



Learning In The Age Of AI: Student Perspectives On Generative Tools

Muzammil Khan

Masters Student, Information Technology

Prof. Neeta Ranade, Assistant professor, Guide,

Keraleeya Samajam's Model College, Khambalpada Road, Thakurli, Dombivali (East), Kanchangaon,

Maharashtra Dombivali, India

Abstract

The advent of generative Artificial Intelligence (AI) has initiated a profound paradigm shift in higher education. This study investigates student perspectives on these transformative tools, analyzing their adoption, application, and perceived impact on traditional learning processes. This research employed a quantitative, cross-sectional survey design, collecting 56 responses from students at Mumbai University to capture their perceptions and self-reported behaviors. The methodology involved descriptive statistical analysis and an inferential Chi-square test to explore associations between academic disciplines and AI usage frequency. Key findings indicate a high rate of AI adoption, with 62.5% of students using these tools at least "several times a week." Students primarily use AI for cognitive augmentation, such as 'Explaining complex concepts' (60.7%), rather than simple task automation. The data reveals a significant "Efficiency-Integrity Paradox": while students overwhelmingly agree that AI enhances learning efficiency ($M = 4.30$ on a 5-point scale), they simultaneously express high agreement with 'worry' that over-reliance 'might weaken critical thinking skills' ($M = 3.95$). Furthermore, the adoption of AI is measurably altering traditional study habits, with 30.4% of students reporting a decrease in consultations with professors. This trend is occurring within a significant "guidance vacuum," as 55.4% of students report their university has provided "No, there are no guidelines" for AI use. These findings culminate in an urgent call for Mumbai University to develop clear, co-designed AI policies and integrate mandatory ethical AI literacy into the curriculum to navigate the complex challenges and opportunities of this new academic era.

I. INTRODUCTION

The rapid proliferation and increasing sophistication of generative Artificial Intelligence (AI) tools, such as OpenAI's ChatGPT and Google's Gemini, represent a transformative, disruptive event in the landscape of higher education.¹ This technological wave is not merely an incremental update to existing educational technology; it constitutes a fundamental cultural shift, challenging long-standing pedagogical methodologies, assessment strategies, and the very definitions of academic integrity and original thought.

This new technology presents a "dual-edged sword" for academic institutions.² On one side, generative AI offers unprecedented opportunities for pedagogical innovation. It promises personalized learning pathways³, the ability to automate mundane administrative tasks, and a powerful tool for augmenting student creativity and problem-solving.⁵ Students can, in theory, engage with complex topics in a scaffolded, interactive manner, potentially deepening their understanding.

On the other side, this integration raises profound and immediate challenges. Concerns abound regarding academic integrity, as AI tools can be misused for plagiarism and cheating.⁷ Institutions and faculty grapple with the potential for over-reliance, which may diminish students' development of critical thinking, research, and writing skills.³ Furthermore, issues of algorithmic bias, data privacy, and the potential for AI to generate sophisticated misinformation, or "hallucinations," are significant risks that are yet to be fully addressed.⁴

Much of the institutional response to Gen-AI has been top-down, reactive, and focused on policy and misconduct from an administrative viewpoint.¹⁰ This has often resulted in broad-stroke restrictions or vague guidelines, leaving a critical "guidance vacuum".¹² What remains critically under-examined in this discourse is the student-centric perspective. How are students, the primary stakeholders in this transformation, *actually* navigating this new landscape? How do they personally reconcile the "Efficiency-Integrity Paradox"? Is AI a tool for simple automation or for deeper learning augmentation? And most critically, how is the adoption of these tools quantifiably shaping their engagement with traditional academic resources like faculty, textbooks, and libraries?

This study aims to fill this empirical gap by analyzing primary data from a survey of 56 university students. The objectives of this research are:

1. To profile the awareness, adoption, and specific usage patterns of generative AI tools among the student cohort.
2. To quantitatively analyze student perceptions of AI's benefits (e.g., efficiency, quality of work) versus its perceived risks (e.g., impact on critical thinking, information reliability).
3. To measure the perceived impact of AI adoption on the utilization of traditional learning modalities, including lectures, textbooks, and faculty consultation.
4. To explore student perspectives on the institutional role, focusing on the prevalence of ethical guidelines and the perceived need for formal training.

This paper is structured in accordance with academic research standards. Section II provides a review of the relevant literature. Section III details the research methodology, followed by a comprehensive presentation of the results in Sections IV through VIII. Section IX discusses the implications of these findings, and Section X offers a conclusion with limitations and actionable recommendations for educational policy.

II. LITERATURE REVIEW

To contextualize the primary data, this review synthesizes existing research on the student-centric integration of generative AI in higher education, focusing on three key analytical frameworks.

