

4th International Conference on Media, Economy, Communication and Intelligence Management (MECI 2026)

Article

University student preferences for generative AI subscriptions

Yubin Liu^{1,*}

¹ School of Economics, Xiamen University, Xiamen, China

* Correspondence: Yubin Liu, School of Economics, Xiamen University, Xiamen, China

Abstract: Subscription design for generative artificial intelligence requires an understanding of which capabilities users value and which trade-offs they accept. This paper examines university student preferences for model capability, context capacity, privacy protection, and monthly price. The empirical material comes from a Chinese competition research report describing ten exploratory interviews and a choice-based conjoint survey with 828 retained respondents. The report applies conditional logit and latent class models to subscription choices that include a free alternative. Its reported willingness-to-pay estimates suggest that expert-level capability and enhanced privacy are important sources of perceived value, whereas context expansion does not show a stable aggregate premium. Three reported preference segments further suggest that capability demand, budget sensitivity, and privacy concerns should be considered jointly. Historical product examples from Anthropic and MiniMax illustrate alternative ways to translate technical capacity into service differentiation. Because individual choice records are unavailable and some statistical entries in the source report show discrepancies, the findings are treated as exploratory reported evidence; precise effect sizes and segment membership require replication. The paper develops implications for capability-based product tiers, credible privacy commitments, and task-specific context provision, while distinguishing candidate subscription prices from empirically optimized prices. Ultimately, this study provides actionable insights for technology developers aiming to align their subscription models with the specific needs and financial constraints of the academic user demographic.

Keywords: artificial intelligence; higher education; conjoint analysis; pricing; privacy; product design

Received: 07 August 2026

Revised: 13 September 2026

Accepted: 25 September 2026

Published: 01 October 2026



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Generative artificial intelligence (AI) services integrate multiple functional attributes—such as reliable reasoning, document processing, persistent interaction, and user control over submitted information—that users may value heterogeneously. University students routinely engage with these services in academic contexts including learning, writing, programming, and research [1]. A technically superior product does not automatically justify a higher subscription price if its enhancements are misaligned with users' actual tasks or exceed their willingness to incur recurring costs. For service providers, the central challenge is therefore to identify attribute combinations that deliver measurable value to target users, rather than treating all technical advances as inherently premium-worthy.

Empirical evidence indicates that generative AI can enhance professional writing performance, underscoring the importance of task productivity in shaping perceived service value [2]. However, productivity improvements alone do not determine willingness to pay. Students may continue relying on freely available tools, adopt multi-service workflows, or tolerate reduced output quality to avoid subscription fees.

Moreover, uncertainty regarding how unpublished text or original research data is handled—particularly concerning storage, usage, and retention—may deter adoption even among users who recognize a service’s functional strengths. Subscription decisions thus reflect a dual evaluation: anticipated gains in task efficiency and the terms under which those gains are realized.

This paper investigates three interrelated questions: Which subscription attributes are most influential in university students’ service selection? How do preference patterns vary across distinct user subgroups? And what implications arise for product design when observed preferences are contrasted with current industry strategies? The empirical foundation is drawn from a nationally administered survey conducted as part of the 2026 Chinese competition report *Zhi Jia Ji He*, which employed interviews, choice-based conjoint tasks, conditional logit modeling, and latent class analysis [3]. This paper restructures that evidence into a coherent academic argument and interprets aggregate findings through the lens of established theories of consumer choice and pricing behavior.

The contribution lies in offering a focused analysis of how perceived capability, data privacy assurance, and subscription value interact within a student user population. It also distinguishes three conceptually distinct constructs that are frequently conflated: the estimated monetary premium associated with a given attribute, the preference profile characterizing a specific user segment, and the profit-maximizing subscription price a provider would set under full market knowledge. The available evidence robustly supports exploratory interpretation of the first two constructs [4]. It does not permit estimation of the third, as cost structures, competitive dynamics, and real-world subscription uptake data remain unavailable.

2. Literature Review and Research Expectations

2.1. Digital Service Pricing and Task Value

Research on pricing information goods examines how pricing structures align with consumer heterogeneity and usage patterns. Nonlinear pricing models have been applied to analyze how differentiated service offerings can accommodate diverse user needs [5]. In the context of generative AI services, differentiation may encompass model access levels, usage allowances, task-specific support, and governance functionalities. Consequently, subscription value should be evaluated holistically at the service bundle level, while accounting for the availability and functionality of free alternatives.

