reported with effect sizes, and corrected for multiple comparisons using the Holm procedure. After correction, one difference among the ten Likert-type items remained statistically significant: students in 2026 reported significantly lower agreement that AI helps manage time more efficiently ($d = -0.30$), a moderation of early productivity optimism that runs counter to prevailing productivity-gain narratives. The distribution of reported emotion toward AI also shifted significantly between waves ($\chi^2(3) = 17.52, p < .001$): curiosity remained most common but declined from 54% to 40% of respondents, while fear nearly doubled from 18% to 33%. Students reported frequent AI use, high dependency on digital devices, moderate verification behavior, and a gap between understanding institutional AI policies and deriving confidence from them. The findings suggest that student interaction with AI has shifted from curiosity and experimentation toward normalized operational integration, accompanied by more tempered efficiency expectations and emerging governance and AI-literacy challenges.

Keywords: artificial intelligence (AI), generative artificial intelligence, AI in education, ChatGPT, student perceptions, AI governance

Introduction

The term and concepts of artificial intelligence (AI) have existed for decades, although widespread public awareness and practical adoption have accelerated substantially over the last several years. The release of generative AI technologies such as ChatGPT in late 2022 significantly altered both public and academic discussions surrounding artificial intelligence. What was once viewed primarily as an emerging or

experimental technology has increasingly become integrated into educational, workplace, and consumer environments.

The rapid rise of generative AI systems has fundamentally altered how students interact with information, complete assignments, conduct research, and engage with digital technologies. Large language models (LLMs) are now capable of producing human-like written responses, summarizing information, generating computer code, and assisting with a wide variety of academic and professional tasks. Unlike earlier educational technologies that often required specialized technical knowledge, modern generative AI systems are accessible through conversational interfaces that can be used by students with little or no technical background. Research on artificial intelligence in education (AIED) anticipated substantial expansion of computer-assisted education well before the generative AI era (Zawacki-Richter et al., 2019; Zhang & Aslan, 2021), and market analyses conducted at the time of the original 2024 study projected rapid growth in AI-related educational technology.

Figure 1 places the two waves of the present study in that market context. Rather than reproducing successive third-party charts, the figure redraws the headline projections from two vintages of the same industry analysis: the projection available at the time of the 2024 baseline study and the projection available at the time of the 2026 replication. Between the two study waves, the projected size of the global AI-in-education market was revised substantially upward (Grand View Research, 2022, 2026), indicating that the growth anticipated during the baseline period was not merely speculative.




Figure 1. Global AI-in-Education Market Projections at the Time of Each Study Wave.

Note. Values redrawn by the authors from Grand View Research report vintages published in 2022 and 2026 (retrieved July 2026). The 2022-vintage series reflects a projection of USD 2.75 billion (2022) growing to USD 32.3 billion (2030); the 2026-vintage series reflects USD 8.3 billion (2025) and USD 11.4 billion (2026) growing to USD 57.2 billion (2033).

Figure 2 situates the two survey waves against the adoption curve of the most widely used generative AI application. The 2024 wave was administered when ChatGPT reported on the order of one to two hundred million weekly active users; by the 2026 wave, that figure had grown to approximately 900 million (Hu, 2023; OpenAI disclosures as reported by Reuters and TechCrunch, 2023–2026). The two cohorts in this study therefore encountered generative AI under markedly different conditions of social saturation, which is precisely the change the repeated cross-sectional design is intended to capture.




Figure 2. ChatGPT Weekly Active User Milestones with the Two Survey Waves Overlaid.

Note. Milestones redrawn by the authors from OpenAI public disclosures as reported by Reuters and TechCrunch (retrieved July 2026). Shaded bands indicate the approximate administration windows of the 2024 and 2026 survey waves.

The rapid adoption of generative AI has also forced universities to reconsider long-standing assumptions regarding academic integrity, assessment design, information literacy, and acceptable technology use. Institutions that initially responded to generative AI with uncertainty or restrictive policies have increasingly shifted toward governance frameworks focused on responsible use, disclosure, and AI literacy. At the same time, public opinion remains cautious. As Figure 3 shows, education is among the domains in which U.S. adults are least optimistic about AI's effects: only about one quarter of adults expect a positive effect of AI on education, roughly half the share who expect a positive effect on medical care (Pew Research Center, 2025).




Figure 3. Share of U.S. Adults Expecting a Positive Effect of AI, by Domain.

Note. Redrawn by the authors from Pew Research Center (2025) survey data (retrieved July 2026). Related polling indicates that half of U.S. adults report feeling more concerned than excited about the growing use of AI in daily life.

The original 2024 study examined student perceptions regarding artificial intelligence during the early stages of widespread generative AI adoption. At that time, many students expressed curiosity regarding AI technologies while also reporting ethical concerns and uncertainty regarding future impacts. Since then, generative AI systems have become substantially more embedded in everyday student workflows and educational environments. The present study revisits and expands upon the original research by re-administering the 2024 survey instrument to a new cohort of students and introducing additional questions related to AI usage behaviors, dependency on digital devices, verification practices, and institutional AI policies.

