Journal of Learning and Development Studies
ISSN: 2752-9541
DOI: 10.32996/jlds
Journal Homepage: www.al-kindipublisher.com/index.php/jlds

Institutional Silence and Informal Governance of Generative AI in Moroccan Higher Education

Safae Kharmach ✉[1] **and Abderrahim El Karfa**[2]

[12] *Faculty of Letters and Human Sciences Dhar El Mahraz, Sidi Mohamed Ben Abdellah University, Fez, Morocco*

Corresponding Author: Safae Kharmach**, E-mail:** safaekharmach00@gmail.com

ABSTRACT

Moroccan higher education has made the teaching of artificial intelligence compulsory in parts of the system while leaving its use by students in assessed work unaddressed. Where no rule is written, the working judgement falls to individual professors, and students assemble their understanding from whatever reaches them. This study establishes what undergraduates at one Moroccan university believe is permitted and where that belief originates. A trilingual questionnaire completed by 173 students across three faculties covered institutional awareness, sources of understanding, and what professors had said. Nine specified uses of generative AI were judged twice, once for permissibility and once for personal acceptability, and one open question invited an account of a moment of uncertainty. More than half did not know whether their university has any written rule, 9.2% had read an institutional document mentioning the technology, and institutional text was the main source for at most 5.2%. Own judgement and the internet accounted for 52.6%, and in every faculty the most commonly reported message from professors was held by fewer than two students in five. Acceptability ratings nonetheless separated supportive from substitutive uses, with no variation by faculty, level or gender. Confidence was unrelated to how strictly students judged substitution, and lower only among those whose professors had contradicted one another, while silence showed no such association. The absence of a rule has produced uncertainty and not permissiveness, so a statement at the level of the act, delivered through professors, would codify a distinction students already hold.

KEYWORDS

Generative artificial intelligence, institutional policy, student perceptions, academic regulation, Moroccan higher education

ARTICLE INFORMATION

ACCEPTED: 09 October 2026 **PUBLISHED:** 09 November 2026 **DOI:** 10.32996/jlds.2026.6.10.4

1. Introduction

Student use of generative AI has become routine across university systems, and the rules meant to govern it remain uneven where they exist at all. A survey of 1,041 undergraduates in the United Kingdom found that 92% had used such tools in some capacity and 88% for assessed work, against 53% a year earlier, with 18% placing AI-generated text directly into submitted work (Freeman, 2025). Figures of this order now appear across a range of national systems, and the practice they describe has outpaced whatever guidance reaches the students producing it. Students have nonetheless arrived at a working sense of what they are permitted to do. Where that sense came from, and whether anyone in authority put it there, are the questions this study takes up.

Where institutions have acted, they have often acted late and passed the decision downward. Analyses of the world's highest-ranked universities find published guidance at fewer than half of them, and most of those documents ask teaching staff to set expectations in their own syllabi (Moorhouse et al., 2023). What counts as permissible therefore varies from classroom to classroom, and the student has no reliable way of knowing that it does.

In Moroccan higher education, the rules have not been written at all. No national framework governs students' use of generative AI in assessed work, and nothing at the institutional level reaches students in the form of a published rule. Radouane Mrabet, former president of Sidi Mohamed Ben Abdellah University, told Médias24 in February 2026 that what exists is scattered and uncoordinated for want of a legal framework, and called for a national strategy and a framework law. He had also seen no comprehensive study of how these tools are used in Moroccan universities (Ibriz, 2026).

What the system has regulated is instruction. A new cahier des normes pédagogiques, published in the Bulletin officiel, requires every Master des sciences et techniques programme to include at least one module in artificial intelligence from 2026–2027,

Copyright: © 2026 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (https://creativecommons.org/licenses/by/4.0/). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

whatever the student's specialism (Hespress, 2026). Teaching about these tools has therefore become compulsory in one part of the system, while their use by students remains unaddressed in all of it.

Students form their own understanding of what is allowed, and where that understanding comes from is an empirical question. If no institutional rule reaches them, whatever they believe about permissible use has come from somewhere else: from a professor who mentioned it in class, from a friend, from the internet, or from their own sense of what is fair. Each of those sources gives a different answer and none carries institutional authority. Students in the same programme can therefore hold incompatible beliefs about the same act, and that inconsistency belongs to the institution that left the question open.

This study examines that situation at a single Moroccan university. It establishes what students believe is permitted across nine specified uses of generative AI, where that belief came from, what their professors have told them, and how their reading of the rules relates to their own sense of what is acceptable. Each use is judged twice, once for whether it is allowed and once for whether it is right, which separates the official rule, if any, from the student's own judgement and makes the distance between them measurable. Disagreement is treated as a measure in its own right, since scatter within a single department indicates where the rule is actually being set.

Three contributions follow. The study supplies the student-level evidence that Moroccan work on artificial intelligence in higher education has not yet produced. It offers a way of measuring regulatory incoherence from survey data, through the dispersion of belief within an institutional unit. And it identifies where students locate the boundary of acceptable practice when no policy defines it, which is what any institution would need before writing one.

2. Literature Review

2.1 From detection to permission

The early literature on generative AI in higher education was preoccupied with detection and prohibition. Cotton et al. (2024) set out the terms, describing the opportunities the technology offered and the threat it posed to conventional assessment, and concluding that institutions needed to revise their academic integrity provisions before the technology settled into student practice. Their argument has aged in an instructive way, since the revision they called for has been slower than the settling.

Subsequent work has moved the question from whether students use these tools to which uses are legitimate, and the acts in question are not equivalent. Asking a model to explain a concept a lecture left obscure, asking it to summarise an assigned reading instead of reading it, and submitting its output as original work carry different consequences for learning. A rule that treats them as one act will not be followed. Freeman (2025) finds a gradient of this kind in reported practice, with explanation, summary and the generation of research ideas the commonest uses. Regulation that does not distinguish among them is therefore aimed at a uniform practice that students do not have.

