Ruwad Al-Mustaqbal System: A Knowledge-Behavior Gap-Aware Framework (KBG-Ruwad) for Personalized Responsible AI Education among Omani Undergraduates

Samah Alzidi, Aseel Abdul Nassor Ahmed Al Shizawi and Arunadevi Karuppasamy
Department of Computing Sciences, Gulf College, Muscat, Oman

2310149@gulfcollege.edu.om
2310179@gulfcollege.edu.om
arunadevi@gulfcollege.edu.om

Abstract: The use of generative AI among undergraduate students has surpassed structured instruction in responsible use of generative AI, and this gap is documented, with students' awareness of and knowledge of AI ethics not matching with their ability to apply it in academic contexts. This has been identified as a critical issue in Oman, with Vision 2040 stating AI adoption as a visionary goal, and most Omani HEIs lack formal AI governance, along with faculty AI literacy falling short of visionary goals. AI ethics are structured into progressive levels in competency frameworks like UNESCO's, and adaptive systems already personalize content. This paper proposes a conceptual framework, KBG-Ruwad, that introduces a novel measure of the knowledge-behavior gap (KBG). The framework assesses, for each student and across each dimension, the discrepancy between self-reported ethical knowledge and behavior demonstrated in scenario-based tasks. In addition, it measures students' level of ethical knowledge, their confidence in that knowledge, and the misconceptions they hold. The framework further extends the existing seven-dimensional model. Building on these measures, KBG-Ruwad extends the seven-dimension Responsible AI Competency Vector with discipline-specific, misconception-aware adaptive scenarios, and with a guidance function sensitive to gap, confidence, discipline and learning history; it further introduces a student-owned Learning Passport and a privacy-protected institutional dashboard. The paper reviews the AI-competency, AI-ethics pedagogy and Omani AI-readiness literature; argues for integration over invention; critically analyzes the components and flow of the proposed framework; and identifies gaps in the current specification and challenges for its operationalization

Keywords: AI Ethics Education, Knowledge-Behavior Gap, Awareness Profiling, Personalized Learning, Responsible AI, AI Competency Assessment

1. Introduction

While using generative AI tools in the undergraduate classroom has become commonplace, the use of ethical judgement in working with generative AI tools, what information can be provided to a generative AI tool, who can claim credit for the work created, and when is the work generated by the tool not the work of the student, is not taught as a distinct, practiced skill. In general, universities have responded with citation policies and brief orientation modules instead of ongoing teaching and learning and have not determined which students have which misconceptions or whether a student who can state a principle is using it. This deficit however, in Oman, does not happen in isolation but rather is a part of a specific national agenda. Digital transformation and AI adoption are strategic objectives of Oman Vision 2040, pursued in part through the country's workforce development objectives (Oman Vision 2040, 2019). The national AI and advanced-technologies programme, in turn, focuses explicitly on strengthening Oman's position on international measures of AI-readiness (MTCIT, 2024). Benayoune, Slimi and Al Habsi (2026) conducted a mixed-methods study on the 28 Omani higher education institutions, finding only 5.1% had a formal policy on AI and 74.4% did not, while 59% of the faculty surveyed were concerned about students relying heavily on generative tools. The results of the survey conducted among the 92 faculty members of a private Omani university revealed moderate awareness and usage of AI along with significantly higher acceptance of the technology with positive attitudes, which indicates that the gap between awareness and usage of the technology cannot be addressed by such general awareness campaigns (Natarajan et al., 2026). This paper presents a conceptual model, termed KBG-Ruwad, that measures this gap, dimension-by-dimension, and combines this measurement with confidence, discipline and learning history, to guide towards non-punitive, personalized advice by filling the gap. The elements of its individual components are not new to the literature, as sections 3 explain, but its novelty is the way in which they are integrated around a single measurable discrepancy. The related work is reviewed in Section 2, the gap and objectives are stated in Section 3; the framework is formalized in Sections 4 and 5, a critical, structural analysis of the work is provided in Section 6, Section 7-9 discuss ethics, significance and limitations, and Sections 10 concludes.

