

ACCEPTABLE USES OF GENAI: POST-SECONDARY STUDENTS' JUDGEMENTS

RAHUL KUMAR
BROCK UNIVERSITY

ROBERT MCGRAY
BROCK UNIVERSITY

Abstract

Generative artificial intelligence (GenAI) is reshaping academic practice in higher education, yet evidence on how students evaluate its acceptable uses remains limited. This study surveyed 988 post-secondary students across 39 countries, recruited through the Prolific online research platform, on the acceptability of GenAI use across 22 academic tasks using a binary Yes/No format. Using an interpretive framework grounded in normative boundary-setting, academic integrity values, and academic socialization, the analysis reveals differentiated acceptability patterns: Students endorse GenAI for instrumental and research-support tasks while maintaining normative boundaries around evaluative and judgement-intensive activities. Gender differences are minimal; level of study is the primary source of variation, with graduate students reporting higher acceptability for research- and professionally oriented uses. Respondents residing in Canada ($n = 58$) reported significantly lower acceptance of four text-focused tasks, including academic composition and summarization, than the non-Canadian cohort. Findings have implications for institutional policy, pedagogy, and academic integrity frameworks.

Keywords: GenAI, academic integrity, student perceptions, higher education, normative boundary-setting

Résumé

L'intelligence artificielle générative (IA générative ou GenAI) transforme les pratiques pédagogiques dans l'enseignement supérieur, mais les données probantes sur la manière dont les étudiants évaluent les usages acceptables de cette technologie demeurent limitées. Cette étude a sondé 988 étudiant[e]s de l'enseignement postsecondaire provenant de 39 pays, recruté[e]s par l'entremise de la plateforme de recherche en ligne Prolific, afin d'évaluer l'acceptabilité de l'utilisation de l'IA générative dans 22 tâches scolaires au moyen d'un format de réponse binaire (oui/non). En s'appuyant sur un cadre interprétatif fondé sur l'établissement de frontières normatives, les valeurs d'intégrité scolaire

et de socialisation scolaire, l'analyse met en évidence des profils différenciés d'acceptabilité : les étudiant[e]s approuvent l'utilisation de l'IA générative pour les tâches instrumentales et de soutien à la recherche, tout en maintenant des frontières normatives autour des activités évaluatives et de celles exigeant un jugement approfondi. Les différences selon le genre sont minimales; le niveau d'études constitue la principale source de variation, les étudiant[e]s aux cycles supérieurs rapportant une acceptabilité plus élevée des usages liés à la recherche et aux activités à visée professionnelle. Les répondant[e]s résidant au Canada ($n = 58$) ont déclaré un niveau d'acceptation significativement plus faible pour quatre tâches axées sur la production de textes, notamment la rédaction et la synthèse de travaux universitaires, comparativement aux répondant[e]s ne résidant pas au Canada. Ces résultats ont des implications pour les politiques institutionnelles, les pratiques pédagogiques et les cadres de référence en matière d'intégrité scolaire.

Mots-clés: IA générative; intégrité scolaire; perceptions étudiantes; enseignement supérieur; établissement de frontières normatives

INTRODUCTION: PERCEIVED PERMISSIBILITY OF GENAI USAGE

The rapid diffusion of GenAI in higher education has prompted sustained debate about its appropriate role in academic work. Scholarly and institutional discussions have been marked by polarized positions, ranging from calls to permit and integrate GenAI into teaching and learning practices (Chan & Hu, 2023; Christ-Brendemühl, 2024; Halaweh, 2023; Humble, 2025; McGuire, 2023) to arguments advocating restriction or prohibition based on concerns about academic integrity, authorship, and assessment validity (Alrayes et al., 2024; Duah & McGivern, 2024; Sullivan et al., 2023). These debates reflect broader uncertainty about how emerging AI technologies intersect with long-standing academic norms.

Despite the intensity of this discourse, empirical attention to students' perspectives has been comparatively limited. Existing studies document widespread student engagement with GenAI tools and explore motivations for use through small-scale surveys and qualitative research. However, large-scale empirical

evidence on how students distinguish between acceptable and unacceptable uses of GenAI across specific academic tasks remains scarce. In particular, little is known about whether students' judgements reflect indiscriminate acceptance, ethical ambiguity, or more differentiated normative reasoning.

This study addresses this gap by reporting findings from a large-scale survey of post-secondary students' judgements of acceptable GenAI use across a wide range of academic activities. Guided by an interpretive framework grounded in normative boundary-setting, academic integrity values, and academic socialization, the study maps task-level acceptability patterns. It examines differences by gender and level of study. A supplementary comparative analysis of Canadian respondents further situates the empirical contributions within the Canadian post-secondary institutional context. By foregrounding student perspectives, the study contributes empirical clarity to ongoing debates about GenAI use in higher education and informs institutional responses that align with students' existing normative reasoning.

THEORETICAL FRAMEWORK

This study employs an interpretive framework to examine post-secondary students' normative judgements regarding acceptable uses of GenAI across academic practices. While the rapid integration of GenAI into higher education has provoked extensive debate about appropriate boundaries for these technologies (Chan & Hu, 2023; Sullivan et al., 2023), empirical evidence concerning how students themselves conceptualize these boundaries remains limited. Rather than approaching acceptability through technology adoption models (Davis, 1989; Venkatesh et al., 2003), this framework positions students' judgements as instances of normative boundary-setting within higher education communities (Gieryn, 1983; Lamont & Molnár, 2002).

Normative Boundary-Setting as Core Construct

Normative boundary-setting refers to the social processes through which academic communities delineate acceptable and unacceptable practices based on shared understandings of academic work (Gieryn, 1983; Lamont & Molnár, 2002). These boundaries are socially constructed and institutionally negotiated rather than determined solely by formal rules or technological affordances. In higher education, such boundaries shape interpretations of authorship, intellectual responsibility, and appropriate tool use across assessed and non-assessed activities (Gellai, 2022).

GenAI disrupts established boundaries by blurring distinctions between human cognition and machine assistance. As Bearman et al. (2022) demonstrate, debates surrounding AI in higher education reflect deep uncertainty about epistemic authority and academic legitimacy. Questions of acceptability, therefore, extend beyond technical competence or awareness and instead probe where legitimate academic practice ends and impermissible substitution begins.