2.2 How far institutional policy reaches

Where institutional guidance exists, its content determines how far it reaches. In the documents Moorhouse et al. (2023) examined, the integrity provisions concentrated on acknowledgement and conceded that detection software could not be relied on, which left what counted as a declarable use to be settled case by case. The assessment provisions ran in the same direction, recommending that staff test their own tasks against the tools and revise whatever the tools could complete. Both place the working judgement with the individual teacher, and neither gives the student a rule that holds from one course to the next.

That delegation is the mechanism this study examines. Chan (2023) surveyed students and teaching staff at a Hong Kong university and found their views closer together than the debate would suggest, with both groups optimistic about the technology and both concerned that misuse would confer an unfair advantage. Neither group was confident that AI use could be detected, and teaching staff were slightly less confident than students. The staff expected to enforce the rule doubted they could detect a breach of it, which leaves it standing as an expectation a teacher states rather than a provision an institution applies. Governance of permissible use then falls to whoever is nearest to the student.

Freeman's (2025) figures supply a useful contrast with the Moroccan case. Around 80% of the British undergraduates surveyed agreed that their institution's policy on AI was clear, and 76% expected their institution would spot the use of AI in assessed work. Those beliefs need not be accurate to matter, since four students in five could name the existence of a rule whether or not detection works as they expect. Only 36% reported receiving institutional support in developing AI competence, which suggests those institutions have been quicker to regulate than to teach. The Moroccan situation described below inverts the order.

2.3 Where students learn what is allowed

The guidance examined so far exists as documents, and a document governs only the student who has encountered it. A rule therefore reaches a student through a channel, and the channel shapes both its content and the authority it carries. A statement in a syllabus is durable and can be checked. A remark made once in a lecture is neither, and a classmate's account of what another professor said carries no authority at all while still being the version a student acts on. Selwyn (2016) argues that the distance between what a technology is expected to do and what it does once deployed is the normal condition, and the same distance opens between a rule as written and a rule as received.

A rule learned through any of these channels cannot be verified by the student who holds it. The remark in the lecture cannot be replayed, and nothing in it marks whether it stated an institutional requirement or one teacher's preference. The student acts on it regardless, and carries the risk of having read it wrongly. Caution and licence are equally defensible responses to the same instruction, which is why students in one programme arrive at different positions on a single act. The degree of that divergence is measurable, and Section 3 sets out how this study measures it.

2.4 Generative AI and ICT in Moroccan higher education

Moroccan research on information and communication technologies in education has consistently identified a distance between what institutions procure and what practitioners use. Messaoudi and Talbi (2012) assessed GENIE, the national programme that equipped public schools with multimedia rooms and trained their teachers to use them, and found the provision in place well ahead of its classroom use. Riyami (2018) found the same in higher education, in a Moroccan institution running a programme jointly with a French university. There a model developed elsewhere produced friction that had little to do with the technology itself. El Ferouali and Ouhadi (2023) identify the absence of locally grounded provision as a recurring limitation, and Ben Attou (2021) sets out the structural constraints under which open-access faculties have attempted digitalisation. E-learning adoption shows the same (Nouib et al., 2022), which suggests the gap is a feature of how provision reaches practice rather than of any one technology.

Work specifically on artificial intelligence is more recent and has approached the subject from the institutional and technical sides. Ihya et al. (2020) used machine learning to predict whether Moroccan students would accept an e-orientation system, meaning a tool that guides the choice of field of study. El Mrabet and Ait Moussa (2023) developed a supervised machine learning framework for predicting academic orientation. Both treat artificial intelligence as something the institution deploys and ask whether a system of that kind would be accepted. The tool arrives from the institution and the student receives it, which reverses the direction generative AI travels. Didi (2025) examined what institutions have written rather than what they have built. A review of the conduct codes at Mohammed V University, ENSA Kenitra and UM6P found none of the three carrying rules on AI-authored work, attribution or disclosure, and the administrators interviewed attributed the inaction to an absence of national guidance.

The study closest to the present one is Mjahad et al. (2025), who surveyed 103 academics at the College of Arts and Humanities of Sidi Mohamed Ben Abdellah University in Fez. They found a predominantly positive outlook, with 69.9% agreeing that AI integration can enhance educational practice, alongside concern about academic integrity, reduced face-to-face interaction and unequal access. Age and academic speciality significantly shaped those perceptions. That survey covers the staff side of the institution examined here. Academics at the faculty report concern about students' academic integrity, and the present study asks whether any rule governing that conduct has reached the students themselves.

2.5 The gap this study addresses

Institutional guidance on generative AI is incomplete even in well-resourced systems, and where it exists, it is frequently delegated to the individual teacher (Chan, 2023; Moorhouse et al., 2023). Moroccan higher education has produced none at the national or institutional level. Didi (2025) found no rules on AI-authored work in the conduct codes of three Moroccan institutions, and the country's own senior figures describe the situation as uncoordinated and without legal footing (Ibriz, 2026). Moroccan research has examined how institutions deploy these systems, how academics regard them, and how students use them (Errifaiy et al., 2026; Ihya et al., 2020; Mamoune et al., 2025; Mjahad et al., 2025). None of it asks what those students believe they are permitted to do.

This study addresses that omission at a single Moroccan university. It establishes what students believe the rules to be in a system that has not written any, and treats the spread of those beliefs as a measure of how far the absence has reached them. It asks:

RQ1. To what extent are students aware of any institutional regulation governing the use of generative AI in academic work?

