Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI Risk Awareness and Reported Rule Following among Law Students in Bangladesh

Dr Kazi Abdul Mannan

drkaziabdulmannan@gmail.com

Center for Academic & Professional Career Development and Research (CAPCDR)
https://orcid.org/0000-0002-7123-132X

Md. Mafuz

Shanto-Mariam University of Creative Technology https://orcid.org/0009-0005-9056-3046

Research Article

Keywords: generative AI, AI literacy, regulatory awareness, law students, responsible AI, Bangladesh

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11636662/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

AI Risk Awareness and Reported Rule Following among Law Students in Bangladesh

Md. Mafuz
Department of Law, Shanto-Mariam University of Creative Technology, Dhaka, Bangladesh
Email: mafuzmorul@gmail.com ORCID: 0009-0005-9056-3046

Prof. Dr Kazi Abdul Mannan
Department of Business Administration, Shanto-Mariam University of Creative Technology, Dhaka, Bangladesh
Corresponding author: drkaziabdulmannan@gmail.com ORCID: 0000-0002-7123-132X

Abstract

Understanding the risks of generative artificial intelligence (AI) and believing that regulatory awareness helps responsible use are distinct from reporting adherence to applicable rules. This cross-sectional study examined those distinctions among 133 respondents recorded as law students in a 385-record survey on AI governance in Bangladesh. We report full item distributions, joint endorsement and paired contrasts, and fit exploratory models of the original rule-following rating. Among law students, 108 (81.2%) endorsed avoiding unprotected disclosure of confidential or sensitive information, 69 (51.9%) endorsed verification before reliance or sharing and 64 (48.1%) endorsed following applicable institutional or professional rules. Confidentiality endorsement exceeded rule following by 33.1 percentage points (pointwise 95% CI 21.8 to 44.4; Holm-adjusted $p < 0.001$), while the verification difference did not meet the adjusted threshold. Of the 103 respondents endorsing the benefit of regulatory awareness, 51 also endorsed rule following. Self-reported understanding of ethical and legal risks, general AI awareness and perceived benefit of regulatory awareness did not meet the adjusted threshold in exploratory regression. The conclusion was unchanged in binary and full-sample sensitivity analyses. The findings distinguish a specific information-handling precaution from reported adherence to unspecified rules. They support evaluating whether students can locate, interpret and apply relevant rules in concrete tasks. The study does not measure demonstrated AI literacy, knowledge of regulations, audited conduct or the effect of education on compliance.

Keywords: generative AI; AI literacy; regulatory awareness; law students; responsible AI; Bangladesh.

1. Introduction

A law student using generative AI may need to decide whether an output is accurate, whether information can safely be entered and whether a particular use is permitted. These questions can arise in the same activity but call for different forms of judgement. Checking a legal proposition requires a reliable source. Deciding whether to disclose information requires attention to sensitivity and safeguards. Determining whether a use is permitted requires identifying the relevant institution, course, placement or professional setting and the rules that apply there. Familiarity with one task does not demonstrate competence in the others.

This distinction has practical consequences for teaching. A university may explain that generative systems can produce unsupported information without showing students how to verify an answer. It may warn against disclosure without giving students a usable way to identify protected information. It may publish a policy without explaining how different requirements apply to coursework, research and client-related activities. Students can agree with the principles and still need help translating them into decisions. Guidance on generative AI in education accordingly treats governance and capability development as separate parts of responsible use.[1]

1.1 Research problem

The assigned Group 17 topic concerns AI literacy, regulatory awareness and responsible AI behaviour among university students in Bangladesh. The supplied questionnaire does not test a student's knowledge of a particular regulation. Instead, it asks whether respondents understand ethical and legal risks and whether they believe greater regulatory awareness improves their behaviour. It separately asks respondents to report verification, confidentiality precautions and rule following. This distinction changes the empirical question. We can examine whether reported awareness or a belief about its benefits covaries with reported rule following, but cannot determine whether respondents know the rules or comply with them in practice.