If acceptability judgements constitute instances of normative boundary-setting rather than purely individual preferences, they should exhibit systematic patterning by task function. Research on educational technology suggests a plausible differentiation: tools perceived as supporting learning processes are accepted more readily than those that encroach upon evaluative or authoritative functions (Kang et al., 2022; Or, 2023). Extending this logic to GenAI, the framework anticipates that students will distinguish between GenAI as an academic tool and GenAI as an academic agent. Whether, and how sharply, this distinction structures students' judgements is an empirical question the present study addresses.

Academic Integrity Values as Interpretive Anchors

To interpret these boundary-setting judgements, the framework employs academic integrity values as interpretive anchors rather than as compliance metrics. The International Center for Academic Integrity (ICAI; 2021) identifies honesty, trust, fairness, respect, responsibility, and courage as foundational values underpinning academic practice. These values function as shared reference points through which academic communities evaluate the legitimacy of practices (McCabe et al., 2012). High acceptance of AI-assisted summarization or proofreading likely reflects judgements that such practices enhance efficiency without compromising authorial honesty or fairness. Conversely, resistance to AI involvement in grading, feedback provision, or assessment design suggests that students associate these activities with trust relationships, accountability structures, and institutional authority that exceed task completion alone.

This interpretation aligns with findings showing consistent caution toward AI use in evaluative contexts across demographic groups. Activities involving academic judgement, such as composing feedback or examination questions, appear to remain symbolically and practically reserved for human actors. Framing results through academic integrity values situates student judgements within established academic

norms rather than reducing them to compliance concerns or misconduct frameworks. This approach recognizes students as legitimate participants in ethical deliberation, consistent with Valerio's (2024) emphasis on students as key stakeholders in AI ethics discussions.

Academic Socialization as an Analytical Lens for Level Differences

Differences between undergraduate and graduate students are interpreted through academic socialization theory, which examines how students internalize the norms, expectations, and responsibilities of academic participation over time (Weidman et al., 2001). This process intensifies as students advance through educational levels and assume differentiated academic roles (Austin, 2002; Gardner, 2008).

Graduate education involves a shift from structured knowledge consumption toward independent research, scholarly writing, and, in some cases, instructional or evaluative responsibilities. This role differentiation provides a lens for interpreting findings showing that graduate students report higher acceptability across multiple AI-supported practices, particularly those related to research production and scholarly work. Greater acceptance of AI for tasks such as composing research reports or data analysis and interpretation reflects graduate students' engagement with knowledge production rather than diminished integrity concerns.

LITERATURE REVIEW

GenAI in Higher Education: Normative Uncertainty and Policy Fragmentation

GenAI technologies have been rapidly adopted across higher education, prompting renewed scrutiny of academic practices related to learning, assessment, and authorship. Empirical studies and sector reports consistently indicate widespread student engagement with GenAI tools, alongside heightened concern about their

implications for academic integrity and educational quality (Kumar & McGray, 2024). Scholarly work has increasingly framed the resulting debates as normative rather than purely technical: Bearman et al. (2022), in a critical discourse analysis of how AI is invoked across prominent higher education journals, reveal that references to AI are consistently entangled with contested constructions of authority and agency—redistributed across institutions, teachers, students, and machines—raising questions that extend well beyond technical implementation.

At the institutional level, analyses of AI governance reveal considerable variation in how these normative questions are addressed. Studies of university policies demonstrate fragmented and uneven approaches to regulating AI use in learning, teaching, and research, ranging from restrictive guidance to more permissive or ambiguous frameworks (Atkinson-Toal & Guo, 2024; Gellai, 2022; Spivakovsky et al., 2023). In the Canadian post-secondary context specifically, this policy fragmentation is well-documented. Da Mota (2024) found no existing or proposed Canadian legislation addressing AI use in research institutions, and his survey of university-level policies revealed a multi-tiered landscape ranging from basic classroom guidelines to formal task forces.

Positioning students as normative agents rather than passive recipients of policy supports the examination of acceptability judgements as indicators of how academic norms are internalized and applied (Bai et al., 2023; Deckker & Sumanasekara, 2025). From this standpoint, student responses to acceptability questions reflect situated evaluations informed by academic expectations, task demands, and integrity considerations, rather than mere personal preference.

Task-Based Differentiation in Educational Technology Use

Research on educational technology has long recognized that acceptance and use vary by task type (Kang et al., 2022; Or, 2023). Studies examining digital tools in higher education consistently show that technologies perceived as

supporting learning processes are more readily accepted than those that appear to encroach on evaluative or authoritative functions. This distinction is particularly salient in discussions of AI, where the delegation of cognitive tasks raises questions about authorship, judgement, and responsibility.

However, much of the existing literature addresses task differentiation implicitly or within narrow contexts. Few studies have systematically examined acceptability across a broad range of academic activities, spanning writing support, research practices, professional communication, content creation, and assessment-related functions (Kamalov et al., 2023). As a result, the contours of students' task-based normative boundaries around GenAI remain insufficiently mapped.

Positioning of the Present Study

Existing research highlights significant gaps in the empirical literature on GenAI in higher education. While institutional policies and ethical frameworks continue to evolve unevenly, and while adoption-focused studies document widespread engagement with AI tools, there remains limited empirical evidence on how students evaluate the acceptability of specific AI-supported academic practices. Moreover, task-based distinctions and level-of-study differences have rarely been examined systematically through a normative lens.

The present study addresses these gaps by empirically mapping post-secondary students' acceptability judgements across a wide range of academic tasks, with particular attention to differences by gender and level of study. By situating these judgements within an interpretive framework grounded in normative boundary-setting, academic integrity, and academic socialization, the study contributes empirical clarity to ongoing debates about GenAI use in higher education.

METHODOLOGY

This study employed a cross-sectional survey design to examine post-secondary students' perceptions of acceptable uses of GenAI in academic contexts. A survey instrument, developed from an earlier study by the authors, was adapted and expanded to reflect the growing diversity of GenAI-supported academic practices. The final survey included 22 items assessing whether specific uses of GenAI were considered acceptable in post-secondary education contexts. Each acceptable use item required a binary response (Yes/No) indicating whether the respondent considered the use acceptable. In addition to acceptability judgements, participants reported whether they had used GenAI for academic work within the past year. The items were presented in a fixed order through Qualtrics survey software. The items spanned a range of academic activities, including writing support, research-related tasks, professional communication, content creation, and evaluative or judgement-intensive practices. Internal consistency across the acceptability items was high (Cronbach's $\alpha = .897$), indicating coherent response patterns.