RQ2. From which sources do students derive their understanding of what is permitted, and how do they resolve uncertainty when it arises?

RQ3. What do students report their professors telling them about generative AI use, and how consistent are those accounts within and between faculties?

RQ4. Where do students locate the boundary of acceptable generative AI use, how far do they diverge from one another, and how far does that boundary sit from what they believe is permitted?

RQ5. Is students' confidence in knowing what is permitted associated with the strictness of their own ethical judgements and with how sharply they distinguish supportive from substitutive uses?

3. Method

3.1 Design

This study examines what undergraduates at a Moroccan university believe is permitted in their use of generative AI, where they get that belief, and how it relates to their sense of what is acceptable. A quantitative approach was therefore adopted, as it allows beliefs to be measured across a large group and compared systematically between faculties (Creswell & Creswell, 2018). Specifically, the study used a cross-sectional survey design, with data collected at a single point in time. One open-ended question was added to capture students' own accounts of how they resolve uncertainty and to help interpret the quantitative results.

3.2 Setting and participants

Data were collected between 30 September and 13 October 2026, and only students enrolled in the Licence cycle were eligible. A non-probability convenience sampling technique was used, as the questionnaire was shared through professors and through faculty and departmental WhatsApp groups. This approach suited a student population spread across several faculties with no central sampling frame. The final sample consisted of 173 students, of whom 115 (66.5%) answered in Arabic, 30 (17.3%) in English and 28 (16.2%) in French. Three faculties were represented in numbers large enough for the comparison the third research question requires, with Letters and Human Sciences at 72 (41.6%), Sciences at 56 (32.4%) and Law, Economics and Social Sciences at 35 (20.2%). Table 1 sets out the sample.

Table 1

Characteristics of the Sample (N = 173)

	Category	n	%
Language of response	Arabic	115	66.5
	French	28	16.2
	English	30	17.3
Faculty	Letters and Human Sciences	72	41.6
	Sciences	56	32.4
	Law, Economics and Social Sciences	35	20.2
	A school (ENS, ENSA, EST, ENCG)	10	5.8
Level	Semesters 1 or 2	81	46.8
	Semesters 3 or 4	28	16.2
	Semesters 5 or 6	64	37.0
Gender	Female	105	60.7

Note. Categories are the response options offered on the questionnaire, harmonised across its English, French and Arabic versions.

3.3 Instrument

The questionnaire ran to nine sections and roughly forty response points and was estimated to take eight to ten minutes, opening with consent and background items and closing with an open-response question. Sections 2 and 3 covered students' use of generative AI and their awareness of any institutional rules, and Sections 4 and 5 covered where their understanding came from and what their professors had said. Section 6 presented the nine acts described below, and an attention check sat on its own page between it and Section 8. The last section asked what students expected to happen if AI use were detected, whether they knew of any case, and how useful clear written rules would be. The full instrument appears in the Appendix.

The questionnaire was designed in Google Forms and developed for this study, with the Arabic and French versions carrying identical items in identical order to the English. It comprised nine sections: consent and background, the frequency and duration of AI use, awareness of institutional rules (three items and a five-point confidence scale), where students had learned what is permitted (seven channels by multiple response and a single-choice main source), what their professors had said (four items), the two rating grids described below, an attention check on its own page, expected consequences and the value of written rules (three items), and a final open-ended question. Completion was estimated at eight to ten minutes (see Appendix).

3.4 Procedure and screening

The questionnaire was open online from 30 September to 13 October 2026. Screening rules were fixed before any result was examined, and a response was excluded if it failed the attention check or if all eighteen Section 6 ratings were identical. Of the 218 students who consented, 27 were not enrolled in the Licence cycle, three abandoned the instrument before completing both grids, and 15 failed the attention check. No respondent gave eighteen identical ratings, so that rule removed nobody. The analyses below rest on the remaining 173.

3.5 Analysis

Analyses were conducted in IBM SPSS Statistics 28, following a plan fixed before collection and reported here in the order of the research questions. Proportions carry 95% confidence intervals throughout. The disagreement index used for RQ3 is one minus the proportion of respondents in a faculty selecting the most common message, so a value near zero describes a faculty whose students report the same thing. Paired comparisons between the two grids use the Wilcoxon signed-rank test, since the ratings are ordinal and their distributions are not symmetric. Group comparisons use independent-samples t-tests with Welch's correction, and associations between the confidence item and the derived Grid B measures are reported as both Pearson and Spearman coefficients. The open responses were coded thematically following Braun and Clarke (2006), for the origin of the uncertainty and for how the respondent resolved it. Harman's single-factor test is also reported, since all measures are self-reported on one

occasion. Disagreement within a faculty is summarised by an index equal to one minus the proportion of its students selecting the most commonly reported message, so a value near zero indicates a faculty whose students report the same thing.

3.6 Ethics

Formal ethical review is not an established practice for anonymous survey research at the institution. Participation was voluntary and anonymous. The information page stated that no name, email address or student number was collected and that responses could not be traced to an individual. It also stated that nothing reported would reach any professor or member of administrative staff and that only overall results would be published. Recruitment ran through shared links rather than in class, so no student was approached in a setting where declining would have been visible.

4. Results

4.1 What students know of the rules

Most students do not know whether a rule exists, and of those with a view, nearly all believe there is none. Asked whether the university has official written rules on AI use in academic work, 96 respondents (55.5%) did not know and 69 (39.9%) said there were none. Only 8 (4.6%) said there were. The same pattern holds for the respondents' own establishment, where 96 were again unable to say, 65 reported none and 12 reported that rules exist. Only 16 students (9.2%) had read a document from the university or their establishment mentioning artificial intelligence, against 110 (63.6%) who had not and 47 (27.2%) who could not remember.