Research Gap and Research Questions

While prior studies have explored student perceptions regarding artificial intelligence in countries such as Estonia, Turkey, Romania, England, and India, there remains limited research comparing how student perceptions and usage behaviors have changed across the period of rapid mainstream generative AI adoption in the United States using the same instrument at the same institution. Much of the earlier research focused on initial awareness, curiosity, and broad perceptions of AI rather than operational integration, dependency behaviors, verification practices, institutional governance, and policy clarity. This study addresses that gap through a repeated cross-sectional comparison of two independent student cohorts surveyed with the same instrument in 2024 and 2026, extended with new measures of AI usage frequency, digital dependency, information verification, and institutional AI policy awareness. The study addresses the following research questions:

RQ1: *How did student perceptions of artificial intelligence differ between the 2024 and 2026 survey waves?*

RQ2: *How frequently do students report using AI technologies for information retrieval and problem solving, and how dependent do they report being on digital devices?*

RQ3: *To what extent do students report verifying AI-generated information using alternative sources?*

RQ4: *How do students perceive institutional AI policies, and is understanding of those policies accompanied by greater confidence in using AI?*

RQ5: *Do student perceptions of AI differ by gender, GPA, or year of study?*

Literature Review

Artificial intelligence in education continues to represent one of the most rapidly evolving areas of educational technology research. Earlier studies examining AI in higher education often focused on intelligent tutoring systems, adaptive learning systems, and computer-assisted instruction (Zawacki-Richter et al., 2019; Zhang & Aslan, 2021). However, the introduction of generative AI technologies dramatically expanded both the accessibility and functionality of AI systems within educational environments.

Keles and Aydin (2021) examined university student perceptions regarding artificial intelligence and found that students often focused on negative aspects of AI applications despite recognizing potential educational benefits. Similarly, Ardelean and Veres (2023) surveyed university students in Romania and found that respondents expressed both optimism and concern regarding AI technologies. Their findings suggested that curiosity remained one of the dominant emotional responses toward AI while concerns regarding workforce replacement and societal impact persisted. Kumar and Raman (2022) found that students generally perceived AI positively within educational contexts, particularly when AI was used to support learning and instructional activities, and Leoste et al. (2021) found that students and instructors frequently expressed optimistic attitudes toward emerging technologies and AI applications in education. These findings were

important during the early adoption stage because they suggested that students were not necessarily opposed to AI; rather, many were curious, cautiously optimistic, and still attempting to understand its implications.

More recent literature has shifted toward operational AI use, productivity implications, and dependency concerns. Noy and Zhang (2023) reported substantial productivity improvements associated with generative AI-assisted writing tasks in a controlled experiment, finding both reduced completion time and improvements in output quality. Firaina and Sulisworo (2023) and Fauzi et al. (2023) similarly found that students frequently perceived generative AI systems as beneficial for improving efficiency and assisting with information retrieval.

At the same time, researchers have expressed concern regarding the long-term implications of dependency on generative AI technologies. Nah et al. (2023) and Meyer et al. (2023) discussed concerns regarding excessive trust in AI-generated responses despite the known risks associated with hallucinations and misinformation, and Rahman and Watanobe (2023) argued that AI systems provide substantial opportunities for learning support while also creating potential risks associated with overreliance and reduced critical thinking engagement. Institutional governance has likewise become an increasingly important research theme: Chan (2023) proposed a comprehensive framework for AI policy development within higher education institutions and argued that universities must provide clearer guidance regarding acceptable AI usage.

Since the baseline wave of the present study, several large-scale student surveys have documented the pace of generative AI normalization. In the United Kingdom, the Higher Education Policy Institute found that the share of undergraduate students using AI tools in some form rose from 66% in 2024 to 92% in 2025, with use in assessment preparation rising from 53% to 88% over the same year (Freeman, 2024, 2025). The Digital Education Council's (2024) global survey similarly reported that 86% of students use AI in their studies, with 54% using it at least weekly. U.S.-focused reporting, including the EDUCAUSE Students and Technology research (EDUCAUSE, 2024) and the Tyton Partners Time for Class series (Tyton Partners, 2025), documents comparable normalization alongside persistent institutional-policy ambiguity.

Provider-side usage research, including the Anthropic Education Report on student use of Claude (Anthropic, 2025) and large-scale analyses of ChatGPT usage (Chatterji et al., 2025), further confirms that information seeking, explanation, and writing assistance dominate student AI use. These studies are, however, predominantly single-wave usage snapshots or track usage rates rather than perceptions. The present study complements them by comparing perceptions measured with the same instrument, at the same institution, at two points spanning the normalization period, and by pairing those perception measures with new items on dependency, verification, and institutional policy.

Methodology

Research Design

This quantitative study employed a repeated cross-sectional design: the survey instrument used in the original 2024 study was re-administered in 2026 to a new, independent cohort of students at the same institution. The design is explicitly not a longitudinal panel; the same individuals were not tracked over time. Consequently, between-wave differences reflect a combination of period effects (the general saturation of generative AI between 2024 and 2026), cohort composition, and instructional changes, and cannot be attributed to within-person attitude change. This constraint is revisited in the limitations section. The repeated cross-sectional design nonetheless permits a direct comparison of how successive cohorts at the same institution, in the same courses, perceived AI at two points spanning a period of rapid normalization.

