Babylonian Journal of Artificial Intelligence
Vol.2026, **pp.** 132-165
DOI: https://doi.org/10.58496/BJAI/2026/016; ISSN: 3006-5437
https://mesopotamian.press/journals/index.php/BJAI/index

Research Article

University Artificial Intelligence Policies and Students' Perceived Career Readiness: A Qualitative Study of the Policy–Experience Gap in Higher Education

Mohammed Alalwan[1], Sam Ansari[2,3], Elton Xhetani[1], Soliman Mahmoud[2,3,4], Dhafar Hamed Abd[5], Ayad Turky[6], Hissam Tawfik[2,7], Abir Hussain[2,3,8,*]

[1]*Royal Holloway Business School, Royal Holloway, University of London, London, United Kingdom*
[2]*Department of Electrical Engineering, College of Engineering, University of Sharjah, Sharjah, United Arab Emirates*
[3]*Research Institute of Science and Engineering, University of Sharjah, Sharjah, United Arab Emirates*
[4]*Department of Electrical Engineering, Fayoum University, Faiyum, Egypt*
[5]*College of Computer Science and Information Technology, University of Anbar, Ramadi, Iraq*
[6]*College of Computing and Informatics, University of Sharjah, Sharjah, United Arab Emirates*
[7]*School of Built Environment, Engineering and Computing, Leeds Beckett University, Leeds, United Kingdom*
[8]*School of Computer Science and Mathematics, Liverpool John Moores University, Liverpool, United Kingdom*

A R T I C L E I N F O

Article History
Received 08 Sep 2026
Revised 11 Nov 2026
Accepted 12 Nov 2026
Published 27 Nov 2026

Keywords
academic policymaking
artificial intelligence literacy
career readiness
generative artificial intelligence
higher education governance

ABSTRACT

The rapid expansion of generative artificial intelligence has presented universities with the need to balance safeguarding academic integrity and the originality of student work with preparing students for AI-influenced professional environments. This study examined the relationship between university policies on artificial intelligence and students' reported experiences of AI use and perceived career readiness. The research used an exploratory qualitative design with descriptive quantification and drew on three official university policy documents and semi-structured interviews with 10 students at undergraduate, graduate, and doctoral levels. The data were analysed using hybrid thematic analysis, combining deductive coding based on the affective, behavioural, cognitive, and ethical dimensions of the ABCE model with inductive identification of organisational patterns. In total, five main themes and 14 subthemes were identified. The analysed policies primarily emphasised academic integrity, student responsibility, risk management, and data protection, whereas student narratives more often emphasised efficiency, workload reduction, practical AI use, and preparation for future work environments. Although participants recognised AI-related errors, risks of over-reliance, and ethical concerns, their validation practices were often situational and self-directed. Participants also reported concerns about false detection or misclassification of AI use, fragmented institutional communication, limited student participation, and insufficient integration of AI education into the curriculum. Based on these findings, an exploratory AI-readiness governance framework is proposed around course-level policy clarity, systematic training in AI-output validation, ethics and data-security education, and structured mechanisms for student feedback during policy review. In this study, career readiness refers to students' perceived preparation for AI-influenced professional contexts rather than objectively measured employment outcomes, workplace performance, or employer-assessed competence. The findings therefore support a shift from purely regulatory approaches toward educational, participatory, and skills-oriented AI governance.

* *Corresponding author. Email: abir.Hussain@sharjah.ac.ae*

1. INTRODUCTION

The rapid emergence of generative artificial intelligence (AI) has presented higher education with one of the most important challenges of digital transformation [1]. On the one hand, universities must maintain academic integrity, the authenticity of student work, data security, and the validity of academic qualifications, while on the other hand, they must adapt their educational programmes to the needs of an AI-driven job market. This is important because organisational change, particularly in technology-intensive settings, is rarely a simple linear process; rather, it is shaped by interactions among institutional structures, stakeholder expectations, implementation practices, and emergent adaptation [2–4]. In such a context, AI is not only a technological tool but also a structural driver for redefining educational policies, assessment methods, the role of academic staff, student responsibility, and the meaning of professional preparation.

The importance of this issue becomes even more pronounced when the nature of generative AI is considered. Large language models are evolving rapidly and introduce educational opportunities alongside challenges involving output reliability, academic integrity, critical evaluation, and responsible use [5, 6]. These characteristics make it difficult for universities to rely on static guidance alone, because policy and curriculum must respond to changing tool capabilities, patterns of student use, and evolving academic expectations [7–9]. From this perspective, AI-related curriculum and policy reform is appropriately understood as an ongoing organisational adaptation process requiring institutional feedback and stakeholder participation [2, 3].

At the labour-market level, professional reports indicate growing employer attention to AI-related skills, while academic research increasingly examines how AI is reshaping workplace skill requirements [10, 11]. Such evidence provides contextual justification for examining how higher education supports students' preparation for AI-influenced professional environments, without treating industry reports as direct measures of career readiness. In this context, the concept of "AI literacy" is emerging as a key component of perceived professional preparation; a concept that includes the ability to recognise, use, evaluate, and responsibly deploy AI systems [12].

However, AI literacy is not a one-dimensional concept and cannot be limited solely to the ability to work with tools such as ChatGPT or other generative systems. In this study, the four-dimensional affective–behavioural–cognitive–ethical (ABCE) model is used as the main theoretical framework. This model explains AI literacy in four dimensions: affective, behavioural, cognitive, and ethical [13]. The affective dimension relates to students' attitudes, motivation, trust, and anxiety toward AI; the behavioural dimension refers to the practical and everyday use of AI tools; the cognitive dimension focuses on the ability to evaluate, validate, and understand the limitations of AI outputs; and the ethical dimension includes responsibility, privacy, scientific integrity, and prevention of technological abuse [12, 14]. Accordingly, the AI-literacy literature suggests that stronger preparation for AI-influenced professional contexts may require the integrated development of these four dimensions rather than behavioural tool use alone.

Research on AI literacy similarly indicates that meaningful engagement with AI extends beyond technical tool use to include understanding, evaluation, responsible application, and the contextual conditions under which AI is used [12, 14]. Accordingly, the present study considers students' reported experiences and adaptations alongside their practical use of AI tools.

However, one important gap in the existing literature is that many university policies on AI are focused more on control, limitation, identification of unauthorised AI use, and protection of academic reputation, while the reported experience of students is often shaped around productivity, workload reduction, time saving and self-directed use of AI tools. From the perspective of organisational change theory, this situation can create a kind of conflict between institutional priorities and the needs of the recipients of change [15]. As the main recipients of educational changes, students are not only affected by university policies, but how these policies are interpreted, accepted or circumvented can determine the success or failure of AI-based educational reforms.

Accordingly, this paper examines the relationship between university policies on artificial intelligence and students' reported experiences of AI use and perceived career readiness in higher education. In the present study, career readiness is treated as a perceived and interpretive construct derived from students' accounts of their preparation for AI-influenced professional contexts. The study does not assess employment attainment, employer evaluations, workplace performance, or objectively demonstrated occupational competence. References to career readiness in the empirical findings should therefore be understood as referring to perceived career readiness unless otherwise stated. The research is based on an exploratory qualitative design incorporating descriptive quantification of the identified themes. The research data include three university policy documents on the use of artificial intelligence and semi-structured interviews with 10 students. During the analysis phase, five main themes were extracted, including emotional, behavioural, cognitive, ethical responses, and organisational adaptation to artificial intelligence. To enhance analytical transparency, the 14 identified sub-themes were quantified descriptively by recording the number of interview participants whose accounts reflected each sub-theme. These frequencies were calculated exclusively from the student interview data and were not extended to the institutional policy documents. Instead, the three university policies were examined qualitatively to identify direct substantive corre-

spondence, indirect policy relevance, or the absence of construct-equivalent policy evidence in relation to the student-reported issues. This source-specific analytical approach enables a systematic interpretation of the relationships between institutional AI governance and students' reported experiences without treating policy provisions and student-experience constructs as numerically equivalent.

The main contribution of this paper is threefold. First, it applies the ABCE model as an analytical framework for examining students' AI-literacy experiences in relation to perceived preparation for AI-influenced professional contexts. Second, by combining policy-document analysis with student narratives, it identifies thematic correspondences and divergences between institutional policy orientations and student-reported experience. Third, descriptive quantification is used to show how frequently particular student-experience subthemes appeared across the interviews, while policy provisions are interpreted separately at the conceptual level. The findings are intended to inform universities, educational policymakers, and curriculum designers by identifying areas in which regulatory approaches could be complemented by educational, participatory, and skills-oriented support.

The remainder of the paper is organized as follows. Section 2 reviews the theoretical foundations and research background, and explains the links between university AI policies, AI multidimensional literacy, organisational change, and students' perceived career readiness. Section 3 describes the research methodology, including an exploratory qualitative design, data sources, data collection process, coding framework based on the ABCE model, hybrid thematic analysis, descriptive quantification, and strategies to strengthen the trustworthiness of the analysis. Section 4 presents the findings in the form of five main themes including students' emotional, behavioural, cognitive, and ethical responses, as well as organisational adaptation, and examines the gap between policy coverage and students' reported experience. Section 5 interprets the findings based on the ABCE model and the organisational change perspective, and presents a proposed framework for governing AI readiness in higher education. In Section 6, limitations of the study and suggested directions for future research are discussed, and finally, Section 7 summarises the most important results, theoretical and practical contributions of the study, and its implications for academic policymaking and students' career preparation.

2. RELATED WORK AND CONCEPTUAL BACKGROUND

The existing literature shows that the integration of AI in higher education cannot be considered simply as the addition of a technological tool to the teaching and learning process, but rather as a complex organisational change that affects academic policy-making, curriculum design, student performance assessment, and students' perceived professional preparation. In this context, two main approaches to understanding organisational change can be identified: planned change and emergent change. Planned change is usually defined as a linear, purposeful, top-down process in which the organisation's management determines and implements the change path [4], [16]. In contrast, emergent change emphasises the continuous, nonlinear, and adaptive nature of organisational transformation and considers change to be the result of ongoing interactions between technology, people, structures, and environmental conditions [3], [17]. Because generative AI is continuously evolving and its applications in educational and professional settings are constantly being redefined, an emergent change approach is more appropriate for analysing AI-based curriculum reforms.

In higher education, universities must balance the protection of academic standards, authorship, and assessment credibility with the pedagogical imperative to develop students' capacity for responsible and critical AI use. Recent policy research shows that institutional guidance frequently prioritises academic integrity and assessment control while also recognising the need for clearer educational integration and course-level guidance [7–9]. This situation is consistent with the organisational change literature, as change is often accompanied by conflicts between the values, interests, and priorities of different groups within the organisation [18], [19]. In this article, students are considered as the main recipients of change, as AI policies, assessment practices, and curriculum reforms directly affect their learning experience, academic behaviours, and perceptions of professional readiness [15]. Accordingly, understanding how institutional AI governance relates to student experience requires attention to both formal policy provisions and students' reported encounters with AI-related guidance, expectations, and support. This rationale supports the study's comparative focus on university policy documents and student narratives.

Alongside the organisational-change perspective, AI literacy represents a central component of students' preparation for future workplaces. AI literacy generally refers to the ability to understand, use, evaluate, and responsibly deploy AI systems [12]. However, this general definition is not sufficient for a detailed analysis of students' readiness, because the effective use of AI is not limited to technical skills, but also includes attitude, behaviour, critical thinking, and ethical responsibility. Accordingly, the ABCE model explains AI literacy in four dimensions: emotional, behavioural, cognitive, and ethical [13]. The emotional dimension is related to students' motivation, trust, concern, and psychological readiness to use AI; the behavioural dimension concerns the practical use of AI tools in academic and everyday activities; the cognitive dimension focuses on the ability to evaluate, validate, and understand the limitations of AI outputs; and the ethical dimension is related to issues such as academic integrity, privacy, data security, and responsibility in using this

technology [12, 14].

The ABCE model is useful because it frames AI literacy as a multidimensional construct rather than equating frequent tool use with professional preparation. A student who uses AI to generate text, summarise resources, or reduce workload may be behaviourally engaged with the technology, but such engagement does not by itself establish broader competence. Stronger preparation is more plausibly associated with behavioural capability alongside motivation to learn, critical evaluation, and ethical responsibility [12, 13]. Therefore, if universities focus only on controlling abuse, detecting unauthorised use, or limiting AI applications, they may not be able to develop AI literacy in students in a comprehensive manner.

The existing literature also emphasises the importance of critical evaluation when students work with generative AI. Large language models can generate inaccurate, biased, or fabricated information, and their educational use therefore requires fact checking, critical reasoning, and continued human oversight [5, 20, 21]. From this perspective, students should not only be able to use AI tools, but also be able to examine, compare, challenge, and interpret their outputs in an appropriate academic context. This emphasis is consistent with reflective approaches to learning in which assumptions, evidence quality, and the consequences of AI-assisted decisions are subject to deliberate review [22].

The ethical dimension is also an important component of preparation for AI-influenced professional contexts. In the academic environment, irresponsible generative-AI use can create risks involving authorship, plagiarism, fabricated content, privacy, and the credibility of assessment [6, 8]. In the professional environment, employers also expect graduates to be able to manage the risks associated with data, confidentiality, intellectual property, and accountability in the use of AI systems [23]. Therefore, university policies should go beyond general warnings and explain the boundaries of permissible, impermissible, and conditional use of AI in a clear and enforceable manner for students.

To maintain evidential precision, sources were used according to the type of claim being supported. Peer-reviewed academic literature provides the primary basis for theoretical, methodological, and empirical propositions. Official institutional documents provide the primary evidence for statements about the content of specific university AI policies. In addition, professional or industry reports are used only to contextualise contemporary labour-market developments. Contextual reports are therefore not treated as substitutes for peer-reviewed evidence concerning AI literacy, educational outcomes, or perceived career readiness.

TABLE I. Synthesis of the Literature on AI Policy, AI Literacy, and Perceived Career Readiness in Higher Education

Research Stream	Core Argument	Key References	Relevance to This Study	Research Gap Addressed in This Paper
Organisational change in higher education	Organisational change is complex, non-linear, and shaped by tensions between institutional control, stakeholder needs, and external pressures. Planned change provides structure, whereas emergent change better explains continuous adaptation under uncertainty.	[3, 4, 16, 17]	AI-driven curriculum reform is treated as an emergent change process because generative AI evolves continuously and requires ongoing institutional adaptation.	Existing studies discuss organisational change broadly, but provide limited empirical evidence on how students experience AI-related policy changes as direct recipients of institutional reform.
AI unpredictability and institutional policy response	Generative AI creates uncertainty for universities because its capabilities, risks, and acceptable uses are continuously changing. Institutions therefore attempt to balance innovation with academic integrity and risk management.	[7–9]	This stream explains why universities tend to develop policies focused on control, monitoring, responsible use, and protection of academic standards.	Prior work provides limited evidence on how institutional responses to generative AI correspond with students' reported practical needs and perceived professional preparation.
AI literacy as a multidimensional competence	AI literacy involves more than technical usage; it includes the ability to understand, use, evaluate, and apply AI responsibly. The ABCE model conceptualises AI literacy through affective, behavioural, cognitive, and ethical dimensions.	[12–14]	The ABCE model provides the theoretical basis for analysing students' affective responses, practical AI use, evaluative reasoning, and ethical awareness.	AI-literacy scholarship increasingly defines the required competencies, but fewer studies examine how these competencies relate simultaneously to institutional AI policy and students' reported experiences.
AI skills and perceived professional preparation	AI-related competencies are increasingly relevant to contemporary work, while professional reports and academic research point to growing expectations around effective and responsible use of AI-enabled tools.	[10, 11]	This stream supports examining how students perceive their preparation for AI-influenced professional environments.	The connection between AI policy in universities and students' perceived professional preparation remains underexplored, particularly from the student perspective.
Cognitive risks of AI use	Uncritical dependence on AI-generated outputs may weaken evaluative reasoning, reduce critical thinking, and increase the risk of accepting inaccurate or fabricated information.	[5, 20, 21]	This stream supports the analysis of students' verification practices, awareness of AI limitations, and ability to critically assess AI-generated content.	Existing research highlights cognitive risks, but limited work investigates how students actually manage these risks in daily academic practice.
Ethical and academic integrity concerns	AI use raises ethical concerns related to plagiarism, fabricated references, data privacy, confidentiality, and accountability. Responsible AI use is therefore essential for both academic and professional readiness.	[6, 8, 23]	This stream informs the analysis of students' ethical boundaries, privacy concerns, and understanding of acceptable AI use.	Although ethical risks are widely discussed, there is limited evidence on whether institutional guidance provides students with sufficiently clear and actionable boundaries.
Student voice and policy–experience alignment	Change recipients are directly affected by institutional reforms, and their interpretation of policy influences the success of change implementation. In higher education, students are central stakeholders in AI-driven curriculum reform.	[15, 19]	This stream frames students as active interpreters of AI policy rather than passive recipients of institutional rules.	Integrated analyses that connect formal AI guidance with students' reported perceptions, concerns, and self-directed AI practices remain comparatively limited.