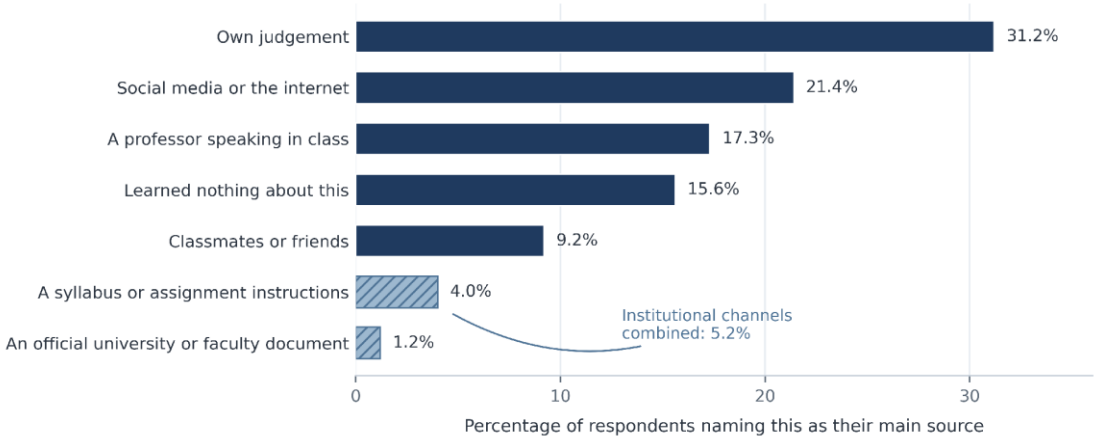
Students are unsure what is permitted, expect to be punished for getting it wrong, and have seen almost no one punished. Mean confidence in knowing what is and is not permitted was 2.68 on a five-point scale (SD = 1.29, n = 154), with 67 respondents (43.5%) in the bottom two points and 34 (22.1%) in the top two. Asked what would happen to a student found to have submitted AI-generated work as their own, 35 respondents (20.2%) named a reduced mark, 76 (43.9%) a zero on the work and 7 (4.0%) formal disciplinary action. Only 8 (4.6%) expected nothing at all. Yet 112 (64.7%) knew of no student who had faced consequences, and 43 (24.9%) had heard of cases they could not verify. Asked how useful clear written rules would be, 121 respondents (74.2%) chose the top two points of the scale (M = 4.18, SD = 1.12, n = 163).

4.2 The sources students rely on

The understanding students hold comes overwhelmingly from themselves and from the internet. Figure 1 gives the main source each respondent named, with own judgement leading at 54 respondents (31.2%) and social media or the internet next at 37 (21.4%). A professor speaking in class accounts for 30 (17.3%) and classmates for 16 (9.2%), while 27 students (15.6%) report having learned nothing about the matter. The same ordering appears in a separate multiple-response item asking which channels students had learned from, where own judgement was selected 80 times and the internet 64, against 5 selections for a syllabus and 2 for an official document. A student's main source should also appear among the channels they selected, yet for 40 respondents it did not, so the two items are reported separately.

Figure 1

Main Source of Students' Understanding of What Is Permitted



Note. N = 173. Single-choice item. Hatched bars mark the two channels the institution controls in writing.

The institution's own writing reaches almost nobody. Grouped by kind, the main sources divide into 91 respondents (52.6%) relying on themselves or the internet, 46 (26.6%) on another person, 27 (15.6%) on nothing at all, and 9 (5.2%) on a document. They are the whole of the institution's written reach into the sample, made up of 7 students (4.0%) naming a syllabus or set of assignment instructions and 2 (1.2%) naming an official university or faculty document. Three times as many students had learned nothing as had learned from anything the university has written. Only one of those nine reported having read an institutional document mentioning artificial intelligence, so even 5.2% is best read as an upper bound.

The source a student relies on makes no detectable difference to what they conclude. Crossing the four source groups against the three answers on whether the university has written rules gives twelve cells, six of which had expected counts below five and one as low as 0.42. The association was therefore tested by Monte Carlo simulation with 20,000 resamples rather than by the chi-square approximation, returning p = .87. That is, students who learned from a professor were no more likely to believe a rule exists than those who learned from social media. With only eight respondents believing a rule exists, the test has little power to detect a small association, so the result is an absence of evidence rather than evidence of equivalence.