Participants were recruited using the Prolific online research platform. Eligibility criteria required participants to be currently enrolled in a post-secondary education (PSE) institution and to be fluent in English, as the survey was administered in English. Participants completed the survey in a single sitting and were compensated in accordance with Prolific's standard compensation guidelines.

A total of 988 participants completed the survey. Respondents represented 39 countries, reflecting an international post-secondary student population. Of the sample, 53.2% ($n = 526$) identified as male and 46.8% ($n = 462$) identified as female. In addition, 29.4% of participants reported using a paid version of a GenAI tool at the time of data collection. Participants self-reported their level of study as undergraduate or graduate at the time of survey completion. Data were collected in February 2025.

Descriptive statistics were used to summarize overall acceptability patterns across the 22 GenAI use categories. Group comparisons were conducted to examine differences in acceptability judgements by gender and level of study (undergraduate vs. graduate). Given the categorical nature of the outcome variables, chi-square tests of independence were used to assess statistical significance.

A secondary comparative analysis examined Canadian respondents ($n = 58$) against the non-Canadian cohort ($n = 930$). Canadian status was operationalized by current country of residence, as reported in response to the survey item “In which country do you currently reside?”; no nationality-based screening was applied during recruitment. The category therefore reflects residence in Canada rather than Canadian nationality and encompasses both domestic and international students living in the country. Given the eligibility requirement of current post-secondary enrolment, residence serves as a close proxy for enrolment at a Canadian institution, although institutional location was not verified directly.

Given the exploratory and descriptive purpose of the analysis, no adjustment for multiple comparisons was applied. Instead, emphasis was placed on effect sizes and the consistency of observed patterns across related tasks to support substantive interpretation.

Because the outcome variables were binary (acceptable vs. not acceptable), effect sizes were reported as odds ratios (OR) and risk differences (RD), expressed as percentage point differences. Odds ratios indicate the relative likelihood that one group endorsed a given AI use as acceptable compared to another and are standard for dichotomous outcomes; however, as they can be less intuitive for applied interpretation, risk differences were also calculated to convey the absolute magnitude of group differences.

Cohen’s d was not used, as it is designed for continuous outcomes and assumes normally distributed variables. In the context of binary acceptability judgements, odds ratios and risk differences provide more appropriate and transparent indicators of effect magnitude.

FINDINGS

The findings are organized as follows. Primary descriptive and comparative results are reported for the full international sample ($N = 988$), including overall acceptability patterns across 22 GenAI use categories and group-level differences by gender and level of study. Subsequently, a secondary comparative analysis examines Canadian respondents ($n = 58$) against the non-Canadian cohort ($n = 930$), situating the empirical findings within the Canadian post-secondary institutional context.

Participant Demographics and Overall AI Usage

The study sample comprised 988 post-secondary students representing 39 countries, with a balanced gender distribution (53.2% male, 46.8% female) and near-equal representation of undergraduate (52.0%) and graduate (48.0%) students. In addition, 29.4% of participants reported using paid versions of GenAI tools, suggesting meaningful engagement with these technologies for academic work.

Overall adoption of GenAI was high, with 93.7% of respondents reporting use for academic purposes within the past year. A statistically significant gender difference was observed ($\chi^2 = 9.35$, $df = 2$, $p = .009$), with males reporting slightly higher usage rates (95.8%) than females (91.6%), though the magnitude of this difference was modest. Most participants (83.9%) also indicated an intention to continue using GenAI for academic work in the following year, pointing to sustained rather than transient engagement.

Acceptability Patterns Across AI Applications

Participants evaluated the acceptability of GenAI use across 22 distinct academic application categories (see Table 1). Analysis revealed clear differentiation between task types, with distinct clustering patterns emerging across functional domains.

Table 1
Twenty-Two GenAI Use Categories: Respondent Agreement Frequencies and Proportions

Variable	Meaning	<i>n</i> (Yes)	Yes %
AUSE1	Email to students	710	71.86
AUSE2	Email to teachers	683	69.13
AUSE3	Emails to administrators	663	67.11
AUSE4	Emails to external parties	725	73.38
AUSE5	Composing reference letters	588	59.51
AUSE6	Composing research reports	526	53.24
AUSE7	Composing grant applications	505	51.11
AUSE8	Composing recommendation letters	519	52.53
AUSE9	Proofreading and editing documents	832	84.21
AUSE10	Summarizing long documents	900	91.09
AUSE11	Composing teaching materials	657	66.50
AUSE12	Composing website posts	703	71.15
AUSE13	Composing blog posts	682	69.03
AUSE14	Composing social media posts	723	73.18
AUSE15	Composing assignments and essays	526	53.24
AUSE16	Composing examination questions	486	49.19
AUSE17	Composing student feedback	408	41.30
AUSE18	Reviewing academic articles	583	59.01
AUSE19	Assessing the peer review process	478	48.38
AUSE20	Generating a research hypothesis	693	70.14
AUSE21	Creating survey questions or interview protocols	766	77.53
AUSE22	Data analysis and interpretation	769	77.83

Note. *N* = 988 for all items. All respondents provided a response to every item; Opposing (No) values may be derived by subtracting Yes *n* from 988.

The highest levels of acceptability were concentrated around instrumental support functions. Tasks such as summarizing long documents, proofreading and editing written work, and data analysis and interpretation received the strongest endorsement. These applications involve AI augmenting rather than replacing human cognitive effort, preserving student agency over core intellectual work.

Communication tasks demonstrated moderate to high acceptability, with variability based on stakeholder relationships. “Emails to external parties” (AUSE4) showed 73.4% acceptance ($n = 725$), while “Emails to students” (AUSE1) garnered 71.9% acceptance ($n = 710$). Interestingly, “Composing social media posts” (AUSE14) achieved 73.2% acceptance ($n = 723$), suggesting students view public-facing communication as appropriate for AI assistance.