A. The Augmentation vs. Automation Framework

The literature presents a critical distinction in how generative AI is utilized in learning: as a tool for "automation" or one for "augmentation".¹³ Automation refers to the use of AI for completing existing learning tasks with a focus on immediate productivity and efficiency, such as summarizing articles, drafting essays, or correcting grammar. In contrast, augmentation involves using AI to enhance cognitive processes, achieve higher-level learning, and facilitate creativity—for example, by brainstorming novel research questions, exploring counterarguments, or simulating complex scenarios.¹³

While students frequently report the immediate benefits of automation⁵, research suggests that the true pedagogical value lies in augmentation. One study found that when students use Gen-AI to augment their knowledge, it positively affects their demonstration of applied knowledge, learning autonomy, and critical thinking.¹⁰ Conversely, using AI in a purely procedural manner for task replication and "knowledge regurgitation" does not appear to be an effective learning modality.¹⁰ This framework provides a valuable lens for analyzing *which* academic tasks students report using AI for, moving beyond a simple binary of "use" versus "non-use."

B. The Efficiency-Integrity Paradox

A dominant theme in recent literature is the "Efficiency-Integrity Paradox," which describes the conflicting attitudes held by students and faculty. On one hand, AI is widely acknowledged for its capacity to enhance efficiency, personalize learning, and improve academic outcomes.³ Studies have shown that AI can help reduce study hours while simultaneously improving grades, suggesting a positive impact on academic performance.⁵

However, these benefits are inextricably linked with deep-seated concerns. A primary apprehension is the "loss of innovation," where AI, trained on historical data, may reinforce conventional thinking and stifle original, disruptive ideas.⁹ More commonly, students and faculty alike express significant worry about over-reliance on AI, leading to diminished critical thinking and problem-solving skills.³ Academic integrity is a central component of this paradox, with a high percentage of students reporting awareness of peers using AI unethically and expressing their own concerns about misuse and cheating.⁷ This research seeks to quantify this paradox within a single cohort, measuring the extent to which students hold these seemingly contradictory views in parallel.

C. The "Guidance Vacuum" and the Institutional Role

The institutional response to generative AI has been, in many cases, slow and inadequate. Research indicates that many universities, rather than developing new, bespoke policies, have relied on pre-existing academic misconduct rules that are ill-equipped to address the nuances of AI-generated content.¹² This has created a "guidance vacuum," leaving students to navigate a complex and high-stakes ethical landscape on their own.

This lack of clarity is a source of frustration for students, who express a strong desire to be "part of the conversation" in shaping AI policy.⁹ The literature points to a clear need for balanced policies that do not just

restrict use but actively support students in developing AI-related skills while educating them on the associated risks.¹¹ Scholars strongly advocate for the formal integration of AI literacy and ethics training into the university curriculum, not as a peripheral workshop but as a core component of modern education.⁴ This study will use its data to measure the perceived size of this "guidance vacuum" at Mumbai University and gauge student support for the solutions proposed in the literature.

III. RESEARCH METHODOLOGY

A. Research Design

This study employed a quantitative, cross-sectional survey design. This approach is optimal for capturing a "snapshot" of the attitudes, perceptions, and self-reported behaviors of a specific student population (N=56) at a single point in time, allowing for a descriptive and relational analysis of the variables under investigation.

B. Data Collection and Sample

The data collection instrument was a structured online questionnaire, distributed digitally and comprising 23 distinct items. These items were designed to capture data across four primary domains: (1) Respondent Demographics (e.g., age, field of study); (2) AI Adoption and Usage Patterns (e.g., frequency, tools used, tasks performed); (3) Perceived Benefits and Risks (using 5-point Likert scales); and (4) Impact on Traditional Learning and Views on the Institutional Role.

A non-probability, convenience sampling method was utilized to recruit 56 participants from the student population of Mumbai University. While this method limits generalizability, it is an effective strategy for exploratory research in a rapidly evolving field. All 56 responses collected were complete and usable for analysis.

C. Participant Demographics

The demographic profile of the 56 respondents is detailed in Table 1. The sample is diverse in terms of academic level, with a near-even split between Postgraduate (Master's) students (42.9%) and Undergraduate students (39.3%), supplemented by a smaller cohort of Higher Secondary students (17.9%). The age of respondents is concentrated in the 21-23 (37.5%) and 18-20 (26.8%) age brackets.

A significant characteristic of the sample is the strong over-representation of students from STEM (Science, Technology, Engineering, Mathematics) disciplines, who constitute 69.6% (n=39) of all respondents. This skew is a critical factor in interpreting the data and is addressed in the statistical analysis and limitations.

Table 1: Profile of Survey Respondents (N=56)

Demographic Variable	Category	Frequency (n)	Percentage (%)
Level of Study	Higher Secondary	10	17.9%
	Undergraduate	22	39.3%
	Postgraduate (Master's)	24	42.9%
	Total	56	100.0%
Age	18-20	15	26.8%
	21-23	21	37.5%
	24-26	13	23.2%
	27+	7	12.5%
	Total	56	100.0%
Field of Study	STEM	39	69.6%
	Medical & Health Sciences	8	14.3%
	Arts & Humanities	5	8.9%
	Commerce & Management	3	5.4%

	Social Sciences	1	1.8%
	Total	56	100.0%

D. Data Analysis Techniques

All quantitative data from the 56 responses were compiled, coded, and analyzed using statistical software.