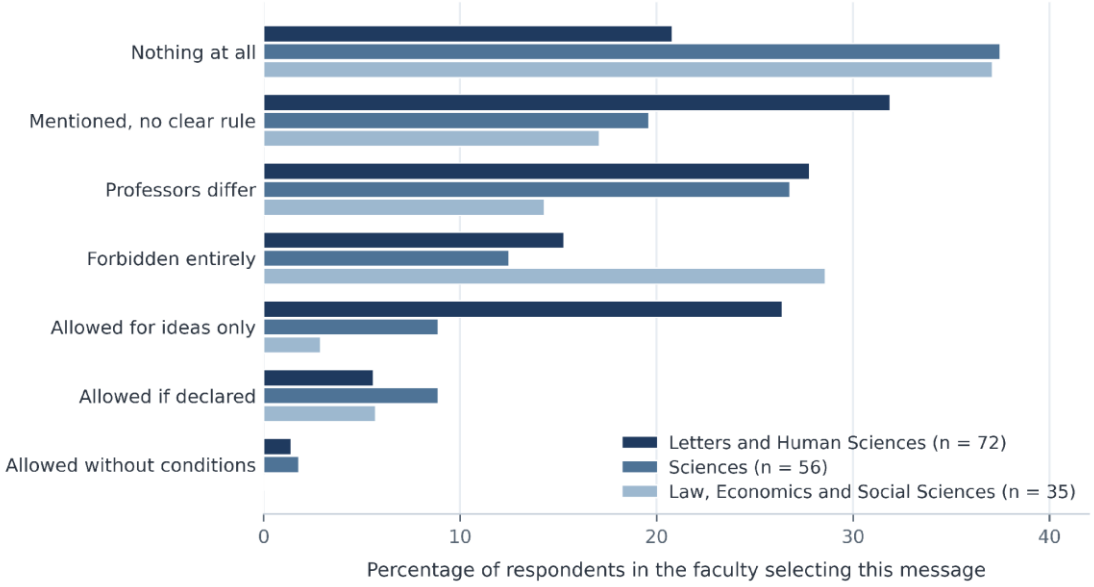
4.3 Reported guidance from Professors

Professors are largely silent on the question, and where they speak, they do not agree. Of the 164 students who answered, 43 (26.2%) reported that no professor had ever addressed AI use in their work, and only 10 (6.1%) reported that most of their professors had. Asked what their professors had said, 51 of the 173 (29.5%) selected nothing at all, 43 (24.9%) that the matter had been mentioned without a clear rule, and 41 (23.7%) that different professors said different things. Definite rules were reported less often and pointed in different directions. Outright prohibition was reported by 33 students (19.1%), permission for ideas but not for writing by 26 (15.0%), permission conditional on declaring the use by 12 (6.9%), and permission without conditions by 2 (1.2%). Asked directly whether different professors had given different rules for the same kind of use, 49 (28.3%) said they had, 74 (42.8%) said they had not, and 50 (28.9%) could not say.

Figure 2 breaks this down by faculty. The index runs from 0.62 in Sciences through 0.63 in Law, Economics and Social Sciences to 0.68 in Letters and Human Sciences, so the most frequently reported answer is selected by between 32% and 38% of a faculty's students. In two of the three, that answer is that nothing was said. The incoherence runs between faculties as well as inside them. Outright prohibition was reported by 28.6% in Law, Economics and Social Sciences against 12.5% in Sciences, and permission for ideas but not for writing by 26.4% in Letters and Human Sciences against 2.9% in Law.

Figure 2

What Students in Each Faculty Report Being Told



Note. Multiple-response item, so percentages within a faculty exceed 100. The fourth group, 10 students not classified by faculty, is omitted.

4.4 Permissibility and acceptability across the nine acts

Students rank the nine acts in the order the design anticipated, with the four supportive acts taking the top four places and the five substitutive ones the bottom five. Table 2 and Figure 3 give the item statistics for both grids. Believed permissibility runs from 4.65 for asking a tool to explain a concept down to 2.28 for submitting generated text with little or no change. Three acts fall below the scale midpoint: summarising an assigned reading instead of reading it, using a tool in a take-home set, and submitting generated text. The six above it are not uniformly endorsed, since fewer than half of students rate asking for a first draft (41.0%) or producing work beyond their own language level (46.2%) as permitted.

Table 2

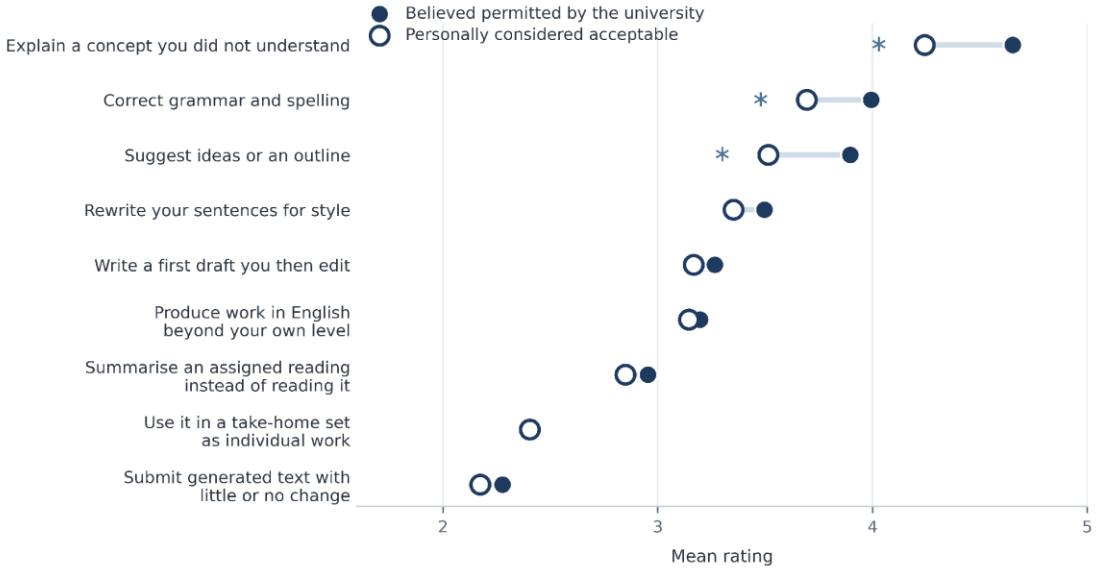
Believed Permissibility and Personal Acceptability of the Nine Acts

Act	A M	A SD	Cannot tell %	B M	B SD	B − A	p
Explain a concept	4.65	0.63	5.2	4.24	1.33	-0.41	< .001
Suggest ideas or an outline	3.90	1.09	16.2	3.51	1.42	-0.38	.005
Correct grammar and spelling	3.99	1.15	12.1	3.69	1.40	-0.30	.032
Rewrite sentences for style	3.50	1.23	22.5	3.35	1.39	-0.14	.388
Summarise a reading instead of reading it	2.95	1.35	23.7	2.85	1.47	-0.10	.507
Write a first draft, then edit it	3.27	1.22	31.8	3.17	1.45	-0.10	.447
Submit generated text unchanged	2.28	1.28	24.9	2.17	1.39	-0.10	.442
Use in an individual take-home task	2.40	1.30	20.2	2.40	1.43	+0.00	.966
Produce English beyond own level	3.20	1.35	20.2	3.14	1.41	-0.05	.829

Note. N = 173. Both grids run 1 to 5. Cannot-tell is the proportion selecting the Grid A midpoint. p values from Wilcoxon signed-rank tests of the paired difference.