The content of the conduct items also differs. Confidentiality concerns whether the respondent avoids entering protected information without safeguards. Verification concerns checking important material before reliance or sharing. Rule following is broader and refers to whatever institutional or professional rules the respondent understands to apply. A respondent can find one practice easier to judge than another. Aggregating these responses would mask which precautions are widely endorsed and which require further inquiry.

[1] Fengchun Miao and Wayne Holmes, *Guidance for Generative AI in Education and Research* (UNESCO 2023).

1.2 Study focus and contribution

We ask how frequently law students endorse each statement; whether verification and confidentiality endorsement differ from rule following within the same students; and whether risk understanding, general AI awareness or belief in the benefit of regulatory awareness are associated with the rule-following rating. We also ask how often students who believe awareness helps report following rules. The achieved law-student sample is 133, while 385 is the size of the mixed-status workbook. This distinction preserves the population focus without relabelling graduates, practitioners or other respondents as students.

The article differs from companion Group 17 studies on trust in governance and support for regulation. It focuses on personal judgements and practices rather than attitudes to national governance. The studies share the same questionnaire records and are not independent samples. The contribution is to show how reported precautions can diverge from reported rule following and to identify questions that can be tested through policy literacy and performance assessment.

2. Literature Review and Theoretical Foundation

The framework distinguishes competence, attitudes about the benefit of awareness and opportunities to follow rules. AI literacy research helps clarify the difference between reported familiarity and performance. The theory of planned behaviour explains why a favourable judgement need not correspond to conduct. Responsible AI in education provides a practical context for assessing what institutions can make clearer or easier to do. These ideas inform interpretation; the questionnaire does not directly measure their complete constructs.

2.1 AI literacy and task performance

AI literacy includes competencies needed to recognise, understand and evaluate AI and use it effectively. A person may know that a model can make errors but still struggle to recognise an error in a particular answer. In legal education, competence may require finding the authoritative source, checking its jurisdiction and assessing whether the passage supports the proposition. This is different from asking whether the person understands that inaccuracy or bias can occur. The study uses the latter as self-reported awareness, not a proxy for verified performance.[2]

Evaluating performance also requires defining the opportunity to act. A verification exercise can provide an answer and several sources, allowing an evaluator to record which

[2] Duri Long and Brian Magerko, ‘What Is AI Literacy? Competencies and Design Considerations’ (Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM 2020) doi:10.1145/3313831.3376727.

claims a learner checks. A confidentiality exercise can ask a learner to identify sensitive details and choose a safer alternative. A policy exercise can provide the applicable rule and require the learner to apply it to a particular use. Using several tasks reveals where difficulties arise; a single rating of responsible behaviour cannot provide the same evidence.

2.2 Attitudes, perceived control and conduct

The theory of planned behaviour distinguishes attitudes towards conduct, subjective norms, perceived control and intentions. A favourable view of a practice is only one part of explaining whether a person intends or is able to perform it. In this setting, a student can believe that knowing rules is helpful and still have difficulty finding the relevant policy or interpreting it. The views of instructors and peers could also influence behaviour. The questionnaire does not directly measure these mechanisms. The theory is used to interpret why belief in awareness and reports of conduct need not coincide.[3]

Rule following depends on what rules are understood to apply. A student could answer with a course policy in mind, while another thinks about a professional placement or a general ethical expectation. The instrument does not identify the applicable institution or rule, whether it was communicated or whether participants faced a recent decision to which it applied. The rule-following response is therefore not an audit of compliance. A neutral response could reflect uncertainty about the question or a lack of applicable experience rather than non-compliance.