Sample and Procedure

The survey was administered online via Qualtrics to students enrolled in Information Systems courses at a private university in the western United States during the 2024 and 2026 terms in which the participating courses were offered. All students enrolled in participating course sections were invited to complete the survey; participation was voluntary and anonymous, no personally identifying information was collected, and students received a small amount of extra credit for completion. After screening (described below), the 2024 wave yielded 233 usable responses and the 2026 wave yielded 217, for a combined analytic sample of 450. Across both waves, 242 respondents identified as male and 198 as female, with the remaining 10 selecting another option or not responding; gender-based analyses use the male and female groups. Although the participating courses primarily enroll freshman and sophomore students, upper-division students also participated: of the 364 respondents reporting year of study, 68% were freshmen or sophomores and 32% were juniors, seniors, or other.

Instrument

The instrument consisted of the eleven original items used in the 2024 study (Table 1) plus six new items added in 2026 (Table 2). Items primarily used 5-point Likert scales ranging from Strongly Disagree to Strongly Agree; the informedness and digital dependency items used 10-point scales. The original items examined student awareness of AI, emotional responses toward AI, ethical concerns, perceptions regarding productivity and societal impact, and trust in AI-generated information. The new 2026 items examined frequency of AI usage, dependency on digital devices, comparison of AI-generated information against alternative sources, and institutional AI policy awareness. Items were derived from prior AI-perception and educational technology studies including Jeffrey (2020), Keles and Aydin (2021), Ardelean and Veres (2023), Chan (2023), Rahman and Watanobe (2023), Noy and Zhang (2023), and Pesovski et al. (2024).

Table 1. Items Administered in Both the 2024 and 2026 Waves

#	Item	Source
1	On a scale of 1 to 10, how informed are you about the concept of artificial intelligence?	Jeffrey, 2020
2	What sources do you use to learn about the concept of artificial intelligence?	Ardelean & Veres, 2023
3	AI encourages dehumanization.	Ardelean & Veres, 2023
4	Robots will replace people at work.	Ardelean & Veres, 2023
5	The ethical use of AI is a major concern for me.	Jeffrey, 2020
6	AI helps with managing time more efficiently.	Chan, 2023
7	AI will impact the presidential election.	Created by authors
8	AI helps to solve many problems in society more efficiently.	Keles & Aydin, 2021
9	When you think about AI do you feel Curiosity, Fear, Indifference, Trust?	Ardelean & Veres, 2023
10	AI can contribute positively to societal issues.	Ardelean & Veres, 2023
11	I trust AI algorithms to provide accurate information.	Ardelean & Veres, 2023

Table 2. New Items Administered in the 2026 Wave Only

#	Item	Source
1	How often do you use AI tools (e.g., virtual assistants, recommendation algorithms) to find information or solve problems?	Noy & Zhang, 2023
2	On a scale of 1 to 10, how dependent are you on digital devices for day-to-day tasks and information retrieval?	Rahman & Watanobe, 2023
3	I compare multiple sources of information before forming an opinion based on AI recommendations.	Pesovski et al., 2024
4	My institution has clearly communicated its policies regarding the acceptable use of AI tools.	Chan, 2023
5	I understand when the use of AI is permitted or prohibited in my institution.	Chan, 2023
6	Institutional AI policies make me more confident using AI technologies.	Chan, 2023

Data Screening and Analysis

Raw Qualtrics exports were screened to remove test entries and erroneous submissions; internet-capture metadata (e.g., IP addresses and coordinates) was removed prior to analysis. Item-level missing responses were retained and handled on a per-test basis. Two 2026 responses on the digital dependency item fell outside the 1–10 scale range and the corresponding respondents were excluded from analyses of the 2026-only items. Statistical analyses were conducted in R (t.test, psych, openxlsx, officer); figures were produced in Python (matplotlib). Between-group comparisons used Welch (unequal-variances) independent-samples t-tests, reported with the test statistic, Welch degrees of freedom, exact p-values, and Cohen’s d computed with the pooled-standard-deviation denominator. Because each family of comparisons involved multiple simultaneous tests, family-wise error was controlled with the Holm procedure applied within each family (ten between-wave tests; ten gender tests; eight GPA tests; eight year-of-study tests). Associations with GPA, recorded as an ordinal category, were tested with Spearman rank correlations. Year-of-study differences were examined with Welch one-way ANOVA across reported categories and a planned contrast comparing underclass (freshman/sophomore) with upperclass (junior and above) respondents. Between-wave differences in the distributions of the categorical items (sources for learning about AI and reported emotion toward AI) were tested with chi-square tests of independence, with Cramér’s V reported as the effect size. For the six new 2026 items, internal consistency was estimated with Cronbach’s alpha and the factor structure was examined with a principal-axis exploratory factor analysis. A complete reproduction script covering every statistic reported in this paper accompanies the manuscript and is available from the authors.