Figure 3

Believed Permissibility and Personal Acceptability of Nine Uses of Generative AI



Note. N = 173. Both scales run from 1 to 5, where 1 is clearly not permitted or completely unacceptable and 5 is clearly permitted or completely acceptable. Acts are ordered by believed permissibility. An asterisk marks a paired difference significant at $p < .05$ (Wilcoxon signed-rank).

The cannot-tell midpoint was used at least once by 125 of the 173 students (72.3%). Across the nine acts it accounts for 19.7% of all Grid A responses, running from 5.2% for asking a tool to explain a concept up to 31.8% for asking it to write a first draft the student then edits. Uncertainty is therefore lowest on the act students most often rate as permitted and highest on the one closest to the boundary they draw.

Dispersion runs against the prediction, since ratings of what is permitted were less spread out than ratings of what is acceptable, at a mean item standard deviation of 1.18 on Grid A against 1.41 on Grid B. The pattern holds on every one of the nine acts, so it does not rest on a single item. H4 anticipated the reverse, on the reasoning that an unsupplied rule would produce wider disagreement than a personal judgement. The cannot-tell option is part of the explanation, since it gives uncertain respondents a midpoint to settle on and compresses the Grid A distribution. The rest is that personal ethical judgement varies more between students than their guesses about institutional policy do.

Comparing the grids act by act produces a small asymmetry at the permissive end, with three differences reaching significance, all in the same direction and all among the supportive acts. Students believe the university permits asking a tool to explain a concept ($M = 4.65$) more than they personally find it acceptable ($M = 4.24$, $p < .001$). The same holds for suggesting ideas or an outline (3.90 against 3.51, $p = .005$) and correcting grammar and spelling (3.99 against 3.69, $p = .032$). None of the substitutive acts differs, and across all nine, the two grids do not separate ($p = .12$). Where students draw a line of their own, it sits inside the line they attribute to the institution.

4.5 Reliability, confidence and contradiction

Students' ethical judgements are internally consistent even though their picture of the rules is not. Cronbach's alpha was .87 for the four supportive acts and .77 for the five substitutive acts, both adequate for a scale of this length. The exploratory factor analysis of all nine Grid B items, reported here as a separate check, found two dimensions underlying the ratings rather than one. The four supportive acts load together on the first, while submitting generated text, using a tool in a take-home set and summarising an assigned reading define the second. The distinction fixed before collection is recovered in outline, though asking for a first draft and producing work beyond one's own language level sit between the two clusters rather than with the substitutive acts.

On the scale itself, students are markedly stricter about substitution than about support. Mean acceptability was 3.70 ($SD = 1.18$) for the four supportive acts and 2.75 ($SD = 1.03$) for the five substitutive ones, a gap of 0.95 points on a five-point scale. Confidence explains none of it, since students who felt sure they knew the rules drew the line in the same place as students who did not. Confidence correlates at $r = .03$ ($p = .72$) with acceptability of substitution and $r = .00$ ($p = .98$) with the size of the gap, and Spearman coefficients give the same result, so H5 is not supported. Faculty, level of study and gender make no difference to it either, with one-way comparisons returning $p = .52$, .88 and .54. The distinction holds at the same sharpness across the institution.

Confidence has one correlate, and it lies with the professors rather than with the students. Students who reported that different professors had given different rules for the same kind of use were less confident of knowing what is permitted than those who

reported no such contradiction. Their mean was 2.30 against 2.91, $t(105) = -2.58$, $p = .012$, on the 107 students who answered both items. Harman's single-factor test returned a first unrotated factor accounting for 22.2% of variance across the eighteen rating items, well below the conventional threshold, so a single response tendency is unlikely to account for the pattern. Silence shows no such association, since students whose professors had never addressed AI use were slightly more confident than those whose professors had, at 2.85 against 2.63 ($p = .42$).

4.6 How students describe the uncertainty

Sixty-eight students answered the open question, four of them only to say they had never faced such a situation. The remaining accounts were coded for the origin of the uncertainty and for how the respondent resolved it. The patterns that recur concern where a private rule comes from, how the boundary is named, when the rule becomes visible, and why the obvious question goes unasked. Quotations are translated from the French and Arabic originals.

These accounts most often describe a rule the student has built alone and applies consistently. The line is almost always drawn between content and form. One respondent writing up an organic chemistry practical put it as "the scientific content was 100% mine, the form was smoothed by the tool". Others name the boundary as unavailable rather than contested, reasoning outward from what the absent rule does not mention. A student who fed financial data into a model for a management case study observed that "the rules only ever talk about writing, so I assumed data analysis was fine, but I'm honestly not sure".

Where a rule does become visible, it is often through its enforcement. One student prepared for an examination using a tool, sat it expecting a good result, and "was surprised that the mark had been reduced because I had used AI". The rule stays invisible until that moment partly because the question never gets asked. For instance, a second-year economics student, unsure about using a tool for an inflation report, wrote that the professor "had never spoken about it". The same student added, "I didn't dare ask him, I was afraid he would think I wanted to cheat". Against this, four respondents did ask and received an immediate and workable answer, in one case that a tool could be used to understand a novel while the writing had to be the student's own.