2. Literature Review

The power of AI-driven personalization has been proven. A global review of 25 studies found that adaptive content delivery, real-time feedback and progress tracking are already common features of AI-driven personalization (Fortuna et al., 2025). AI-powered virtual assistants are used in a similar way to enhance student engagement, and interest in this application has grown rapidly, although mainly from the perspective of educators rather than students (Sajja et al., 2024). UNESCO's AI competency framework for students already defines what students should know in four aspects: human-centred mindset, ethics, techniques and system design, organized as three development phases (UNESCO, Miao and Shiohira, 2024). Personalization and progression of competency levels are not new, as this literature focuses on the competency of subject-matter or general AI competency, rather than if the behavior correlates with the stated knowledge; and also as the technical and ethical competencies evolve as progressions of their own, and this could be measured separately from the learning of the content, as demonstrated in the early research on the foundation of K-12 education (Casal-Otero et al., 2023). While the studies above focus on the technology of content personalization, a separate strand of literature treats AI ethics as a pedagogical problem: it asks how ethical reasoning about AI use should be taught, rather than how content delivery can be automated. A 2025 EDUCAUSE framework provides an AI ethics foundation based on scenario analysis in real world settings instead of compliance (Georgieva and Stuart, 2025). A systematic review of AI-ethics interventions revealed that several case studies and group projects are already common, but the effects on behavior, as opposed to declarative knowledge, is inconsistent and often missing (Wiese et al., 2025). A review of 68 K-12 studies reached a similar conclusion, finding that ethics-based approaches also support AI-literacy development (Ma et al., 2025). A separate review of 34 studies found limited incorporation of cognitive, behavioral and attitudinal indicators, with most designs favoring dilemma-based tasks (Yan, Liu and Chau, 2025). A mixed-studies review similarly found that very few interventions used an ethics framework as a diagnostic tool for individual learners (Fu and Weng, 2024). Indeed, scenario-based AI-ethics teaching is already common, and the open question is how scenario performance should relate to an individual profile that can be updated, reinforced by a coordination-game analysis which shows that responsible norms rely on the expectations and assessment design of peers and that reflective tasks change norms over time (Ogbo et al., 2026). A research paper addressing AI's role in supporting Oman's knowledge-based economy discusses personalized and adaptive learning as an approach to aligning higher education with Vision 2040 (Sujee, Othmani and Dalwai, 2024). A review suggests that measures for reducing bias should be implemented on an institutional level, not solely within the classroom, and that it is necessary for institutional leadership to take responsibility for an institution's relationship to responsible adoption (Barnes and Hutson, 2024; Tarisayi, 2024). Skill development in general has been proposed to be adaptive and AI-driven, with a personalized-pathway approach (Akintola, Akintayo and Kadri, 2024), while studies underline both the promise and the academic-integrity concerns of generative AI in education have been documented (Adel, Ahsan and Davison, 2024). There are no studies that suggest there is a platform that would be specifically designed to identify and address the AI ethics knowledge and practice gaps that exist in the individual student population.

3. Research Gap and Objectives

Existing frameworks and adaptive learning systems for AI competency already establish several features, including AI-literacy assessment, awareness of ethics and privacy, academic-integrity instruction, case-based scenarios, and personalized recommendations and progress dashboards (Fortuna et al., 2025; Sajja et al., 2024; UNESCO, Miao and Shiohira, 2024). Similarly, AI-ethics research already draws on case studies and group reasoning rather than relying solely on declarative teaching (Wiese et al., 2025). What the literature does not depict is these features brought together to form a mechanism that compares the information a student reports he/she knows to behavior, dimension by dimension, confidence level by confidence level, and guides personalized, discipline-specific guidance from that comparison. While the UNESCO framework outlines knowledge expectations, it does not have a clear pathway on how to consistently measure such knowledge; while adaptive learning systems personalize content, they lack a focus on multidimensional AI-ethics competencies; and while the AI-ethics pedagogy documents scenario-based teaching, they rarely connect performance to an individual, updatable profile (UNESCO, Miao and Shiohira, 2024; Fortuna et al., 2025; Wiese et al., 2025; Yan, Liu and Chau, 2025). This gap is part of an Omani environment where the policy intent of Vision 2040 has been clearly stated and where there is a recognized lack of AI governance (Benayoune, Slimi and Al Habsi, 2026; Natarajan et al., 2026). KBG-Ruwad's main contribution is an integration of existing components, rather than a wholly new component, into a discipline-sensitive, non-punitive framework for undergraduate AI-ethics education in Oman. Within this framework, the knowledge-behavior gap is formalized as a measurable