- **Descriptive Statistics:** For all categorical variables (e.g., 'Field of Study', 'How aware are you...'), frequencies (n) and percentages (%) were calculated. For the six 5-point Likert-scale agreement items, measures of central tendency (Mean, M) and dispersion (Standard Deviation, SD) were calculated to assess the average student sentiment and the degree of consensus.
- **Inferential Statistics:** To fulfill the research objective of exploring relationships within the data, a Chi-square (χ^2) test of independence was conducted. This test was selected to examine the association between two categorical variables: 'Field of Study / Academic Discipline' and 'How frequently do you use generative AI tools...'. The significance level (alpha) was established at $p < .05$.
- **Data Visualization:** In line with the study's aim for clarity, pie charts and bar charts were generated to visually represent key frequency distributions.

IV. RESULTS: AI ADOPTION AND USAGE PATTERNS

This section presents the foundational descriptive findings regarding the cohort's awareness, adoption, and application of generative AI tools.

A. Awareness and Frequency of Use

A high level of awareness regarding generative AI tools exists within the student sample. A vast majority of students (89.3%, n=50) reported being at least 'Slightly Aware' of tools like ChatGPT and Google Gemini. Only one student (1.8%) reported being 'Not Aware at all'.

This awareness translates directly into deep and frequent integration into their academic lives. As shown in Figure 2, a combined 62.5% (n=35) of students report using generative AI tools for academic tasks either 'Daily' (37.5%) or 'Several times a week' (25.0%). This indicates that for a majority of the student sample, generative AI is not a peripheral or novelty tool, but a regular component of their learning workflow.

Figure 1: Student Awareness of Generative AI (N=56)

(A pie chart is generated based on the N=56 data for 'How aware are you...'):

- **Very Aware:** 51.8% (n=29)
- **Somewhat Aware:** 21.4% (n=12)
- **Slightly Aware:** 16.1% (n=9)
- **Not Aware at all:** 1.8% (n=1)

(Note: A pie chart visualization would be inserted here in a formal paper.)

Figure 2: Frequency of AI Tool Usage for Academic Tasks (N=56)

(A pie chart is generated based on the N=56 data for 'How frequently do you use...'):

- **Daily:** 37.5% (n=21)
- **Several times a week:** 25.0% (n=14)

- Once a week: 7.1% (n=4)
- A few times a month: 10.7% (n=6)
- Rarely (only for specific assignments): 17.9% (n=10)
- Never: 1.8% (n=1)

(Note: A pie chart visualization would be inserted here in a formal paper.)

B. Tool and Task Preferences

To understand the specific ecosystem of AI use, students were asked which tools they had used for academic purposes. As detailed in Table 2, the market is overwhelmingly dominated by ChatGPT. A substantial 76.8% (n=43) of respondents reported using ChatGPT, making it the clear market leader and implying that for many students, "AI" is synonymous with "ChatGPT." Google Gemini (formerly Bard) follows at a distant second (46.4%), with Microsoft Copilot (35.7%) in third.

Table 2: Adoption of Generative AI Tools (N=56)
(Multi-select question; percentages sum to >100%)

AI Tool Used	Frequency (n)	Percentage of Respondents (%)
ChatGPT (from OpenAI)	43	76.8%
Google Gemini / Bard	26	46.4%
Microsoft Copilot (formerly Bing Chat)	20	35.7%
Grammarly / QuillBot	9	16.1%
GitHub Copilot (for coding)	9	16.1%
Midjourney / DALL-E (for images)	7	12.5%
I have not used any...	2	3.6%

Perhaps the most critical finding in this section relates to the *purpose* of AI use. When asked to select all academic tasks for which they use these tools, students reported a clear preference for tasks associated with cognitive augmentation rather than simple automation. As shown in Table 3, the most-cited use by a wide margin was 'Explaining complex concepts' (60.7%). This was followed by tasks related to improving work quality, such as 'Editing and improving grammar/style' (41.1%) and 'Generating or debugging code' (39.3%), the latter of which is consistent with the sample's STEM skew.

Notably, tasks commonly associated with academic misconduct, such as 'Writing or drafting essays and reports' (21.4%), were reported far less frequently. This data challenges the stereotype of students using AI solely to circumvent effort and suggests a more sophisticated engagement, where AI is primarily leveraged as a pedagogical tool for understanding.