5. Discussion

5.1 A settled ethics and an unsettled map

The clearest result of this study is that the two things usually treated as one come apart. That is, students at this university hold a coherent and fairly strict ethical position on generative AI, and they have almost no information about what their institution expects. The ethical position is measurable and stable, since acceptability ratings separate the acts that support a student's own work from those that replace it, with the substitutive ones rated almost a full point lower. That gap holds across faculties, years and genders. The informational position is the opposite, with more than half unable to say whether a rule exists and fewer than one in ten having read a document mentioning the technology. Nearly three quarters could not say whether a given everyday act is permitted.

That separation changes how the problem should be framed. Debate about students and generative AI has largely been conducted as a question of integrity, with Cotton et al. (2024) calling for revised academic integrity provisions and much of what followed concerned with detection. The students surveyed here are not waiting for a rule in order to decide what is acceptable, since they have already decided. They lack any means of discovering whether the institution agrees with them, which makes this a problem of communication rather than of conduct.

5.2 Contradiction rather than silence

Lower confidence is associated with one variable alone, which is whether a student's professors have contradicted one another, while how often they use these tools, their faculty and their year of study make no difference to it. Silence leaves a student guessing from a blank page, while conflicting instruction leaves them unable to treat any instruction as authoritative. The second is the worse of the two, since it supplies information that cancels itself.

Moorhouse et al. (2023) found that where institutional guidance exists, it hands the working judgement to the individual teacher. Yet this university has not written the guidance, so nothing was handed anywhere. The judgement sits with each professor all the same, arriving there by default rather than by decision, which means no document records it and no one is answerable for the result. Thus the disagreement index measures what that produces, since in every faculty surveyed the most commonly reported message is held by fewer than two students in five. That is what an institution looks like when each professor legislates independently and no one reconciles the outcomes.

5.3 A penalty without a rule

Students expect to be punished for breaking a rule that most of them believe has never been written, with nearly half anticipating a zero on the work and two-thirds knowing of no case in which anyone was penalised. That combination is costly in a particular way, since the cost falls on the cautious. The open responses show students holding back from uses they believed were probably permitted, and above all declining to ask. One did not dare approach a professor for fear the question itself would be read as an admission. Another withheld from a supervisor that a thesis title had come from a model, to avoid a misunderstanding that had not yet arisen.

This also answers a worry that runs through the policy literature, which is that an absent rule will be exploited. The present data point the other way, since on the three acts where believed permissibility and personal acceptability diverge significantly, students

judge the university to be more permissive than they themselves are. No act produces the reverse, though across all nine the two grids do not separate. They are not reading the silence as licence.

5.4 Implications for policy and practice

Written guidance asks less of an institution than the detection and assessment redesign the literature has concentrated on, and the students want it, with 74.2% rating clear written rules at the top two points of the scale. A university could simply announce that its existing academic integrity rules already cover AI use, but that would change nothing here, since fewer than one in ten students have read an institutional document mentioning the technology. A statement at the level of the act would reach them, naming which of the ordinary things students do with these tools are acceptable and which are not. The channel through which such a rule is disseminated matters as much as its wording. An official document is the main source for 1.2% of students and a syllabus for 4.0%, while a professor speaking in class reaches 17.3%, so the professor is the only institutional channel with real reach.

The distinction between support and substitution offers a workable basis for drafting, since students already hold it and apply it consistently. A rule written along that line would codify a norm these students have arrived at on their own rather than impose one on them. The open responses show how little this takes in practice. Four students asked a professor directly and received an answer they could act on. In one case, a tool could be used to understand a novel while the writing had to be their own.

5.5 Limitations

This study covers only a single institution, and its convenience sample was recruited through professors and student networks, so the proportions describe the respondents and not the Moroccan undergraduate population. It measures what students believe without checking it against institutional documents. The absence of a governing rule therefore rests on the published accounts of Didi (2025) and Ibriz (2026), not on verification carried out here. Eligibility was restricted to the Licence cycle, which leaves Master's and doctoral students unexamined even though their work is where questions of authorship are sharpest.

Both rating grids were presented in the same order to every respondent, so the possibility that answering about permissibility influenced the acceptability ratings cannot be excluded. The two source items do not reconcile, since 40 respondents named a main source they had not selected among the channels they had learned from. The confidence measure is a single item, and 19 respondents left it blank. The disagreement index is introduced here for the first time and has no external benchmark, so its values are interpretable against one another and not against any established standard.

Acceptability was self-reported, and students had an obvious interest in presenting themselves as strict, so the gap between supportive and substitutive acts may be wider here than in practice. Anonymity and the absence of any identifying data reduce that pressure without removing it, and Harman's test indicates only that no single response tendency drives the whole pattern. Finally, what students report their professors said is the message received and not the message sent, and the distance between the two is itself worth study.

6. Conclusion

Moroccan higher education has made the teaching of artificial intelligence compulsory while leaving its use by students unaddressed, and this study set out to establish what undergraduates at one university believe in the space that leaves. They believe, in overwhelming proportion, that no rule exists or that they cannot find out. They have assembled a working position from their own judgement and from the internet, with institutional documents reaching at most 5.2% of them. That position is coherent and consistent across faculties, and on the acts where it differs from what they attribute to their university, it is the stricter of the two.

The absence of a rule has therefore not produced the permissiveness institutions tend to fear. It has produced uncertainty, and the uncertainty is sharpest where professors contradict one another. For an institution preparing to write its first guidance, the hard work has already been done by the students, who distinguish between a tool that helps them work and a tool that works in their place. The remaining task is to write that distinction down and deliver it through the professors students actually hear.