construct, as presented in Section 4.2. The aim is to propose and conceptually evaluate a knowledge-behavior gap-aware framework for AI-ethics education for Omani undergraduates. The goals of the work are to develop a multidimensional competency model and a formalized measurement of the gap between knowledge and behavior; to develop discipline-specific, adaptive scenario and guidance mechanisms based on these gaps and on confidence and learning history; and to present a conceptual evaluation and phased implementation plan, including governance safeguards, that can take the model from proposal to a piloted system.

4. The KBG-Ruwad Conceptual Framework

4.1 Framework Dimensions and the Responsible AI Competency Vector

The framework sees AI-ethics competency as seven interrelated dimensions, rather than one score: AI literacy (L), even if it's just the basic understanding users have of the process AI tools use to create their output; ethical awareness (E), the knowledge of core ethical concepts, such as autonomy and informed consent; privacy awareness (P), understanding of privacy concepts and what data is safe to share with the AI tool; bias and fairness awareness (F), the ability to identify bias from AI output; academic-integrity awareness (A), knowing how to use AI tools to cite sources and acknowledge others' contributions, or how to disclose its use; critical evaluation (X), the habit of verifying AI output and not taking it as gospel; and responsible AI-use behavior (Bh), the general pattern of a student's decisions in AI-primitive use. These form a Responsible AI Competency Vector for student i:

$$R_i = [L_i, E_i, P_i, F_i, A_i, X_i, Bh_i]$$

These dimensions reflect the distinction that the K-12 AI-literacy literature has made between technical competency and ethical competency (Casal-Otero et al., 2023) and allow the platform to identify students who demonstrate high ethical awareness but low privacy awareness, a distinction that would not be captured by a single composite score.

4.2 Knowledge-Behavior Gap and Confidence-Aware Misconception Indicator

The framework splits a knowledge score $K_{i,d}$, from the questionnaire and a behavior score $B_{i,d}$, from the choices of scenarios , and defines the knowledge-behavior gap:

$$KBG_{i,d} = K_{i,d} - B_{i,d}$$

The sign and size of $KBG_{i,d}$ distinguish four knowledge-behavior combinations. High knowledge with low behavior indicates an application gap: a student who can correctly describe the privacy principle in the questionnaire but still uploads confidential material in the scenario. These students are expected to benefit more from practice-oriented guidance than from additional theoretical content. Low knowledge with high behavior indicates intuitive but incomplete knowledge, which is worth reinforcing conceptually so that it generalizes beyond the scenarios already encountered. Low knowledge with low behavior is treated as the highest priority for support, since neither the concept nor the corresponding practice is yet in place. High knowledge with high behavior indicates that the two measures agree and that KBG i,d is close to zero; because this combination does not indicate a gap, it does not receive the same targeted guidance as the other three. The second combination, low knowledge with high behavior, is expected to be more common among students who already use AI daily, and a single composite score would not distinguish it from the first combination. The second profile (low knowledge and high behavior) is more common among students who already use AI every day, a single composite score is not enough to differentiate it from the first case. Each scenario is also presented with a question about the student's confidence $Cf_{i,d}$, normalized in [0,1] since he/she could get the response right by chance. This helps to create a misconception-risk indicator:

$$MRI_{i,d} = \left(1 - B_{i,d}\right) \times Cf_{i,d}$$

A high $MRI_{i,d}$ indicates that the student chose the wrong action while reporting high confidence in it, which signals a misconception the student is unlikely to correct without guidance, rather than a result of guessing or acknowledged uncertainty. KBG and MRI together provide two separate signals as opposed to one signal for the profiling in Section 5.3.

5. Framework Components and Conceptual Platform Architecture

The six-layer architecture is summarized in Figure 1, which contains a profiling and knowledge assessment layer, a discipline-specific, adaptive scenario layer which includes the representation of behavior and confidence, the core KBG or MRI computation, a layer of personalized student support combining multidimensional profiling, the guidance function, explanation-based feedback and the Learning Passport, and an institutional insight and refinement layer. The details of each block are found in the corresponding subsection of Section 5.