2.3 Responsible AI and institutional support

Institutions can help students apply rules by giving examples that match the activities students undertake. A statement about academic integrity may be easier to apply if it explains when AI may assist planning, what assistance must be disclosed and what work the student must perform independently. Privacy guidance can likewise distinguish prohibited disclosure from permitted processing under specified safeguards. Where several policies apply, students need to know where to seek clarification and who can resolve conflicting instructions. These are features of the setting; the present survey does not measure their availability or effectiveness.[4]

Institutional trust can be relevant if students are more willing to seek clarification when they expect a fair and useful response. The focal analysis does not test that proposition. Its predictors concern risk understanding, general awareness and the perceived benefit of awareness, not trust in a named institution. The broader questionnaire contains governance

[3] Icek Ajzen, 'The Theory of Planned Behavior' (1991) 50 Organizational Behavior and Human Decision Processes 179.

[4] Miao and Holmes (n 1).

judgements, but using them as though they measured the policy environment would extend the interpretation beyond their wording.

2.4 Analytical expectations

Self-reported risk understanding may coincide with rule following if respondents who recognise ethical and legal risks also attend to relevant requirements. General AI awareness may be related if recognition of inaccuracy and bias accompanies careful use, but the items do not ask whether respondents understand institutional obligations. Belief in the benefit of awareness may track rule following if people who value knowledge also report acting on it. It may also be common among people who are uncertain about their own practice. These are possible explanations for association, not a pre-established causal sequence. We report the overlap and the separate regression estimates to distinguish them.

3. Methodology

The study analyses a cross-sectional questionnaire, with law students as the primary group and all available records used only in a sensitivity model. Reporting distinguishes observed judgements from objective skills and from legally verified conduct.[5]

3.1 Sources and eligibility

The research plan assigns the topic to Md. Mafuz. The Group 17 workbook contains a Data sheet, a codebook and a README, and the supplied questionnaire describes voluntary participation, no requirement to provide a name, aggregate academic reporting and the option to discontinue. The workbook contains 385 records with affirmative consent, adult eligibility and general AI awareness. The recorded current-status category “Law student” defines the primary sample of 133. The remaining 252 records are not treated as students for the main analysis.

We checked sequence uniqueness, valid values for the 20 agreement items, completeness of the analysed variables and duplicate response and profile patterns after excluding the sequence and contextual text. No blank numerical or profile fields or duplicate patterns were found in those checks. The study team supplied the numerical data as empirical and identified the synthetic descriptors attached to the sequence and contextual fields as metadata errors. Contextual text was excluded. The supplied materials do not establish exact recruitment dates, distribution mode, a verified sampling frame or the invitation denominator. The achieved sample is not used to infer a national response rate or probability-sample margin of error.

[5] Erik von Elm and others, ‘The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) Statement’ (2007) 4 PLoS Medicine e296.

3.2 Measures

Table 1 describes the analytical roles. RRB1 measures self-reported risk understanding. RRB2 and RRB3 record reported verification and confidentiality precautions. RRB4 measures reported rule following. RRB5 asks whether awareness helps, rather than measuring regulatory knowledge. The equal-weight mean of AIL1 and AIL3 summarises reported general awareness of inaccuracy and bias. It excludes behaviour and oversight items.

Table 1. Roles of the observed measures

Measure	Item	Interpretation
Risk understanding	RRB1	Self-reported understanding of ethical and legal risks
Verification	RRB2	Reported checking before reliance or sharing
Confidentiality	RRB3	Reported avoidance of unprotected disclosures
Rule following	RRB4	Reported adherence to applicable rules
Awareness benefit	RRB5	Belief that regulatory awareness helps responsible use
General AI awareness	AIL1 and AIL3	Mean of reported recognition of inaccuracy and bias

RRB4 is the regression outcome. RRB2 and RRB3 are used in the paired comparisons.