Table I synthesises the main theoretical streams informing this study and shows how the existing literature leads to the specific policy–experience gap examined in this paper. Based on these foundations, the main research gap addressed in this article is that, although universities are developing policies to manage the use of AI, it remains unclear how such policy orientations correspond with students' practical AI-use experiences, multidimensional AI literacy, and perceptions of preparation for AI-influenced professional contexts. Previous literature has often focused either on institutional policy or on skills associated with the labour market; however, the relationship between formal university policy, students' practical experience, and multiple dimensions of AI literacy has rarely been examined simultaneously. Therefore, by combining policy-document analysis with student interviews, this study examines how institutional policy orientations relate to students' reported experiences and perceived professional preparation. This work does not test whether the policies cause measurable improvements in employment readiness or occupational performance. This orientation is consistent with the structure of the literature and the theoretical framework of the main research, which focuses on emergent change, AI literacy, and perceived professional preparation.

3. METHODOLOGY

This study adopts an exploratory qualitative design, incorporating descriptive quantification of coded interview themes. The analysis investigates the relationship between institutional approaches to AI governance and students' reported experiences with AI use and perceived career readiness. The qualitative component provides an in-depth interpretation of student narratives and institutional policy documents, while descriptive counts are used to indicate the prevalence of identified themes across participants. Policy documents are examined comparatively for relevant institutional provisions rather than treated as numerically equivalent instances of student-experience constructs. These numerical summaries are used to enhance transparency and facilitate structured comparison rather than to support statistical inference or population-level generalisation [24–26]. This redesign builds on the original work methodology, which included policy document analysis and semi-structured interviews with 10 students. The methodological design integrates the analysis of three institutional AI policy documents with semi-structured interviews involving 10 students, enabling a structured examination of the relationship between institutional governance approaches, students' AI-use experiences, and perceived career readiness. Table II summarises the methodological design of the study, while Figure 1 illustrates the workflow used to transform policy documents and student interviews into a structured policy–experience comparison. The main rationale for this design is that the issue of AI integration in higher education cannot be explained by numerical indicators alone, as students' perceptions of university policies, their use of AI tools, ethical concerns, and their perceived preparation for AI-influenced professional contexts are all interpretive, contextual, and multidimensional in nature.

TABLE II. Summary of the Methodological Design

Methodological Component	Procedure	Role in the Analysis
Research design	An exploratory qualitative design was employed, complemented by descriptive quantification of coded themes [25, 26].	To examine student-reported experiences and institutional AI-policy content while providing transparent descriptive summaries of recurring themes.
Document analysis	Three official university AI policy documents were purposively selected and analysed using the identifiers Policy 1–Policy 3. Document analysis followed established qualitative principles [24, 27].	To identify institutional provisions concerning permissible AI use, academic integrity, output validation, data protection, accountability, and student responsibility.
Semi-structured interviews	Ten undergraduate, master's, and doctoral students participated in semi-structured interviews conducted via Microsoft Teams between 12 and 19 April 2026. Interviews lasted approximately 15–35 minutes, and participants were assigned non-identifying codes (P1–P10). The interview design was informed by established guidance for semi-structured qualitative interviewing and purposive sampling [28, 29].	To capture students' reported experiences of AI use, perceptions of policy clarity, AI-literacy practices, institutional support, and perceived preparation for AI-influenced professional contexts.
Theoretical framework	The ABCE model guided deductive coding across the affective, behavioural, cognitive, and ethical dimensions of AI literacy [12, 13].	To organise the analysis around multidimensional AI-literacy experiences and their relationship with perceived professional preparation.
Hybrid thematic analysis	Deductive coding based on the ABCE dimensions was combined with inductive identification of organisational and context-specific patterns. A structured codebook defined the final 14 sub-themes and specified operational definitions and coding boundaries [30–32].	To preserve theoretical coherence while allowing patterns not predetermined by the ABCE framework to emerge from the interview and policy data.
Analytical trustworthiness	Trustworthiness was supported through explicit operational definitions, inclusion criteria, coding decision rules, direct quotation support, source-level frequency counting, examination of less common cases, and comparison across interviews and policy documents [24, 30, 33].	To maintain a transparent and traceable relationship between the source material, coding decisions, descriptive summaries, and interpretation.

Continued on next page

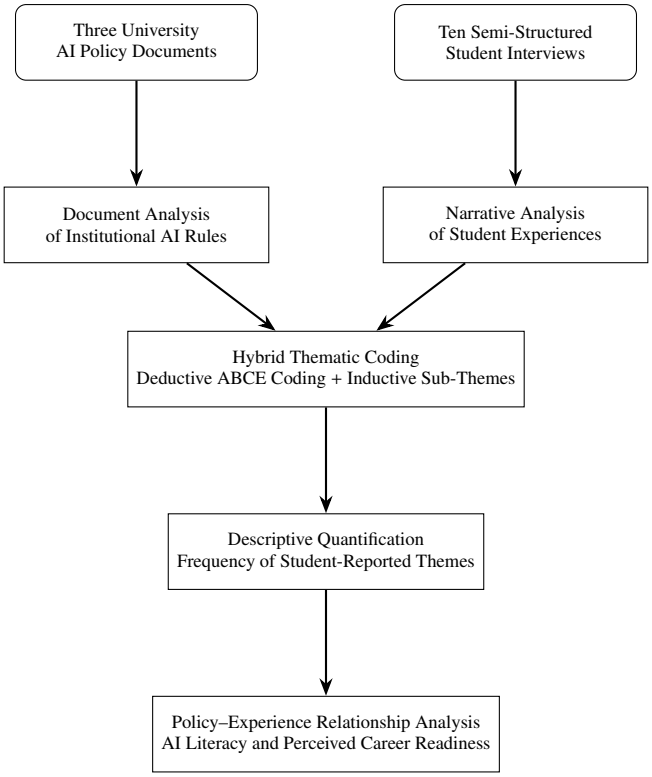

Fig. 1. Research workflow for analysing relationships between university AI policy provisions, student-reported AI-literacy experiences, and perceived career readiness.

Table II – *Summary of the Methodological Design*

Methodological Component	**Procedure**	**Role in the Analysis**
Descriptive quantification and cross-source comparison	For each student-experience sub-theme, the number of interview participants in whose accounts the theme appeared was recorded. Policy documents were analysed separately and classified as showing direct substantive correspondence, indirect policy relevance, or no construct-equivalent policy evidence [25–27].	To describe the distribution of student-reported experiences while preserving the conceptual distinction between interview-derived constructs and institutional policy provisions.
Ethical screening and participant safeguards	The study underwent institutional ethical screening using the Royal Holloway, University of London, Simplified Ethical Approval Form for module MN3311. The form, signed by the supervisor on 2 February 2026, records a determination of "No referral necessary," preceding the interviews conducted in April 2026. This constitutes a documented supervisor-level screening determination rather than formal research ethics committee approval. Participants were informed of the study purpose, confidentiality arrangements, intended use of their responses, and their right to withdraw. Written informed consent was obtained from all ten participants before their interviews. Interview findings were reported using non-identifying codes (P1–P10), without disclosing participant names or other direct personal identifiers.	To document prior ethical screening and participant consent, protect confidentiality, minimise identification risks, and transparently distinguish supervisor-level screening from formal committee approval.

Research data were collected from two main sources. The first source consisted of three official university policy documents on the use of AI in educational settings. These documents were selected using convenience sampling, as the aim of the research was to identify the observable institutional frameworks that students encounter in practice. These documents were purposively selected based on their direct relevance to the institutional governance of AI use in educational settings and their accessibility for systematic analysis. Document analysis allowed the research to extract official university orientations on issues such as the permitted use of AI, academic integrity, student accountability, validation of AI outputs, and data protection [27]. The second source of data consisted of semi-structured interviews with 10 students conducted between 12 and 19 April 2026 via Microsoft Teams. Individual interviews lasted between 15 and 35 minutes, with variation reflecting differences in the amount of detail provided by participants, as is compatible with a semi-structured interview format [28]. The sample included undergraduate, graduate, and doctoral students and was purposively structured

to introduce variation in academic experience rather than to provide statistical representation [29]. The interview protocol comprised four broad areas: participant background and AI-use experience; the affective and behavioural dimensions of the ABCE model; the cognitive and ethical dimensions of AI literacy; and students' perceptions of institutional change and perceived career readiness.

3.1 Participants and Policy Documents

The present study relied on two complementary data sources: semi-structured student interviews and official university policy documents concerning the use of artificial intelligence. The interview sample comprised 10 students drawn from undergraduate, graduate, and doctoral levels, thereby providing variation in academic experience and exposure to university-level educational practices. All interviews were conducted via Microsoft Teams and lasted between 15 and 35 minutes. Variation in interview duration reflected differences in the amount of detail provided by participants and was treated as a normal characteristic of the semi-structured interview process.

The sample included undergraduate, graduate, and doctoral students; however, participant-level information on disciplinary field, age, gender, and exact study level was not available for all participants and is therefore not reported. No demographic, disciplinary, or educational-level subgroup comparisons were conducted. The participant profiles reported in Table III are consequently analytical profiles based only on explicitly documented AI-use patterns, experiences, concerns, and response strategies.

The student interviews and policy documents were treated as two complementary analytical sources rather than as individually matched institution–participant cases. No one-to-one institutional correspondence was established between individual participants and specific policy documents. Accordingly, the policy documents were analysed as institutional cases illustrating approaches to AI governance in higher education, while the interviews were analysed as accounts of students' reported experiences of AI use, institutional guidance, and perceived preparation for AI-influenced professional contexts. The resulting policy–experience comparison therefore identifies thematic correspondences, divergences, and areas of policy silence across the two source types. It should not be interpreted as a direct evaluation of policy implementation within the participants' own institutions.

The study was designed as an exploratory qualitative investigation using a purposive sample. The purpose of recruiting 10 participants was not to obtain statistical representation of a wider student population. The aim was to examine in depth how students described their use of AI, their concerns regarding institutional guidance, and their perceptions of preparation for AI-influenced professional contexts. Inclusion of students from different educational levels was intended to introduce variation in academic experience rather than to support comparisons between degree levels. Accordingly, the findings are interpreted as contextual qualitative evidence and not as demographic or population-level estimates.

The interview participants were recruited through a combination of purposive and convenience sampling, as described in the original study protocol. Selection sought variation in educational level and experience with AI, while existing researcher–participant contacts facilitated recruitment within the available time frame. This recruitment approach, together with the small sample, limits the transferability of the findings and may introduce selection bias.

TABLE III. Analytical Profiles of Interview Participants Based on Explicitly Reported AI-Use Experiences

Participant	Reported Level of AI Engagement	Primary Reported Uses or Experiences	Main Cognitive, Ethical, or Institutional Concern	Reported Response or Adaptation Strategy
P1	Frequent and increasingly confident use	Used AI for a broad range of tasks and for supporting academic activities, including summarising readings and checking study-related information.	Expressed concern about inaccurate AI-detection systems, institutional accusations of unauthorised use, overdependence on AI, and the disclosure of personal or sensitive information.	Attempted to verify study-related outputs using prior knowledge and challenged the model when an answer appeared incorrect.
P2	Regular use for academic and everyday enquiries	Used AI frequently for studies, clarification, and answering general everyday questions.	Reported initial scepticism and acknowledged that AI-generated information was not always systematically verified.	Compared responses generated by different AI models when greater confidence in an answer was required.
P3	Extensive academic use	Reported that AI supported a substantial proportion of academic work and was used for summarisation, clarification, and other study-related activities.	Expressed concern regarding false AI-detection results, reduced independent thinking, incorrect answers, data retention, and unclear boundaries of acceptable academic use.	Re-prompted the system after incorrect responses, although verification was not routinely undertaken unless an answer appeared suspicious.

Continued on next page

TABLE III. Analytical Profiles of Interview Participants Based on Explicitly Reported AI-Use Experiences (continued)

Participant	Reported Level of AI Engagement	Primary Reported Uses or Experiences	Main Cognitive, Ethical, or Institutional Concern	Reported Response or Adaptation Strategy
P4	Task-oriented and efficiency-driven use	Used AI during examination periods and as a faster means of completing tasks and reducing academic workload.	Questioned whether universities were adequately adapting curricula and teaching practices to prepare students for AI-influenced employment.	Relied on AI as an academic support tool while emphasising the need for more professionally oriented university instruction.
P5	Selective use for time-consuming tasks	Used AI primarily when an activity was considered time-consuming and for obtaining information more efficiently.	Reported encountering inaccurate AI-generated information and questioned whether previously received outputs might also have been incorrect.	Reflected more critically on AI accuracy after encountering an incorrect response and advocated discipline-specific AI instruction.
P6	Frequent academic use	Reported using AI almost daily for studying and academic support.	Expressed concern that reliance on AI could reduce independent intellectual effort and recognised that user inputs might be retained or used for model training.	Switched to a different AI tool and reformulated instructions when an output was inaccurate or unsuitable.
P7	Frequent use in academic and everyday contexts	Regularly used tools such as ChatGPT and Claude and regarded them as convenient and efficient sources of information.	Reported a perceived decline in critical-thinking ability and expressed concern about data storage, misuse of personal images, and identity-related risks.	Continued using AI for efficiency while recognising the need for stronger privacy awareness and more structured institutional guidance.
P8	Regular weekly academic use	Used AI several times per week for study-related purposes and recognised its usefulness as an academic support tool.	Expressed concern about false detection, excessive cognitive dependence, unclear ethical boundaries, and the tendency to omit verification when under time pressure.	Distinguished between supportive use and direct copying and recommended greater integration of AI into the university curriculum.
P9	Efficiency-oriented academic use	Used AI mainly to save time, improve structure, and support the preparation of academic work.	Expressed concern that AI-assisted editing or restructuring could be incorrectly identified as misconduct and discussed the ambiguity between legitimate assistance and inappropriate authorship.	Rephrased prompts repeatedly when outputs were unsatisfactory and defined ethical use in terms of retaining personal authorship of the submitted work.
P10	Regular use for time saving and academic support	Used AI to reduce the time required for academic tasks and reported familiarity with inaccurate outputs and fabricated references.	Expressed concerns about unreliable AI-generated sources, inconsistent institutional communication, limited visibility of university guidance, and uncertainty about AI detection. Also described concern that Turnitin might incorrectly flag independently completed coursework as AI-generated, despite not directly copying AI-generated answers.	Reported verifying outputs in most cases, although verification could be omitted when workload was high, and relied on institutional platforms such as Moodle for policy information.

Note: The profiles report only characteristics explicitly supported by the interview data, including patterns of AI engagement, reported uses, cognitive, ethical, or institutional concerns, and response strategies. The available study records indicate that the overall sample included undergraduate, graduate, and doctoral students. However, participant-level study stage, disciplinary field, age, and gender were not available in a sufficiently complete form for reliable reporting. These characteristics were therefore not inferred retrospectively, and no demographic subgroup analyses were conducted.

This reporting decision was adopted to preserve the evidential integrity of the dataset. Contextual clues within individual interviews were not used to reconstruct missing demographic attributes because such reconstruction could introduce misclassification. Consequently, interpretation of the findings is limited to the experiential and analytical characteristics directly supported by the interview records.

The second data source consisted of three official institutional policy documents governing the use of artificial intelligence in higher education. The documents were purposively selected because they addressed complementary aspects of AI governance, including AI literacy, responsible use, academic integrity, assessment, output validation, data privacy, and institutional oversight. For transparency and reproducibility, Table IV reports the official policy document titles, issuing institutions, available version and approval details, and principal scope of application.

Policy 1 was the University of London *Artificial Intelligence (AI) Policy* [34], an institution-wide policy addressing artificial intelligence across education, research, professional services, and wider institutional activities. Policy 2 was

the University of Sharjah *Policy on the Responsible Use of Artificial Intelligence in Academic Texts* [35], which focuses particularly on graduate education, including master's and doctoral theses, graduate-level coursework, research plans, and related scholarly outputs. Policy 3 was the University of North Alabama *Generative AI Policy* [36], which provides institution-wide guidance for students, faculty, and staff concerning responsible generative-AI use in coursework, research, teaching, and administration.

The three documents were analysed as institutional policy cases rather than as individually matched policies governing specific interview participants. Accordingly, their comparison with the student interviews remains thematic and interpretive rather than institution-specific. Institutional provenance is reported to contextualise the policy corpus, but it should not be interpreted as establishing that any specific analysed policy governed a particular interview participant.

Primary institutional policy documents were used as the evidential source for claims concerning specific policy provisions, whereas peer-reviewed academic literature was used to support theoretical and methodological claims. Professional or industry reports were used only for contextual evidence and were not treated as substitutes for peer-reviewed research.

Table IV shows the dominant orientation, main audience, content areas and analytical features of the three policy documents reviewed. A comparison of the documents shows that all three policies address issues such as academic integrity, student responsibility, validation of outputs and data protection; however, the degree of operationalization of the guidelines and their emphasis on teaching applied skills are not the same. While some documents recognise the use of AI in the context of learning and skill development, the dominant orientation of the policies remains based on control, prevention of misuse and transfer of ultimate responsibility to the student or instructor.