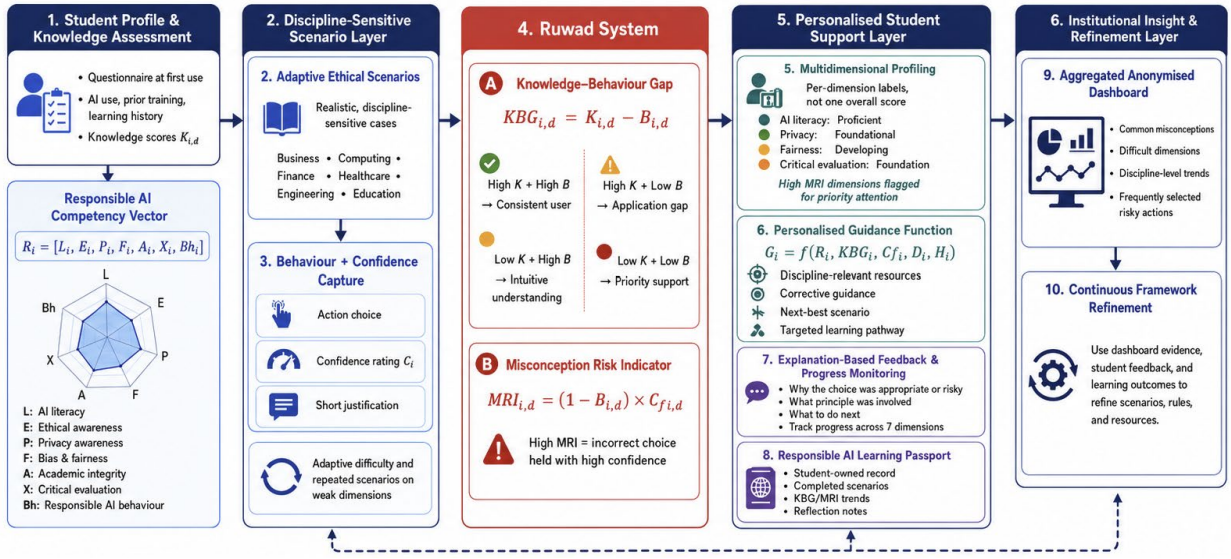


Figure 1. The architecture of the KBG-Ruwad system

5.1 Student Profile and Knowledge Assessment

An initial profile is created from the first- administration of a questionnaire, which asks about the frequency and purpose of using AI tools, understanding of AI hallucination, verification practices, citing AI sources, disclosure practices, processing institutional data, identification of biased outputs, and reliance on answers produced by AI, resulting in $K_{i,d}$ for each dimension. The instrument is conceptualized at this stage which is not used any actual student data but will be proposed to be pilot run and psychometrically validated in the future.

5.2 Discipline-Sensitive, Adaptive Ethical Scenarios

Students are presented with realistic scenarios covering uploading personal information, using AI-generated content without attribution, relying on inaccurate output without verification, acting on biased recommendations, using AI in assessed work, and relying too heavily on AI answers in written responses. Each theme is written as a set of variations, one for each of six sectors: computing, business, finance, healthcare, engineering, and education. Scenario selection is adaptive: if a student struggles on a dimension, they are given a further scenario on that dimension, at an easier or harder level depending on their performance. For each scenario, students choose the action they would take, rate their confidence in that decision, and provide a brief justification. The platform records the consequence of the chosen action, for example that sensitive data becomes exposed and the institution is put at risk, rather than only whether the answer was right or wrong. This branching, adaptive design follows a format that has already been shown to be effective for scenario-based instruction of ethical reasoning (Georgieva and Stuart, 2025; Wiese et al., 2025; Yan, Liu and Chau, 2025).