Results

Between-Wave Comparisons (RQ1)

Table 3 reports the full between-wave results, including means, standard deviations, Welch test statistics, uncorrected and Holm-corrected p-values, and effect sizes. After Holm correction, one difference remained statistically significant: agreement that AI helps with managing time more efficiently declined from 2024 ($M = 3.91$, $SD = 0.92$) to 2026 ($M = 3.60$, $SD = 1.19$), $t(407.2) = 3.12$, corrected $p = .019$, $d = -0.30$. Two differences that were nominally significant before correction did not survive it and should be interpreted only as suggestive: agreement that AI encourages dehumanization was nominally higher in 2026 (uncorrected $p = .024$, $d = 0.22$, corrected $p = .211$), and agreement that AI helps solve problems more efficiently was nominally lower (uncorrected $p = .045$, $d = -0.19$, corrected $p = .359$). Self-reported informedness about AI trended upward ($d = 0.18$) but did not reach significance before or after correction

(uncorrected $p = .059$). No other between-wave differences among the Likert items approached significance. Figure 4 summarizes the means for the 5-point items by wave.

Table 3. Between-Wave Comparisons of AI Perception Items, 2024 vs. 2026 (Welch t-tests)

Variable	2024 M (SD)	2026 M (SD)	t	df	p	p (Holm)	d
INFORMED	6.58 (2.04)	6.95 (2.10)	-1.89	443.7	.059	.413	0.18
DEHUMAN	3.24 (0.94)	3.46 (1.07)	-2.27	431.3	.024	.211	0.22
ROBOREPLACE	3.54 (1.00)	3.54 (1.14)	-0.01	429.5	.995	1.000	0.00
SOLVE_MANY	3.69 (1.00)	3.59 (1.09)	1.06	436.9	.290	1.000	-0.10
TIME_EFFICIENT	3.91 (0.92)	3.60 (1.19)	3.12	407.2	.002	.019*	-0.30
SOLVE_EFFICIENT	3.33 (0.98)	3.12 (1.17)	2.01	422.2	.045	.359	-0.19
VOTE_IMPACT	3.25 (1.01)	3.41 (1.16)	-1.51	429.8	.131	.784	0.14
ETHICAL_CONCERN	3.60 (1.01)	3.71 (1.15)	-1.16	430.3	.248	1.000	0.11
HELP_SOCIETY	3.24 (1.04)	3.34 (1.13)	-0.97	437.1	.333	1.000	0.09
TRUST	3.04 (1.11)	3.14 (1.20)	-0.95	436.5	.342	1.000	0.09

Note. 2024 $n = 230-233$; 2026 $n = 217$ per item. Positive d indicates higher 2026 mean. * significant after Holm correction across the family of ten tests ($\alpha = .05$).




Figure 4. Mean Agreement by Survey Wave for the 5-Point Perception Items.

Note. Error bars are 95% confidence intervals. * significant after Holm correction across the family of ten Welch t-tests.

Sources of AI Information and Reported Emotions

The distribution of reported sources for learning about AI was similar across waves, with the Internet remaining dominant, followed by social media and discussions with family and friends; the between-wave difference in the source distribution was not statistically significant, $\chi^2(4) = 7.96$, $p = .093$ (Figure 5). The distribution of reported emotions toward AI, by contrast, shifted significantly between waves, $\chi^2(3) = 17.52$, $p < .001$, Cramér's $V = .20$ (Figure 6). Curiosity remained the most frequently reported emotion but declined from 54% of respondents in 2024 to 40% in 2026, while the share reporting fear nearly doubled, from 18% to 33%. Reported indifference declined slightly (25% to 22%) and trust remained rare in both waves (3% and 5%).




Figure 5. Reported Sources for Learning about AI, 2024 vs. 2026.

Note. Percentages are of all reported sources within each wave; respondents could select multiple sources.




Figure 6. Reported Emotion toward AI, 2024 vs. 2026.

Gender Differences (RQ5)

Table 4 reports the gender comparisons with full test statistics and Holm correction. Five differences survived correction. Female respondents reported significantly higher agreement that AI helps solve many problems in society ($d = 0.34$), that AI helps solve problems more efficiently ($d = 0.33$), that AI can contribute positively to societal issues ($d = 0.27$), and higher trust in AI algorithms to provide accurate information ($d = 0.36$). Male respondents reported significantly higher agreement that AI encourages dehumanization ($d = -0.30$). Nominal differences in perceived election impact and time efficiency did not survive correction. Figure 7 presents the gender differences as a forest plot with 95% confidence intervals.