For students, the practical utility of enhanced capabilities depends on the academic or project tasks they support. Improved reasoning accuracy may lower verification effort; superior coding assistance may accelerate programming task completion. Such functional benefits suggest a positive capability premium is plausible, though its magnitude requires empirical validation [6]. Accordingly, it is expected that higher prices diminish the appeal of otherwise equivalent plans (H1), and that advanced model capabilities increase perceived attractiveness relative to baseline assistance (H2).

2.2. Privacy Protection and Context Capacity

The economic value of privacy varies according to context, data type, and users’ awareness of how information is used. This consideration is particularly relevant in AI-assisted study environments, where users may submit unpublished arguments, identifiable interview materials, or original research data [7]. Privacy attributes influence perceived control over such materials and expectations regarding disclosure consequences. Consequently, stronger privacy protections are expected to enhance subscription appeal relative to the default data-training configuration (H3).

Context capacity relates differently to task value: increased capacity can aid processing of lengthy documents, yet yields diminishing returns when tasks already fit within baseline limits. The survey’s highest context level integrates expanded token capacity with persistent memory, meaning its estimated effect reflects a composite feature rather than an isolated increase in token allowance. A positive context premium is

therefore hypothesized as H4, while acknowledging that actual task demands and concerns about persistent memory may attenuate or even reverse this effect.

2.3. Attribute Trade-Offs and Preference Heterogeneity

Conjoint analysis assesses how consumers balance trade-offs among product attributes when evaluating alternatives. This approach is grounded in random utility theory, which links observed choices to underlying systematic utility and unobserved factors [7]. Incorporating price as an attribute allows estimated attribute values to be expressed in monetary terms, contingent upon model assumptions. A free alternative holds particular relevance in this context, as student users may decline both paid options without discontinuing use of generative AI.

Aggregate preference estimates may obscure distinct configurations of user needs. Latent class choice models capture such heterogeneity by identifying a finite number of distinct preference profiles [8]. Within this setting, high demand for functional capability does not necessarily coincide with high demand for contextual richness, and heightened privacy concerns do not inherently translate into willingness to select the highest-priced plan. Variations across academic disciplines and latent classes are explored as exploratory research questions. Neither a fixed hierarchy of disciplines nor a pre-specified number of market segments is assumed as a theoretical requirement.

3. Methodology

3.1. Research Design and Empirical Material

The empirical material comprises the source report's questionnaire, interview summaries, aggregate tables, results narrative, and appended analysis code. The underlying project conducted ten online interviews with two junior-college students, four undergraduates, two postgraduate students, and two doctoral students to identify relevant attributes. A pilot survey administered from February 2 to 6, 2026, collected 76 responses, of which 55 were retained [9]. Feedback from this phase informed revisions to the wording and length of the questionnaire.

The formal survey collected 972 questionnaires, with 828 retained after screening. Eligibility required use of generative AI within the preceding three months. Respondents were drawn from all 31 provincial-level administrative regions in mainland China. Although the report states that provincial allocation was intended to be proportional to the square root of student population, it does not provide a realized sampling frame, inclusion probabilities, or survey weights [10]. Consequently, findings are interpreted as descriptive of the retained sample, without claims to national representativeness.

This paper does not re-estimate the models, as individual responses and original estimation files are unavailable. Descriptive percentages are derived from reported counts where those counts are internally consistent. Model outputs are attributed solely to the source report, and inconsistent inferential statements are excluded from presentation as independently verified estimates [11]. Appendix A outlines the principal limitations affecting numerical interpretation.

3.2. Choice Tasks and Attribute Construction

Respondents compared two hypothetical paid subscription options against continuing with a free baseline service. Table 1 summarizes the four attributes used in the design. Monthly paid prices were set at CNY 9, 29, 49, and 79. The questionnaire included 12 displayed choice tasks, one of which served as an instructed-response check in task 5. That check was excluded from analysis, leaving 11 substantive tasks. With full responses, 828 participants would generate 9,108 choice observations; this figure reflects the experimental design's theoretical maximum, not a verified count of usable estimation data.

Table 1. Subscription Attributes and Levels in the Source Questionnaire.