TABLE IV. Profile of the University AI Policy Documents Included in the Analysis

Policy	Official Document Title	Institution	Date / Version Information	Scope and Intended Audience	Principal Analytical Relevance
Policy 1	*Artificial Intelligence (AI) Policy* [34]	University of London, United Kingdom	Reference UOL04; Version 1.2 (2024); approved 11 December 2024; annual review cycle	Institution-wide policy covering AI use across education, research, professional services, and wider institutional activities; relevant to students and staff.	AI literacy; guidance and training; responsible use; critical evaluation of AI outputs; academic integrity; data protection; stakeholder dialogue; skills development; and institutional monitoring.
Policy 2	*Policy on the Responsible Use of Artificial Intelligence in Academic Texts* [35]	University of Sharjah, United Arab Emirates	A formal version identifier and approval/publication date were not stated in the official source used for the analysis.	Graduate-level context, including master's and doctoral theses, thesis proposals, research plans, scientific publications stemming from theses, graduate courses and programmes, assignments, coursework, projects, and academic texts submitted to graduate-studies bodies.	Permitted and prohibited AI use; academic integrity; intellectual originality; disclosure and responsibility; output accuracy; data confidentiality; and graduate-level institutional accountability.
Policy 3	*Generative AI Policy* [36]	University of North Alabama, United States	A formal version identifier and approval/publication date were not stated on the official policy page used for the analysis; the document specifies annual review and updating.	Institution-wide guidance applying to students, faculty, and staff using generative AI in coursework, research, teaching, and administration.	Course-level clarification of acceptable AI use; academic integrity; acknowledgement of AI-generated content; output verification; data privacy and security; responsible use; faculty guidance; and caution regarding AI-detection and AI-grading tools.

Note: Titles, institutional affiliations, scope, and available version/date information are reported from the official institutional sources. Where a formal version identifier or approval/publication date was not explicitly available, no value was inferred. The three documents were analysed as institutional policy cases and were not individually matched to specific interview participants. Consequently, comparisons with student narratives are thematic and interpretive rather than institution-specific evaluations of policy implementation or effectiveness.

Table IV shows that the three documents represent distinct institutional and educational contexts. The University of London policy provides broad institution-wide guidance combining responsible AI adoption, AI literacy, ethical and legal considerations, and skills development. The University of Sharjah policy has a more specialised graduate-education focus and places substantial emphasis on academic integrity, intellectual originality, accountability, and responsible use in academic texts. The University of North Alabama policy provides institution-wide guidance with a strong course-level and teaching-oriented emphasis, including syllabus clarification, responsible use, output verification, data privacy, and caution regarding automated AI-detection and grading tools. These contextual differences were treated as characteristics of the policy corpus rather than as evidence that one policy was more effective than another. The analysis focused on the content and orientation of the documents and on their conceptual relationship to student-reported experiences.

3.2 Hybrid Thematic Analysis and Coding Procedure

Hybrid thematic analysis was used to analyse the data. This approach combines deductive and inductive logic. That is, in the first stage, the four-dimensional model of AI literacy including affective, behavioural, cognitive, and ethical dimensions was used as the initial coding framework, and in the second stage, emerging subthemes were extracted from

the interview data and policy documents [30], [31], [32], [13]. The emotional dimension was used to analyse students' attitudes, trust, concerns, and anxiety towards AI; The affective dimension was used to analyse students' attitudes, trust, motivation, concerns, and anxiety towards AI. The behavioural dimension was used to examine patterns of practical use of AI in academic activities. The cognitive dimension was used to assess students' level of understanding of limitations, errors, and the necessity of validating AI outputs; and the ethical dimension was used to analyse issues such as academic integrity, privacy, data security, and responsibility in the use of AI. In addition, an independent theme titled "Organisational Adaptation to AI" was also extracted to analyse the relationship between university policies, institutional communications, and students' experiences of educational change. In this work, data analysis was also conducted based on a combination of the ABCE model, the organisational change lens, and emerging subthemes.

To preserve consistency across the coding process, the codebook was treated as the primary decision framework throughout the analysis. When new or ambiguous evidence was encountered, the existing operational definitions and distinction rules were consulted before any coding decision was retained. Where the evidence did not satisfy the inclusion criteria for an existing code, it was not assigned merely on the basis of thematic similarity. This rule was particularly important for separating experiential student constructs from institution-level policy provisions.

3.3 Descriptive Quantification and Policy–Experience Comparison

Following completion of the thematic analysis, descriptive quantification was used to improve the transparency and consistency of reporting the student interview findings. For each student-experience sub-theme, the number of participants whose narratives contained the corresponding code was calculated and reported as both a frequency and a percentage of the 10 interview participants. These counts were used solely to describe the distribution and salience of recurring experiences within the interview sample and were not used for statistical inference or population-level generalisation.

A conceptual distinction was maintained between student-experience codes and institutional policy content. Sub-themes describing affective or behavioural states and practices, such as initial scepticism, growing confidence, routine AI use, or efficiency-oriented use, were applied only to student interview data because these constructs describe reported experiences or behaviours rather than properties of policy documents. Policy texts were therefore not coded as exhibiting the same experiential construct merely because they contained statements about awareness, training, acceptable use, educational value, or assessment practices.

Instead, the policy documents were examined for provisions that were conceptually relevant to each student-reported issue. Cross-source comparison distinguished between three analytically conservative categories: (i) *direct substantive correspondence*, where a policy explicitly addressed the same substantive issue identified in the student accounts. (ii) *indirect policy relevance*, where a policy addressed a related institutional concern without representing the same construct. (iii) *no construct-equivalent policy evidence*, where the available policy text did not instantiate the student-reported construct. Accordingly, the policy–experience comparison is interpreted as an analysis of conceptual relationship and institutional response rather than as evidence that identical constructs occur in two different types of data source.

Behavioural sub-themes were subject to the same source-specific coding rule. "Routine integration of AI into academic practice" was assigned only when an interview participant explicitly described recurrent, frequent, or embedded use of AI in academic activities. By contrast, "efficiency and workload reduction" was assigned only when a participant explicitly described time saving, convenience, reduced effort, or workload management as a motivation for AI use. Statements in policy documents promoting AI literacy, encouraging educational use of AI, or modifying assessment practices were not coded as instances of these behavioural constructs. Instead, such statements were analysed as institutional provisions or responses relevant to the broader context in which student AI use occurs.

All cross-source assignments were checked against the operational definitions in the codebook and the supporting policy excerpts. The purpose of this check was to ensure that institutional policy intent was not treated as equivalent to an affective state, behavioural pattern, or perceived organisational experience reported by students. A policy provision was classified as showing *direct substantive correspondence* only when it explicitly addressed the same substantive issue. Where the policy addressed a related institutional concern without representing the same construct, it was classified as *indirect policy relevance*. Where no construct-equivalent provision was identified, the relationship was classified as *no construct-equivalent policy evidence*. These categories were used for cross-source interpretation only and were not converted into policy percentages or treated as inferential measures.

The descriptive frequencies were not used to rank the affective, behavioural, cognitive, ethical, or organisational dimensions in terms of relative strength or weakness. The sub-themes differ in conceptual content and number, and their frequencies represent the occurrence of specific reported experiences rather than scores on a common measurement scale. Interpretations therefore focus on patterns within each theme rather than comparative claims that one dimension is stronger or weaker than another.

The cross-source comparison is presented in Table VI. Student frequencies are reported only for the interview-derived

sub-themes, while the policy columns describe the corresponding institutional provisions and the nature of their relationship to each student-reported issue. The source-level matrix in Figure 2 displays interview-derived student codes only. Policy documents are intentionally excluded because the cross-source analysis does not assume construct equivalence across source types. The policy–experience relationship is summarised separately in Figure 3 using the three qualitative relationship categories defined above.

Data analysis was conducted through a multi-stage process combining deductive and inductive logic. In the first stage, the full text of interviews and policy documents was read several times to familiarize the researchers with the context, structure, and diversity of the data content. In this stage, sections of the text that were relevant to students' attitudes toward AI, practical usage patterns, evaluation and validation of outcomes, ethical responsibility, privacy, organisational communication, and perceived career preparation were identified as initial semantic units. A coding unit could be a sentence, part of a response, or a set of related sentences, provided that it conveyed a distinct and independent analytical concept.

In the second stage, the four affective, behavioural, cognitive, and ethical dimensions of the ABCE model were used as the initial deductive framework. The affective dimension was used to code students' attitudes, trust, scepticism, motivation, and concerns about AI. The behavioural dimension covered actual usage patterns, frequency of use, practical goals, and the role of AI in academic activities. The cognitive dimension was dedicated to the ability to recognise limitations, assess the accuracy of outputs, validate information, and manage incorrect responses. The ethical dimension included issues such as academic integrity, plagiarism, personal responsibility, confidentiality, and data security. Using this theoretical framework provided initial coherence to the coding and allowed the findings to be directly linked to the multidimensional concept of AI literacy.

In parallel with the deductive coding, inductive reasoning was also used to identify patterns that did not fit neatly into the four primary dimensions of the ABCE model. In this way, a separate theme, "organisational Adaptation to AI," was formed, which included issues such as limited student participation in policymaking, fragmented organisational communications, and inadequate integration of AI into the curriculum. This hybrid approach prevented the data from being forced into a predetermined theoretical framework while maintaining the relevance of the analysis to the conceptual foundations of the research.

After identifying the initial codes, codes with similar content were compared and organized into specific subthemes. At this stage, the boundaries between closely related codes were carefully examined. For example, "understanding AI limitations and errors" was distinguished from "inconsistent verification practices": the former captures recognition of inaccurate or unreliable AI-generated information, whereas the latter captures the reported practice of checking that information. Similarly, "routine integration of AI into academic practice" was distinguished from "efficiency and workload reduction", because the former concerns recurrent usage and the latter concerns motivation for use.

In the process of refining the themes, each student-experience sub-theme was compared with the original interview data to ensure that its definition remained consistent with the supporting quotations. Policy excerpts were then reviewed separately to determine whether they showed direct substantive correspondence, indirect policy relevance, or no construct-equivalent policy evidence. Items that could be placed in more than one student sub-theme were categorised according to their dominant meaning and context. More than one interview code could be assigned when a passage conveyed two conceptually distinct meanings. Consequently, the sum of student sub-theme frequencies did not necessarily equal the number of participants.

To increase the transparency and consistency of analytical decisions, a qualitative codebook was developed, including the name of each theme, operational definition, inclusion criteria, rules of distinction, and examples of evidence. The final version of this codebook is presented in Table V. The use of clear operational definitions reduced the likelihood that the meaning of the codes would change during the analysis and allowed the frequency of each subtheme to be calculated based on the same criteria.

The qualitative coding and analytical interpretation were undertaken collaboratively by members of the research team. The analytical process involved applying the established coding framework, examining the identified sub-themes, and evaluating the consistency of coding decisions against the operational definitions and supporting interview evidence. Coding decisions were assessed with reference to the original interview material and the inclusion criteria specified in the codebook. This collaborative approach supported transparency and consistency in thematic interpretation without constituting a formal independent double-coding procedure or a quantitative assessment of inter-coder agreement.

TABLE V. Qualitative Codebook Used for the Analysis of Student Interviews and University AI Policy Documents

Main Theme	Sub-theme	Operational Definition	Inclusion Criteria	Distinction and Decision Rule
Affective responses	Initial scepticism towards AI	Initial distrust, uncertainty, hesitation, or limited understanding experienced when first encountering or considering the use of AI tools.	Statements referring to initial doubt about AI accuracy, usefulness, legitimacy, or relevance to academic work.	Applied only to initial or early-stage attitudes. Ongoing concern about institutional detection was coded separately.
Affective responses	Growing confidence in AI use	An increase in comfort, trust, willingness, or perceived ability to use AI after gaining practical experience.	Statements indicating greater confidence, reduced hesitation, increased reliance, or improved comfort following repeated use.	Confidence was coded separately from frequency of use. Frequent use without an explicit change in confidence was classified under behavioural engagement.
Affective responses	Uncertainty about AI detection	Anxiety or insecurity concerning automated detection, false accusation, misclassification, or unclear enforcement of AI-use rules.	References to AI-detection tools, Turnitin flags, fear of being accused of misconduct, or uncertainty about how instructors identify AI use.	Concerns about plagiarism in principle were classified under academic integrity unless the emphasis was on detection or accusation.
Behavioural responses	Routine integration of AI into academic practice	Regular or recurrent use of AI as part of studying, assignment preparation, clarification, information seeking, or everyday academic activity.	Statements indicating daily, weekly, frequent, or embedded use of tools such as ChatGPT or Claude.	The code captures frequency and integration. Statements focused mainly on saving time were additionally or alternatively coded as efficiency and workload reduction.
Behavioural responses	Efficiency and workload reduction	Use of AI primarily to complete tasks faster, reduce effort, manage workload, or obtain information conveniently.	References to time saving, convenience, faster task completion, reduced workload, or minimisation of effort.	The code reflects motivation for use rather than frequency. Regular use without an efficiency-related rationale was not sufficient for inclusion.
Cognitive responses	Awareness of reduced critical thinking	Recognition that frequent or excessive reliance on AI may weaken independent reasoning, problem solving, intellectual effort, or critical thinking.	Statements expressing concern about cognitive dependence, reduced use of personal knowledge, or allowing AI to perform the user's thinking.	General concerns about inaccurate outputs were coded under understanding AI limitations and errors rather than cognitive dependence.
Cognitive responses	Understanding AI limitations and errors	Recognition or direct experience of inaccurate, incomplete, misleading, fabricated, misquoted, or contextually inappropriate AI-generated outputs.	References to wrong answers, fabricated sources, mathematical limitations, misquotation, hallucination, or unreliable information.	Awareness of an error was coded here. The action taken to verify or correct the error was coded separately.
Cognitive responses	Inconsistent verification practices	Irregular, conditional, time-dependent, incomplete, or absent checking of AI-generated information.	Statements indicating that verification depends on academic importance, prior suspicion, available time, confidence, or workload.	Consistent descriptions of correction through re-prompting or model switching were coded as correction strategies, while checking accuracy was coded here.

Continued on next page

TABLE V. Qualitative Codebook Used for the Analysis of Student Interviews and University AI Policy Documents (continued)

Main Theme	Sub-theme	Operational Definition	Inclusion Criteria	Distinction and Decision Rule
Cognitive responses	Self-directed correction strategies	Informal strategies developed by students to improve, challenge, or replace inaccurate AI-generated responses.	Re-prompting, reformulating a question, requesting another answer, switching models, or comparing outputs across tools.	The code concerns attempts to change or regenerate an output. Independent verification against authoritative evidence was classified under verification practices.
Ethical responses	Academic integrity boundaries	Students' or policies' distinctions between acceptable assistance and unacceptable substitution of original academic work.	References to plagiarism, direct copying, originality, authorship, independent effort, acceptable support, or prohibited generation of assessed content.	Concerns centred on false detection were coded under uncertainty about AI detection rather than integrity boundaries.
Ethical responses	Data privacy and security concerns	Awareness of risks related to uploading personal, confidential, unpublished, institutional, or sensitive information into AI systems.	References to data retention, training use, confidentiality, identity theft, personal images, business information, or sensitive data.	General ethical concerns without a privacy or confidentiality component were not coded in this category.
Organisational adaptation	Absence of student voice	Perception that students have limited influence over the design, revision, interpretation, or implementation of university AI policies.	Student statements describing lack of consultation, compulsory acceptance of rules, limited opportunities for input, or limited feedback channels.	Lack of information about existing rules was coded as fragmented communication unless the emphasis was specifically on participation in policymaking.
Organisational adaptation	Fragmented institutional communication	Inconsistent, dispersed, unclear, course-dependent, or difficult-to-access communication regarding permissible AI use.	References to conflicting module rules, sporadic emails, assignment-specific instructions, Moodle-only information, or dependence on individual instructors.	The code concerns communication and accessibility of policy. Lack of curricular teaching was coded separately as insufficient curriculum integration.
Organisational adaptation	Insufficient curriculum integration	Perceived absence or inadequacy of structured, disciplinary, professional, or practical AI education within university programmes.	Calls for subject-specific AI teaching, formal skill development, curriculum integration, workplace preparation, or practical training.	General awareness campaigns or policy communication were not treated as curriculum integration unless they involved planned educational content or skill development.

Note: The codebook served as the primary decision framework for the analysis. Each student-experience code was applied only when the corresponding operational definition and inclusion criteria were satisfied. Closely related codes were distinguished using the decision rules shown in the final column. Experiential and behavioural constructs were not assigned to policy documents merely because a policy addressed a related institutional issue. Following the coding check, policy provisions were analysed separately according to direct substantive correspondence, indirect policy relevance, or absence of construct-equivalent evidence. A participant was counted once for a student sub-theme when the corresponding concept was identified at least once in that interview; multiple occurrences within the same interview did not increase its descriptive frequency. The operational definitions in this table primarily define student-experience constructs derived from the interview data. Experiential and behavioural sub-themes such as initial scepticism, growing confidence, routine AI integration, and efficiency-oriented use were therefore not attributed directly to policy documents. Policy texts were analysed separately for *direct substantive correspondence*, *indirect policy relevance*, or *no construct-equivalent policy evidence*. Accordingly, awareness, training, encouragement of AI use, assessment-redesign, governance, or communication provisions were not interpreted as evidence of students' affective states, behavioural patterns, or perceived implementation experiences.