Table 4. Gender Comparisons of AI Perception Items (Welch t-tests)

Variable	Male M (SD)	Female M (SD)	t	df	p	p (Holm)	d
INFORMED	6.62 (2.07)	6.92 (2.10)	-1.52	418.6	.129	.387	0.15
DEHUMAN	3.49 (0.97)	3.19 (1.04)	3.07	407.3	.002	.016*	-0.30
ROBOREPLACE	3.52 (1.07)	3.58 (1.08)	-0.55	418.7	.580	1.000	0.05
SOLVE_MANY	3.49 (1.03)	3.84 (1.03)	-3.55	420.2	<.001	.004*	0.34
TIME_EFFICIENT	3.69 (1.07)	3.88 (1.06)	-1.86	423.2	.064	.257	0.18
SOLVE_EFFICIENT	3.09 (1.05)	3.43 (1.08)	-3.40	417.0	.001	.006*	0.33
VOTE_IMPACT	3.42 (1.00)	3.20 (1.17)	2.07	387.4	.039	.197	-0.20
ETHICAL_CONCERN	3.67 (1.05)	3.64 (1.13)	0.21	403.6	.831	1.000	-0.02
HELP_SOCIETY	3.16 (1.06)	3.45 (1.10)	-2.83	414.8	.005	.029*	0.27
TRUST	2.92 (1.15)	3.33 (1.12)	-3.75	423.8	<.001	.002*	0.36

Note. Male $n = 241$ – 242 ; female $n = 196$ – 198 . Positive d indicates higher female mean. * significant after Holm correction across the family of ten tests ($\alpha = .05$). The *HELP_SOCIETY* and *TRUST* items were rebuilt directly from the cleaned per-wave data files so that gender was correctly linked at the respondent level; the values reported here therefore supersede those in the previous draft (which reflected a one-row misalignment) and are in the same direction but slightly stronger.




Figure 7. Gender Differences in AI Perceptions (Female – Male Mean Differences with 95% Welch Confidence Intervals).

Note. Solid black markers and * denote differences significant after Holm correction.

GPA and Year of Study (RQ5)

Because GPA was recorded as an ordinal category, its associations with the perception items were tested with Spearman rank correlations (Table 5). Perceptions were largely independent of GPA. One association survived Holm correction: students reporting higher GPAs expressed slightly higher agreement that AI helps solve many problems in society ($\rho = .13$, corrected $p = .047$); the association between GPA and self-

reported informedness was of similar size but marginal after correction ($\rho = .13$, corrected $p = .051$). All remaining associations were small and non-significant.

Table 5. Spearman Correlations between GPA (Ordinal) and AI Perception Items

Variable	n	ρ	p	p (Holm)
INFORMED	447	.127	.007	.051
DEHUMAN	447	.029	.536	1.000
ROBOREPLACE	446	.080	.092	.462
SOLVE_MANY	446	.130	.006	.047*
TIME_EFFICIENT	447	.091	.053	.320
SOLVE_EFFICIENT	447	-.007	.875	1.000
VOTE_IMPACT	446	-.012	.800	1.000
ETHICAL_CONCERN	444	.039	.413	1.000

Note. * significant after Holm correction across the family of eight tests ($\alpha = .05$).

Year-of-study differences were examined with Welch one-way ANOVA across reported categories and a planned underclass (freshman/sophomore) versus upperclass (junior and above) contrast (Table 6). Two contrasts survived Holm correction: upperclass students reported higher agreement that AI helps with managing time more efficiently ($d = 0.30$, corrected $p = .044$) and that AI will impact the presidential election ($d = 0.34$, corrected $p = .027$). Year-of-study data were missing for 86 respondents, so these results should be interpreted with corresponding caution.

Table 6. Underclass vs. Upperclass Contrast on AI Perception Items (Welch t-tests)

Variable	Underclass M	Upperclass M	t	p	p (Holm)	d
INFORMED	6.66	6.80	-0.62	.537	1.000	0.07
DEHUMAN	3.32	3.34	-0.24	.811	1.000	0.03
ROBOREPLACE	3.52	3.55	-0.19	.849	1.000	0.02
SOLVE_MANY	3.55	3.82	-2.42	.016	.098	0.27
TIME_EFFICIENT	3.68	3.99	-2.75	.006	.044*	0.30
SOLVE_EFFICIENT	3.18	3.42	-1.96	.052	.259	0.22
VOTE_IMPACT	3.17	3.54	-2.96	.003	.027*	0.34
ETHICAL_CONCERN	3.52	3.76	-1.96	.052	.259	0.22

Note. Underclass $n = 244$ – 247 ; upperclass $n = 116$ – 117 . Positive d indicates higher upperclass mean. * significant after Holm correction across the family of eight tests ($\alpha = .05$).

Inter-Item Relationships

Inter-item correlations (Figure 8) indicate that AI perceptions are not isolated attitudes; trust, perceived usefulness, efficiency, social benefit, and concern operate as related components of a broader orientation toward AI.




Figure 8. Inter-Item Pearson Correlations among the Perception Items (Combined Waves, Pairwise n = 444–450).