Attribute	Levels	Model coding
Model capability	Basic assistance; advanced reasoning; expert creation	Two indicators; basic assistance is the reference
Context capacity	About 50 pages; about 500 pages or codebase input; a large knowledge collection with persistent memory	Two indicators; standard capacity is the reference
Privacy protection	Conversations used for training; strict confidentiality described as no retention or disclosure	One indicator; training use is the reference
Monthly price	CNY 9, 29, 49, or 79 for paid plans; CNY 0 for the free baseline	Continuous price variable

The attribute descriptions in Table 1 represent hypothetical service offerings, not empirically validated product capabilities. Specifically, the highest capability level describes enhanced writing and analytical functionality, while the strict-privacy level asserts absolute confidentiality. These formulations aim to elicit respondent perceptions toward aspirational claims, rather than serving as objective benchmarks for model performance or enforceable security assurances [12].

3.3. Conditional Logit and Willingness to Pay

The source report applies conditional logit to model choices as functions of subscription attributes. Within the random utility framework, the utility of alternative j in task t for respondent n consists of systematic utility and an error term. The specification includes a paid-alternative constant, price, and indicator variables for nonbaseline attribute levels:

$$V_{ntj} = \alpha B_{ntj} + \beta_p p_{ntj} + \beta_{s2} S_{2ntj} + \beta_{s3} S_{3ntj} + \beta_{c2} C_{2ntj} + \beta_{c3} C_{3ntj} + \beta_r R_{ntj} \quad (1)$$

Here, B equals one for either paid alternative and zero for the free baseline; S_2 and S_3 identify advanced and expert capability; C_2 and C_3 identify the two context upgrades; and R identifies strict privacy. All attribute indicators and price are zero for the free alternative, whose systematic utility is normalized to zero. This coding follows the appended script and avoids interpreting the paid-alternative constant as a constant for nonpurchase.

Assuming independent type-I extreme-value errors, choice probabilities follow the logit form. An attribute-level willingness to pay (WTP) is obtained by dividing its utility coefficient by the negative price coefficient:

$$P_{ntj} = \frac{\exp(V_{ntj})}{\sum_{k \in J_t} \exp(V_{ntk})}, \quad WTP_a = -\frac{\beta_a}{\beta_p}. \quad (2)$$

These quantities describe hypothetical choices within the specified alternatives. Repeated tasks from one person do not constitute independent observations. Valid uncertainty estimates should account for within-person correlation, for example through respondent-level resampling. The source script groups observations by choice occasion but does not document such an adjustment. Moreover, differing attribute levels do not validate the assumption concerning substitution structure among alternatives.

3.4. Heterogeneity and Case Comparison

The source report estimates separate models for broad academic groups and describes a three-class latent choice model. Conceptually, the latent class likelihood combines a person's sequence of conditional choice probabilities with class-membership probabilities [13]. Class selection should be supported by fit comparisons, convergence checks, and stability across starting values. Because those outputs are not supplied, the three reported profiles are used as an exploratory segmentation rather than an established market taxonomy.

Anthropic and MiniMax are retained as illustrative cases from the original report. The comparison uses dated provider announcements and a technical report that predate

the February 2026 survey. It examines capability differentiation, context provision, and governance. The cases offer contextual interpretation; they do not independently validate student willingness-to-pay estimates or the causal effectiveness of a pricing strategy [14].

4. Results

4.1. Sample Characteristics and Use Patterns

Table 2 presents internally reconcilable sample counts. Undergraduates constitute 580 of the 828 respondents, representing 70.0%. Monthly living expenditure is CNY 2,500 or less for 596 respondents, accounting for 72.0%. Current paid artificial intelligence use is reported by 353 respondents (42.6%), whereas 475 (57.4%) indicate no payment. Because the payment inquiry encompasses both subscriptions and API recharges, this metric is broader than recurring consumer-subscription adoption.

Table 2. Reported Characteristics of the Retained Sample.

Characteristic	Category	Respondents	Percentage
Gender	Female	399	48.2
Gender	Male	429	51.8
Study stage	Lower-year undergraduate	224	27.1
Study stage	Upper-year undergraduate	356	43.0
Study stage	Master's student	186	22.5
Study stage	Doctoral student	33	4.0
Study stage	Junior-college student	29	3.5
Monthly living expenditure	CNY 1,500 or less	134	16.2
Monthly living expenditure	CNY 1,501–2,500	462	55.8
Monthly living expenditure	CNY 2,501–3,500	167	20.2
Monthly living expenditure	Above CNY 3,500	65	7.9
Current paid AI use	Yes	353	42.6
Current paid AI use	No	475	57.4

Note: Counts are from the source report [2]; percentages are recalculated using $N = 828$. Rounding may prevent category percentages from summing to exactly 100.