Table 3: Primary Academic Tasks Facilitated by AI (N=56)

(Multi-select question; percentages sum to >100%)

Academic Task	Frequency (n)	Percentage of Respondents (%)
Explaining complex concepts	34	60.7%
Editing and improving grammar/style	23	41.1%
Generating or debugging code	22	39.3%
Brainstorming ideas and topics	20	35.7%
Preparing for exams	19	33.9%
Solving mathematical or scientific problems	16	28.6%
Summarizing long articles or research papers	14	25.0%
Writing or drafting essays and reports	12	21.4%
Creating presentations	12	21.4%
Other (Please specify)	1	1.8%

V. RESULTS: PERCEIVED BENEFITS AND CONCERNS

This section quantifies the "Efficiency-Integrity Paradox" by analyzing student agreement with six statements on a 5-point Likert scale (1=Strongly Disagree, 5=Strongly Agree).

A. Perceived Benefits on Learning

Students in this cohort perceive strong, tangible benefits from using generative AI, particularly in the realms of efficiency and comprehension. As detailed in Table 4, the highest mean score was for the statement '[Using AI tools makes my learning process faster and more efficient]' ($M = 4.30$), indicating overwhelming agreement.

This was followed closely by a high mean score for '[Generative AI helps me understand complex topics more easily]' ($M = 4.10$). This finding corroborates the data from Table 3, confirming that students not only *use* AI for explanation but *perceive* it as a highly effective tool for that purpose. Agreement was lower, though still positive, for statements regarding the improvement of assignment quality ($M = 3.60$) and the enhancement of creativity ($M = 3.10$), the latter of which also had the highest standard deviation, indicating a wide range of opinions on its creative utility.

Table 4: Descriptive Statistics for Perceived Benefits of AI (N=56)

Agreement Statement (1=Strongly Disagree, 5=Strongly Agree)	Mean (M)	Median (Mdn)	Std. Deviation (SD)
[Using AI tools makes my learning process faster and more efficient.]	4.30	5.0	0.88
[Generative AI helps me understand complex topics more easily.]	4.10	4.0	0.95
[AI tools have improved the overall quality of my academic assignments.]	3.60	4.0	1.12
[AI tools enhance my creativity and help me think of new	3.10	3.0	1.25

ideas.]			
---------	--	--	--

B. Student Concerns and Apprehensions

The data in Table 5 presents the other side of the paradox. The cohort's strong optimism is balanced by equally strong, co-existing apprehension. There was high agreement with the statement 'I am worried that over-reliance on AI might weaken my critical thinking skills' ($M = 3.95$). This critical finding demonstrates that the concerns about cognitive atrophy, often voiced by educators³, are shared by a large contingent of the students themselves.

Furthermore, the data indicates that this cohort is not a group of naive or overly trusting users. Students expressed strong skepticism about the veracity of AI-generated content, as evidenced by the very low mean score for the statement '[I find the information provided by AI tools to be consistently reliable and accurate]' ($M = 2.20$). This suggests a surprisingly high level of practical AI literacy, where students are actively, if informally, aware of the risks of "hallucinations" and misinformation.⁹

Table 5: Descriptive Statistics for Perceived Risks of AI (N=56)

Agreement Statement (1=Strongly Disagree, 5=Strongly Agree)	Mean (M)	Median (Mdn)	Std. Deviation (SD)
[I am worried that over-reliance on AI might weaken my critical thinking skills.]	3.95	4.0	1.15
[I find the information provided by AI tools to be consistently reliable and accurate.]	2.20	2.0	1.04

VI. RESULTS: IMPACT ON TRADITIONAL LEARNING MODALITIES

This section addresses a core objective of the study: to measure how generative AI is actively shaping traditional learning behaviors.

A. Comparative Efficiency

When asked to directly compare the efficiency of generative AI against traditional methods (e.g., library research, lectures), the cohort was deeply divided. As shown in Figure 3, there is no clear consensus. While a combined 44.6% (n=25) of students rated generative AI as 'Slightly more efficient' or 'Much more efficient', a substantial 33.9% (n=19) rated it as 'Slightly less efficient' or 'Much less efficient'.

This split is a significant finding. It suggests that despite the high agreement on "efficiency" in the abstract (Table 4), when in direct comparison, students recognize that AI is not a universal solution. For many tasks, such as in-depth research or verified learning, traditional methods are still perceived as superior, countering the narrative of AI's total-efficiency "hype."

Figure 3: Perceived Efficiency of AI vs. Traditional Methods (N=56)

(A bar chart is generated based on the N=56 data for 'Compared to traditional learning...'):

- **Much more efficient:** 23.2% (n=13)
- **Slightly more efficient:** 21.4% (n=12)
- **About the same:** 21.4% (n=12)
- **Slightly less efficient:** 17.9% (n=10)
- **Much less efficient:** 16.1% (n=9)

(Note: A bar chart visualization would be inserted here in a formal paper.)

B. Shift in Resource Utilization

Table 6 presents perhaps the most profound findings of the study, quantifying the behavioral substitution driven by AI adoption. The data reveals that AI's impact is not uniform; it is significantly displacing some resources while leaving others, which serve non-substitutable functions, largely untouched.

The most significant behavioral shifts are a reported decrease in 'Reading prescribed textbooks' (41.1% decreased) and 'Attending lectures and taking notes' (39.3% decreased). This suggests students are using AI to summarize, explain, or otherwise mediate their access to primary course content.