5.3 Behavioral, Confidence and Multidimensional Profiling

Profiling in this stage draws on a student's scenario choices to build a picture of their behavior, confidence and understanding across dimensions, and to indicate where each of these needs to improve. Scenario choices produce $B_{i,d}$ and $Cf_{i,d}$ for each dimension, from which $KBG_{i,d}$ and $MRI_{i,d}$. are calculated as described in Section 4.2. The labels Foundation, Developing or Proficient are then assigned to each dimension together, to create a multidimensional profile, e.g., AI literacy: Proficient, privacy: Foundation, fairness: Developing, academic integrity: Proficient, critical evaluation: Foundation. If a dimension has a high MRI, regardless of whether it is labeled Foundation/Developing/Proficient, the dimension is flagged for priority attention as a risk due to a

confidently held misconception in the dimension, not because it is labeled Developing/Proficient. This would lose the distinction that was the intent of grouping the seven dimensions in Section 4.1. These are still educational-support categories and not disciplinary designations; the complete profile is not sent to the teacher for the purpose of identifying form, and is not used for the purpose of marking.

5.4 Personalized Guidance and Recommendation Function

Guidance is recommended based on the student competency vector, gap profile, confidence profile, discipline and learning history:

$$G_i = f(R_i, KBG_i, Cf_i, D_i, H_i)$$

A student flagged with high MRI on privacy is prioritized for guidance on that dimension above dimensions the student has already acknowledged as uncertain but has not yet misconceived, and above discipline-relevant privacy dimensions already labeled Proficient.

5.5 Explanation-Based Feedback and Progress Monitoring

In addition to the chosen action, students are asked to submit a brief justification for the action, that will be assessed in a rubric format or by using the NLP model, moving the platform from being a multiple-choice system to an ethical reasoning assessment. Immediate feedback contains information about why a chosen action was right or wrong, which principle the action was based on, and what resources to move to next. Progress monitoring displays activities completed and areas that still need to be developed on the seven dimensions, as a process of self-development and not as a monitoring process.

5.6 Responsible AI Learning Passport

The Responsible AI Learning Passport is proposed to record a student's evolving profile, completed scenarios, dimension-level competencies, KBG and MRI trends, and self-authored reflection notes. The Passport is owned and controlled by the student: it is not visible to instructors by default, and, unlike the aggregated dashboard described below, it is not counted towards grading or used for disciplinary purposes.

5.7 Institutional Insight and Continuous Refinement

An authorized dashboard provides educators with aggregated, anonymized information on common misconceptions, the most challenging dimensions, differences between disciplines, the most frequently chosen risky actions, and changes in awareness over time. Individual-level detail is not routinely visible and remains subject to the safeguards described in Section 7. This evidence, in conjunction with student feedback and the pilot outcomes described in Section 9, is designed not to remain in place after deployment, but to be used to inform rounds of revision of the scenario library, guidance rules and content of the resources, thereby closing the loop in the sense of Figure 1.

6. Findings and Analysis

The findings are organized around four questions: what architecture is structurally built to expose, where it diverges from existing approaches and where gaps remain in its own specification, what practical challenges stand between the specification and a working instrument, and what these findings imply for the next stage of development. Table 1 sets the framework against the established work reviewed in Section 2. The extension in each row is architectural, not conceptual: none of the seven dimensions replaces an existing competency category, but the KBG and MRI measures in Section 4.2 quantify something none of these framework's measures alone, per-dimension discrepancy and confidence, rather than a single composite score. The analysis suggests that diagnosis is not incidental to the diagnostic framework; it is derived from the connections made by the two sections of the framework (Sections 4 and 5). The pipeline is coherent, every downstream component depends on the data it generates, and all components, except for the explanation-based one in Section 5.5, are implementable using methods already described in Section 2, and satisfy feasibility. Consider a business student with moderate privacy familiarity ($K_{i,privacy}$ = high). In the matching scenario, this student uploads a client dataset to a public AI tool to save time, a low-privacy action ($B_{i,privacy}$ = low), and rates their confidence in that choice as high ($Cf_{i,privacy}$ = high). Because $K_{i,privacy}$ is high while $B_{i,privacy}$ is low, the knowledge-behaviour gap ($KBG_{i,privacy}$) is large: the student knew better, so the failure is an application gap rather than a knowledge gap. Because the

action was taken with high confidence despite being wrong, the misconception risk index (MRI $_{i,privacy}$) is also elevated. The rest of this section examines whether that distinction holds up consistently in practice.