In contrast, the lowest acceptability scores clustered around evaluative and authoritative functions. “Composing student feedback” (AUSE17) received the lowest acceptance at 41.3% ($n = 408$), followed by “Assessing the peer review process” (AUSE19) at 48.4% ($n = 478$), and “Composing examination questions” (AUSE16) at 49.2% ($n = 486$). These


activities are closely associated with academic judgement and institutional authority, indicating persistent normative boundaries around assessment integrity that remain largely resistant to AI delegation.

Notably, core academic outputs such as assignments, essays, and research reports received only moderate acceptance, despite high overall levels of GenAI use. This finding suggests that students distinguish between permissible support and impermissible substitution in formally assessed work, reinforcing the view that acceptability judgements reflect ethical reasoning rather than indiscriminate adoption.

A secondary post hoc analysis compared acceptability perceptions between respondents residing in Canada ($n = 58$, 5.9% of the sample) and the remainder of the sample (non-Canadians, $n = 930$). Table 2 displays acceptability rates by category and compares Canadian respondents with non-Canadian respondents across all 22 categories.

Similar to the international cohort, Canadian respondents shared the pattern of endorsing functions relevant to work and research tasks while maintaining normative boundaries for evaluative writing tasks. Canadians perceived four specific tasks as less acceptable than did

Figure 1
Overall Acceptance Rates Across 22 Categories



Note. $N = 988$ for all items.

Table 2
Canadian Subsample Comparison: Acceptability Rates Across 22 GenAI Use Categories

Variable	Meaning	Non-Canadian %	Canadian %	RD (pp)
AUSE1	Email to students	72.04	68.97	-3.07
AUSE2	Email to teachers	69.57	62.07	-7.50
AUSE3	Emails to administrators	67.31	63.79	-3.52
AUSE4	Emails to external parties	73.66	68.97	-4.69
AUSE5	Composing reference letters	60.00	51.72	-8.28
AUSE6*	Composing research reports	54.73	29.31	-25.42
AUSE7	Composing grant applications	51.61	43.10	-8.51
AUSE8	Composing recommendation letters	53.01	44.83	-8.18
AUSE9	Proofreading and editing documents	83.87	89.66	5.79
AUSE10**	Summarizing long documents	91.83	79.31	-12.52
AUSE11***	Composing teaching materials	67.31	53.45	-13.86
AUSE12	Composing website posts	71.61	63.79	-7.82
AUSE13	Composing blog posts	69.25	65.52	-3.73
AUSE14	Composing social media posts	73.44	68.97	-4.47
AUSE15*	Composing assignments and essays	54.62	31.03	-23.59
AUSE16	Composing examination questions	49.35	46.55	-2.80
AUSE17	Composing student feedback	41.72	34.48	-7.24
AUSE18	Reviewing academic articles	58.92	60.34	1.42
AUSE19	Assessing the peer review process	48.82	41.38	-7.44
AUSE20	Generating a research hypothesis	70.22	68.97	-1.25
AUSE21	Creating survey questions or interview protocols	77.53	77.59	0.06
AUSE22	Data analysis and interpretation	77.85	77.59	-0.26

Note. Non-Canadian $n = 930$; Canadian $n = 58$. Non-Canadian percentages are calculated based on the 930 non-Canadian respondents only and therefore differ slightly from the global sample percentages reported in Table 1 ($N = 988$). RD = risk difference, reported in percentage points (pp). Asterisks denote statistically significant differences from chi-square tests of independence: * $p < .001$; ** $p = .001$; *** $p = .030$.

the rest of the sample. Non-Canadians were almost three times more likely to believe an AI is acceptable for composing research reports (AUSE6; 54.73% vs. 29.31%; OR = 2.91, $\chi^2[1] = 14.17$, $p < .001$; RD = -25.42 pp) and for composing assignments and essays (AUSE15; 54.62% vs. 31.03%; OR = 2.68, $\chi^2[1] = 12.20$, $p < .001$; RD = -23.59 pp). Significant differences were also observed for summarizing long

documents (AUSE10; 91.83% vs. 79.31%; OR = 2.93, $\chi^2[1] = 10.54$, $p = .001$; RD = -12.52 pp). Notably, AUSE10 was the category receiving the highest overall endorsement in the international sample. AUSE11 for composing teaching materials also showed divergence from the non-Canadian trends (67.31% vs. 53.45%; OR = 1.79, $\chi^2[1] = 4.71$, $p = .030$; RD = -13.86 pp). In contrast, the two groups

showed near-identical rates across purely instrumental categories, including data analysis and interpretation (AUSE22: 77.85% vs. 77.59%) and creating survey questions or interview protocols (AUSE21: 77.53% vs. 77.59%), confirming that the observed divergences are task-specific rather than a reflection of generalized conservatism among Canadian respondents. Given the small subsample size of the Canadian cohort ($n = 58$), these results should be interpreted as preliminary and exploratory.

Gender-Based Differences in Acceptability (With Effect Sizes)

Analysis of gender differences across the 22 AI use categories revealed limited and narrowly concentrated variation. Only three categories demonstrated statistically significant differences between male and female students, and in each case the magnitude of the difference was modest (see Figure 2).

The largest gender difference was observed for “Composing website posts” (AUSE12), for which male students reported higher acceptability than female students (75.3% vs. 66.5%), corresponding to a risk difference of 8.8 percentage points ($OR = 1.54$, $\chi^2 = 9.38$, $p = .002$). Similar patterns were found for “Composing blog posts” (AUSE13; 72.4% vs. 65.2%), yielding a 7.2 percentage point difference ($OR = 1.41$, $\chi^2 = 6.09$, $p = .014$), and for “Data analysis and interpretation” (AUSE22; 81.4% vs. 73.8%), with a 7.6 percentage point difference ($OR = 1.55$, $\chi^2 = 8.23$, $p = .004$).

Across the remaining categories, no statistically significant gender differences were observed. Female students reported marginally higher acceptability than male students in three categories (AUSE4, AUSE19, AUSE21); however, none reached statistical significance.

Taken together, these findings indicate that gender plays a limited role in shaping students’ normative judgements regarding acceptable AI use. The small number of significant differences and their modest effect sizes suggest that acceptability boundaries are shared mainly across genders.

Given the exploratory and descriptive purpose of the analysis, no adjustment for multiple comparisons was applied; instead, effect sizes were emphasized to support substantive interpretation.