The trustworthiness of the analysis was strengthened through several complementary strategies. First, findings were compared with two different data sources, namely student narratives and official policy documents. This comparison allowed for the identification of alignments, differences, and instances of policy silence. Second, the interpretation of each theme was supported by direct quotations and textual evidence to make the connection between the raw data and the researcher's inference visible. Third, deviant and less common cases were retained rather than restricting the analysis to dominant patterns. Fourth, the frequency of themes was used solely as a complement to the qualitative interpretation and did not replace the contextual and semantic analysis of the data.

In the present study, the descriptive frequency of a student-experience sub-theme was calculated as the number of interview

participants in whose account that sub-theme was identified at least once. Multiple repetitions of the same concept by one participant did not increase that participant's contribution beyond one count for the relevant sub-theme. Policy documents were not assigned parallel frequencies for experiential constructs. Instead, their relationship to each student-reported issue was classified qualitatively using the three cross-source categories defined above.

Given the exploratory qualitative design, analytical trustworthiness was assessed through the transparency of the coding framework, clearly specified operational definitions, systematic examination of supporting evidence, and consistency between source material and thematic interpretation. Although the qualitative analysis involved collaborative coding and interpretation, it did not follow a formal independent double-coding protocol. Consequently, no inter-coder agreement coefficient was calculated. The findings are interpreted as context-dependent qualitative evidence rather than statistically generalisable estimates.

The qualitative codebook presented in Table V identifies the conceptual boundaries between the final 14 student-experience sub-themes. Operational definitions and decision rules were particularly necessary to separate closely related concepts, such as error awareness, response validation, and output-modification strategy. This structure enabled student frequencies to be calculated using consistent criteria while policy excerpts were assessed separately for their substantive relationship to the corresponding student-reported issue.

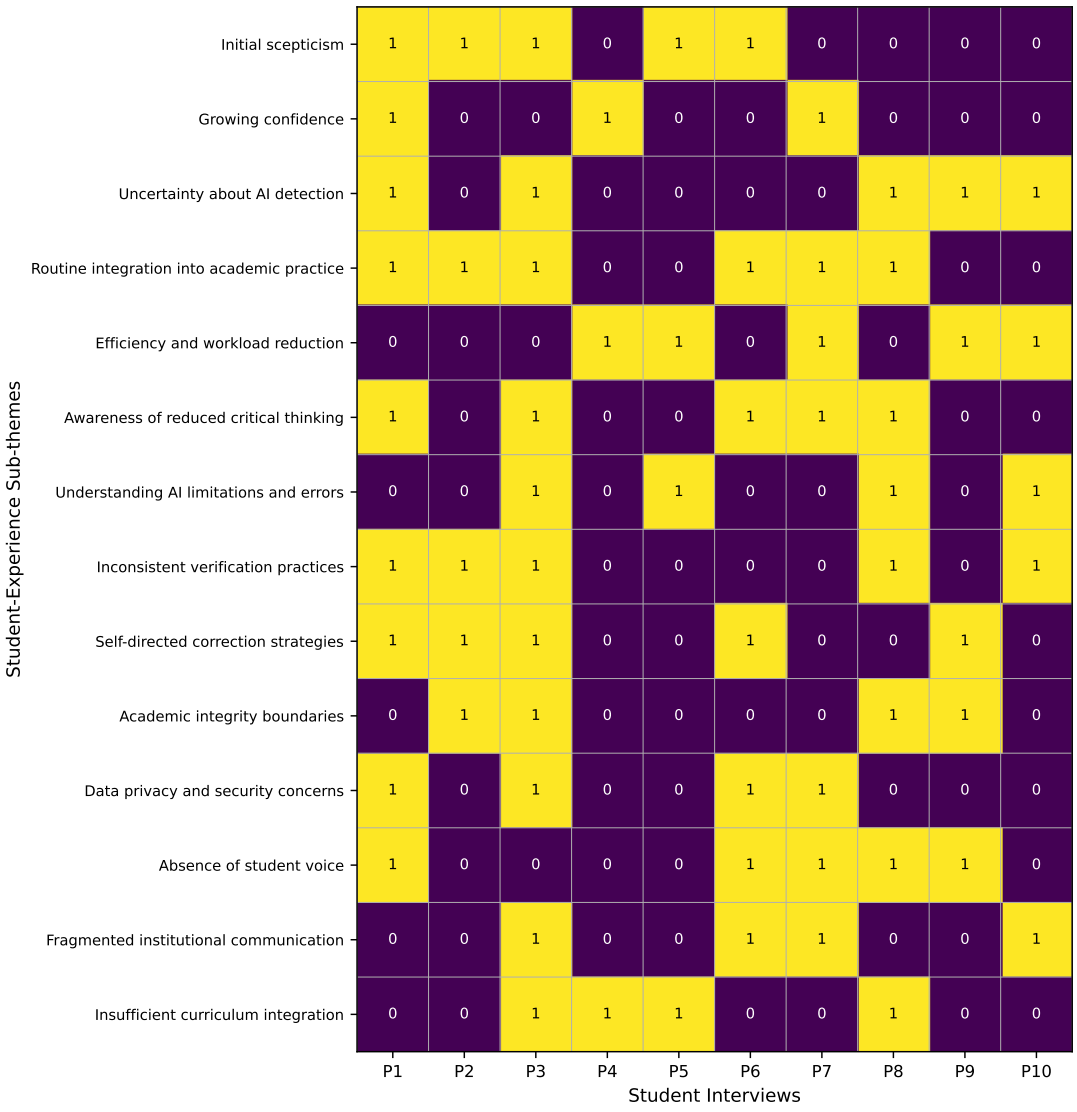

Fig. 2. Binary coding matrix showing the presence or absence of the 14 identified student-experience sub-themes across the 10 interviews. A value of one indicates that the sub-theme was identified at least once in the corresponding interview, whereas a value of zero indicates that it was not identified. Policy documents are excluded because the cross-source analysis does not treat institutional provisions as construct-equivalent occurrences of student-experience codes.

Figure 2 shows the distribution of the 14 student-experience sub-themes across the 10 interviews. The matrix pattern indicates that affective, behavioural, cognitive, ethical, and organisational themes were not uniformly distributed across participants. Rather, each narrative contained a different combination of experiences, concerns, and response strategies. Policy documents are not included in this binary matrix because institutional provisions are analysed as a separate source type and are compared with student-reported issues through qualitative relationship categories rather than construct-equivalent binary coding.

Participant-level coding assignments were cross-checked against the analytical profiles presented in Table III and the operational definitions and inclusion criteria specified in the qualitative codebook before descriptive frequencies were calculated. This consistency check ensured that the coding matrix and frequency summaries were derived using criteria consistent with those applied throughout the qualitative analysis.

3.4 Ethical Considerations and Analytical Trustworthiness

The research underwent a documented ethical screening procedure at Royal Holloway, University of London, using the Simplified Ethical Approval Form for module MN3311. The screening form, completed by the student researcher, Mohammed Alalwan, and signed by the supervisor, Elton Xhetani, on 2 February 2026, records a supervisor-level determination of "No referral necessary." This determination preceded the semi-structured student interviews, which were conducted between 12 and 19 April 2026.

The institutional screening form considered four aspects of the proposed research: whether the study involved covert data collection, the use of research data for purposes outside the study, participation by vulnerable populations, or potential distress or harm to participants. All four questions were answered negatively. Following consideration of the proposed research and its ethical implications, the supervisor selected the option indicating that no referral was necessary. This constitutes a documented supervisor-level ethical screening determination rather than formal approval issued by an institutional research ethics committee. Accordingly, no research ethics committee approval number is claimed. The original signed screening documentation is provided to the journal as supporting evidence.

Before participation, all interviewees were informed of the research objectives, the intended use of their responses, the voluntary nature of participation, confidentiality arrangements, and their right to withdraw. Written informed consent was obtained from all ten participants before the interviews were conducted. To protect participant confidentiality, interview material was reported using non-identifying participant codes (P1–P10). No participant names, contact details, or other direct personal identifiers are disclosed in the manuscript. The findings and supporting quotations were presented with consideration for participant confidentiality and the minimisation of identification risks. The three institutional policy documents were identified analytically as Policy 1–Policy 3, with their official titles and institutional contexts provided separately in Table IV.

Analytical trustworthiness was supported through a structured and traceable coding process rather than through statistical inter-coder agreement. The analysis began with repeated reading of the interview and policy material to establish familiarity with the context and to identify relevant semantic units. A qualitative codebook was then developed for the final sub-themes, with each code defined through an operational definition, explicit inclusion criteria, and a decision rule distinguishing it from conceptually adjacent codes. These rules were used consistently throughout the analysis and were revisited whenever a passage could plausibly be assigned to more than one category.

Coding consistency was further supported through iterative comparison between each assigned code and the underlying source material. After the initial coding stage, sub-theme assignments were rechecked against the original interview extracts and policy statements to ensure that the coded meaning remained consistent with the operational definition. Ambiguous passages were interpreted according to their dominant contextual meaning, while more than one code was retained only when a passage contained conceptually distinct meanings. This procedure was intended to reduce category overlap and minimise interpretive drift across the analysis.

Several additional procedures were used to strengthen transparency and credibility. First, representative interpretations were supported with direct student quotations and policy excerpts, enabling readers to evaluate the link between the source material and the analytical inference. Second, less common and divergent cases were retained rather than restricting the analysis to dominant patterns. Third, the descriptive frequency of each student sub-theme was calculated at the source level: a participant was counted once when the corresponding concept appeared at least once in that interview, irrespective of how many times it was subsequently repeated. This rule reduced the possibility that longer or more detailed interviews would disproportionately influence the descriptive frequencies.

Cross-source comparison provided an additional interpretive check, but it was not treated as independent validation of the interview findings. Policy assignments were checked against the codebook and supporting excerpts. Student-experience constructs and institutional policy provisions were kept analytically distinct, and policy relationships were classified according to direct substantive correspondence, indirect relevance, or absence of construct-equivalent evidence. This

additional audit was used to improve conceptual consistency across the two source types.

Analytical trustworthiness was supported through a collaborative coding process guided by a structured codebook with explicitly defined inclusion criteria and decision rules. The coding framework was applied to the interview material, and analytical consistency was examined by reviewing the correspondence between coded excerpts, operational definitions, and the resulting thematic interpretations. Particular attention was given to maintaining traceability between participant accounts, sub-theme assignments, and the evidence presented in the findings. As the analysis did not employ a formal independent double-coding protocol, inter-coder agreement statistics were not calculated. Instead, analytical rigour was pursued through systematic documentation of coding decisions, evidence-based interpretation, and transparent reporting of the qualitative findings.

4. FINDINGS

The comparative analysis identified several areas of convergence and divergence between the institutional orientations represented in the analysed AI policy documents and the experiences reported by the interviewed students. Because the two data sources were not analysed as individually matched institution–participant cases, these findings should be interpreted as thematic relationships between policy orientations and student experiences rather than as direct evidence of implementation outcomes within specific institutions.

Throughout this section, references to professional or career preparation denote participants' reported perceptions of preparedness for AI-influenced professional contexts. They should not be interpreted as measures of actual employment readiness, workplace performance, employer-assessed competence, or successful transfer of AI-related skills into professional practice. Likewise, the policy analysis identifies relationships between documented policy provisions and student-reported experiences; it does not evaluate institutional policy effectiveness. Across the policy documents, AI governance was commonly framed in terms of academic integrity, responsible use, risk management, accountability, and data protection. By contrast, student narratives placed greater emphasis on practical AI use, productivity, workload management, uncertainty regarding acceptable use, and skills perceived as relevant to future employment. Based on the hybrid thematic analysis, five main themes were identified: affective, behavioural, cognitive, and ethical responses, together with organisational adaptation. Taken together, these themes indicate that students' engagement with AI is multidimensional but uneven. Frequent practical use coexists with variation in affective responses, verification practices, ethical interpretation, and perceptions of institutional support. The findings therefore identify areas requiring more structured educational and organisational support rather than establishing a hierarchical ranking of the individual dimensions of AI literacy.

TABLE VI. Student-Reported Sub-themes and Their Relationship to University AI Policy Provisions

Main Theme	Student-Reported Sub-theme	Students, *n*	Students, %	Relevant Policy Provision(s)	Policy–Experience Relationship
Affective responses	Initial scepticism towards AI	5	50.0	Policies 1–3 contained awareness, guidance, training, or acceptable-use provisions intended to improve understanding of AI.	*Indirect policy relevance.* These provisions may respond to uncertainty or limited understanding, but they are not evidence that the policies themselves contain or demonstrate student scepticism.
Affective responses	Growing confidence in AI use	3	30.0	Policies recognised or encouraged educational engagement with AI and emerging technologies.	*Indirect policy relevance.* Institutional support for AI use does not demonstrate an increase in students' confidence; confidence was coded only when explicitly reported by participants.
Affective responses	Uncertainty about AI detection	5	50.0	Policies 2 and 3 explicitly addressed detection, investigation, or human oversight of suspected AI use.	*Direct substantive correspondence.* The policies explicitly address AI-detection governance, whereas the student code captures anxiety and uncertainty concerning its application.

Continued on next page

TABLE VI. Student-Reported Sub-themes and Their Relationship to University AI Policy Provisions (continued)

Main Theme	Student-Reported Sub-theme	Students, *n*	Students, %	Relevant Policy Provision(s)	Policy–Experience Relationship
Behavioural responses	Routine integration of AI into academic practice	6	60.0	Policies 1 and 2 contain provisions supporting AI literacy or the educational use of technological developments; other policy provisions regulate how AI may be used within teaching and assessment.	*Indirect policy relevance.* These provisions acknowledge or regulate AI use in education but do not constitute evidence that students routinely integrate AI into their academic activities.
Behavioural responses	Efficiency and workload reduction	5	50.0	Policy 2 includes assessment-redesign provisions intended to reduce the risks associated with AI-assisted take-home work; no policy provision explicitly promotes student workload reduction as a behavioural objective.	*No construct-equivalent policy evidence.* Assessment redesign addresses institutional consequences of AI use but does not demonstrate that students use AI to save time, reduce effort, or manage workload.
Cognitive responses	Awareness of reduced critical thinking	5	50.0	Policies included provisions concerning critical analysis, human oversight, and the risks of over-reliance on AI-generated content.	*Direct substantive correspondence.* Policy provisions address preservation of critical judgement, while the student code reflects participants' own awareness of possible cognitive dependence.
Cognitive responses	Understanding AI limitations and errors	4	40.0	The policies addressed inaccuracies, unverifiable information, limitations of AI outputs, or user responsibility for output accuracy.	*Direct substantive correspondence.* Both evidence streams concern limitations and errors, although student accounts reflect lived experience whereas policies state institutional expectations and cautions.
Cognitive responses	Inconsistent verification practices	5	50.0	Policy provisions assigned responsibility for checking or validating AI-generated content and, in some cases, required human oversight.	*Direct substantive correspondence.* Policy requirements explicitly address validation responsibility, whereas the student sub-theme captures irregular or conditional verification behaviour.
Cognitive responses	Self-directed correction strategies	5	50.0	No explicit operational provision was identified for re-prompting, model switching, or other student-developed correction strategies.	*No construct-equivalent policy evidence.* The strategies were reported by students but were not explicitly operationalised in the analysed policy excerpts.
Ethical responses	Academic integrity boundaries	4	40.0	All three policies addressed originality, permissible or prohibited AI use, authorship, or responsibility for submitted work.	*Direct substantive correspondence.* Both evidence streams address boundaries of acceptable academic use, although policies define institutional expectations and students describe how they interpret those boundaries.
Ethical responses	Data privacy and security concerns	4	40.0	The analysed policies addressed confidentiality, data protection, privacy, or risks associated with uploading sensitive information.	*Direct substantive correspondence.* Student concerns and institutional provisions concern the same substantive privacy and data-security risks.

Continued on next page

TABLE VI. Student-Reported Sub-themes and Their Relationship to University AI Policy Provisions (continued)

Main Theme	Student-Reported Sub-theme	Students, *n*	Students, %	Relevant Policy Provision(s)	Policy–Experience Relationship
Organisational adaptation	Absence of student voice	5	50.0	Policies described governance structures, review responsibilities, and, in one case, the importance of staff–student dialogue.	*No construct-equivalent policy evidence.* Formal governance structures do not demonstrate the presence or absence of meaningful student voice; the student code reflects participants' perceived limited involvement in policy development or revision.
Organisational adaptation	Fragmented institutional communication	4	40.0	Policies contained provisions calling for clear communication, dialogue, or course-level explanation of acceptable AI use.	*Direct substantive correspondence.* The policies explicitly address institutional communication as a requirement, whereas the student code captures participants' reported experience of uneven or difficult-to-access communication; this comparison does not establish implementation effectiveness.
Organisational adaptation	Insufficient curriculum integration	4	40.0	Policy 1 explicitly linked AI literacy with study and future-career skills, while other policy excerpts provided more general guidance on educational use.	*Indirect policy relevance.* General commitments to AI literacy and skills development are relevant to curriculum integration but are not equivalent to evidence of systematic, discipline-specific curricular implementation.