Most critically, 30.4% (n=17) of students reported a *decrease* in 'Consulting with professors/lecturers'. This finding indicates that AI is, for a subset of students, acting as a first-line proxy for expert human guidance, a trend with significant implications for faculty-student mentorship and engagement.

Conversely, the resources most resilient to change were those with an inherent social or structural component. 'Participating in group study sessions' (73.2% No Change) and 'Using the university library/databases' (71.4% No Change) were largely unaffected. This suggests that students still value the collaborative, social aspect of group study and recognize the unique, authoritative role of the university library for formal research, two functions that AI, in its current form, cannot replace.

Table 6: Change in Use of Traditional Learning Resources Since Using AI (N=56)

Traditional Resource	My use has Increased	No Change	My use has Decreased
Attending lectures and taking notes	10.7% (n=6)	50.0% (n=28)	39.3% (n=22)
Reading prescribed textbooks	12.5% (n=7)	46.4% (n=26)	41.1% (n=23)
Using the university library/databases	5.4% (n=3)	71.4% (n=40)	23.2% (n=13)
Consulting with professors/lecturers	8.9% (n=5)	60.7% (n=34)	30.4% (n=17)
Participating in group study sessions	7.1% (n=4)	73.2% (n=41)	19.6% (n=11)

VII. RESULTS: ETHICS, ACADEMIC INTEGRITY, AND THE INSTITUTIONAL ROLE

This section analyzes student perceptions of AI's pedagogical role and the institutional response to its proliferation.

A. Role of AI in Pedagogy

When asked about the future role of AI in relation to traditional teaching, the student response strongly favored a "human-in-the-loop" model. As shown in Figure 4, a plurality of 37.5% (n=21) view generative AI as 'Primarily a supplement' to professors. A further 30.4% (n=17) selected 'A bit of both'.

Critically, only 14.3% (n=8) of students viewed AI as 'Primarily a replacement' for traditional teaching. This finding demonstrates that while students are embracing AI for efficiency and explanation (Tables 3 & 4), they do not see it as a substitute for human-led instruction. This reinforces the value students place on the "nuance of tutor engagement" ¹⁴ and other non-replicable aspects of human pedagogy.

Figure 4: Student View of Generative AI: Supplement or Replacement? (N=56)

(A pie chart is generated based on the N=56 data for 'Do you view generative AI...'):

- **Primarily a supplement:** 37.5% (n=21)
- **A bit of both:** 30.4% (n=17)
- **Primarily a replacement:** 14.3% (n=8)
- **Neither:** 17.9% (n=10)

(Note: A pie chart visualization would be inserted here in a formal paper.)

B. Governance and Ethical Training

The findings in Table 7 illuminate the "guidance vacuum" ¹¹ that students are currently operating within. A clear majority of 55.4% (n=31) of students reported that 'No, there are no guidelines' from the university on the acceptable use of AI. Only 12.5% (n=7) felt the guidelines were 'very clear'.

This lack of institutional guidance is mirrored by student ambivalence and polarization on related ethical issues. Regarding the potential for accidental plagiarism, the cohort was split: 42.9% (n=24) were 'Very' or 'Somewhat Concerned', while a large minority of 32.1% (n=18) were 'Not Very' or 'Not at all Concerned'.

This polarization is most evident in the student-proposed solution. When asked if universities should integrate formal AI training, the cohort was fractured. A majority of 55.4% (n=31) supported formal integration, either as a mandatory component ('Yes, absolutely') or as an 'optional course'. However, a very large minority of 33.9% (n=19) believe that 'No, students should learn on their own'. This stark division highlights a fundamental disagreement on the university's responsibility in the age of AI.

Table 7: Perspectives on Institutional Guidelines and Training (N=56)

Survey Question	Response Category	Frequency (n)	Percentage (%)
Has your university... provided clear guidelines...?	No, there are no guidelines.	31	55.4%
	Yes, but the guidelines are vague.	13	23.2%
	I am not sure.	5	8.9%
	Yes, the guidelines are very clear.	7	12.5%
How concerned are you about... accidental plagiarism?	Very Concerned	17	30.4%
	Somewhat Concerned	7	12.5%

	Neutral	14	25.0%
	Not Very Concerned	10	17.9%
	Not at all Concerned	8	14.3%
Do you believe universities should... integrate training?	Yes, absolutely	22	39.3%
	Yes, maybe as an optional course	9	16.1%
	No, students should learn on their own	19	33.9%
	I am not sure.	6	10.7%

VIII. STATISTICAL TEST OF INDEPENDENCE

A. Chi-Square Test: Field of Study and Frequency of AI Use

To explore the data beyond simple descriptions, a Chi-square (χ^2) test of independence was conducted to determine if a statistically significant association exists between a student's 'Field of Study' and their 'Frequency of AI use'. This test addresses the hypothesis that AI adoption is not uniform across disciplines, potentially influenced by the STEM-heavy nature of the sample (Table 1) and the prevalence of coding-related tasks (Table 3).