Level of Study Differences in Acceptability (With Effect Sizes)

Comparisons between undergraduate and graduate students revealed systematic differences in acceptability judgements across a range of AI-supported academic practices (see Figure 3). Graduate students consistently reported higher acceptability than undergraduates, particularly for research-oriented, professional communication, and content-creation tasks. In total, 12 of the 22 categories showed statistically significant differences ($p < .05$), with effect sizes indicating moderate but consistent gaps.


The most pronounced differences emerged in research-related activities. “Composing research reports” (AUSE6) exhibited the largest

Table 3
Significant Difference Categories Across Gender

Variable	Meaning	Yes %		RD (pp)	OR	p-value
		Male	Female			
AUSE12	Composing website posts	75.3	66.5	8.8	1.538	0.002
AUSE13	Composing blog posts	72.4	65.2	7.2	1.406	0.014
AUSE22	Data analysis and interpretation	81.4	73.8	7.6	1.550	0.004

Note. OR = odds ratio. RD = risk difference, reported in percentage points (pp).

Figure 2
Gender-Based Differences in Acceptability of AI Use Across Academic Tasks



Note. $N = 988$ (male $n = 526$; female $n = 462$). Grey shading marks categories with statistically significant gender differences based on chi-square tests of independence (AUSE12, AUSE13, and AUSE22; $p < .05$; see Table 3). Overall patterns indicate substantial similarity in acceptability judgements between male and female students, with modest differences confined to these three categories.

level-of-study gap of any category, with 59.4% of graduate students endorsing this use compared to 47.7% of undergraduates, an 11.7-percentage-point difference ($OR = 1.57, p = .001$). Substantial gaps were also observed for “Reviewing academic articles” (AUSE18; 63.7% vs. 54.5%), a 9.2-percentage-point difference ($OR = 1.48, p = .003$), and “Data analysis and interpretation” (AUSE22; 82.4% vs. 74.0%), an 8.4-percentage-point difference that nonetheless yielded the largest odds ratio among the significant categories ($OR = 1.65, p = .002$). These findings suggest that graduate students view AI-assisted research tasks as more compatible with legitimate academic practice than do undergraduates.

Professional communication and academic production tasks showed similarly consistent patterns. Graduate students were more likely to endorse AI use for “Emailing administrators” (AUSE3; 71.2% vs. 63.1%; $RD = 8.1$

pp; $OR = 1.45, p = .008$) and “Emailing instructors” (AUSE2; 73.0% vs. 65.0%; $RD = 8.0$ pp; $OR = 1.45, p = .001$). Differences also appeared for “Grant applications” (AUSE7; $RD = 7.4$ pp; $OR = 1.35, p = .022$) and “Composing recommendation letters” (AUSE8; $RD = 8.2$ pp; $OR = 1.39, p = .013$), reinforcing the interpretation that graduate students’ broader academic responsibilities shape more expansive acceptability boundaries.

Content-creation tasks followed the same directional pattern, though with slightly smaller effect sizes. Graduate students reported higher acceptability for “Website posts,” “Blog posts,” and “Social media posts” (AUSE12–AUSE14), with risk differences ranging from 6.0 to 6.3 percentage points and odds ratios between 1.35 and 1.50. While modest, these differences were consistent and statistically significant, indicating that graduate students are more inclined to view AI-assisted content generation as appropriate within academic contexts.

Despite these differences, several important areas of convergence remained. “Proofreading and editing documents” (AUSE9) was the only category in which undergraduate acceptance marginally exceeded graduate acceptance (85.0% vs. 83.3%), though this difference was not statistically significant ($p = .471$). More importantly, both groups expressed comparable skepticism toward AI use in core evaluative and judgement-intensive activities. No significant level-of-study differences were observed for “Composing examination questions” (AUSE16; $p = .243$) or “Assessing the peer-review process” (AUSE19; $p = .184$), with acceptance rates remaining below 50% for both undergraduate and graduate students.

These findings indicate that the level of study represents a meaningful dimension of differentiation in students’ normative judgements of acceptable AI use, particularly for tasks associated with research production and professional academic practice. At the same time, the

persistence of shared boundaries around evaluative functions suggests enduring norms regarding academic judgement and integrity that transcend educational level.


Summary of Key Findings

Three interrelated patterns emerge from the findings. First, the near-universal adoption of GenAI for academic purposes (93.7%) and high internal consistency across acceptability items ($\alpha = .897$) indicate that GenAI is a normalized feature of post-secondary academic practice.

Second, acceptability judgements vary systematically by task type. High levels of acceptance cluster around instrumental and supportive activities, such as summarization, proofreading, and data analysis, whereas substantially lower acceptance is observed for evaluative and judgement-intensive functions, including assessment design, feedback provision, and examination-related tasks. These bound-

Figure 3

Percentage of Students Endorsing Acceptable AI Uses, by Level of Study



Note. $N = 988$ (undergraduate 52.0%; graduate 48.0%). Grey shading marks the 12 categories with statistically significant differences by level of study based on chi-square tests of independence ($p < .05$; see Table 4). Graduate students reported higher acceptability in every significantly different category.

Table 4
Significant Difference Categories Across Level of Study

Variable	Meaning	UG Yes %	Grad Yes %	RD (pp)	OR	p-value
AUSE2	Email teachers	65.0	73.0	8.0	1.453	0.001
AUSE3	Email administrators	63.1	71.2	8.1	1.447	0.008
AUSE6	Composing research reports	47.7	59.4	11.7	1.566	0.001
AUSE7	Composing grant applications	47.3	54.7	7.4	1.347	0.022
AUSE8	Composing recommendation letters	48.5	56.7	8.2	1.385	0.013
AUSE11	Composing teaching materials	62.6	70.6	8.0	1.441	0.008
AUSE12	Composing website posts	67.5	73.8	6.3	1.504	0.033
AUSE13	Composing blog posts	65.8	72.1	6.3	1.358	0.037
AUSE14	Composing social media posts	70.0	76.0	6.0	1.355	0.039
AUSE17	Composing student feedback	37.5	45.5	8.0	1.391	0.013
AUSE18	Reviewing academic articles	54.5	63.7	9.2	1.475	0.003
AUSE22	Data analysis and interpretation	74.0	82.4	8.4	1.649	0.002

Note. OR = odds ratio. RD = risk difference, reported in percentage points (pp).

aries remain stable across gender and level of study, indicating widely shared norms regarding academic authority and responsibility.