Appendix A reproduces the wording and anchors. We report coefficient alpha and rank correlations as diagnostics, not validation. Different item content is the principal reason to avoid a single composite for responsible AI. The study does not estimate latent variables or mediation.[6]

3.3 Statistical analysis

Endorsement combines ratings of 4 and 5. Neutrality is included in non-endorsement and is not interpreted as opposition. We report full response distributions and cross-tabulate the perceived benefit of awareness with rule following. Two paired contrasts compare verification and confidentiality separately with rule following. The paired estimate is the comparator-only count minus the rule-following-only count, divided by 133. Exact McNemar tests use the discordant pairs. Pointwise percentile intervals use 10,000 whole-record resamples with replacement and seed 20261104. Holm adjustment covers both tests.[7]

The exploratory linear model uses the original 1–5 rule-following rating. The three focal terms enter together. The adjusted student model adds categorical location, training and

[6] Klaas Sijtsma, 'On the Use, the Misuse, and the Very Limited Usefulness of Cronbach's Alpha' (2009) 74 Psychometrika 107.

[7] Quinn McNemar, 'Note on the Sampling Error of the Difference between Correlated Proportions or Percentages' (1947) 12 Psychometrika 153; Sture Holm, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65.

AI-use frequency; status is constant. The full-sample sensitivity model also adjusts for status. Unstandardised coefficients use HC3 standard errors and pointwise normal-approximation intervals, and Holm adjustment covers the three focal terms within each model. A binary sensitivity model treats endorsement as a 0–1 outcome and reports linear probability differences. The full-sample model overlaps the student model and is not external validation or a test of differences between groups.[8]

3.4 Ethics and data protection

Ethical approval for this study was obtained from the Center for Academic & Professional Career Development and Research (CAPCDR), an external institutional body, under approval reference CAPCDR: SMUCT/LLB/SS-0020260908-13. External ethical review was sought because Shanto-Mariam University of Creative Technology does not have an institutional ethics committee responsible for reviewing this type of research.

4. Results

The results describe the law-student sample, item distributions and joint endorsements before presenting paired differences and adjusted models. Unless otherwise indicated, the denominator is 133.

4.1 Student sample

Table 2 shows age, location and training. Forty-eight respondents (36.1%) were aged 21–23. Dhaka accounted for 58 records (43.6%), district towns for 33 and other city corporations for 30. Seventy-three respondents reported no formal training, 41 reported a short orientation and 19 reported course or certified training.

Table 2. Respondent profile

Characteristic	Category	n (%)
Age	21-23	48 (36.1%)
Age	24-26	30 (22.6%)
Age	18-20	26 (19.5%)
Age	31+	15 (11.3%)
Age	27-30	14 (10.5%)
Status	Law student	133 (100.0%)
Location	Dhaka	58 (43.6%)
Location	District town	33 (24.8%)
Location	Other city corporation	30 (22.6%)

[8] James G MacKinnon and Halbert White, ‘Some Heteroskedasticity-Consistent Covariance Matrix Estimators with Improved Finite Sample Properties’ (1985) 29 Journal of Econometrics 305; Holm (n 7).

Characteristic	Category	n (%)
Location	Upazila-rural area	10 (7.5%)
Location	Other	2 (1.5%)
Training	No	73 (54.9%)
Training	Yes, short orientation	41 (30.8%)
Training	Yes, course or certified training	19 (14.3%)

Recorded status, rather than age or educational level, defined student eligibility. The location field does not identify a campus or recruitment site. Training categories describe respondents' reports and do not document course content.

4.2 Distribution of responses

Table 3 gives the full ratings. Risk understanding was endorsed by 77 respondents (57.9%) and verification by 69 (51.9%). Confidentiality was endorsed by 108 (81.2%). Reported rule following was endorsed by 64 (48.1%), with 56 neutral responses and 13 disagreements. The benefit of regulatory awareness was endorsed by 103 (77.4%).

Table 3. Distribution of responsible AI responses

Item	1	2	3	4	5	Endorsed
RRB1	0	7	49	60	17	77 (57.9%)
RRB2	2	10	52	53	16	69 (51.9%)
RRB3	0	3	22	72	36	108 (81.2%)
RRB4	0	13	56	55	9	64 (48.1%)
RRB5	0	4	26	65	38	103 (77.4%)

Endorsement combines categories 4 and 5.