Table 1: Comparative Analysis of Existing AI Competency, Ethics, and Adaptive Learning Frameworks and the Proposed KBG-Ruwad Extension

Existing Framework	Existing Contribution	KBG-Ruwad Extension
UNESCO AI Competency Framework for Students	Combines ethics and technical competence in a progressive manner (UNESCO, Miao and Shiohira, 2024)	Incorporates a dimension of measured knowledge-behavior gap and confidence-weighted misconception detection.
AI-ethics pedagogy research	The use of case studies and as everyday practice the use of scenarios for teaching. (Wiese et al., 2025)	Rather than just an exercise, links scenario performance to an individual, updatable profile and adaptive scenario selection
Adaptive learning systems	Personalize content delivery and monitor progress (Fortuna et al., 2025; Sajja et al., 2024)	Focus on personalizing around competency vector and knowledge/behavior gap rather than overall subject progress
Academic-integrity guidance	Citation and disclosure norms for AI-assisted work	Embeds those norms within one of seven dimensions (confidence weighted)

Comparing the framework against Table 1 and against its own specification surfaces two groups of issues. The first concerns severity, not validity. If the confident/uncertain contrast were scored on a composite AI-ethics measure, the design argued for in earlier sections, it would fall into one of the "moderate" bands, and the same is true of the knowledge-only and behavior-only designs. So, while the framework distinguishes these cases correctly, it does not obviously rank them as more severe than a composite score would. The second group covers three assumptions that Section 6.1 relies on but does not itself examine. First, the seven dimensions in Section 4.1 were not derived empirically; they were argued for from the literature and from the problem statement in Section 1. Their coverage is assumed to be complete, and they are treated as independent of one another, neither claim is tested. Second, KBG $_{i,d}$ = K $_{i,d}$ – B $_{i,d}$ treats a given numeric gap as meaning the same thing on a common scale across dimensions and disciplines. Sections 5 establish only that the instruments share a similar convention at a given scale point, not that the gaps themselves are comparable in magnitude or meaning across dimensions. Third, the multiplicative form MRI $_{i,d}$ = (1 – B $_{i,d}$) × Cf $_{i,d}$ is consistent with the modelling approach in Section 4.2, but it is a design choice, not a form derived from or validated against a theory of misconception risk; a different functional form would weigh the same underlying data differently. Taken together, these results support the framework's diagnostic advantage over traditional methods, but they also show that its internal assumptions are untested by the comparative argument itself. They need to be examined separately, through how the framework performs in operation.

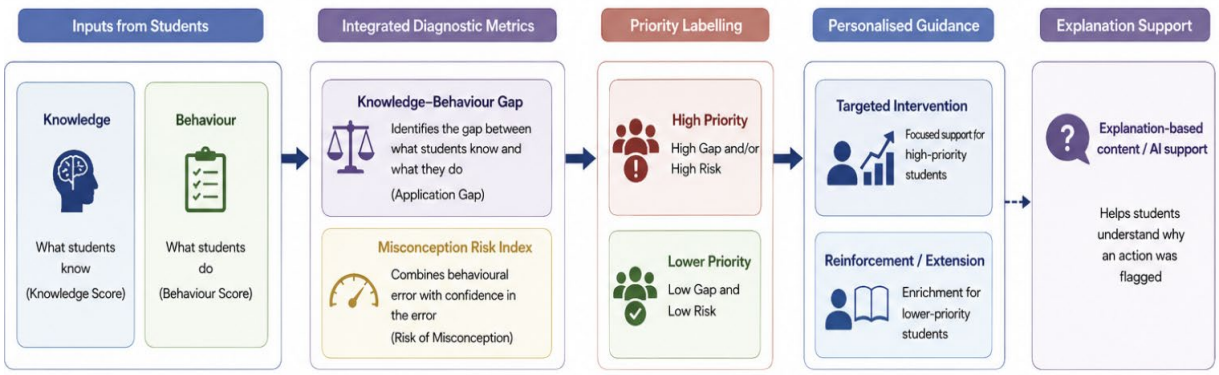


Figure 2: Findings from the Proposed Framework for Diagnostic Assessment and Personalized Student Guidance