Third, the level of study constitutes the primary source of variation in acceptability judgements. Graduate students consistently report higher acceptance of AI use for research-oriented, professional communication, and content-creation tasks. These differences align with processes of academic socialization, as students' roles shift from structured learning toward independent scholarly and research practice.

In addition, the secondary Canadian comparison indicates broad convergence with the international pattern, alongside significantly lower acceptance among respondents in Canada for four tasks: composing research reports (AUSE6), summarizing long documents (AUSE10), composing teaching materials (AUSE11), and composing assignments and essays (AUSE15). These near-identical rates in purely instrumental categories suggest a task-specific pattern rather than a generalized reluctance to use AI.

Taken together, the persistence of low acceptability for assessment-related uses across all groups points to enduring normative commitments associated with academic integrity. Rather than reflecting uncertainty or permissiveness, students' judgements indicate a nuanced differentiation between AI uses that support academic work and those that risk substituting for human judgement in ways that remain normatively constrained.

DISCUSSION AND IMPLICATIONS

Theoretical Contributions

This study provides empirical evidence for reconceptualizing how normative boundaries around technology use emerge within academic communities. The finding that 93.7% of students actively use GenAI while simultaneously maintaining clear boundaries around specific applications challenges traditional technology

adoption models (Davis, 1989; Venkatesh et al., 2003) that focus on acceptance vs. rejection binaries. Instead, students engage in sophisticated boundary-setting processes, actively negotiating which uses align with academic values, consistent with Gieryn's (1983) boundary-work framework.

The theoretical framework of normative boundary-setting proves particularly robust in explaining the observed patterns. Students consistently distinguish between AI as an instrument for efficiency vs. AI as a substitute for human judgement. This distinction transcends demographic variables, suggesting that boundary-setting operates at the level of shared academic values rather than individual characteristics, aligning with Lamont and Molnár's (2002) analysis of boundary processes. The persistence of these boundaries despite near-universal adoption indicates that normative frameworks can remain stable even as practices evolve, supporting Bearman et al.'s (2022) observation that AI discourse in higher education involves complex negotiations of epistemic authority.

Academic socialization primarily enables boundaries to become permeable. Graduate students' expanded acceptance of AI for research-related tasks does not represent abandonment of academic integrity but rather reflects their evolving position within the academic hierarchy, consistent with established socialization theory (Austin, 2002; Gardner, 2008; Weidman et al., 2001). As students transition from knowledge consumers to knowledge producers, their normative boundaries shift to accommodate the demands of scholarly productivity while maintaining skepticism toward AI in evaluative functions.

Future research might examine whether the pattern of increasing boundary permeability observed across levels of study extends to faculty, or whether the dual pressures of productivity and academic gatekeeping produce a more complex or even reversed dynamic.

The data also reveal an implicit values hierarchy within academic integrity frameworks (ICAI, 2021). While efficiency and productivity gain acceptance as legitimate reasons for AI

use, authenticity and human judgement remain non-negotiable in assessment contexts, reflecting McCabe et al.'s (2012) findings about core academic values. This hierarchy becomes evident through the stark contrast between 91.1% acceptance of summarizing documents and 41.3% acceptance of composing student feedback.

The Paradoxes of Acceptable Use

The findings reveal several paradoxes that complicate straightforward policy responses to GenAI in higher education. The most striking feature is the *producer-recipient asymmetry*: students broadly accept using AI to generate content for others, while simultaneously rejecting AI-generated content directed at themselves. This paradox reveals tensions similar to those identified by Sullivan et al. (2023) regarding academic integrity in the AI era. Students appear to operate under an implicit double standard—expecting human engagement with their potentially AI-assisted work while accepting AI assistance in their own production.

The substitution vs. support dilemma presents another critical paradox. The highest acceptance rate (91.1%) for summarizing long documents raises concerns about cognitive outsourcing, echoing McGuire's (2023) warnings about erosion in digital literacy. If students consistently use AI to summarize academic texts rather than engaging directly with primary sources, they may never develop critical reading skills essential for scholarly work. Similarly, the high acceptance of AI for proofreading (84.2%) and data analysis (77.8%) creates uncertainty about skill development, a concern raised by Humble (2025) in analyzing university AI guidelines.

The transparency challenge represents a third paradox. With 93.7% of students using AI for academic work, traditional detection-based approaches to academic integrity become increasingly futile, as Duah and McGivern (2024) argue. Yet only moderate acceptance of AI for core academic outputs suggests students recognize ethical boundaries even while actively using these tools. This creates what Chan and

Hu (2023) describe as an environment of strategic ambiguity, where actual practice likely exceeds stated acceptance.

Policy Implications

These findings require reconsideration of institutional approaches to GenAI. The data strongly argue against both blanket prohibitions and unrestricted permission, supporting Christ-Brendemühl's (2024) call for nuanced institutional guidelines. The significant differences between undergraduate and graduate students suggest that one-size-fits-all policies will inevitably misalign with student realities, consistent with findings about differentiated technology adoption patterns in higher education (Caner-Yildirim, 2025; Torbaghan et al., 2023).

The concentration of low acceptance around assessment-related tasks validates concerns raised by Spivakovsky et al. (2023) and supports maintaining strict institutional boundaries around examination questions, student feedback, and peer review, even where flexible AI use in drafting and communication is permitted. These findings align with Halaweh's (2023) framework for responsible AI implementation in education.

The emergence of transparency and disclosure as policy imperatives reflects broader discussions about AI governance in higher education (Gellai, 2022). If students expect human evaluation of their potentially AI-assisted work, reciprocal transparency requires instructors to disclose any AI use in course design, content creation, or feedback provision. This bidirectional transparency could foster more honest dialogue about appropriate use while acknowledging the technology's pervasive presence, as Valerio (2024) suggests regarding ethical AI integration.

Canadian Context

The secondary comparison of Canadian respondents against the non-Canadian cohort reinforces and extends the broader policy argument. That Canadian respondents diverged significantly from their non-Canadian counterparts

specifically on AI-assisted academic composition—while converging on instrumental tasks—suggests that normative boundary-setting around assessed academic writing may be particularly salient within the Canadian post-secondary context. This task-specific conservatism, if confirmed in future research with nationally representative Canadian samples, would argue for differentiated institutional policy guidance that explicitly addresses AI use in formally assessed writing rather than relying on generalized permission or prohibition frameworks.