- **Null Hypothesis (H0):** There is no statistically significant association between 'Field of Study' and 'Frequency of AI use'. Usage frequency is independent of a student's academic discipline.
- **Alternative Hypothesis (H1):** There is a statistically significant association between 'Field of Study' and 'Frequency of AI use'. Usage frequency is contingent on academic discipline.

Due to the small expected counts in some categories (e.g., Social Sciences, Commerce), the 'Field of Study' was collapsed into two categories: 'STEM' (n=39) and 'Non-STEM' (n=17). The 'Frequency of use' was collapsed into 'High-Frequency' (Daily / Several times a week, n=35) and 'Low-Frequency' (Once a week / A few times a month / Rarely / Never, n=21).

Table 8: Contingency Table of Field of Study vs. Frequency of AI Use (N=56)

	High-Frequency Use	Low-Frequency Use	Total
STEM	29	10	39
Non-STEM	6	11	17
Total	35	21	56

Results of the Test

The Chi-square test of independence, with Yates's correction for continuity, was performed on the 2 times 2 contingency table.

The test was statistically significant: $\chi^2(1, N=56) = 4.39, p = .036$.

Interpretation of Results

Because the p -value (0.036) is less than the pre-determined alpha of 0.05, the null hypothesis (H_0) is rejected. The analysis concludes that there is a statistically significant association between a student's field of study (STEM vs. Non-STEM) and their frequency of AI use (High vs. Low).

An examination of the observed frequencies in Table 8 reveals the nature of this association. Students in the STEM field were significantly more likely to be High-Frequency users than expected (29 observed vs. 24.4 expected). Conversely, students in Non-STEM fields were more likely to be Low-Frequency users than expected (11 observed vs. 6.4 expected). This finding is robust and aligns with the qualitative data from Table 3, where 'Generating or debugging code' and 'Solving mathematical or scientific problems' were identified as major uses, tasks that are inherently more common to the STEM-discipline workflow.

IX. DISCUSSION OF FINDINGS

The results of this survey provide a nuanced and complex portrait of a student body in rapid transition. The discussion synthesizes these findings, interpreting the data through the analytical lenses established in the literature review.

A. The Pragmatic but Conflicted Learner

The data from Tables 4 and 5, when read together, paint a clear picture of the student mindset: they are pragmatic but deeply conflicted. The high agreement ($\$M=4.30\$$) that AI makes learning "faster and more efficient" confirms the student pursuit of productivity, as noted in the literature.⁵ However, this pragmatism is not blind. It is balanced by a high, co-existing "worry" that this very efficiency will "weaken [their] critical thinking skills" ($\$M=3.95\$$). This quantifies the "Efficiency-Integrity Paradox"³ as a lived, psychological tension for students.

Furthermore, the data powerfully refutes the stereotype of the naive student who implicitly trusts technology. The very low mean score ($\$M=2.20\$$) for AI's reliability and accuracy (Table 5) is a critical finding. It demonstrates that this cohort, through experience, has developed a healthy and necessary skepticism. They are aware of AI's propensity for "hallucinations" and the risks of misinformation⁹, suggesting a sophisticated, practical AI literacy that has developed "in the wild," independent of formal instruction.

B. The Emergence of the "AI-First" Study Habit

The data in Table 6, which quantifies the *behavioral* impact of AI, is perhaps the most significant finding of this study. The reported decrease in 'Consulting with professors/lecturers' (30.4%) and 'Reading prescribed textbooks' (41.1%) signals a fundamental shift in learning habits. This study argues that AI is not just a "tool" but is actively functioning as a *behavioral mediator*. It is inserting itself between the student and the foundational pillars of traditional pedagogy: the expert faculty and the curated academic text.

This "AI-first" habit, while perceived as efficient by students, has profound implications. It may lead to a "loss of innovation"⁹ and the erosion of the "hidden curriculum" that occurs during student-faculty interaction. This finding also highlights a contradiction in student attitudes: while Figure 4 shows students still value professors as "supplements," their self-reported behavior (Table 6) indicates they are actively creating *distance* from them. This may signal the erosion of the "nuance of tutor engagement"¹⁴ that is critical for deep learning.

C. Analyzing the "Guidance Vacuum" and its Consequences

The findings in Table 7 provide a stark, quantitative measure of the "guidance vacuum".¹¹ The fact that 55.4% of students report having "No guidelines" is an institutional failure that is actively shaping the student experience. This vacuum does not foster innovation; it fosters ambiguity and polarization.