Note: Student frequencies describe the prevalence of interview-derived sub-themes in the purposive sample ($N = 10$). Policy statements are not counted as parallel occurrences of student-experience constructs. "Direct substantive correspondence" indicates that the policy explicitly addresses the same substantive issue; "indirect policy relevance" indicates a related institutional provision without construct equivalence; and "no construct-equivalent policy evidence" indicates that the available policy text does not instantiate the student-reported construct. These categories are qualitative and should not be interpreted as measures of policy effectiveness or implementation.

The cross-source comparison distinguishes the prevalence of student-reported experiences from the content of institutional policy. Student frequencies therefore describe only the interview sample, whereas policy analysis identifies whether the available provisions show direct substantive correspondence, indirect policy relevance, or no construct-equivalent evidence. This distinction is important because institutional intentions cannot be treated as observations of student affect, behaviour, or implementation experience. Across the five themes, several substantive issues were addressed in both sources, including AI-detection governance, critical evaluation of AI output, AI limitations and inaccuracies, validation responsibility, academic integrity, data privacy, and institutional communication. Other student experiences, including initial scepticism, growing confidence, routine integration, and perceived curriculum insufficiency, were only indirectly related to broader policy provisions. No construct-equivalent policy evidence was identified for students' efficiency and workload-reduction motivations, self-directed correction strategies, or perceived absence of student voice. These distinctions provide a more conservative account of the policy–experience relationship and avoid interpreting institutional intentions as evidence of student states, behaviours, or implementation outcomes. The cognitive evidence nevertheless identifies an important difference between institutional expectations and reported practice. Participants described awareness of AI errors, fabricated sources, potential reductions in critical thinking, and irregular validation practices, while relevant policy provisions emphasised responsibility, critical analysis, verification, or human oversight. No explicit operational guidance was identified in the analysed excerpts for the self-directed correction strategies reported by students. The comparison therefore concerns the relationship between policy expectations and reported student practices rather than a numerical difference between equivalent constructs.

4.1 Students' Affective Responses to AI

The affective findings indicate that students' relationships with AI developed through experience rather than following a uniform pattern. Five of the ten participants described initial scepticism, distrust, uncertainty, or limited understanding

when they first encountered AI tools. These responses were associated with unfamiliarity with the technology, uncertainty regarding the reliability of AI-generated outputs, and ambiguity concerning the acceptable place of AI within academic work. By contrast, three participants explicitly described increased confidence or greater comfort after gaining practical experience with AI. These findings indicate a transition for some students from initial uncertainty towards greater acceptance; however, such a transition was not reported by all participants and should not be inferred from frequency of use alone.

Importantly, the policy documents were not coded as containing either "initial scepticism" or "growing confidence," because these sub-themes represent affective states reported by students rather than properties of institutional policy documents. Instead, policy statements concerning awareness, training, acceptable AI use, or the educational potential of AI were interpreted as institutional provisions that may be conceptually relevant to students' reported affective experiences.

A more direct policy–experience relationship was observed for uncertainty about AI detection. Five participants (P1, P3, P8, P9, and P10) expressed concern about false identification, misclassification, or possible accusations of unauthorised AI use. These accounts indicate that students were not only concerned about the boundaries of acceptable AI use, but also about how AI-assisted work might be interpreted by institutional detection systems. The findings therefore highlight the importance of transparent AI-detection procedures, human oversight, and clear institutional guidance regarding the distinction between legitimate AI-supported academic work and academic misconduct.

TABLE VII. Student-Reported Affective Responses and Corresponding University AI Policy Provisions

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Initial scepticism towards AI	"Initially when I heard about it I was quite skeptical to be honest I didn't really understand what it was" – P1. "I was a bit late to use it because I remember at first it came out and I was sceptical" – P2. "At first I was very dismissive, I thought it was unreliable" – P3. "I was sceptical at first, and I remember when I first got introduced to it I didn't really think much of it" – P5. "At first I was a bit sceptical using it" – P6.	Policy 1 provides guidance and training intended to help students and staff understand generative AI tools. Policy 2 promotes awareness of responsible AI use. Policy 3 requires instructors to explain acceptable AI use at the course level.	Initial scepticism is an affective student-experience construct and was therefore coded only from interview data. The cited policy provisions are indirectly relevant because they seek to improve understanding and awareness, but they do not constitute evidence that the policies themselves contain or demonstrate scepticism.
Growing confidence in AI use	"Quite confident using it for any sort of task [because] I know that the tool can give me what I need so I'm quite happy to use it" – P1. "When I had my exam season in university I felt a lot more confident because I realised I could rely on AI to offer assistance" – P4. "I've got more comfortable with it and I feel like it really is an efficient tool that you can use in your day to day life" – P7.	Policy 1 recognises the potential of generative AI to enhance teaching and learning. Policy 2 promotes the effective use of technological developments in education and research. Policy 3 encourages engagement with emerging technologies to support learning and skill development.	Growing confidence was identified only when participants explicitly described an increase in comfort, trust, or perceived ability following experience with AI. Institutional encouragement of AI-supported learning is relevant to the broader environment of adoption but is not treated as evidence of increased student confidence.
Uncertainty about AI detection	"I've seen TikTok's of Turnitin and their AI detector that is flagging students incorrectly [...] I'm scared they could use this feature to try and claim AI use" – P1. "I'm only [using AI] to summarise readings and things like that so I don't want them to use this as evidence of AI use against me" – P3. "Sometimes even if you're doing stuff in your own words it could be detected as plagiarism so it could be a bit risky" – P8. "I'm worried that the advice and tweaks it gives me could be detected as AI even if it's me writing it" – P9. "When I'm doing coursework and I have to submit it on Turnitin, even though I never actually use AI to copy and paste answers, I still worry that the system might incorrectly flag my work as AI-generated and that I could be accused of academic misconduct despite completing the work myself." – P10.	Policy 2 gives instructors and supervisors responsibility and authority to examine suspected AI-generated work. Policy 3 cautions against using AI-detection software as the sole basis for determining originality and recommends human oversight.	The student accounts reveal explicit concerns about false identification, misclassification, and possible accusations of unauthorised AI use. These concerns were expressed by P1, P3, P8, P9, and P10, including apprehension that independently completed coursework could be incorrectly flagged as AI-generated. The corresponding policy documents address AI detection, academic integrity, and oversight procedures, indicating direct substantive relevance to the source of students' reported uncertainty. However, the existence of these provisions does not in itself demonstrate that students experience the institutional safeguards as clear, consistent, or sufficiently reassuring in practice.

Table VII distinguishes the student-reported affective constructs from the corresponding institutional provisions. This distinction is analytically important: policy commitments to awareness, training, or responsible AI use may represent institutional responses to emerging challenges, but they should not be interpreted as direct evidence of students' affective states. From a career-readiness perspective, these findings indicate that confidence in AI should be understood as informed and critical confidence rather than simple acceptance. Universities can support such confidence by clarifying permissible use, communicating the limitations of AI-detection tools, and establishing transparent procedures for reviewing suspected misuse. However, the present data support conclusions about students' reported experiences and relevant policy provisions rather than the effectiveness of those provisions in changing affective outcomes.

4.2 Behavioural Responses and Patterns of AI Use

The behavioural findings describe how participants incorporated AI into their academic activities and the motivations underlying that use. All participants reported at least some experience with AI tools. However, occasional exposure was not treated as equivalent to routine integration. Six of the ten participants explicitly described recurrent or frequent use of AI for studying, clarification, information seeking, summarisation, or other academic activities. These accounts indicate

that AI had become embedded in the academic practice of a substantial proportion of the interview sample, although the frequency and depth of use were not uniform across participants.

A distinct behavioural pattern concerned the motivation for using AI. Five participants explicitly associated their use with saving time, reducing effort, managing workload, obtaining information more quickly, or completing time-consuming tasks more efficiently. This sub-theme was analytically separated from routine integration because frequency of use and motivation for use represent different constructs. A participant could use AI regularly without identifying efficiency as the primary motivation, while another could use it selectively specifically to reduce the time or effort required for particular activities.

The policy documents were not coded as exhibiting either routine AI integration or efficiency-oriented behaviour. Statements promoting AI literacy or encouraging the educational use of technological developments were treated as institutional provisions that acknowledge or support AI engagement but do not establish how frequently students actually use AI. Likewise, provisions concerning assessment redesign were interpreted as institutional responses to the consequences of AI-assisted work rather than as evidence that students use AI to reduce workload.

The comparison nevertheless reveals an important difference in emphasis. Student accounts describe AI largely in terms of practical use, convenience, and immediate task-related benefits. By contrast, the policy provisions reviewed focus more strongly on responsible use, educational governance, academic integrity, and assessment design. This difference should not be interpreted as direct evidence that the policies are ineffective. Rather, it indicates that students' reported motivations and institutional policy objectives operate at different analytical levels.

Frequent use or efficiency-oriented use should also not be equated with professional competence or career readiness. The interview data demonstrate patterns of use and students' perceived benefits, but they do not directly assess workplace performance or employment outcomes. From an educational perspective, routine AI use may contribute to perceived preparedness when it is accompanied by the ability to select appropriate tools, evaluate output quality, recognise limitations, document use, and retain independent judgement. The findings therefore identify an opportunity for universities to connect students' existing patterns of AI use with more structured, discipline-specific learning activities.

TABLE VIII. Student-Reported Behavioural Patterns and Corresponding University AI Policy Provisions

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Routine integration of AI into academic practice	"I use it very often in my daily life just to ask random questions and a lot for my studies as well for clarification" – P2. "In terms of daily life not that much but in terms of studies, almost all of my work is in some form helped with AI" – P3. "I would say I pretty much use it every day for studying" – P6. "I always refer back to AI models like ChatGPT and Claude in my day to day life" – P7. "I would say I'd mostly use it for my studies a couple of times every week" – P8.	Policy 1 states that the institution will support AI literacy through learning and skills development for students and staff. Policy 2 promotes the effective use of technological developments in support of education and research.	Routine integration was coded only when students explicitly described recurrent or frequent AI use. The policy provisions acknowledge and support educational engagement with AI but do not provide evidence of the frequency or embeddedness of students' actual use. Their relationship to this behavioural theme is therefore indirect.
Efficiency and workload reduction	"I consider it a faster route to getting things done and minimising workload" – P4. "I use AI for personal uses if something is kind of time consuming" – P5. "It's a fast and convenient way for me to gather any information I want" – P7. "My own personal use of AI tends to be in situations where it's more so about saving time" – P9. "It saves me quite a bit of time" – P10.	Policy 2 recommends reducing the weighting of some take-home assignments and placing greater emphasis on face-to-face quizzes, examinations, and projects as part of the institutional response to AI-assisted assessment.	Efficiency and workload reduction represent students' motivations for using AI. The assessment provision does not represent the same behavioural construct. Rather, it reflects an institutional response to the implications of AI use for assessment design. No explicit policy provision was identified that framed student workload reduction as an intended objective of AI use.

Table VIII presents the student-reported behavioural patterns alongside the corresponding policy provisions while maintaining the distinction between individual behaviour and institutional intent. P1 was classified as a positive case for routine integration because Table III describes frequent AI use across a broad range of academic activities, satisfying the codebook criterion for recurrent or embedded academic use. The behavioural rows in Table VI follow the same source-specific coding rule. Accordingly, policy statements encouraging AI literacy or modifying assessment practices are interpreted as institutional provisions or responses and not as evidence of routine integration or efficiency-oriented student behaviour.

4.3 Cognitive Responses and the Issue of Validating AI Outputs

The cognitive findings reveal a notable distinction between awareness of AI-related risks and the systematic application of practices for evaluating and validating AI-generated outputs. Five of the ten participants expressed concern that extensive or habitual reliance on AI could reduce independent thinking, intellectual effort, or critical reasoning. These accounts

indicate awareness of potential cognitive dependence. However, they represent participants' reported concerns rather than objectively measured cognitive decline. Such awareness also did not consistently translate into structured practices for evaluating, verifying, or challenging AI-generated information.

A similar pattern was observed in participants' understanding of AI limitations. Four participants explicitly described inaccurate responses, weaknesses in complex tasks, misquotation, or fabricated sources. These accounts demonstrate recognition of important limitations in AI-generated content, but recognition of an error is analytically distinct from the ability to validate an output systematically. Five participants reported verification practices that were conditional, time-dependent, or otherwise inconsistent. Some checked outputs primarily when an answer appeared suspicious, whereas others indicated that verification could be omitted under time pressure. The central cognitive issue identified in the data is therefore not an absence of awareness, but a gap between recognising potential problems and applying consistent, methodologically robust validation procedures.

The most consequential issue within the cognitive findings concerned the validation stage. Although relevant policy provisions emphasised user responsibility for the accuracy of AI-generated content, the interview data indicated that students did not always apply consistent or structured procedures for checking that accuracy. Responsibility for validation should therefore be distinguished from demonstrated validation capability: requiring students to verify AI-generated content does not, by itself, establish that they have been provided with systematic methods for doing so.

Strategies for correcting incorrect outputs were also largely self-taught and informal. Five participants mentioned methods such as rewriting the prompt, asking for the answer again, comparing across multiple AI tools, or changing the model. These behaviours represent a practical adaptation to the limitations of AI, but are not sufficient from an educational perspective, as these strategies rely more on trial and error and are not necessarily designed around valid frameworks for evaluating information, source validity, or critical thinking. This finding suggests that one of the most important educational needs for universities is to teach rigorous methods for evaluating and validating AI outputs. Table IX reports the cognitive responses to AI, including students' concerns about reduced critical thinking, recognition of AI limitations, inconsistent verification practices, and self-directed correction strategies.

TABLE IX. Student-Reported Cognitive Responses and Corresponding University AI Policy Provisions

Student-Reported Sub-theme	**Representative Student Evidence**	**Corresponding Policy Provision(s)**	**Analytical Interpretation**
Awareness of reduced critical thinking	"That's definitely something I worry about I'm hoping not to get too dependent on it because it's not good" (P1). "I feel like in recent times I always need it to guide me or to give me steps or to structure my papers" (P3). "I was a bit sceptical using it because of the fact that I felt that I would use less of my own personal brain power, or my own actual knowledge" (P6). "My ability to think critically has actually decreased over the past years" (P7). "You're going to let it do your thinking for you, which is probably quite dangerous in the long term" (P8).	Policy 1 requires critical analysis and oversight of generative-AI outputs. Policy 3 cautions that AI systems cannot substitute for the nuanced human judgement involved in critical thinking, creativity, and individual learning needs.	The student accounts of P1, P3, P6, P7, and P8 describe perceived risks of cognitive dependence and reduced independent reasoning. The corresponding policy provisions address critical oversight and human judgement but do not demonstrate whether these capacities are maintained in practice.
Understanding AI limitations and errors	"AI is almost always pretty bad at complex mathematical questions" (P3). "I remember when I was searching up something for studies and the answer came out wrong. It made me kind of think what other stuff have I been getting wrong" (P5). "It had times where things were misquoted a little bit" (P8). "Sometimes ChatGPT creates fake sources and references so you have to be careful" (P10).	Policy 1 assigns users responsibility for the accuracy of AI-generated information. Policy 2 addresses inaccurate, unverifiable, or non-existent references. Policy 3 acknowledges the possibility of AI-generated inaccuracies.	Student experiences of inaccurate answers and fabricated references correspond directly to policy provisions addressing output reliability. Such correspondence does not establish that students have been taught systematic verification procedures.

Continued on next page

Table IX – *Continued from previous page*

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Inconsistent verification practices	"For everyday things [...] I typically won't verify it, but if it's to do with my studies [...] I try to verify that with my own knowledge" (P1). "I guess I don't typically verify the information. I kind of trust it a lot of the time" (P2). "No, quite naive of me I usually don't verify the answer. Unless I know that the answer is wrong" (P3). "Maybe if I'm in a rush or I don't really have time to cross-check everything that would be a situation I wouldn't [verify]" (P8). "Most of the time, probably 80% of the time [I verify], if I'm behind on work in general, I might just take it as the answer" (P10).	Policy 2 makes students responsible for work submitted under their names. Policy 3 encourages users to recognise AI limitations and verify accuracy before academic or research use.	The policies address responsibility for accurate use, whereas students report conditional, inconsistent checking. The existence of verification guidance does not establish consistent implementation in student practice.
Self-directed correction strategies	"I tend to get angry at the model and say [...] this is incorrect" (P1). "I use multiple different models and compare the answers" (P2). "I would prompt it and say [...] please try again" (P3). "I would switch to a different AI tool and guide that one" (P6). "I'll rephrase what I'm trying to ask it a few times" (P9).	Policies 1–3 provide guidance on responsible use, output accuracy, and verification, but the cited provisions do not prescribe these particular student-generated prompting or model-switching practices.	These accounts document informal output-correction strategies. General institutional guidance is indirectly relevant, but should not be interpreted as direct evidence that these specific strategies are taught or required by university policy.