Acknowledgements

Author contributions (CRediT): Safae Kharmach was responsible for conceptualisation, methodology, investigation, formal analysis, data curation, visualisation, and the writing of the original draft. Abderrahim El Karfa provided supervision and contributed to the methodology. Both authors contributed to the review and editing of the manuscript and approved the final version.

Declarations

Ethical standards: Participation was voluntary and anonymous, consent was recorded from every respondent, and no identifying data were collected. Formal ethical review is not an established practice for anonymous survey research at the institution.

Use of Generative AI: Artificial Intelligence was used in the analysis of the survey data and to assist with editing text for clarity. The research questions, theoretical framing, interpretation of findings and conclusions are the authors' own. The final manuscript was authored and reviewed by the listed authors, who take full responsibility for its content.

Funding: This research received no specific grant from any funding agency.

Conflicts of Interest: The authors declare none.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Baba, N. E., Anasse, K., & Slimani, N. (2025). Exploring the relationship between artificial intelligence awareness and its use among students: Cadi Ayyad University as a case study. *European Journal of Education and Pedagogy*, *6*(4), 22–33. https://doi.org/10.24018/ejedu.2025.6.4.960

[2] Ben Attou, A. (2021). La digitalisation de l'enseignement supérieur: Limites, enjeux et perspectives. *Espace Géographique et Société Marocaine*, *1*(53). https://doi.org/10.34874/IMIST.PRSM/EGSM/28421

[3] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

[4] Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, *20*, Article 38. https://doi.org/10.1186/s41239-023-00408-3

[5] Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*(2), 228–239. https://doi.org/10.1080/14703297.2023.2190148

[6] Didi, N. (2025). Safeguarding educational integrity: Ethical and regulatory perspectives on generative AI in Moroccan higher education. *CONFLUENCE: Journal of Multidisciplinary Studies*, *1*(1), 209–221. https://doi.org/10.34874/PRSM.cjms-vol1iss1.5027

[7] El Ferouali, S., & Ouhadi, S. (2023). Digital transformation in Moroccan higher education: A literature review. *African Journal of Management, Engineering and Technology*, *1*(2), 135–148. https://revues.imist.ma/index.php/AJMET/article/view/44921

[8] El Mrabet, H., & Ait Moussa, A. (2023). A framework for predicting academic orientation using supervised machine learning. *Journal of Ambient Intelligence and Humanized Computing*, *14*(12), 16539–16549. https://doi.org/10.1007/s12652-022-03909-7

[9] Errifaiy, H., Aergoub, C., Rhezali, M., Nejmi, H., & Abouelhassan, T. (2026). Perceptions and attitudes toward artificial intelligence among Moroccan public medical school students: A cross-sectional study. *Frontiers in Education*, *11*, Article 1800337. https://doi.org/10.3389/feduc.2026.1800337

[10] Freeman, J. (2025). *Student generative AI survey 2025 (HEPI Policy Note 61)*. Higher Education Policy Institute. https://www.hepi.ac.uk/wp-content/uploads/2025/02/HEPI-Policy-Note-61.pdf

[11] Ibriz, S. (2026, February 22). *Universités marocaines: Face à l'IA, l'appel urgent à une stratégie nationale*. Médias24. https://medias24.com/2026/02/22/universites-marocaines-face-a-lia-lappel-urgent-a-une-strategie-nationale-1631604/

[12] Ihya, R., Aitdaoud, M., Namir, A., Guerss, F. Z., & Haddani, H. (2020). Using machine learning algorithms to predict the e-orientation systems acceptancy. In M. Ben Ahmed, A. A. Boudhir, D. Santos, M. El Aroussi, & İ. R. Karas (Eds.), *Innovations in smart cities applications edition 3* (pp. 117–130). Springer. https://doi.org/10.1007/978-3-030-37629-1_10

[13] Mamoune, L., Kharmach, S., El Karfa, A., & Bennis, M. (2025). AI tools in academic research: Perceptions, adoption, and challenges among Moroccan student researchers. *Arab World English Journal*, *16*(4), 479–492. https://dx.doi.org/10.24093/awej/vol16no4.26

[14] Messaoudi, F., & Talbi, M. (2012). *Réussir l'intégration des TICE au Maroc: Regard sur le déploiement de la stratégie nationale GENIE*. EpiNet, (143). https://www.epi.asso.fr/revue/articles/a1203e.htm

[15] Mjahad, R., Boukranaa, A., El Karfa, A., & Sandy, K. (2025). The implementation of artificial intelligence in Moroccan higher education: Benefits and challenges from the perception of academics. *SAGE Open*, *15*(3). https://doi.org/10.1177/21582440251381402

[16] Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open*, *5*, Article 100151. https://doi.org/10.1016/j.caeo.2023.100151

[17] Moustghfir, S. (2026). Global lessons on AI: Challenges and strategies for tertiary education in Morocco. *International Journal of Innovative Research in Multidisciplinary Education*, *5*(1), 52–68. https://doi.org/10.5281/zenodo.18516889

[18] Nouib, A., El Ganich, S., & Moubarik, R. (2022). Le comportement d'adoption des technologies d'e-learning dans les universités marocaines: Une étude empirique exploratoire. *Revue de l'Entrepreneuriat et de l'Innovation*, *5*(17). https://doi.org/10.34874/IMIST.PRSM/reinnova-v5i17.36209

[19] Riyami, B. (2018). *Analyse des effets des TIC sur l'enseignement supérieur au Maroc dans un contexte de formation en collaboration avec une université française [Thèse de doctorat, Université Bretagne Sud]*. HAL Thèses en ligne. https://theses.hal.science/tel-02004288

[20] Selwyn, N. (2016). Is technology good for education? Polity Press.