The RRB block alpha was 0.289. Pairwise Spearman correlations ranged from approximately-0.021 to 0.153. The relatively weak covariation is consistent with interpreting the statements individually. Only ten students (7.5%) endorsed all five, which describes a response pattern rather than a competence threshold.

4.3 Joint and paired results

Table 4 shows that 51 students (38.3%) endorsed the benefit of regulatory awareness and rule following together. Another 52 endorsed the benefit without endorsing rule following, 13 endorsed rule following without endorsing the benefit, and 17 endorsed neither.

Table 4. Joint endorsement of perceived awareness benefit and reported rule following

Awareness benefit	Rule following not endorsed	Rule following endorsed	Total
Not endorsed	17 (12.8%)	13 (9.8%)	30 (22.6%)
Endorsed	52 (39.1%)	51 (38.3%)	103 (77.4%)

Awareness benefit	Rule following not endorsed	Rule following endorsed	Total
Total	69 (51.9%)	64 (48.1%)	133 (100.0%)

Percentages are of all 133 law-student records.

Of those endorsing the benefit, 49.5% endorsed rule following, compared with 43.3% of the remainder. These are descriptive proportions. The 52 students endorsing the benefit alone should not be described as violating rules: the category includes neutrality, and no conduct was audited.

Table 5 compares reported verification and confidentiality with rule following. Confidentiality exceeded rule following by 33.1 percentage points (95% CI 21.8 to 44.4), based on 58 versus 14 discordant responses. Verification exceeded rule following by 3.8 points (95% CI-7.5 to 15.8), based on 34 versus 29 discordant responses.

Table 5. Paired contrasts with reported rule following

Contrast	Comparator only	Rule following only	Gap (pp)	95% CI (pp)	Holm p
RRB2- RRB4	34	29	3.8	-7.5 to 15.8	0.615
RRB3- RRB4	58	14	33.1	21.8 to 44.4	<0.001

Pointwise percentile intervals and Holm-adjusted exact McNemar tests.

The confidentiality contrast retained an adjusted p value below 0.001. The verification contrast did not (adjusted p = 0.615). The latter does not establish equivalent endorsement. The findings distinguish a specific information-handling precaution from a broad report of rule adherence.

4.4 Exploratory models

Table 6 reports the regression estimates. Adjusted coefficients were-0.074 for risk understanding (95% CI-0.277 to 0.130),-0.060 for general AI awareness (−0.315 to 0.195) and 0.124 for perceived awareness benefit (−0.082 to 0.330). None met the multiplicity-adjusted threshold. The adjusted student model explained 9.7% of the observed variation.

Table 6. Associations with the reported rule following response

Model (n)	Term	Estimate	95% CI	Holm p	R^2
Unadjusted (133)	Risk understanding	-0.004	-0.194 to 0.187	1.000	0.016
Unadjusted (133)	General AI awareness	-0.064	-0.309 to 0.180	1.000	0.016
Unadjusted (133)	Awareness benefit	0.127	-0.068 to 0.322	0.605	0.016
Adjusted (133)	Risk understanding	-0.074	-0.277 to 0.130	0.957	0.097
Adjusted (133)	General AI awareness	-0.060	-0.315 to 0.195	0.957	0.097
Adjusted (133)	Awareness benefit	0.124	-0.082 to 0.330	0.718	0.097
Binary endorsement (133)	Risk understanding	-0.034	-0.166 to 0.097	1.000	0.093

Model (n)	Term	Estimate	95% CI	Holm p	R^2
Binary endorsement (133)	General AI awareness	-0.034	-0.205 to 0.136	1.000	0.093
Binary endorsement (133)	Awareness benefit	0.037	-0.088 to 0.162	1.000	0.093
All respondents (385)	Risk understanding	-0.001	-0.115 to 0.112	0.982	0.078
All respondents (385)	General AI awareness	0.057	-0.078 to 0.192	0.818	0.078
All respondents (385)	Awareness benefit	0.093	-0.015 to 0.202	0.278	0.078

HC3 intervals are pointwise. Binary endorsement coefficients are probability differences; other coefficients are on the 1–5 scale.