New 2026 Items: Usage, Dependency, Verification, and Policy (RQ2–RQ4)

Table 7 reports descriptive statistics for the six new 2026 items. Students reported moderate-to-high frequency of AI use for information retrieval and problem solving ($M = 3.24$ on a 5-point scale) and high dependency on digital devices ($M = 7.02$ on a 10-point scale). Verification behavior was moderate: agreement with comparing multiple sources before forming opinions based on AI recommendations averaged 3.52. Students reported relatively high agreement that they understand when AI use is permitted or prohibited ($M = 3.81$) and moderate agreement that their institution has clearly communicated its AI policies ($M = 3.62$); however, agreement that institutional policies make them more confident using AI was notably lower ($M = 3.28$). This gap between policy understanding and policy-derived confidence is shown in Figure 10.

Table 7. Descriptive Statistics for the New 2026 Items

Variable	Content	M	SD	n
USE_AI	Frequency of AI tool use (1–5)	3.24	0.95	206
DIGITAL_DEPEND	Dependency on digital devices (1–10)	7.02	1.78	206
COMPARE_MULT	Compares multiple sources before forming opinion (1–5)	3.52	1.01	205
INSTITUTE_AI_POLICY	Institution clearly communicated AI policies (1–5)	3.62	1.09	203
AI_PERMITTED	Understands when AI is permitted/prohibited (1–5)	3.81	1.09	203
INSTITUTE_POLICIES	Policies increase confidence using AI (1–5)	3.28	1.00	203

Note. Two respondents with out-of-range values on the dependency item were excluded from these analyses.

Because the six new items had not previously been piloted together, their measurement properties were examined. The three institutional-policy items exhibited acceptable internal consistency (Cronbach's $\alpha = .76$). Sampling adequacy supported factor analysis ($KMO = .72$; Bartlett's test $\chi^2 = 265.0$, $p < .001$), and a principal-axis exploratory factor analysis of the six items yielded a coherent two-factor solution (Table 8):

the three policy items and the verification item loaded on a first factor reflecting institutional governance and verification, while the usage-frequency and digital-dependency items loaded on a second factor reflecting personal usage and dependency. These results provide initial psychometric support for the new items while confirming that governance perceptions and personal usage behavior are related but distinct constructs.

Table 8. Exploratory Factor Analysis Loadings for the New 2026 Items (Principal Axis, Rotated)

Item	F1: Policy & verification	F2: Usage & dependency	h ²
INSTITUTE_AI_POLICY	.82	.04	.67
AI_PERMITTED	.81	.04	.66
COMPARE_MULT	.54	.13	.31
INSTITUTE_POLICIES	.52	.18	.31
DIGITAL_DEPEND	.04	.87	.76
USE_AI	.16	.36	.16

Note. $n = 202$ complete cases. Loadings shown with signs oriented for readability; loadings $\geq .30$ are conventionally treated as salient.




Figure 9. Descriptive Results for the 2026-Only Survey Items ($n = 203\text{--}206$).

Note. Error bars are one standard deviation. The digital dependency item (right panel) uses a 10-point scale; all other items use 5-point scales.




Figure 10. Institutional Policy Understanding vs. Confidence Gap (2026 Wave).

Discussion

The most robust and most surprising finding of this study is a decline, not an increase, in students' perceptions of AI's efficiency benefits. Agreement that AI helps with managing time more efficiently fell from 3.91 to 3.60 between the 2024 and 2026 waves ($d = -0.30$), and this was the only difference among the ten Likert-type items to survive correction for multiple comparisons. A nominally significant decline in agreement that AI helps solve problems more efficiently pointed in the same direction. This pattern runs counter to the dominant productivity-gain narrative in the experimental literature (Noy & Zhang, 2023) and to satisfaction trends reported in large usage surveys. It is, however, consistent with an interpretation in which early enthusiasm is recalibrated through sustained practical experience: as generative AI moved from novelty to routine infrastructure, students appear to have developed more tempered, experience-based assessments of what these tools actually save them, including exposure to hallucinations, inconsistent outputs, and the overhead of verifying AI-generated content. External evidence is compatible with this reading; the most recent UK survey wave reports that while overall AI use continued to rise, the share of students using AI to generate text declined year over year (Freeman, 2026), suggesting more selective, workflow-conscious use rather than indiscriminate enthusiasm. If the normalization thesis is correct, the central question facing institutions is no longer whether students will adopt AI but how their expectations of it mature.

A second robust between-wave finding emerged from the categorical emotion item: the distribution of reported emotions toward AI shifted significantly between waves, with curiosity declining from 54% to 40% of respondents and fear nearly doubling from 18% to 33%. This shift is consistent with the efficiency-recalibration interpretation — as AI moved from novelty to infrastructure, the affective mix moved away from open-ended curiosity toward apprehension — and it directly reflects the “from curiosity” framing of this study's title. Against this backdrop, the remaining between-wave results are best characterized as stability with suggestive movement. Self-reported informedness trended upward ($d = 0.18$), consistent with two additional years of public discussion and classroom exposure, and agreement that AI encourages dehumanization was nominally higher in 2026 ($d = 0.22$); neither survived correction, and both should be treated as hypotheses for future waves rather than established changes. Notably, core attitudes about robots replacing workers, AI's societal contribution, ethical concern, and trust were essentially unchanged,

indicating that the period effect between 2024 and 2026 operated narrowly on efficiency perceptions and emotional response rather than broadly on AI attitudes.