Knowledge search and learning represent the most frequently reported applications, followed by writing, programming, and academic research (Figure 1). These applications underscore the relevance of task performance to subscription design, although they do not establish that every student requires identical advanced functionality [14]. The coexistence of routine learning and specialized tasks motivates a detailed examination of differences in attribute preferences.




Figure 1. Reported Generative AI Use Cases. Percentages Use 828 Respondents as the Denominator. Multiple Responses Were Allowed, So the Percentages Are Not Market Shares and Do Not Sum to 100. Counts Are Taken from the Source Report [2].

4.2. Reported Attribute Preferences

The report consistently describes a negative price association and positive associations for enhanced capability and privacy. Its narrative reports monthly premiums

of CNY 10.04 for advanced reasoning, CNY 25.80 for expert creation, and CNY 24.15 for strict privacy (Table 3). These are source-reported values, not estimates reproduced in this paper. Their numerical precision exceeds what can presently be verified from the inconsistent regression entries.

Table 3. Attribute-level Willingness to Pay as Reported in the Source Narrative.

Upgrade from the baseline	Reported premium in CNY per month	Interpretation permitted by the available material
Advanced reasoning	10.04	Reported positive capability premium
Expert creation	25.80	Reported premium above basic assistance
Strict privacy	24.15	Reported privacy premium close to the expert-capability premium
Context expansion	Not reproduced	Numerical entries conflict; the aggregate narrative reports no stable premium

Note: These are attributed report values [2], not newly estimated results. Standard errors, confidence intervals, and significance markers are omitted because the supplied statistical entries cannot be reconciled. The premiums refer to individual upgrades from baseline, not complete subscription prices.

The reported directions are consistent with H1–H3, but the available record does not support renewed significance testing. Expert capability and strict privacy have similar reported monetary values. Consequently, the evidence does not justify describing privacy as a minor consideration or claiming that capability is statistically more valuable. Comparing the two would require an uncertainty estimate for their difference.

For context clarification, the report's aggregate interpretation provides insufficient support for H4. This finding should be read as an absence of a stable reported premium under the questionnaire's descriptions. It does not show that long-context capability has no value, that the true effect is zero, or that industry investment is excessive. The baseline already permits roughly 50 pages, and the highest level also introduces persistent memory, which may alter the perceived privacy implications of the offer.

4.3. Exploratory Preference Segments

The estimated latent class shares are 65.58%, 7.73%, and 26.69% (Table 4). The largest group is characterized as capability-oriented and exhibits relatively low price sensitivity. A smaller group values multiple functions but also shows pronounced responsiveness to price changes. The third group prioritizes privacy protection above other attributes [3]. These distinct preference profiles demonstrate that a single average willingness-to-pay estimate may inadequately inform product design decisions.

Table 4. Exploratory Latent Class Profiles Reported in the Source.

Profile	Reported sample share	Reported preference pattern	Product-design implication
Capability-oriented	65.58%	Strong interest in expert capability; relatively weak price response	Test a capability-focused offer with transparent usage limits
Broad-function and value-conscious	7.73%	Values several features while remaining price responsive	Test useful bundles rather than presume acceptance of a high price
Privacy-oriented	26.69%	Strong privacy preference; limited interest in some feature upgrades	Test a privacy-focused offer without unnecessary feature bundling

Note: Shares and profile descriptions originate in the report [2]. Class-selection statistics, membership uncertainty, and a complete estimation output are unavailable. These figures are not national population proportions or predicted market shares.

Disciplinary comparisons further indicate variation in attribute emphasis: users from economics, management, and law backgrounds tend to be more capability-oriented, whereas those in humanities-related fields place comparatively greater importance on privacy. However, the academic categories used are nonstandard, and their reported frequencies lack internal consistency. Moreover, logit coefficients estimated separately across groups are not directly comparable due to potential differences in scale. Collectively, these findings support investigating preference heterogeneity, though they do not establish causal links between academic discipline and preference structure.

5. Discussion

5.1. Interpreting Capability and Privacy as Subscription Value

The central implication is that capability and privacy can be complementary dimensions of a service's usefulness. A student may value accurate assistance but avoid submitting the material needed to obtain it. Privacy commitments can therefore affect whether capability is usable in a particular workflow. This interpretation aligns with established understandings of privacy valuation in context, although the current study does not directly test an interaction between capability and privacy.