The binary sensitivity analysis and the 385-record full-sample analysis likewise yielded no focal association meeting the adjusted threshold. The estimates do not establish equivalence or rule out practically relevant associations. They provide no basis for claiming that increasing awareness improves compliance.

5. Discussion

The most clearly supported finding is the within-student difference between confidentiality and rule following. More than four fifths endorsed avoiding unprotected disclosure, while fewer than half endorsed following applicable rules. Verification was closer to rule following and did not differ clearly after adjustment. Together, these results show why a generic responsible-use score can obscure which conduct respondents report.

5.1 A specific precaution and a broad obligation

The confidentiality statement describes a recognisable action and identifies a condition under which disclosure may be inappropriate. The rule item requires respondents to decide which rules are relevant and consider their broader practice. A student may find it easier to say how they handle sensitive information than to make a general statement about compliance. The difference may therefore reflect specificity, opportunity to act or uncertainty about the governing rules. The survey cannot distinguish these explanations.

The 56 neutral responses on rule following are especially important. Treating neutrality as a negative answer in a paired analysis is a way to define endorsement, not evidence that those students broke a rule. A follow-up could ask whether the student can identify a relevant policy, has used a tool in an activity covered by it and can describe how they interpreted a specific provision. That would distinguish lack of knowledge from lack of opportunity or uncertainty about the wording.

5.2 Belief in awareness and evidence of conduct

The joint results show that belief in the benefit of awareness is common but does not identify the same group as reported rule following. The adjusted estimates do not clarify whether awareness is positively associated with rule following in this sample. A person may value knowledge of rules and still lack the resources, instruction or opportunity needed

to apply them. Another may report adherence because an instructor gave a clear task-specific instruction, even without strong general awareness. The theory of planned behaviour provides a vocabulary for separating attitudes, control and conduct, but the survey does not measure all of these mechanisms.[9]

Likewise, a lack of adjusted evidence is not evidence that instruction has no benefit. Neither the content nor timing of training is documented in enough detail to identify an intervention. Respondents may differ in their confidence about their own abilities and in the standards they use to judge them. A study of the effect of instruction should specify what was taught, compare performance on common tasks and document whether learners can apply what they learned in a new situation.[10]

5.3 Implications for legal education

The findings support assessing verification, confidentiality and policy application separately. For example, a student could correctly avoid entering client material and still rely on an unsupported legal proposition. Another might verify the proposition but fail to disclose assistance where a course requires it. A common score would not show which correction is needed. A useful assessment provides a realistic task and evaluates the decision, the sources used and the student's explanation of what remains uncertain.[11]

6. Theoretical and Practical Implications

The study separates awareness beliefs from reported conduct. This helps identify what an institution should measure when it evaluates responsible-use guidance. Better familiarity with risks may be a valuable intermediate outcome; the ability to apply a rule remains a separate question.

6.1 Teaching and assessment

A short sequence of cases could ask students to check an AI-generated legal claim, remove or avoid confidential information, and decide whether a proposed use is permitted under a supplied course policy. Assessment could record the authority checked, the reasoning for the privacy decision and the exact policy language applied. Cases should include an ambiguous situation in which the appropriate response is to request clarification. This tests whether students can recognise the limits of their knowledge rather than rewarding a confident answer to every question.

[9] Ajzen (n 3).

[10] Long and Magerko (n 2).

[11] Miao and Holmes (n 1); Long and Magerko (n 2).

6.2 Institutional practice

Universities can make rules easier to find and interpret by linking each rule to concrete examples. Guidance should distinguish AI assistance in preparatory work from requirements governing submission, research or professional placements. Staff can identify whom students should contact when the context is unclear. Evaluations can ask whether learners locate the current policy and reach an appropriate decision, rather than measuring only whether they report having heard of it. Such evaluations would test a mechanism the present study cannot observe.[12]

7. Conclusion Limitations and Future Research

Among 133 law students identified in a 385-record Bangladesh survey, confidentiality was endorsed substantially more often than reported rule following. Verification did not differ clearly from rule following in the paired comparison. Belief in the benefit of regulatory awareness was common, but the exploratory models did not establish a reliable adjusted association between the measured awareness judgements and reported rule adherence. The evidence supports assessing specific precautions and policy application separately.