Students' cognitive responses indicated that their awareness of the limitations and risks of AI was relatively well-developed, but this awareness had not yet translated into a coherent set of assessment and validation skills. Five out of 10 participants expressed concerns about the potential for reduced critical thinking, overreliance on automated responses, or diminished independent judgement. This concern suggests that some students distinguish between supporting the use of AI and handing over the thinking process to the tool, but lack structured training to practically manage this boundary.

This finding highlights the importance of distinguishing between "using AI" and "cognitively responsible use of AI." Merely using a tool repeatedly does not indicate the ability to assess response quality, identify bias, detect fabricated information, or understand model limitations. In relation to perceived professional preparation, participants' accounts suggest that responsible AI use involves maintaining independent judgement, critical analysis, and ultimate human responsibility rather than delegating these functions to the tool.

Four participants directly mentioned the errors and limitations of AI. These experiences included incorrect answers, weakness in complex problems, incorrect citations, and the generation of fictitious sources or references. Policy documents also emphasised at various levels the user's responsibility for the accuracy of information and the need to check the outputs. However, alignment in risk recognition did not translate into alignment in the ability to manage it. Policies mostly stated that the student was responsible for the accuracy of the content provided, but did not teach practical validation methods in a clear and step-by-step manner.

A recurring issue within the cognitive findings concerned the consistency of verification practices. Five participants stated that checking the accuracy of AI answers depended on the importance of the topic, the level of prior familiarity, suspicion that an output might be incorrect, time pressure, or workload. Some checked an answer only when they suspected an error, whereas others reported omitting verification when time was limited. This pattern suggests that validation had not yet become a stable and systematic practice among all participants.

Inconsistency in validation is particularly important from a professional perspective. In the workplace, incorrect AI output can affect decision-making, service quality, information confidentiality, professional credibility, and organisational liability. Therefore, relying on intuition, personal suspicion, or time to check the accuracy of an answer is not enough. Students must be able to recognise the type of claim, select the appropriate source to check it, compare independent evidence, and adjust their level of confidence based on the quality of the data.

Five participants also mentioned strategies for correcting inappropriate answers. These strategies included rewriting the request, requesting a second answer, changing the model, comparing multiple tools, or guiding the system more precisely. These behaviours represent a form of practical adaptation and self-directed learning, but should not be confused with systematic validation. Producing a different answer or obtaining agreement between multiple models does not necessarily establish the accuracy of the information because different tools may rely on similar patterns, overlapping data sources, or convergent errors.

Therefore, students' remediation strategies can be considered useful but insufficient. Redesigning the prompt may improve

the quality, clarity, or appropriateness of the response, but it is not a substitute for checking authoritative sources, comparing with independent evidence, and making expert judgments. Universities should explicitly teach the difference between output remediation and output validation. In such training, students should understand that remediation refers to improving the way the response is produced, while validation refers to assessing the accuracy, adequacy, source, and trustworthiness of the response.

Based on these findings, the cognitive dimension of AI literacy should be included in the curriculum as a set of teachable and assessable skills. This training could include recognizing signs of fabricated information, checking primary sources, comparing output with authoritative scientific literature, assessing model limitations, documenting validation steps, and determining the appropriate level of human supervision. Such an approach can transform students' experimental and trial-and-error use into a professional, systematic, and defensible process.

The cognitive findings should not be interpreted as indicating that the cognitive dimension is weaker than the other dimensions of AI literacy. Rather, they identify a specific within-theme pattern: participants demonstrated meaningful awareness of AI limitations and possible cognitive risks, while systematic validation practices were less consistently reported. This distinction between awareness and operational validation capability provides the empirical basis for recommending more explicit instruction in source checking, comparison with independent evidence, error detection, and appropriate human oversight.

The student frequencies and qualitative relationships between student-reported sub-themes and the analysed policy provisions are summarised in Table VI and Figure 3. Direct participant evidence concerning cognitive dependence, AI limitations and errors, verification practices, and self-directed correction strategies is presented in Table IX. The findings reflect participants' reported experiences rather than measured changes in critical-thinking performance.

4.4 Ethical Responses and Boundaries of Responsible AI Use

The fourth theme relates to the ethical dimension of AI literacy. The findings indicate that students are generally aware of the ethical boundaries of AI use. Four out of ten participants clearly distinguished between using AI to assist and leaving the production of scholarly content to AI. In their view, using AI to structure, explain, summarise, or support can be acceptable, but directly copying the output or presenting AI-generated text as one's own work constitutes unethical behaviour or plagiarism. This finding is consistent with university policies emphasizing scholarly originality, individual effort, and student responsibility.

However, general awareness does not mean that the boundaries of permissible use are completely clear. Many students still face ambiguity about exactly what level of AI assistance is acceptable and when it constitutes academic misconduct. The main problem, therefore, is not a lack of ethical awareness, but a lack of clear, practical, and situational guidelines. This is particularly relevant to students' perceived preparation for future professional contexts, where AI use may need to be balanced with confidentiality, intellectual property, originality, and professional accountability.

In the area of data security and privacy, four respondents also raised concerns about data storage, the use of input data to train models, the theft of personal information, or the misuse of sensitive images and data. University policy documents also explicitly warn against the inclusion of confidential, personal, or unpublished information in public AI tools. Thus, there is some alignment between policy and student experience on this front. However, this alignment has remained more of a warning than a practical training in data management, sensitive information recognition, or the safe use of AI tools. Table X illustrates the ethical dimension of AI use, particularly students' understanding of academic integrity boundaries and their concerns regarding data privacy and security.

TABLE X. Student-Reported Ethical Responses and Corresponding University AI Policy Provisions

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Academic integrity boundaries	"I don't want to have AI generated work. I want to have my own original authentic work" (P2). "If you just copy whatever it says that's bad because its plagiarism, but other than that I think it's a very useful tool to use" (P3). "If you're trying to use it to write a whole paragraph or trying to copy it then that's what I would call immoral" (P8). "I'd tend to draw the line if I claim the work that AI is generating as my own. If it's something that I wrote and AI helped me structuralise [...] I feel like that's ok because it's the same as having someone else check your work" (P9).	Policy 1 identifies plagiarism and copyright concerns associated with generated content. Policy 2 treats AI as a support tool rather than a substitute for scholarly reasoning. Policy 3 emphasises original thought, critical analysis, and personal academic effort.	Student accounts distinguish acceptable assistance from replacement of personal authorship. The policies directly address academic integrity boundaries, although formal rules alone do not demonstrate consistent student interpretation or compliance.

Continued on next page

Table X – *Continued from previous page*

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Data privacy and security concerns	"I try not to put my personal information or sensitive stuff because I'm scared that it could steal my data" (P1). "I think it's important to understand that the AI actually retains all of the information that you upload onto it" (P3). "I'm aware of the fact that some of the data is stored and is used for training purposes" (P6). "I don't know what they would do with the information or the images. I see so much news about certain platforms using your face and it gets stolen for identity theft" (P7).	Policy 1 discourages entry of personal and confidential business information into AI tools. Policy 2 prohibits uploading confidential, unpublished, or sensitive data into publicly accessible AI tools. Policy 3 warns of potential exposure of content entered into non-enterprise AI services.	Both sources recognise risks involving sensitive information and external AI tools. This is direct substantive correspondence in the issues addressed, but it does not establish that students possess operational data-protection skills or consistently follow these precautions.

The students' ethical responses showed that they generally distinguish between supporting the use of AI and completely leaving the production of academic content to this tool. Four out of 10 participants explicitly stated that using AI to explain concepts, summarise, organize text, or improve the structure of work can be acceptable. But they considered it unethical to directly copy the output or present the produced content as independent student work. This finding suggests that there is a basic awareness of scholarly originality and the responsibility of authorship among a subset of participants.

However, the existence of a general awareness does not mean that the boundaries of permissible use are clear in all situations. Students may experience ambiguity in distinguishing between linguistic editing, restructuring, idea generation, summarizing, rewriting, and original content production. This ambiguity is increased when each course, instructor, or assignment has different rules about the use of AI. The main issue, therefore, is not simply the moral acceptance or rejection of AI, but the lack of situation-specific and enforceable criteria for determining the permissible level of assistance from this tool.

The policy documents also emphasised academic integrity, individual responsibility, and the need to preserve the intellectual contribution of the student. All three documents did not see the use of AI as a substitute for reasoning, scholarly effort, or independent writing. In this sense, a certain alignment was observed between the policy direction and the general understanding of students. However, the policies mainly stated general principles and, in many cases, did not provide clear and staged examples of permissible, conditional, and impermissible uses.

This distinction between moral principle and practical guidance is important. Phrases such as "responsible use," "maintaining originality," or "avoiding substitution of scholarly effort" are useful for setting a general direction, but are not sufficient for making decisions in real situations. Students need to know when it is permissible to use AI for ideation, language editing, text structuring, data analysis, code generation, summarisation, or resource preparation, and at what point does this assistance become an undue interference in the scientific production process.

The second issue in the ethical dimension was data protection and privacy. Four participants expressed concerns about data being stored, using input data to train models, misusing personal images, identity theft, or exposing sensitive information. These concerns indicate that some students are aware of the risks associated with entering personal or confidential data into public AI tools.

On this axis, greater alignment was observed between students and policies. All three documents warned against entering personal, confidential, unpublished, or organisational data into public tools. However, this alignment remained more at the warning level and did not necessarily lead to practical training on identifying sensitive data, choosing safe tools, anonymizing data, managing consent, or using organisational systems.

From the perspective of perceived preparation for AI-influenced professional contexts, this limitation is particularly important. In professional settings, work with AI may involve customer information, research data, organisational documents, intellectual property, health information, or confidential business content. Therefore, general awareness of the risk is not enough, and students must be able to recognise the level of data sensitivity, assess the risk of disclosure, and review organisational and contractual constraints before using AI tools.

The ethical findings of this study suggest that universities should move away from purely warning policies toward scenario-based training. Providing real-world examples of permitted, conditional, and prohibited use, teaching how to declare the use of AI, explaining authorship responsibility, recognizing confidential data, and practicing ethical decision-making can help transform students' general awareness into actionable, professional competence.

The student frequencies and their corresponding relationships to the analysed policy provisions are summarised in Table VI and Figure 3. Detailed evidence concerning academic-integrity boundaries and data-privacy concerns is reported in Table X.

4.5 Organisational/Institutional Adaptation and the Policy–Experience Gap

The fifth theme concerns students' perceptions of organisational adaptation to AI. Five of the ten participants reported limited student voice, limited participation in decision-making, or dependence on decisions made by faculty members and university administration. The policy documents identified formal governance responsibilities, councils, committees, or review mechanisms; however, the existence of these structures does not by itself demonstrate either the presence or the absence of meaningful student participation. The comparison therefore distinguishes students' perceived involvement from the formal governance arrangements described in policy.

Four participants reported fragmented or inconsistent institutional communication about permissible AI use, describing reliance on assignment briefs, Moodle, sporadic emails, or individual course instructors. Relevant policy provisions called for clear communication, staff–student dialogue, or course-level explanation of acceptable AI use. These provisions show direct substantive correspondence with the issue of communication, but the present data do not establish whether the stated communication requirements were implemented effectively. The comparison therefore identifies a difference between policy intent and reported experience without treating student accounts as a direct evaluation of policy effectiveness.

Finally, four participants perceived AI integration in the curriculum as insufficient and called for more structured, professional, and discipline-specific AI education. Some policy provisions promoted AI literacy, skills development, or preparation for future professional use, but these general commitments are not equivalent to evidence of systematic curricular implementation. The relationship is therefore interpreted as indirect policy relevance rather than as a direct measure of curricular adequacy.

The results suggest that students' reported preparation for AI-influenced professional contexts is uneven across behavioural, cognitive, ethical, and institutional dimensions. Students use AI, but often without structured training; they are aware of important risks, but their validation methods remain inconsistent; they recognise general ethical boundaries, but face ambiguity in practical situations; and some participants perceived limited opportunities to contribute to AI-related policy development or review. These organisational findings are interpreted as student-reported experiences rather than as direct evidence about the composition, operation, or effectiveness of institutional governance structures. Accordingly, universities may benefit from complementing regulatory approaches with educational, participatory, and skills-based support that responds to the preparation needs reported by students. Table XI presents the organisational dimension of AI adaptation, showing students' perceptions of limited voice, fragmented communication, and insufficient curriculum integration.

TABLE XI. Student-Reported Organisational Adaptation and Corresponding University AI Policy Provisions

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Absence of student voice	"Whatever the university says you have to go along with that" (P1). "There's not much of a student input to the AI rules the university uses [...] it is very much based on what the lecturers and what the university administration sees would be good for the students" (P6). "I don't think there is a meeting with students about whether we can or can't use it. I feel like you have to accept what you're told" (P7). "They don't necessarily reach out to us and ask for feedback" (P8). "I feel like it's mostly up to the professor's discretion whether we're allowed to use AI and what capacity we're encouraged to do that" (P9).	Policy 1 assigns institutional accountability to a governing board. Policy 2 provides for council oversight of policy compliance. Policy 3 specifies annual review by an AI policy committee.	These provisions describe governance and review structures, but do not establish whether students have meaningful input into policymaking. The student-reported experience of limited voice has no demonstrated construct-equivalent policy evidence in these cited governance provisions.
Fragmented institutional communication	"More of my assignment briefs contained policies on AI, but there were no open announcements on AI in general" (P3). "We'd only get emails about the fact that you're not supposed to use AI for certain assignments or if we are able to use AI for certain assignments" (P6). "I haven't really been communicated about it unless it's from a module teacher and that's not really official communication" (P7). "It's not directly notified; you have to go through Moodle to find it" (P10).	Policy 1 calls for dialogue between staff and students. Policy 2 calls for clear communication on whether and how AI tools may be used. Policy 3 asks faculty to explain acceptable AI use in course syllabi.	The policy provisions address communication and course-level clarity directly. Student accounts nevertheless describe inconsistent or difficult-to-find guidance. This comparison does not independently assess implementation at participants' institutions.

Continued on next page

Table XI – *Student-Reported Organisational Adaptation and Corresponding University AI Policy Provisions*

Student-Reported Sub-theme	Representative Student Evidence	Corresponding Policy Provision(s)	Analytical Interpretation
Insufficient curriculum integration	"None of my modules have any sort of AI components into them, which I think is pretty inadequate for a business economics world" (P3). "They're just trying to teach you the principles of business management. I think that is a flaw on the university's part. They are not looking to consider new ways of educating people about how to enter the workforce" (P4). "The university should be teaching you how to use it for specific subjects because AI is a global thing everyone needs to learn how to use" (P5). "Universities could try and integrate AI more into the syllabus. I feel like they are still a little bit behind" (P8).	Policy 1 identifies AI literacy and future-oriented skills development as institutional objectives.	This policy commitment is indirectly relevant to students' calls for formal, discipline-specific AI education. It does not by itself demonstrate that AI skills have been systematically integrated into the curriculum.

The organisational findings concern participants' reported experiences of participation, communication, and curriculum integration. Five of the ten participants described limited opportunities to contribute to the design, review, interpretation, or refinement of AI-related rules. These statements constitute the empirical basis for the sub-theme *absence of student voice*; they should not be interpreted as direct evidence about the formal composition or operation of institutional governance bodies.

The analysed policy documents identify boards, councils, committees, oversight responsibilities, and review mechanisms. However, the presence of such structures does not, by itself, establish whether students are included in policy formulation or revision, nor whether any participation is substantively meaningful. The policy evidence was therefore not coded as demonstrating either the presence or the absence of student voice. Instead, the cross-source relationship for this sub-theme is classified as *no construct-equivalent policy evidence*.

Four participants reported that communication concerning AI use was fragmented, course-dependent, sporadic, or difficult to locate. Their accounts referred to assignment-specific instructions, Moodle-based information, emails, and individual instructor guidance. These reports constitute the empirical basis for the student-experience code *fragmented institutional communication*.

By contrast, the analysed policy documents contain provisions calling for clear communication, dialogue, or explicit course-level explanation of acceptable AI use. These provisions demonstrate that communication is recognised as an institutional policy concern. They do not demonstrate that communication was either clear or fragmented in implementation. The cross-source relationship is therefore one of *direct substantive correspondence* with the issue of communication rather than construct equivalence or evidence of implementation effectiveness.

Four participants also perceived AI integration in the curriculum as insufficient and called for more structured, professional, and discipline-specific AI education. Some policy provisions promote AI literacy, skills development, or preparation for future professional use; however, these general commitments do not establish systematic curriculum implementation. The relationship is therefore interpreted as *indirect policy relevance* rather than as evidence of curricular adequacy or inadequacy.

Taken together, these organisational findings identify student-reported governance concerns and areas of policy relevance without converting formal policy structures or stated institutional intentions into evidence of students' actual participation or communication experience. This distinction is maintained throughout the cross-source analysis and the subsequent discussion.

From an organisational change perspective, the findings support the need to move towards an adaptive and participatory approach. AI policies should not be viewed as static, final documents, but should be reviewed at regular intervals, based on changing technology, faculty experience, student feedback, and labour market needs. Establishing formal channels for receiving feedback, having student representatives on relevant committees, and clearly disseminating the reasons for policy changes can increase acceptance and enforceability.