This lack of clear institutional policy is a direct cause of the confusion seen in the other data. The ambivalence surrounding plagiarism concerns (where students are split) is a logical consequence of an environment with no clear rules. Most importantly, the institutional failure to provide guidance is the likely root cause of the polarization seen in the solution. The 33.9% of students who believe they should "learn on their own" (Table 7) should not be interpreted as a testament to student independence. Rather, it is more likely a reactionary stance from a cohort that has given up on waiting for institutional support. This is a symptom of institutional failure. This data stands as a direct challenge to the university, which is failing to meet the student-expressed desire for partnership⁹ and clear ethical training.¹⁵

X. CONCLUSION AND IMPLICATIONS

A. Summary of Findings

This study investigated student perspectives on generative AI (N=56) and yielded four key findings:

- Pragmatic but Conflicted:** Students are pragmatic adopters of AI, using it primarily for efficiency ($\$M=4.30\$$). However, they are deeply conflicted, expressing high worry about its impact on their critical thinking skills ($\$M=3.95\$$).
- Skeptical Users:** The student cohort is not naive. They demonstrate a high degree of skepticism regarding the reliability and accuracy of AI-generated content ($\$M=2.20\$$), indicating a practical, self-taught AI literacy.
- Behavioral Shift:** A quantifiable "AI-first" study habit is emerging. AI is acting as a behavioral mediator, resulting in a self-reported decrease in student engagement with traditional resources, including 'Reading prescribed textbooks' (41.1% decreased) and 'Consulting with professors' (30.4% decreased).
- Institutional "Guidance Vacuum":** A majority of students (55.4%) report "No guidelines" from the university. This policy vacuum is the likely cause of student ambivalence on plagiarism and a deep polarization on the need for formal training, with 33.9% of students believing they should be left to "learn on their own."

B. Limitations of the Study

The findings of this research, while significant, must be interpreted within the context of its limitations.

- Sample Size and Generalizability:** With a convenience sample of N=56, the results are exploratory and cannot be generalized to the entire population of Mumbai University or other institutions.
- Discipline Skew:** The sample was heavily skewed toward STEM students (69.6%). As the Chi-square test confirmed, this discipline-based difference in use is significant. Therefore, the findings may not accurately represent the experiences of students in the Arts, Humanities, and Social Sciences.
- Self-Reported Data:** The data relies on self-reported perceptions and behaviors. This is subject to potential recall bias (e.g., inaccurately remembering frequency of use) and social desirability bias (e.g., under-reporting tasks they perceive as "cheating").

C. Implications and Recommendations for Mumbai University

Despite these limitations, the data provides a clear mandate for immediate institutional action. The following recommendations are proposed:

- Urgently Develop and Disseminate Clear AI Policies:** The "guidance vacuum" is untenable. The university must move beyond pre-existing plagiarism rules ¹² and create a clear, specific, and nuanced policy that distinguishes between acceptable (e.g., augmentation) and unacceptable (e.g., misconduct) uses of generative AI.
- Co-Design Policy with Students:** The student voice is fractured (Table 7). To build consensus and ensure buy-in, the university must actively invite students to "be part of the conversation".⁹ A working group comprising faculty, administrators, and, crucially, students from diverse disciplines should be formed to co-design these new policies.
- Integrate Ethical AI Literacy into the Curriculum:** The 55.4% of students asking for formal training are correct. The university should develop and integrate modules on ethical AI literacy, not as a one-off workshop, but as a core component of the curriculum.⁴ This training must focus on ethics, verification, source attribution, data privacy, and understanding AI's limitations, thereby empowering students to be

critical and responsible users.

4. **Faculty Development and Assessment Re-Design:** Students are changing their habits (Table 6). Faculty must be supported with training on how to "stress-test" their assessments against AI.¹¹ More importantly, faculty should be trained to design *AI-inclusive* assignments (e.g., tasks that require students to critique an AI's output) that leverage AI to build, rather than atrophy, the critical thinking skills that students themselves are afraid of losing.

D. Future Research

This exploratory study opens several avenues for future research.

- **Replication:** This study should be replicated with a larger, stratified random sample to ensure generalizability and to allow for a more robust analysis of differences across all academic disciplines.
- **Qualitative Inquiry:** The quantitative data reveals *what* students are doing, but not always *why*. Qualitative methods, such as in-depth interviews or focus groups, are needed to explore the "why" behind these findings. For example, *why* do 30.4% of students feel AI is a valid substitute for professor consultation?
- **Longitudinal Study:** A longitudinal study is needed to track how student perceptions and behaviors (especially the data in Table 6) evolve over time, as they become more enmeshed with AI throughout their academic careers.