The study does not identify a representative sample of all university students or verify actual compliance. Recruitment details are incomplete, the measures are self-reports and the applicable rules were not specified. Cross-sectional adjustment does not establish direction of association. Full-sample sensitivity analysis overlaps the primary sample, and companion papers share the workbook.

Future research could document the rules participants are expected to follow, recruit from more than one programme and compare policy interpretation before and after defined instruction. Performance tasks could measure verification, confidentiality decisions and rule application separately. A follow-up could then examine whether skills transfer to an unfamiliar task and whether students use available routes for clarification.

Declarations

Ethical approval and oversight

Ethical approval for this study was obtained from the Center for Academic & Professional Career Development and Research (CAPCDR), an external institutional body, under approval reference CAPCDR: SMUCT/LLB/SS-0020260908-13. External ethical review was sought because Shanto-Mariam University of Creative Technology does not have an institutional ethics committee responsible for reviewing this type of research.

[12] Miao and Holmes (n 1).

Informed consent

All analysed records contain affirmative consent and adult eligibility responses. The questionnaire described voluntary participation and the option to discontinue.

Data availability

Requests for de-identified numerical data and analytical specifications may be directed to the corresponding author, subject to confidentiality and applicable institutional requirements.

Funding and competing interests

No funding or competing interest was disclosed in the available materials.

Author contributions

Md. Mafuz is the lead student author. Prof. Dr Kazi Abdul Mannan is the supervising co-author and corresponding author.

Generative AI disclosure

Generative AI assisted drafting, editing and preparation of analytical code. Numerical results were computed programmatically from the supplied workbook. The authors are responsible for verification and interpretation.

Bibliography

Ajzen I, 'The Theory of Planned Behavior' (1991) 50 Organizational Behavior and Human Decision Processes 179

Holm S, 'A Simple Sequentially Rejective Multiple Test Procedure' (1979) 6 Scandinavian Journal of Statistics 65

Long D and Magerko B, 'What Is AI Literacy? Competencies and Design Considerations' (Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM 2020) doi:10.1145/3313831.3376727

MacKinnon JG and White H, 'Some Heteroskedasticity-Consistent Covariance Matrix Estimators with Improved Finite Sample Properties' (1985) 29 Journal of Econometrics 305

McNemar Q, 'Note on the Sampling Error of the Difference between Correlated Proportions or Percentages' (1947) 12 Psychometrika 153

Miao F and Holmes W, *Guidance for Generative AI in Education and Research* (UNESCO 2023)

Sijtsma K, ‘On the Use, the Misuse, and the Very Limited Usefulness of Cronbach’s Alpha’ (2009) 74 Psychometrika 107

von Elm E and others, ‘The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) Statement’ (2007) 4 PLoS Medicine e296

Appendix A Questionnaire wording

Table A1 gives the focal items and the general AI awareness statements.

Table A1. Exact questionnaire wording

Code	Questionnaire statement
AIL1	I understand that generative AI can produce plausible but inaccurate information.
AIL3	I understand that AI outputs may reflect bias or incomplete data.
RRB1	I understand major ethical and legal risks associated with generative AI.
RRB2	I check important AI-generated information before relying on or sharing it.
RRB3	I avoid entering confidential or sensitive information into AI tools without appropriate safeguards.
RRB4	I follow applicable institutional or professional rules when using AI.
RRB5	Greater regulatory awareness improves my responsible AI behaviour.

Source: Group 17 Codebook. Ratings range from 1 strongly disagree to 5 strongly agree.

RRB4 is the outcome and does not enter any predictor.