Future Pedagogical Considerations

The widespread adoption documented in this study demands immediate curricular adaptation. Thus, AI literacy must become a core competency alongside traditional information literacy and academic writing skills, supporting arguments by George and Reddy (2024) about digital empowerment in education. This includes not merely technical training in AI tools but also critical engagement with questions of cognitive development, skill acquisition, and intellectual authenticity.

Assessment design faces particular pressure for innovation. If students routinely have access to AI to complete assignments, traditional take-home essays and problem sets may no longer effectively measure individual learning. Alternative assessment methods may become necessary to maintain evaluation integrity while acknowledging AI's role in professional practice, as suggested by recent pedagogical research (Estaiteyeh & Kumar, 2025; Mittal et al., 2022; Sidi et al., 2022).

The preparation gap between undergraduate restrictions and graduate expectations requires intentional bridging. Institutions might develop transitional courses or workshops that explicitly teach responsible AI use, moving students from prohibited use to guided practice to autonomous integration. The higher acceptance among graduate students across multiple categories suggests they already navigate this transition independently; formal support could make this progression more equitable and effective.

Finally, the international composition of our sample, spanning 39 countries, raises the possibility of cultural variation in AI acceptability judgements that the present study was not designed to examine systematically. The secondary analysis of Canadian respondents reported in the Findings section offers a preliminary indication that such variation may exist for specific task categories; however, cross-cultural differences in acceptability norms warrant dedicated investigation in future research, including studies with larger nationally representative samples and culturally validated instruments, as suggested in cross-cultural studies of academic technology adoption (Al-Rahmi et al., 2022; Alrayes et al., 2024).

LIMITATIONS

This study has several limitations that should be considered when interpreting the findings. First, the data represent a snapshot in time, collected in February 2025, during a period of rapid technological and institutional change. Second, the acceptability items reflect AI-supported practices available at the time of data collection; emerging developments, including agentic AI systems, may introduce new forms of use that alter existing normative boundaries. Third, the study relies on self-reported judgements of acceptability, which reflect normative evaluations rather than observed behaviour. Fourth, the binary Yes/No response format, while suited to the large-scale cross-national scope of the study, does not capture gradations of acceptability; respondents with qualified or context-dependent endorsements are indistinguishable from those holding unambiguous positions, which may obscure the extent of normative ambivalence in the data. Future research employing Likert-type scales would allow for more nuanced measurement of acceptability judgements. Relatedly, AUSE22 did not distinguish quantitative from qualitative analysis; outsourcing statistical computation may be judged differently than qualitative interpretation, and future research should disaggregate these analytic modes. Finally, although the sample was large and internationally diverse, participants were

recruited through the Prolific online research platform, whose users tend to be self-selected and may overrepresent technology-engaged, English-proficient students with prior familiarity with online survey participation. Moreover, because the survey was administered exclusively in English, the findings cannot fully capture the Canadian post-secondary context, in which francophone institutions and French-speaking students constitute a substantial component of the sector. This may inflate acceptance rates relative to the broader post-secondary population and should be considered when generalizing the findings across linguistic, cultural, or institutional contexts.

CONCLUSION

This study provides empirical evidence that post-secondary students hold coherent and differentiated judgements regarding acceptable uses of GenAI in academic contexts. Rather than treating AI use as uniformly permissible or problematic, students draw clear normative boundaries that vary by task type and academic role. Evaluative and judgement-intensive functions remain widely constrained, while instrumental and research-oriented uses are more readily accepted.

Importantly, the level of study emerges as the primary axis of differentiation, reflecting processes of academic socialization rather than demographic variation. These findings challenge approaches to AI governance that assume either student permissiveness or uniform norms across educational stages. Instead, they point to the need for guidance and policy frameworks that acknowledge students' existing normative reasoning and the differentiated nature of academic work.

By foregrounding student perspectives on acceptable AI use, this study contributes to ongoing debates about academic integrity, pedagogical design, and institutional responses to GenAI. It offers an empirical foundation for future research examining how normative boundaries evolve as technologies, roles, and expectations in higher education continue to change.

AI USE STATEMENT

Generative AI tools were used to improve linguistic flow, condense selected passages, and review the manuscript for potential APA style inconsistencies. The authors conducted all analyses and developed all interpretations, arguments, and conclusions. They reviewed critically, verified, and substantially revised all AI-assisted output and retain full responsibility for the accuracy, integrity, and final content of the manuscript.

REFERENCES

- Al-Rahmi, A. M., Shamsuddin, A., Wahab, E., Al-Rahmi, W. M., Alismaiel, O. A., & Crawford, J. (2022). Social media usage and acceptance in higher education: A structural equation model. *Frontiers in Education*, 7, 1–14. <https://doi.org/10.3389/educ.2022.964456>
- Alrayes, A., Henari, T., & Ahmed, D. A. (2024). ChatGPT in education—Understanding the Bahraini academics perspective. *Electronic Journal of E-Learning*, 22(2), 112–134. <https://doi.org/10.34190/ejel.22.2.3250>
- Atkinson-Toal, A., & Guo, C. (2024). Generative artificial intelligence (AI) education policies of UK universities. *Enhanced Teaching and Learning in Higher Education*, 2, 70–94. <https://doi.org/10.62512/etlhc.20>
- Austin, A. E. (2002). Preparing the next generation of faculty: Graduate school as socialization to the academic career. *The Journal of Higher Education*, 73(1), 94–122. <https://doi.org/10.1080/00221546.2002.11777132>
- Bai, J. Y. H., Marín, V. I., Jung, I., & Zawacki-Richter, O. (2023). Future prospects of artificial intelligence in education: A preliminary analysis of educator perspectives from focus groups in Japan, Spain, and Germany. *Ubiquity Proceedings*, 3(1), 96–102. <https://doi.org/10.5334/uproc.74>
- Bearman, M., Ryan, J., & Ajjawi, R. (2022). Discourses of artificial intelligence in higher education: A critical literature review. *Higher Education*, 86(2), 369–385. <https://doi.org/10.1007/s10734-022-00937-2>
- Caner-Yildirim, S. (2025). Modeling ChatGPT adoption among undergraduates: An integrated UTAUT2 and digital competence framework. *SAGE Open*, 15(2). <https://doi.org/10.1177/21582440251343340>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20, Article 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Christ-Brendemühl, S. (2024). Leveraging generative AI in higher education: An analysis of opportunities and challenges addressed in university guidelines. *European Journal of Education*, 60(1), Article e12891. <https://doi.org/10.1111/ejed.12891>
- Da Mota, M. (2024). *Toward an AI policy framework for research institutions* [Working paper]. Centre for International Governance Innovation Digital Policy Hub.
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Deckker, D., & Sumanasekara, S. (2025). Student-centric ethical frameworks for AI-driven education: Participatory consent and data ownership in a global perspective. *International Journal of Science and Research Archive*, 15(3), 1491–1503. <https://doi.org/10.30574/ijsra.2025.15.3.1920>