Context capacity requires a similarly task-specific interpretation. For a short assignment, increasing an already sufficient window may add little visible benefit. For a code repository or a collection of research documents, the same increase may be consequential. A provider should therefore connect context features to observable improvements in task completion and retrieval quality. The survey does not isolate the contribution of capacity from persistent memory, making a blanket recommendation to reduce long-context investment unwarranted.

5.2. Anthropic and MiniMax as Contextual Cases

In March 2024, a major AI developer introduced a new model family, distinguishing variants by performance characteristics, inference speed, and pricing, while initially specifying a uniform 200K token context window across all versions. This illustrates how consistent context capacity can coexist with differentiation along other capability dimensions. Later that year, an enterprise-focused offering extended the context window to 500K tokens, incorporated organizational access controls, and included a commitment not to use customer-provided conversations or content for model training. Collectively, these developments indicate that both functional capabilities and governance features contribute to product differentiation, and that extended context remains positioned within premium-tier offerings [8]. However, they do not isolate context length as a decisive factor in purchasing decisions.

Another firm's January 2025 technical report describes a model architecture supporting a one-million-token training context and demonstrating inference-time extrapolation to four million tokens. This reflects substantial technical investment in long-context functionality. While no student subscription premium is documented, the case suggests that context value may be realized in application domains beyond those typically assessed in academic surveys. Consequently, it serves primarily to delineate the scope of student-oriented findings rather than to support broad industry generalizations.

These examples involve distinct product types and deployment contexts—API services, enterprise solutions, and consumer subscriptions—which are not directly comparable units of analysis. Their shared relevance lies in product architecture: firms may decouple capabilities, integrate them with workflow tools, or modulate access conditions. Whether any such configuration yields higher revenue specifically within university markets remains an empirical question requiring direct evaluation.

5.3. Implications for Product Design and Pricing

A practical starting point is a small set of task-relevant service offerings. An entry-level offering can provide reliable everyday assistance with clearly defined usage allowances and transparent data-use terms. A research-oriented offering can emphasize reasoning quality, coding support, or document workflow capabilities. An enhanced governance offering may be introduced for users requiring additional administrative controls or granular data management functions, while preserving a consistent baseline of privacy protection across all user tiers. These represent design considerations rather than definitive conclusions regarding the optimal number of service plans.

The source report identifies CNY 29, 49, and 79 as candidate monthly price points. These figures serve as initial test values derived from the survey's stated price range. Attribute-based willingness-to-pay analysis cannot determine the optimal price for a fully specified plan, and reported segment preference shares do not reflect actual conversion rates. A rigorous model-based pricing exercise would require internally consistent parameter estimates, explicitly defined plan features, a credible outside option, usage-dependent cost structures, and assumptions about competitive alternatives. Subsequent field testing would be necessary to measure real-world conversion, retention, and usage behavior across alternative pricing configurations.

For universities, the findings underscore the importance of affordability and equitable access conditions alongside technical capability. Institutional provision may help ensure students have consistent access to essential tools; however, the current evidence does not permit quantification of its educational impact or distributional effects. Student status and living expenditure should not be interpreted as direct proxies for financial need. Decisions regarding subsidized access must therefore rely on information beyond the observed preference profiles.

5.4. Limitations and Future Research

A key limitation lies in the empirical foundation. Aggregate reporting enables examination of broad trends but prevents verification of model estimates or reconstruction of individual respondent choices [10]. Statistical inconsistencies preclude definitive hypothesis testing, precise cross-group comparisons, or validated forecasting. Consequently, the study remains exploratory in nature.

The choice experiment captures stated preferences under hypothetical scenarios. Respondents may interpret expert capability, confidentiality, and memory differently, and their stated choices may not align with actual payment behavior. Future designs should disentangle context size from persistent memory, articulate privacy safeguards using concrete and credible language, and randomize or systematically document task presentation order. Additionally, an outside option should reflect alternatives realistically available to students.

A replication effort should implement independent quality-screening criteria, retain respondent-level dependence structures, compare multiple preference modeling approaches, and assess predictive performance on held-out respondents. Longitudinal or field-based evidence would be essential to examine subscription retention, price adaptation over time, and evolving perceptions of AI service quality—enabling a shift from exploratory design toward rigorous analysis of dynamic pricing mechanisms [2, 12].