The demographic analyses (RQ5) revealed that gender was the most consequential background variable. After correction, female respondents reported more favorable views of AI's problem-solving capability, efficiency, and societal contribution, and higher trust in AI-provided information, while male respondents expressed greater concern about dehumanization. These robust differences indicate that AI perceptions are not uniform across student groups, and they add nuance to the mixed gender findings in the survey literature, where male students have often been reported as more enthusiastic adopters (Freeman, 2025); perceived benefit and adoption enthusiasm are evidently not the same construct, and their relationship warrants direct study. By contrast, GPA was largely unrelated to AI perceptions, with only a small positive association with perceived societal problem solving surviving correction, and year-of-study effects were confined to upperclass students rating AI's time efficiency and election impact higher than underclass students. Taken together, these results suggest that demographic differences in perception are real but selective, concentrated in gender rather than academic standing.

The new 2026 items address the operational side of normalization (RQ2–RQ4). Students reported frequent AI use and high dependency on digital devices for day-to-day tasks, consistent with the emerging literature on cognitive offloading and reliance on digital systems (Rahman & Watanobe, 2023). Verification behavior, however, was only moderate. Given continued concerns regarding hallucinations, misinformation, and biased outputs (Nah et al., 2023; Meyer et al., 2023), inconsistent verification is arguably the central AI-literacy gap: if AI tools are now a routine channel for information retrieval, comparing sources becomes a core academic skill rather than a secondary one.

Institutional policy perceptions produced the study's most actionable finding. Students reported comparatively high understanding of when AI use is permitted or prohibited, yet markedly lower agreement that institutional policies make them more confident using AI. The exploratory factor analysis reinforces the point: governance-related perceptions form a coherent construct that is distinct from personal usage and dependency. Policy communication alone therefore appears insufficient. Institutions may need to move beyond rule statements toward concrete examples, course-level guidance, and AI-literacy support so that students understand not only what is allowed but how to use AI responsibly and effectively (Chan, 2023).

Table 9 summarizes the broader interpretive framework suggested by these results: a shift from an early-adoption orientation characterized by curiosity, experimentation, and optional use toward a normalization orientation characterized by workflow integration, operational dependency, and governance and literacy concerns. This framework does not imply that all students have moved uniformly from one stage to another; students still report curiosity and generally positive perceptions. Rather, the 2026 findings suggest that AI is increasingly treated as a practical workflow tool whose limitations are better understood, and whose responsible use requires verification habits and institutional support.

Table 9. Evolution of Student AI Interaction from 2024 to 2026

2024 Early-Adoption Orientation	2026 Normalization Orientation
Curiosity	Operational dependency
Experimentation	Workflow integration
Optional tool	Embedded infrastructure
Awareness	AI literacy and governance

Conclusion, Limitations, and Future Research

This study re-administered a 2024 survey of student perceptions of artificial intelligence to a new cohort in 2026, providing a repeated cross-sectional comparison spanning a period of rapid generative AI normalization, and extended the instrument with new measures of usage, dependency, verification, and institutional policy. The findings indicate that students continue to perceive substantial benefits associated with AI technologies while reporting more tempered efficiency expectations than the 2024 cohort, a significant shift in reported emotion away from curiosity and toward fear, robust gender differences in perceived benefit and trust, high operational dependency, moderate verification behavior, and a gap between understanding institutional AI policies and deriving confidence from them. The central issue for higher education is accordingly no longer whether students are aware of or curious about AI, but how AI is integrated into academic work and whether institutions provide sufficient structure to support responsible, verified, and confident use.

Several limitations should be acknowledged. First, the design is repeated cross-sectional, not longitudinal: the 2024 and 2026 samples are different cohorts, so between-wave differences are confounded with cohort composition, instructional changes, and the general saturation of AI over the period, and no within-person change can be inferred. Second, the sample is a convenience sample of students enrolled in Information Systems courses at a single private university in the western United States, recruited with an extra-credit incentive, which introduces self-selection and limits generalizability; although the courses primarily enroll underclass students, upper-division students also participated. Third, the study relies on self-reported perceptions and behaviors rather than observed AI usage or verification practices. Fourth, while the new 2026 items showed acceptable internal consistency and a coherent factor structure, they remain newly assembled measures whose validity should be established further in future administrations. Finally, the rapid evolution of generative AI means that student attitudes and institutional policies may continue changing substantially over short periods, so the present comparisons capture a moving target.

Future research should continue examining the long-term implications of dependency on generative AI systems, changes in critical thinking and verification behavior, and how institutions can balance AI integration with academic integrity and skill development. Additional waves using the same instrument would strengthen the temporal comparison, and within-person panel designs would permit inferences about individual attitude change that the present design cannot support. Future work should also test whether the efficiency-perception decline observed here replicates in other institutional contexts, and should examine the relationship between perceived benefit and adoption enthusiasm across demographic groups.