In addition to the specification level gaps described earlier, this analysis suggests that there are some practical issues in making the framework a working instrument. It is not an easy task to write scenarios that are valid in the respective fields, without any scenario-writer bias as to what is the "responsible" action: The professional norms vary by field. A longstanding measurement challenge is to elicit reliable confidence ratings, where students may anchor on a fixed scale point, even if the rating is not a true reflection of their confidence. Adaptive difficulty, as suggested in Section 5.2, can become confused with under-competence if based solely on behavior, without taking confidence into account. The explanation-based module of section 5.5 relies on scoring unvalidated for this domain and therefore its reliability against human judgement is not presumed adequate. Creating the Learning Passport and dashboard to be non-punitive, then interpreting the flag as an indicator of risk is both an institutional and cultural challenge, as well as a technical one. These results show overall that operationalizing the framework poses challenges that are not covered by the gaps identified but are specific and additional. The gaps and problems highlighted above suggest five methodological priorities for the validation stage, beyond the general validation from expert review to pilot, described above. The score-scaling and equating procedures for both the $K_{i,d}$ and $B_{i,d}$ should be standardized and reported before any KBG value should be considered significant. The MRI formulation should be regarded as one of several and not the only formulation that should be tested prior to being accepted. Empirical pilot distributions should be used to determine KBG/MRI thresholds, and thresholds should be reported transparently and not determined a priori. A single subject-matter expert should not be used for more than one scenario in each discipline; cross-checks should be made to minimize single author bias. The explanation-based module needs to be tested with independent human raters before going into operation, and if reliability is not achieved, the guidance function does not require the use of the explanation-based module. Taken together, these findings suggest the institutional value put forward in Section 8 depends on addressing the gaps, the value that this framework provides at this point is the set of measures against which the addressing of the gaps earlier can be tested.

7. Ethical, Privacy and Responsible-AI Considerations

A platform profiling students' ethical understanding, confidence and misconception risk carries its own obligations, treated here as a first-order design requirement. Proposed safeguards include that informed consent is obtained before assessment, that data is minimized to items relevant to competency modelling, that data is anonymized before aggregate reporting, that access to identifying profiles and the Learning Passport is restricted (and should not be routinely visible to instructors), that assessment criteria are transparent, that guidance is supported by human oversight, that dashboards are interpreted by humans, that bias is monitored on an ongoing basis (as profiling that was tuned on one cohort may perform differently across linguistic backgrounds, disciplines or prior AI exposure), and that a student right to review recommendations is established, and that data is appropriately stored.

8. Conclusion

Students' use of AI has progressed faster than formal teaching on responsible AI use. Interventions are frequently generic and fail to pinpoint the student, discipline, and learning areas that need assistance. They also do not differentiate between students who do not know how to explain an ethical principle and students who know the ethical principle but are not able to make it a part of their practice. KBG-Ruwad tackles this by creating a clear diagnostic framework of the knowledge–behavior gap, adding confidence into a misconception-risk indicator, and offering discipline-specific and non-punitive guidance. The novelty of it is that the components are integrated in a systematic way around a measurable discrepancy between what the students know and what they do. The results and findings are articulate as a testable specification and state the measurement assumptions, gaps in implementation and challenges to be overcome in operation that require empirical assessment. The proposed order of expert review, prototyping, ethical approval and pilot testing is therefore proposed as a way of moving KBG-Ruwad from conceptual specification to evidence based validation. If the framework is validated, it can equip higher education institutions in Oman with a concrete guide to enhance and build the competence of responsible-AI and lay the groundwork for the workforce development goals of Oman Vision 2040.

Acknowledgment

The research leading to these results from the undergraduate research grant approved by Research and Innovation Authority (RIA) of the Sultanate of Oman under the Block Funding Programme (BFP), Agreement No. BFP/URG/ICT/25/390.

Ethics Declaration

This study proposes a conceptual framework and does not involve the collection, use, or analysis of primary or secondary data. No human participants, personal information, clinical records, or identifiable data were used. Therefore, institutional ethical approval was not required.

AI Declaration

AI-assisted tools were utilized only to improve linguistic quality and readability and creating the figure from my proposed work.

References

Adel, A., Ahsan, A. and Davison, C. (2024) "ChatGPT promises and challenges in education: computational and ethical perspectives", Education Sciences, 14(8), p.814.

Akintola, A.S., Akintayo, M. and Kadri, T. (2024) "Adaptive AI systems in education: real-time personalised learning pathways for skill development", Journal of Artificial Intelligence, Machine Learning and Data Science, 1, pp.2489-2494.