6. Conclusion

University student demand for generative AI subscriptions is best understood as a set of trade-offs among task capability, privacy, and price. Evidence indicates meaningful interest in enhanced capabilities and stronger privacy protection, while revealing no stable aggregate willingness to pay for the bundled contextual enhancement feature. Segment-level patterns further suggest that users prioritize distinct combinations of functionality and affordability. Case examples from Anthropic and MiniMax illustrate how these dimensions can inform product design, while also underscoring that student preference findings are not broadly generalizable across AI markets. The resulting recommendation is to test differentiated, task-relevant offers supported by transparent

data practices. Precise willingness-to-pay estimates, segment sizes, and optimal pricing require replication using a consistent respondent-level dataset.

7. Appendix Numerical Provenance

The descriptive tables use the report's main sample of 828, consistent with gender, education, expenditure, and payment counts. A concluding passage cites 825, while a model-screening section reports 497. Discipline counts sum to 831, precluding reproduction of the disciplinary distribution.

The regression table lists a price coefficient of -0.01513 , whereas the narrative states -0.0185 . Several coefficients, standard errors, z statistics, and confidence limits also differ. For instance, the reported paid-alternative coefficient lies outside its listed confidence interval. Consequently, the WTP values in Table 3 are quoted from the narrative without inferential statistics or claim of recoverability from that table.

The appended script applies an upper cap of CNY 12 to privacy WTP draws in one calculation and later generates uncapped results. The origin of the narrative's CNY 24.15 figure remains unidentified. The script also excludes the 40% of respondents with the lowest fitted choice likelihood prior to re-estimation; improved fit following this selection does not constitute independent validation of data quality. The reported pseudo- R^2 of 0.2894 and hit rate of 64.19% are therefore not interpreted as evidence of out-of-sample performance. No individual responses have been reconstructed or simulated.

References

1. S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023.
2. A. Sundararajan, "Nonlinear pricing of information goods," *Management Science*, vol. 50, no. 12, pp. 1660–1673, 2004.
3. A. Acquisti, C. Taylor, and L. Wagman, "The economics of privacy," *Journal of Economic Literature*, vol. 54, no. 2, pp. 442–492, 2016.
4. P. E. Green and V. Srinivasan, "Conjoint analysis in marketing: new developments with implications for research and practice," *Journal of Marketing*, vol. 54, no. 4, pp. 3–19, 1990.
5. D. McFadden, "Conditional logit analysis of qualitative choice behavior," 1972.
6. W. A. Kamakura and G. J. Russell, "A probabilistic choice model for market segmentation and elasticity structure," *Journal of Marketing Research*, vol. 26, no. 4, pp. 379–390, 1989.
7. K. Train, *Discrete Choice Methods with Simulation*, 2nd ed. Cambridge, UK: Cambridge University Press, 2003.
8. S. Joshi, "Open-source vs. commercial coding assistants: A 2025 comparison of deepseek r1, qwen 2.5 and claude 3.7," *International Journal of Computer Applications Technology and Research*, vol. 14, no. 09, pp. 6–18, 2025.
9. A. Li et al., "Minimax-01: Scaling foundation models with lightning attention," *arXiv preprint arXiv:2501.08313*, 2025.
10. P. C. Hong, Y. B. Choi, and Y. S. Park, "AI diffusion and the new triad of supply chain transformation: Productivity, perspective, and power in the era of Claude, ChatGPT, Gemini, LLaMA, and Mistral," *Logistics*, vol. 10, no. 2, p. 40, 2026.
11. O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, "Systematic review of research on artificial intelligence applications in higher education—where are the educators?," *International Journal of Educational Technology in Higher Education*, vol. 16, no. 1, p. 39, 2019.
12. H. Crompton and D. Burke, "Artificial intelligence in higher education: The state of the field," *International Journal of Educational Technology in Higher Education*, vol. 20, no. 1, p. 22, 2023.
13. F. Jafari and A. Keykha, "Identifying the opportunities and challenges of artificial intelligence in higher education: a qualitative study," *Journal of Applied Research in Higher Education*, vol. 16, no. 4, pp. 1228–1245, 2024.
14. M. Bearman, J. Ryan, and R. Ajajawi, "Discourses of artificial intelligence in higher education: A critical literature review," *Higher Education*, vol. 86, no. 2, pp. 369–385, 2023.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.