Institutional communication should also shift from a fragmented model to a multi-level, coordinated framework. In this framework, the university could set general principles at the institutional level, faculties could explain course requirements, and each instructor could specify the permissible use of AI at the course and assignment level. Coordination across these levels can avoid inconsistencies, ambiguity, and excessive reliance on instructor interpretation.

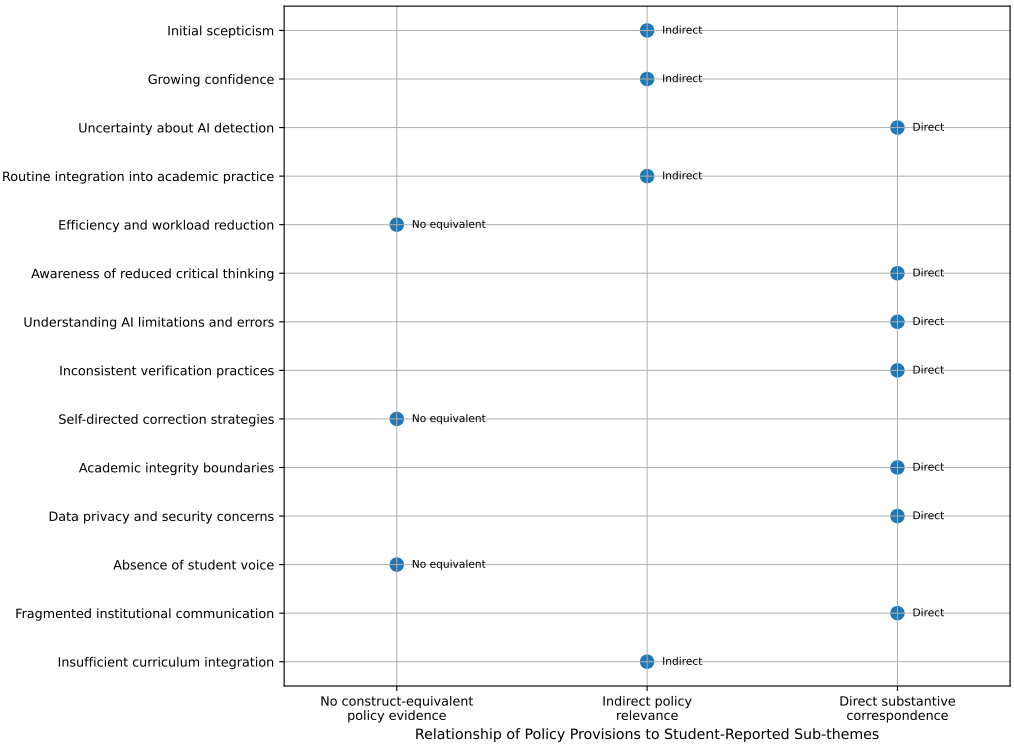

Fig. 3. Policy–experience relationship across the 14 student-reported sub-themes. Each sub-theme is explicitly labelled and assigned to one of three qualitative relationship categories: direct substantive correspondence, indirect policy relevance, or no construct-equivalent policy evidence. The categories describe conceptual relationships between the interview and policy data and should not be interpreted as quantitative measures of policy effectiveness or implementation quality.

Finally, integrating AI into the curriculum must go beyond general training in tools. Students need activities in which they practice tool selection, application design, output validation, documentation of usage, data protection, and professional responsibility. Such training should be tailored to the discipline, as the types of risks, applications, and acceptable levels of use of AI vary across disciplines.

The institutional findings suggest that readiness for AI is not only dependent on individual student knowledge and behaviour, but is also influenced by the quality of university governance, transparency of communication, the level of stakeholder engagement, and how skills are integrated into the curriculum. Accordingly, supporting students' perceived preparation for AI-influenced professional contexts may require coordinated attention at the student, course, curriculum, and institutional policy levels.

The student frequencies and their corresponding relationships to the analysed policy provisions are summarised in Table VI and Figure 3. Detailed evidence concerning perceived limited student voice, fragmented institutional communication, and insufficient curriculum integration is reported in Table XI.

Figure 3 summarises the relationship between the 14 student-reported sub-themes and the corresponding policy provisions. It does not compare student frequencies with policy percentages because the two sources capture different analytical constructs. Instead, it distinguishes direct substantive correspondence, indirect policy relevance, and the absence of construct-equivalent policy evidence.

Direct substantive correspondence was observed for issues explicitly addressed in both sources, including AI-detection governance, critical evaluation of AI output, AI limitations and inaccuracies, verification responsibility, academic integrity, data privacy, and institutional communication. Other student experiences, such as initial scepticism, growing confidence, routine use, and curriculum-related concerns, were only indirectly related to broader policy provisions. In contrast, no construct-equivalent policy evidence was identified for students' efficiency and workload-reduction motivations, self-directed correction strategies, or perceived absence of student voice. These distinctions provide a more conservative representation of policy–experience relationships and avoid interpreting institutional intentions as evidence of student states, behaviours, or implementation outcomes.

5. DISCUSSION: TOWARD AN AI-READINESS GOVERNANCE FRAMEWORK

The findings suggest a broader divergence between the institutional priorities represented in the analysed AI policies and the experiences and expectations reported by students. The analysed policies commonly frame AI in terms of academic integrity, risk of misuse, data privacy, individual responsibility, and responsible use. In contrast, the student narratives more often describe AI as a tool for accelerating learning, reducing workload, accessing information efficiently, and preparing for AI-influenced professional contexts. Because the policy documents were not individually matched to the interviewed participants, this divergence should not be interpreted as evidence that a particular institution failed to implement its own policy. Rather, it identifies recurring tensions between common policy orientations and student-reported needs that warrant further institution-specific investigation.

The comparison establishes an important distinction between policy content and implementation outcomes. The analysed documents address output validation, academic integrity, privacy, critical thinking, and transparent communication, but the presence of these provisions cannot by itself demonstrate how consistently they are implemented, experienced, or translated into educational practice. Conversely, student-reported uncertainty regarding validation or acceptable use cannot be attributed to failure of a particular analysed policy because the policy documents and participants were not matched institutionally. The evidential contribution of the study is therefore to identify where formal policy provisions correspond directly, indirectly, or not at all with the issues raised in student accounts, while leaving implementation quality and causal effectiveness for institution-specific evaluation.

From an organisational change theory perspective, the findings illustrate a tension between formal policy structures and the rapidly changing conditions surrounding generative-AI use. The analysed documents describe formal governance, oversight, and review arrangements, while the student interviews reveal perceptions of changing tools, ambiguous rules, and uneven access to practical guidance. The present study does not establish whether the institutions themselves operate predominantly through a top-down or participatory governance model. Rather, it suggests that AI governance may benefit from iterative, flexible, and feedback-based processes capable of responding to emerging circumstances.

Participants' perceptions of limited involvement in policy development and review highlight the potential value of more visible and systematic student-feedback mechanisms. Because the policy documents do not establish the actual extent or quality of student participation, this finding is interpreted as a student-reported governance concern rather than as a verified institutional characteristic. Student participation need not transfer policymaking authority to students; instead, student experience can be incorporated alongside the perspectives of academic staff, administrators, technology specialists, and relevant professional stakeholders when policies are reviewed.

Viewed through the ABCE framework, the student accounts illustrate distinct, context-dependent patterns within the affective, behavioural, cognitive, and ethical dimensions of AI literacy. Within the affective dimension, some participants described increasing confidence in using AI, while others expressed concerns about inaccurate AI detection and potential institutional consequences. Within the behavioural dimension, participants reported incorporating AI tools into academic activities, particularly for efficiency, information seeking, and workload management. Within the cognitive dimension, participants recognised the possibility of inaccurate or fabricated outputs and the risks of over-reliance, although their verification practices were often conditional and self-directed. Within the ethical dimension, participants raised concerns about academic integrity, data privacy, and uncertainty regarding the boundaries of permissible AI use. These findings describe the characteristics of student-reported experiences within each dimension, without implying a comparative ranking of the dimensions or treating their respective sub-theme frequencies as measurements of a common underlying scale. Collectively, the findings identify areas in which participants perceived a need for more structured institutional support, without establishing objectively demonstrated AI competence.

Taken together, these within-theme findings highlight the importance of distinguishing frequent AI use from comprehensive AI literacy and perceived career readiness. Regular engagement with AI tools does not, by itself, demonstrate that students can select appropriate technologies, critically evaluate output accuracy, recognise potential bias, protect sensitive data, or exercise independent professional judgement. The conceptual literature emphasises the complementary roles of practical capability, critical evaluation, ethical responsibility, and transferable skills in preparation for AI-influenced professional contexts. The present findings identify student-reported experiences and perceived educational support needs across these areas, but do not establish comparative strengths or weaknesses among the ABCE dimensions, nor do they provide objective evidence of workplace competence or employment readiness.

Based on these findings, the present study proposes an evidence-informed governance framework for strengthening institutional support for AI literacy and perceived professional preparation in higher education. The framework was derived directly from the recurring policy–experience relationships identified across the five analytical themes rather than introduced independently of the empirical analysis. The first pillar, course-level policy clarity, responds to reported uncertainty regarding permissible AI use and uneven access to institutional guidance. The second pillar, cognitive validation capability, reflects participants' reported experiences of AI inaccuracies, inconsistent verification practices, and concerns about

reduced independent reasoning. The third pillar, ethical and data-security practice, synthesises recurring issues concerning academic integrity, privacy, authorship, and user accountability. The fourth pillar, participatory and adaptive governance, arises from student-reported concerns regarding limited opportunities for input, fragmented communication, and the need for policies to remain responsive to changing AI practices. Accordingly, the framework represents a structured synthesis of the principal educational and governance needs identified across the interview and policy data.

The first component is to clarify the boundaries of permitted, conditional, and prohibited use at the level of each course and each type of assessment activity. This emphasis is consistent with recent higher-education policy research calling for explicit, context-sensitive guidance at institutional and course levels [7, 9]. Students should know under what circumstances the use of AI for ideation, language editing, text structuring, code generation, data analysis, or summarisation is permissible and when it can be considered an unjustified substitute for scholarly effort.

The second component is to develop cognitive abilities to validate AI outputs. This training should go beyond warning about the possibility of error and familiarize students with practical methods for checking sources, comparing independent evidence, detecting spurious references, assessing confidence levels, and determining the need for human oversight. Training should also clearly distinguish between improving an output through prompt revision and validating the accuracy of the resulting content.

The third component is to teach ethics and data security based on real-world scenarios. Rather than relying solely on general terms, students should be exposed to situations in which they have to make decisions about privacy, intellectual property, consent, data disclosure, authorship, and disclosure of AI use. Discipline-specific scenarios can help students understand the implications of using AI in workplace-like situations.

The fourth component is to establish participatory governance and a continuous feedback mechanism that can make policy review more responsive to recurring student-reported ambiguities and emerging forms of AI use. Policies should be reviewed at regular intervals, using evidence from students, academic staff, administrators, professional bodies, and evolving labour-market requirements. Potential mechanisms include student representation on relevant committees, periodic surveys of policy ambiguities, and transparent dissemination of policy changes. Their effectiveness in reducing policy–experience tensions would require empirical evaluation.

These four components are designed to operate in an integrated manner across three levels: institutional policy, curriculum, and instructional activities. Institutional policy should set out general principles and ethical requirements; curriculum should define expected competencies in a discipline-specific manner; and each course should specify how AI is permitted to be used in specific activities and assessments. Lack of coordination between these levels can lead to inconsistent rules and a fragmented experience for students.

Figure 4 integrates these four empirically derived pillars into a single governance architecture linking institutional policy, curriculum, instructional practice, and stakeholder feedback. The framework is intended to translate the recurring needs identified in the data into actionable institutional mechanisms. Its proposed outcomes are balanced AI literacy, responsible academic use, and stronger perceived preparation for AI-influenced professional contexts; these pathways remain conceptual propositions requiring future empirical validation.

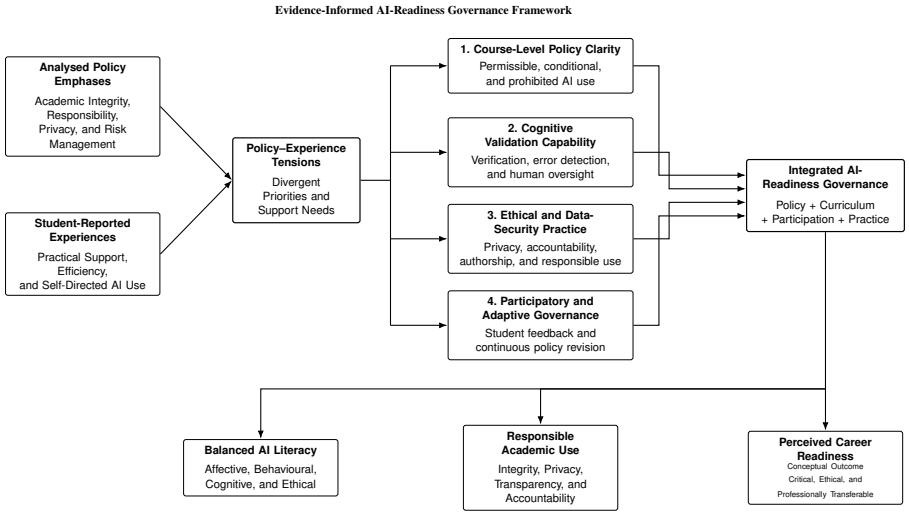

Fig. 4. Evidence-informed AI-readiness governance framework derived from the policy–experience relationships identified in this study. The four pillars represent recurring needs concerning course-level policy clarity, cognitive validation capability, ethical and data-security practice, and participatory and adaptive governance. Perceived career readiness is represented as a conceptual and prospective outcome rather than an empirically measured employment outcome.

The framework's practical value lies in connecting policy design with curriculum, instructional practice, and stakeholder feedback rather than treating these as separate institutional functions. It therefore provides a coherent basis for designing and testing interventions that respond to the support needs identified in the present data. Whether these mechanisms improve objectively measured AI competence, employment readiness, or workplace performance remains an empirical question for future longitudinal and institution-specific research.

The proposed framework is based on the principle that policy should not aim to choose between complete prohibition and unrestricted freedom. A broad prohibition can lead students to covert or self-taught use, while unguided encouragement increases the risk of cognitive dependence, error, privacy violations, and undermines scholarly originality. The appropriate approach is to regulate use with education, transparency, and accountability.

In practice, universities can implement this framework by developing a graded guide to AI use, designing short training modules on validation and privacy, adding AI use statements to assignments, creating templates for declaring how to use it, and providing a feedback channel for students. Also, the assessment of AI proficiency should not focus solely on the ability to produce output, but also include the ability to justify the choice of tool, examine the quality of the response, identify risk, and accept ultimate responsibility.

The theoretical contribution of this study lies in integrating three areas that are often examined separately: institutional AI policy, multidimensional AI literacy, and perceived professional preparation. The findings extend the ABCE perspective by showing that student-reported affective, behavioural, cognitive, and ethical experiences are embedded within an organisational context of policy clarity, educational provision, communication, and participation. This integration provides a more complete basis for theorising AI readiness in higher education while remaining appropriately bounded to perceived preparation rather than demonstrated employment outcomes.

The career-readiness element of the framework therefore represents a conceptual and prospective outcome. The present data identify perceived preparation needs and governance priorities with sufficient consistency to motivate the framework, but they do not establish that implementing it will improve graduate employment, employer-assessed competence, workplace performance, or longitudinal skill transfer. Those relationships constitute explicit hypotheses for subsequent evaluation.

The proposed framework should therefore be interpreted as a conceptually grounded, exploratory model whose effectiveness has not yet been empirically established. The model is derived from the analysis of 10 interviews and three policy documents, and to empirically evaluate it, future studies should include larger samples, different disciplines, multiple universities, and the perspectives of faculty members, administrators, and employers. Despite this limitation, the framework provides a coherent and actionable basis for designing and empirically testing more educational, participatory, and skills-oriented approaches to AI governance.

6. STUDY LIMITATIONS AND FUTURE RESEARCH

The present study has several limitations that should be considered when interpreting the findings. First, the interview data were based on a limited sample of 10 students. The aim of the study was not to achieve statistical generalizations, but rather to provide an in-depth and contextualized interpretation of students' experiences of encountering university AI policies. However, the limited sample size makes it difficult to generalize the prevalence of the identified perspectives and patterns to a broader student population.

Second, a limitation concerns participant-characteristic reporting. Although the available records establish that the interview sample included students across undergraduate, graduate, and doctoral levels, complete participant-level information on disciplinary field, age, gender, and study stage was not available for reliable reporting. The study therefore cannot examine whether the identified experiences vary systematically according to demographic characteristics, discipline, or educational level. Future studies should collect and report a structured participant profile at recruitment and use more diverse samples to examine potential variation across disciplines, degree levels, age groups, gender, institutional settings, and cultural contexts.