Barnes, E. and Hutson, J. (2024) "Navigating the ethical terrain of AI in higher education: strategies for mitigating bias and promoting fairness", Forum for Education Studies, 2(2).

Benayoune, A., Slimi, Z. and Al Habsi, A. (2026) "Artificial intelligence policy challenges and institutional readiness in Omani higher education", Discover Education, 5, p.86.

Casal-Otero, L., Catala, A., Fernandez-Morante, C., Taboada, M., Cebreiro, B. and Barro, S. (2023) "AI literacy in K-12: a systematic literature review", International Journal of STEM Education, 10(1), pp.1-17.

Fortuna, A., Prasetya, F., Samala, A.D., Rawas, S., Criollo-C, S., Kaya, D., Raihan, M., Andriani, W., Safitri, D. and Nabawi, R.A. (2025) "Artificial intelligence in personalized learning: a global systematic review of current advancements and shaping future opportunities", Social Sciences and Humanities Open, 12, p.102114.

Fu, Y. and Weng, Z. (2024) "Navigating the ethical terrain of AI in education: a systematic review on framing responsible human-centered AI practices", Computers and Education: Artificial Intelligence, 7, p.100306.

Furman, G.D. (2024) "Enhancing engineering education: the role of artificial intelligence in personalizing learning and outcomes", in Proceedings of the International Conference on Big Data Engineering and Education (BDEE). IEEE, pp.61-65.

Georgieva, M. and Stuart, J. (2025) "Ethics is the edge: the future of AI in higher education", EDUCAUSE Review, 24 June.

Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) "Design science in information systems research", MIS Quarterly, 28(1), pp.75-105.

Ma, M., Ng, D.T.K., Liu, Z. and Wong, G.K.W. (2025) "Fostering responsible AI literacy: a systematic review of K-12 AI ethics education", Computers and Education: Artificial Intelligence.

MTCIT (2024) Oman AI and Digital Future Program (2024-2026). Muscat: Ministry of Transport, Communications and Information Technology, Sultanate of Oman.

Natarajan, J.R., Ayyappan, J.P., Sharma, K., Krishnasamy, T. and Seshan, V. (2026) "Faculty readiness for AI-driven digital transformation in Omani higher education: examining the attitudes-usage gap", Frontiers in Education, 11, p.1774840.

Ogbo, N.B., Song, Z., Ghareeb, S. and Han, T.A. (2026) Mathematical modelling of ethical AI use in higher education: a coordination game framework for future-facing learning. Preprint. Teesside University.

Oman Vision 2040 (2019) Oman Vision 2040. Muscat: Ministry of Energy and Minerals, Sultanate of Oman.

Sajja, R., Sermet, Y., Cikmaz, M., Cwiertny, D. and Demir, I. (2024) "Artificial intelligence-enabled intelligent assistant for personalized and adaptive learning in higher education", Information, 15(10), p.596.

Sujee, S., Othmani, S.H. and Dalwai, T. (2024) "Toward a knowledge-based economy: empowering higher education in Oman with AI for Vision 2040 achievement", in Mansour, N. and Bujosa Vadell, L.M. (eds) Finance and Law in the Metaverse World. Cham: Springer, pp.133-142.

Tarisayi, K.S. (2024) "Strategic leadership for responsible artificial intelligence adoption in higher education", CTE Workshop Proceedings, 11, pp.4-14.

UNESCO, Miao, F. and Shiohira, K. (2024) AI Competency Framework for Students. Paris: UNESCO.

Wiese, L.J., Patil, I., Schiff, D.S. and Magana, A.J. (2025) "AI ethics education: a systematic literature review", Computers and Education: Artificial Intelligence, 8, p.100405.

Yan, Y., Liu, H. and Chau, T. (2025) "A systematic review of AI ethics in education: challenges, policy gaps, and future directions", Journal of Global Information Management, 33(1), pp.1-50.

Zawacki-Richter, O., Marin, V.I., Bond, M. and Gouverneur, F. (2019) "Systematic review of research on artificial intelligence applications in higher education - where are the educators?", International Journal of Educational Technology in Higher Education, 16(1), p.39.