XI. REFERENCES

1. 10 Luo (2024), cited in Tandfonline.
2. 7 Vinyard, M., & Roosa, M. (2024), cited in JHU Press.
3. 1 MDPI. (2024). Generative AI in Education: Perspectives Through an Academic Lens. Signals.
4. 17 Abbasi, M. U. R., et al. (2024). Impact of AI on the inclusion of learners with special needs. PMC.
5. 13 Razmerita, L., et al. (2025). Student Perspectives on Generative AI Use in Higher Education: The Automation and Augmentation of Learning. Springer.
6. 10 Tandfonline. (2025). Logistic regressions on GenAI use and learning outcomes.
7. 5 Stanford University. (2024). Analyzing the impact of AI tools on student study habits and academic performance.
8. 6 Arxiv. (2024). AI systems assist students in organizing study routines.
9. 3 MDPI. (2025). Investigating the impact of AI technologies on students' learning processes.
10. 18 Emerald. (2024). Evaluating the impact of students' generative AI.
11. 11 HEPI. (2025). Student generative AI survey 2025.
12. 8 Campbell University. (2025). AI in Higher Education: A Summary of Recent Surveys.
13. 9 Jisc. (2025). Student perceptions of AI 2025.
14. 14 Jisc. (2025). Student perceptions of AI 2025. National Centre for AI.
15. 4 MDPI. (2024). University students' perceptions of the use of generative artificial intelligence.
16. 2 Harvard University. (2024). Academic integrity and teaching without AI.
17. 15 Frontiers in Education. (2024). Awareness of the tools and the consequences of their incorrect usage.
18. 16 PMC. (2021). Students' understanding of AI and ethics in K-12 settings.
19. 12 University of Kansas. (2024). Maintaining academic integrity in the AI era.

Works cited

1. Generative AI in Education: Perspectives Through an Academic Lens - MDPI, accessed on November 5, 2025, <https://www.mdpi.com/2079-9292/14/5/1053>
2. Academic Integrity and Teaching With(out) AI, accessed on November 5, 2025, <https://oaisc.fas.harvard.edu/academic-integrity-and-teaching-without-ai/>
3. The Impact of Artificial Intelligence (AI) on Students' Academic Development - MDPI, accessed on November 5, 2025, <https://www.mdpi.com/2227-7102/15/3/343>
4. Student Perceptions of the Use of Gen-AI in a Higher Education Program in Spain - MDPI, accessed on November 5, 2025, <https://www.mdpi.com/2673-6470/5/3/29>
5. Analyzing the Impact of AI Tools on Student Study Habits and Academic Performance - Stanford SCALE, accessed on November 5, 2025, <https://scale.stanford.edu/ai/repository/analyzing-impact-ai-tools-student-study-habits-and-academic-performance>
6. Analyzing the Impact of AI Tools on Student Study Habits and Academic Performance - arXiv, accessed on November 5, 2025, <https://arxiv.org/html/2412.02166v1>
7. Student Perspectives on Using Generative Artificial Intelligence for Research: A Qualitative Approach - Preprint, accessed on November 5, 2025, https://preprint.press.jhu.edu/portal/sites/default/files/08_25.4vinyard.pdf
8. AI in Higher Education: A Meta Summary of Recent Surveys of Students and Faculty, accessed on November 5, 2025, <https://sites.campbell.edu/academictechnology/2025/03/06/ai-in-higher-education-a-summary-of-recent-surveys-of-students-and-faculty/>
9. Student perceptions of AI 2025 - Jisc, accessed on November 5, 2025, <https://www.jisc.ac.uk/reports/student-perceptions-of-ai-2025>
10. Full article: Mastering knowledge: the impact of generative AI on student learning outcomes, accessed on November 5, 2025, <https://www.tandfonline.com/doi/full/10.1080/03075079.2025.2487570>
11. Student Generative AI Survey 2025 - HEPI - Higher Education Policy Institute, accessed on November 5, 2025, <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/>
12. Maintaining academic integrity in the AI era | Center for Teaching Excellence, accessed on November 5, 2025, <https://cte.ku.edu/maintaining-academic-integrity-ai-era>
13. Student Perspectives on Generative AI Use in Higher Education: The Automation and Augmentation of Learning - CBS Research Portal, accessed on November 5, 2025, <https://research.cbs.dk/en/publications/student-perspectives-on-generative-ai-use-in-higher-education-the>
14. Student Perceptions of AI 2025 - Artificial intelligence, accessed on November 5, 2025, <https://nationalcentreforai.jiscinvolve.org/wp/2025/05/21/student-perceptions-of-ai-2025/>
15. Ethical use of ChatGPT in education—Best practices to combat AI-induced plagiarism, accessed on November 5, 2025, <https://www.frontiersin.org/journals/education/articles/10.3389/feduc.2024.1465703/full>
16. Artificial intelligence in education: Addressing ethical challenges in K-12 settings - NIH, accessed on November 5, 2025, <https://pmc.ncbi.nlm.nih.gov/articles/PMC8455229/>
17. Exploring the effects of artificial intelligence on student and academic well-being in higher education: a mini-review - PubMed Central, accessed on November 5, 2025, <https://pmc.ncbi.nlm.nih.gov/articles/PMC11830699/>
18. Evaluating the impact of students' generative AI use in educational contexts | Journal of Research in Innovative Teaching & Learning | Emerald Publishing, accessed on November 5, 2025, <https://www.emerald.com/jrit/article/17/2/152/1226727/Evaluating-the-impact-of-students-generative-AI>