Third, the study's policy analysis was limited to three official documents, selected based on direct relevance to the topic and accessibility. These documents are not representative of all policy approaches in higher education, and there may be important differences between universities, countries, and accreditation systems. Furthermore, the document analysis reflects the official content of policies and does not necessarily reflect the quality of their implementation at the faculty, course, or assessment activity level.

A further limitation concerns the relationship between the two data sources. The policy documents and interview participants were not analysed as individually matched institutional cases. Consequently, the study cannot determine whether a specific participant's reported experience resulted directly from the implementation of a particular analysed policy. The policy–experience comparison should therefore be understood as thematic and exploratory. Future research should employ matched institutional designs in which students are interviewed within the same universities from which the

analysed policies are obtained, allowing more direct assessment of policy interpretation, implementation, and experienced effects.

Fourth, the data were largely based on students' perceptions and self-reports. Responses may have been influenced by memory, personal perceptions, a desire to provide a socially acceptable response, or concerns about academic judgement. The present study did not directly observe students' actual behaviour when using AI, the quality of their AI-related performance, or the extent to which their performance aligned with policies. The use of observation, task analysis, tool-use recording, and scenario-based studies could help future research to more accurately assess actual behaviour.

Fifth, because the analysis was not independently replicated by a second coder, interpretive consistency rests on the transparency of the codebook, coding rules, source-level coding checks, and explicit linkage between evidence and interpretation. No inter-coder agreement coefficient was therefore calculated. Although the structured codebook, operational definitions, decision rules, direct quotation support, and systematic checking of source assignments strengthened the dependability and transparency of the analysis, they do not substitute for independent inter-rater assessment. Future studies could strengthen this aspect further through independent double coding of a defined proportion of the dataset, followed by documented consensus review and, where methodologically appropriate, reporting of an agreement indicator.

Sixth, generative AI and its related policies are evolving rapidly. Study findings should be interpreted within the specific time context in which the data were collected. Changing tool capabilities, evolving organisational systems, reforming university regulations, and evolving labour market expectations can change student usage patterns and concerns in a short period of time. Therefore, longitudinal studies are necessary to examine the evolution of AI literacy and the impact of policy revisions.

Seventh, career readiness was examined through students' self-reported perceptions of their preparation for AI-influenced professional contexts and through its conceptual relationship with dimensions of AI literacy. The study did not measure graduate employment outcomes, workplace performance, employer evaluations, professional task performance, or the longitudinal transfer of AI-related competencies from university to employment. Consequently, the findings should not be interpreted as evidence that participants were objectively more or less career-ready. Future research should incorporate graduate tracking, employer perspectives, longitudinal designs, and performance-based assessment using authentic or simulated professional AI tasks to determine whether the competencies identified in this study transfer effectively to workplace settings.

Given these limitations, future research directions could include validating the proposed framework across multiple universities, comparing across disciplines, assessing faculty and administrators' perspectives, longitudinally studying student attitude change, and testing educational interventions. Also, the effectiveness of scenario-based training, graded AI guidelines, and student engagement mechanisms in improving validation skills, policy transparency, and perceived professional preparation could be explored.

7. CONCLUSION

This study comparatively examined institutional approaches to AI governance, students' reported experiences of AI use, and their perceived preparation for AI-influenced professional contexts through the analysis of three policy documents and semi-structured interviews with 10 students. Its empirical contribution lies in identifying recurring patterns across student narratives and systematically relating those patterns to institutional policy provisions. The findings should therefore be interpreted as evidence concerning perceptions, experiences, and thematic policy–experience relationships rather than as direct measures of employment readiness, workplace performance, or institution-specific policy effectiveness.

The analysis revealed a substantive difference in emphasis between the two source types. The analysed policies foregrounded academic integrity, risk management, student responsibility, privacy, and responsible use, whereas participants more often described AI in relation to practical academic support, efficiency, workload management, information access, and preparation for AI-influenced professional contexts. Because the policy corpus and interview sample were not individually matched by institution, these differences do not establish that a particular policy succeeded or failed in implementation. Instead, they identify concrete areas in which institutional provisions and student-reported needs converge, diverge, or remain only indirectly related.

Viewed through the ABCE framework, the student accounts revealed distinct, context-dependent patterns within the affective, behavioural, cognitive, and ethical dimensions of AI literacy. Frequent use was not accompanied uniformly by systematic validation, stable ethical interpretation, or confidence in institutional guidance. Participants also reported concerns regarding false AI detection, fragmented communication, and limited opportunities for participation in policy-related processes. These findings provide a strong basis for identifying educational and governance priorities while remaining bounded to the experiences reported by this sample.

A central conclusion of the study is therefore that policy coverage and student experience should not be treated as interchangeable evidence. A policy may address validation, privacy, communication, or academic integrity without demonstrat-

ing how effectively those provisions are implemented; likewise, a student-reported concern does not establish institutional policy failure. The comparative value of the study lies in making these relationships visible and analytically explicit.

On this basis, the proposed AI-readiness governance framework integrates course-level policy clarity, systematic training in output validation, ethics and data-security education, and participatory feedback mechanisms. The framework is designed to connect institutional policy, curriculum, and instructional practice with the areas of support identified in the student accounts. Perceived career readiness is represented as a conceptual and prospective outcome rather than an empirically demonstrated employment outcome.

The study therefore supports a shift from treating AI governance primarily as a matter of control towards a more integrated educational, participatory, and competency-oriented approach. This conclusion remains strong but appropriately delimited: the present evidence identifies what students report needing and how those needs relate to formal policy provisions; it does not establish causal policy effects or objective employment readiness. Future longitudinal and institution-specific studies should test whether the proposed mechanisms improve measurable AI competence, employer-assessed capability, workplace performance, and graduate employment outcomes.

Conflicts of Interest

The authors declare no conflicts of interest.

Funding

This research received no external funding.

CRediT Author Contribution Statement

Mohammed Alalwan: Conceptualization, Investigation, Data curation, Writing – original draft. **Sam Ansari:** Methodology, Investigation, Data curation, Formal analysis, Validation, Visualization, Writing – original draft, Writing – review & editing. **Elton Xhetani:** Conceptualization, Supervision, Writing – review & editing. **Soliman Mahmoud:** Methodology, Validation, Supervision, Writing – review & editing. **Dhafar Hamed Abd:** Formal analysis, Writing – review & editing. **Ayad Turky:** Validation, Writing – review & editing. **Hissam Tawfik:** Formal analysis, Visualization, Writing – review & editing. **Abir Hussain:** Conceptualization, Methodology, Formal analysis, Supervision, Project administration, Writing – review & editing.

All authors have read and approved the final version of the manuscript and agree to be accountable for their respective contributions.

Data Availability

The qualitative data supporting this study were obtained from semi-structured interviews with 10 students conducted between 12 and 19 April 2026 and from three institutional artificial intelligence policy documents. The original interview data have been retained by the research team.

Owing to participant confidentiality, privacy considerations, and the potential risk of indirect identification, the interview data are not publicly available. De-identified interview materials can be provided to the Editor upon reasonable request for confidential editorial review, subject to participant consent, the necessary institutional permissions, and applicable ethical and data-protection requirements. Any authorised disclosure will be preceded by an assessment of potential identification risks and the application of appropriate confidentiality safeguards.

The institutional policy documents analysed in this study are identified and cited in the manuscript.

Acknowledgment

The authors sincerely thank the students who voluntarily participated in this study and shared their experiences and perspectives on artificial intelligence in higher education. Their contributions provided the qualitative evidence underpinning the research.

REFERENCES

[1] L. Yan, S. Greiff, Z. Teuber, and D. Gašević, "Promises and challenges of generative artificial intelligence for human learning," *Nature Human Behaviour*, vol. 8, no. 10, pp. 1839–1850, 2024. [Online]. Available: https://doi.org/10.1038/s41562-024-02004-5

[2] B. Burnes, "Complexity theories and organizational change," *International Journal of Management Reviews*, vol. 7, no. 2, pp. 73–90, 2005. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-2370.2005.00107.x

[3] K. E. Weick and R. E. Quinn, "Organizational change and development," *Annual Review of Psychology*, vol. 50, no. 1, pp. 361–386, 1999.

[4] J. Phillips and J. D. Klein, "Change management: From theory to practice," *TechTrends: For Leaders in Education and Training*, vol. 67, no. 1, pp. 189–197, January 2023, published online September 9, 2022. [Online]. Available: https://doi.org/10.1007/s11528-022-00775-0

[5] E. Kasneci, K. Sessler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günnemann, E. Hüllermeier, S. Krusche, G. Kutyniok, T. Michaeli, C. Nerdel, J. Pfeffer, O. Poquet, M. Sailer, A. Schmidt, T. Seidel, M. Stadler, J. Weller, J. Kuhn, and G. Kasneci, "Chatgpt for good? on opportunities and challenges of large language models for education," *Learning and Individual Differences*, vol. 103, p. 102274, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1041608023000195

[6] D. R. E. Cotton, P. A. Cotton, and J. R. Shipway, "Chatting and cheating: Ensuring academic integrity in the era of chatgpt," *Innovations in Education and Teaching International*, vol. 61, no. 2, pp. 228–239, 2024. [Online]. Available: https://doi.org/10.1080/14703297.2023.2190148

[7] C. K. Y. Chan, "A comprehensive ai policy education framework for university teaching and learning," *International Journal of Educational Technology in Higher Education*, vol. 20, no. 1, p. 38, 2023. [Online]. Available: https://doi.org/10.1186/s41239-023-00408-3

[8] J. L. (Jess), "A critical review of genai policies in higher education assessment: a call to reconsider the "originality" of students' work," *Assessment & Evaluation in Higher Education*, vol. 49, no. 5, pp. 651–664, 2024. [Online]. Available: https://doi.org/10.1080/02602938.2024.2309963

[9] Y. An, J. H. Yu, and S. James, "Investigating the higher education institutions' guidelines and policies regarding the use of generative ai in teaching, learning, research, and administration," *International Journal of Educational Technology in Higher Education*, vol. 22, no. 1, p. 10, 2025. [Online]. Available: https://doi.org/10.1186/s41239-025-00507-3

[10] Microsoft and LinkedIn, "Ai at work is here. now comes the hard part," Microsoft and LinkedIn, Work Trend Index Annual Report, 2024, accessed: 18 April 2026. [Online]. Available: https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part/

[11] A. Margaryan, "Artificial intelligence and skills in the workplace: An integrative research agenda," *Big Data & Society*, vol. 10, no. 2, p. 20539517231206804, 2023. [Online]. Available: https://doi.org/10.1177/20539517231206804

[12] T. K. F. Chiu, "Ai literacy and competency: definitions, frameworks, development and future research directions," *Interactive Learning Environments*, vol. 33, no. 5, pp. 3225–3229, 2025. [Online]. Available: https://doi.org/10.1080/10494820.2025.2514372

[13] M. Z. Asghar, K. A. Duah, J. Iqbal, and H. Järvenoja, "Evidence from west africa on the interplay of affective behavioral cognitive and ethical dimensions of ai literacy in ghanaian and nigerian universities," *Discover Computing*, vol. 28, no. 1, p. 169, 2025. [Online]. Available: https://doi.org/10.1007/s10791-025-09691-2

[14] A. Benlian and M. Pinski, "The ai literacy development canvas: Assessing and building ai literacy in organizations," *Business Horizons*, vol. 69, no. 6, pp. 759–775, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0007681325001673

[15] Y. M. R. Alhezzani, "Change recipients' resistance and salience to organizational re-creation: the effects of participation and coercion strategies on change derailment," *Organization Management Journal*, vol. 18, no. 1, pp. 2–18, 07 2020. [Online]. Available: https://doi.org/10.1108/OMJ-10-2018-0608

[16] D. R. Bamford and P. L. Forrester, "Managing planned and emergent change within an operations management environment," *International Journal of Operations & Production Management*, vol. 23, no. 5, pp. 546–564, 05 2003. [Online]. Available: https://doi.org/10.1108/01443570310471857

[17] R. T. By, "Organisational change management: A critical review," *Journal of Change Management*, vol. 5, no. 4, pp. 369–380, 2005. [Online]. Available: https://doi.org/10.1080/14697010500359250

[18] S. Karasvirta and S. Teerikangas, "Change organizations in planned change – a closer look," *Journal of Change Management*, vol. 22, no. 2, pp. 163–201, 2022. [Online]. Available: https://doi.org/10.1080/14697017.2021.2018722

[19] B. Burnes, "Kurt lewin and the planned approach to change: A re-appraisal," *Journal of Management Studies*, vol. 41, no. 6, pp. 977–1002, 2004. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-6486.2004.00463.x

[20] M. Bearman, J. Tai, P. Dawson, D. Boud, and R. Ajjawi, "Developing evaluative judgement for a time of generative artificial intelligence," *Assessment & Evaluation in Higher Education*, vol. 49, no. 6, pp. 893–905, 2024. [Online]. Available: https://doi.org/10.1080/02602938.2024.2335321

[21] A. Salido, I. Syarif, M. S. Sitepu, Suparjan, P. R. Wana, R. Taufika, and R. Melisa, "Integrating critical thinking and artificial intelligence in higher education: A bibliometric and systematic review of skills and strategies," *Social Sciences & Humanities Open*, vol. 12, p. 101924, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2590291125006527

[22] C. Argyris and D. A. Schön, *Organizational Learning II: Theory, Method, and Practice*, 2nd ed. Reading, MA: Addison-Wesley, 1996.

[23] L. J. Wiese, I. Patil, D. S. Schiff, and A. J. Magana, "Ai ethics education: A systematic literature review," *Computers and Education: Artificial Intelligence*, vol. 8, p. 100405, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666920X25000451

[24] V. Braun and V. Clarke, *Successful Qualitative Research: A Practical Guide for Beginners*, 1st ed. London: SAGE Publications, 03 2013. [Online]. Available: https://www.sagepub.com/shop/buy-a-book/successful-qualitative-research-1-233059

[25] J. A. Maxwell, "Using numbers in qualitative research," *Qualitative Inquiry*, vol. 16, no. 6, pp. 475–482, 2010. [Online]. Available: https://doi.org/10.1177/1077800410364740

[26] M. Sandelowski, "Real qualitative researchers do not count: The use of numbers in qualitative research," *Research in Nursing & Health*, vol. 24, no. 3, pp. 230–240, 2001. [Online]. Available: https://doi.org/10.1002/nur.1025

[27] G. A. Bowen, "Document analysis as a qualitative research method," *Qualitative Research Journal*, vol. 9, no. 2, pp. 27–40, 08 2009. [Online]. Available: https://doi.org/10.3316/QRJ0902027

[28] H. Kallio, A.-M. Pietilä, M. Johnson, and M. Kangasniemi, "Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide," *Journal of Advanced Nursing*, vol. 72, no. 12, pp. 2954–2965, 2016. [Online]. Available: https://doi.org/10.1111/jan.13031

[29] L. A. Palinkas, S. M. Horwitz, C. A. Green, J. P. Wisdom, N. Duan, and K. Hoagwood, "Purposeful sampling for qualitative data collection and analysis in mixed method implementation research," *Administration and Policy in Mental Health and Mental Health Services Research*, vol. 42, no. 5, pp. 533–544, 2015. [Online]. Available: https://doi.org/10.1007/s10488-013-0528-y

[30] K. Proudfoot, "Inductive/deductive hybrid thematic analysis in mixed methods research," *Journal of Mixed Methods Research*, vol. 17, no. 3, pp. 308–326, 2023. [Online]. Available: https://doi.org/10.1177/15586898221126816

[31] O. C. Robinson and D. Bailey-Rodriguez, "Deductive qualitative research: an integrative approach to designing studies," *Qualitative Research in Psychology*, vol. 23, no. 2, pp. 369–384, 2026. [Online]. Available: https://doi.org/10.1080/14780887.2025.2604773

[32] D. R. Thomas, "A general inductive approach for analyzing qualitative evaluation data," *American Journal of Evaluation*, vol. 27, no. 2, pp. 237–246, 2006. [Online]. Available: https://doi.org/10.1177/1098214005283748

[33] Y. S. Lincoln, "Emerging criteria for quality in qualitative and interpretive research," *Qualitative Inquiry*, vol. 1, no. 3, pp. 275–289, 1995. [Online]. Available: https://doi.org/10.1177/107780049500100301

[34] University of London, "Artificial intelligence (ai) policy," University of London, 2024, reference UOL04, Version 1.2, approved 11 December 2024; annual review cycle. [Online]. Available: https://www.london.ac.uk/sites/default/files/artificial-intelligence-policy-uol.pdf

[35] University of Sharjah, "Policy on the responsible use of artificial intelligence in academic texts," University of Sharjah, Graduate Studies, policy covering graduate theses, doctoral and master's courses, research plans, related scholarly outputs, and other graduate academic texts. [Online]. Available: https://www.sharjah.ac.ae/-/media/project/uos/sites/uos/graduate-studies/bylaws-and-regulations/policy-on-the-responsible-use-of-artificial-intelligence-in-academic-texts.pdf

[36] University of North Alabama, "Generative ai policy: Guidelines for responsible generative ai use," University of North Alabama, institution-wide guidance for students, faculty, and staff; reviewed and updated annually. [Online]. Available: https://www.una.edu/academics/generative-ai-policy.html