- Duah, J. E., & McGivern, P. (2024). How generative artificial intelligence has blurred notions of authorial identity and academic norms in higher education, necessitating clear university usage policies. *International Journal of Information and Learning Technology*, 41(2), 180–193. <https://doi.org/10.1108/ijilt-11-2023-0213>
- Estaiteyeh, M., & Kumar, R. (2025, November 4). *How AI is challenging the credibility of some online courses*. The Conversation. <https://doi.org/10.64628/AAM.rv597ntsq>
- Gardner, S. K. (2008). Fitting the mold of graduate school: A qualitative study of socialization in doctoral education. *Innovative Higher Education*, 33(2), 125–138. <https://doi.org/10.1007/s10755-008-9068-x>
- Gellai, D. B. (2022). Enterprising academics: Heterarchical policy networks for artificial intelligence in British higher education. *ECNU Review of Education*, 6(4), 568–596. <https://doi.org/10.1177/20965311221143798>
- George, G., & Reddy, R. (2024). Digital empowerment in education: Exploring the dynamics and significance of ICT—A review. *International Journal of Scientific Research in Engineering and Management*, 8(1), 1–10. <https://doi.org/10.55041/ijssrem28245>
- Gieryn, T. F. (1983). Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *American Sociological Review*, 48(6), 781–795. <https://doi.org/10.2307/2095325>
- Halaweh, M. (2023). ChatGPT in education: Strategies for responsible implementation. *Contemporary Educational Technology*, 15(2), Article ep421. <https://doi.org/10.30935/cedtech/13036>
- Humble, N. (2025). Higher education AI policies—A document analysis of university guidelines. *European Journal of Education*, 60(3), Article e70214. <https://doi.org/10.1111/ejed.70214>
- International Center for Academic Integrity. (2021). *The fundamental values of academic integrity* (3rd ed.). https://academicintegrity.org/aws/ICAI/asset_manager/get_file/911282?ver=1
- Kamalov, F., Calonge, D. S., & Gurrib, I. (2023). New era of artificial intelligence in education: Towards a sustainable multifaceted revolution. *Sustainability*, 15(16), Article 12451. <https://doi.org/10.3390/su151612451>
- Kang, H., Han, J., & Kwon, G. H. (2022). The acceptance behavior of smart home health care services in South Korea: An integrated model of UTAUT and TTF. *International Journal of Environmental Research and Public Health*, 19(20), Article 13279. <https://doi.org/10.3390/ijerph192013279>
- Kumar, R., & McGray, R. (2024, August 15). *Global trends in generative artificial intelligence adoption* [Paper presentation]. Digital Pedagogy Institute 2024, Waterloo, ON (Online). <http://hdl.handle.net/10464/18633>
- Lamont, M., & Molnár, V. (2002). The study of boundaries in the social sciences. *Annual Review of Sociology*, 28, 167–195. <https://doi.org/10.1146/annurev.soc.28.110601.141107>
- McCabe, D. L., Butterfield, K. D., & Treviño, L. K. (2012). *Cheating in college: Why students do it and what educators can do about it*. Johns Hopkins University Press.
- McGuire, A. (2023). Leveraging ChatGPT for rethinking plagiarism, digital literacy, and the ethics of co-authorship in higher education. *Irish Journal of Technology Enhanced Learning*, 7(2), 21–31. <https://doi.org/10.22554/ijtel.v7i2.131>
- Mittal, P., Kaur, A., & Jain, R. (2022). Online learning for enhancing employability skills in higher education students: The mediating role of learning analytics. *TEM Journal*, 11(4), 1469–1476. <https://doi.org/10.18421/tem114-06>

- Or, C. C. P. (2023). Towards an integrated model: Task-technology fit in Unified Theory of Acceptance and Use of Technology 2 in education contexts. *Journal of Applied Learning & Teaching*, 6(1), 151–163. <https://doi.org/10.37074/jalt.2023.6.1.8>
- Sidi, Y., Blau, I., & Shamir-Inbal, T. (2022). Mapping active and collaborative learning in higher education through annotations in hyper-video by learning analytics. *Journal of Computer Assisted Learning*, 38(6), 1752–1764. <https://doi.org/10.1111/jcal.12714>
- Spivakovsky, O., Omelchuk, S., Kobets, V., Valko, N., & Malchykova, D. (2023). Institutional policies on artificial intelligence in university learning, teaching and research. *Information Technologies and Learning Tools*, 97(5), 181–202. <https://doi.org/10.33407/itlt.v97i5.5395>
- Sullivan, M., Kelly, A., & McLaughlan, P. (2023). ChatGPT in higher education: Considerations for academic integrity and student learning. *Journal of Applied Learning & Teaching*, 6(1), 31–40. <https://doi.org/10.37074/jalt.2023.6.1.17>
- Torbaghan, M. E., Sasidharan, M., Jefferson, I., & Watkins, J. S. (2023). Preparing students for a digitized future. *IEEE Transactions on Education*, 66(1), 20–29. <https://doi.org/10.1109/te.2022.3174263>
- Valerio, A. S. (2024). Anticipating the impact of artificial intelligence in higher education: Student awareness and ethical concerns in Zamboanga City, Philippines. *Cognizance Journal of Multidisciplinary Studies*, 4(6), 408–418. <https://doi.org/10.47760/cognizance.2024.v04i06.024>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Weidman, J. C., Twale, D. J., & Stein, E. L. (2001). Socialization of graduate and professional students in higher education: A perilous passage? *ASHE-ERIC Higher Education Report*, 28(3). <https://files.eric.ed.gov/fulltext/ED457710.pdf>