Acknowledgements

Generative AI tools were used to assist with portions of the survey data cleaning process and with copyediting for grammar, punctuation, and writing clarity. The authors also acknowledge that research is both an analytical and creative endeavor, and that AI-curated music recommendations, including artists such as Pearl Jam and AC/DC, occasionally provided welcome inspiration during the preparation of this manuscript. All study design decisions, statistical analyses, interpretations, and conclusions are solely those of the authors.

References

- Anthropic. (2025, April). Anthropic Education Report: How university students use Claude.
<https://www.anthropic.com/news/anthropic-education-report-how-university-students-use-claude>

- Ardelean, T. K., & Veres, E. (2023). Students' perceptions of artificial intelligence in higher education. *SWS International Scientific Conferences on Social Sciences*, 1–11. <https://doi.org/10.35603/sws.iscss.2023/s08.38>
- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, 38. <https://doi.org/10.1186/s41239-023-00408-3>
- Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). How people use ChatGPT (NBER Working Paper No. 34255). National Bureau of Economic Research. <https://doi.org/10.3386/w34255>
- Digital Education Council. (2024). Digital Education Council Global AI Student Survey 2024. <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-student-survey-2024>
- EDUCAUSE. (2024). 2024 EDUCAUSE Students and Technology Report. EDUCAUSE. <https://www.educause.edu/ecar/research-publications/2024/students-and-technology-report>
- Fauzi, F., Tuhuteru, L., Sampe, F., Ausat, A. M. A., & Hatta, H. R. (2023). Analysing the role of ChatGPT in improving student productivity in higher education. *Journal on Education*, 5(4), 14886–14891. <https://doi.org/10.31004/joe.v5i4.2563>
- Firaina, R., & Sulisworo, D. (2023). Exploring the usage of ChatGPT in higher education: Frequency and impact on productivity. *Buletin Edukasi Indonesia*, 2(1), 39–46. <https://doi.org/10.56741/bei.v2i01.310>
- Freeman, J. (2024). Provide or punish? Students' views on generative AI in higher education (HEPI Policy Note 51). Higher Education Policy Institute.
- Freeman, J. (2025). Student Generative AI Survey 2025 (HEPI Policy Note 61). Higher Education Policy Institute. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/>
- Freeman, J. (2026). Student Generative AI Survey 2026 (HEPI Report 199). Higher Education Policy Institute. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2026/>
- Grand View Research. (2022). AI in education market size & share report, 2022–2030. <https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-education-market-report>
- Grand View Research. (2026). AI in education market size, share and growth report, 2026–2033. <https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-education-market-report>
- Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base – analyst note. Reuters. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Jeffrey, T. (2020). Understanding college student perceptions of artificial intelligence. *Systemics, Cybernetics and Informatics*, 18(2), 8–13.
- Keles, P. U., & Aydin, S. (2021). University students' perceptions about artificial intelligence. *Shanlax International Journal of Education*, 9(S1), 212–220. <https://doi.org/10.34293/education.v9iS1-May.4014>
- Kumar, V. R., & Raman, R. (2022, March). Student perceptions on artificial intelligence (AI) in higher education. In 2022 IEEE Integrated STEM Education Conference (ISEC) (pp. 450–454). IEEE.

- Leoste, J., Jögi, L., Õun, T., Pastor, L., San Martín López, J., & Grauberg, I. (2021). Perceptions about the future of integrating emerging technologies into higher education—the case of robotics with artificial intelligence. *Computers*, 10(9), 110. <https://doi.org/10.3390/computers10090110>
- Meyer, J. G., Urbanowicz, R. J., Martin, P. C. N., O'Connor, K., Li, R., Peng, P.-C., Bright, T. J., Tatonetti, N., Won, K. J., Gonzalez-Hernandez, G., & Moore, J. H. (2023). ChatGPT and large language models in academia: Opportunities and challenges. *BioData Mining*, 16, 20. <https://doi.org/10.1186/s13040-023-00339-9>
- Nah, F. F.-H., Zheng, R., Cai, J., Siau, K., & Chen, L. (2023). Generative AI and ChatGPT: Applications, challenges, and AI-human collaboration. *Journal of Information Technology Case and Application Research*, 25(3), 277–304. <https://doi.org/10.1080/15228053.2023.2233814>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Pesovski, I., Santos, R., Henriques, R., & Trajkovik, V. (2024). Generative AI for customizable learning experiences. *Sustainability*, 16(7), 3034. <https://doi.org/10.3390/su16073034>
- Pew Research Center. (2025, April 3). How the U.S. public and AI experts view artificial intelligence. <https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/>
- Rahman, M. M., & Watanobe, Y. (2023). ChatGPT for education and research: Opportunities, threats, and strategies. *Applied Sciences*, 13(9), 5783. <https://doi.org/10.3390/app13095783>
- Tyton Partners. (2025). Time for Class 2025. Tyton Partners. <https://tytonpartners.com/time-for-class-2025/>
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1), 1–27. <https://doi.org/10.1186/s41239-019-0171-0>
- Zhang, K., & Aslan, A. B. (2021). AI technologies for education: Recent research & future directions. *Computers and Education: Artificial Intelligence*, 2, 100025. <https://doi.org/10.1016/j.caeai.2021.100025>