[Work in Progress] A Framework for AI-Situated Ethical Reasoning in Engineering Education

Dr. Douglas Yung, Syracuse University

Doug Yung is a Teaching Professor and Director of the Undergraduate Program in Biomedical Engineering at Syracuse University. His teaching and scholarship focus on experiential learning, engineering design education, and ethics in emerging technologies. He has developed and implemented several innovative, hands-on assessment models that integrate computation, teamwork, and professional competencies into early engineering courses. Dr. Yung's recent work explores collaborative and competency-based assessment in first-year engineering, embedding Kolb's experiential learning cycle and learner engagement frameworks into design-focused exams and lab experiences.

Beyond the classroom, Dr. Yung is active in K–12 STEM outreach, diversity and inclusion initiatives, and faculty development for evidence-based pedagogy. He serves as an ASEE Campus Representative, a Meredith Professorship Faculty Fellow, and a member of multiple university and national committees dedicated to advancing engineering education. His work reflects a commitment to cultivating curiosity, creativity, and reflective practice in the next generation of engineers.

A Framework for AI-Situated Ethical Reasoning in Engineering Education (Work in Progress)

Abstract

As artificial intelligence becomes more deeply embedded in engineering practice, engineering ethics education must move beyond general moral awareness toward forms of reasoning that address the distinctive socio-technical demands of AI-enabled systems. Existing approaches to ethics assessment in engineering education, including the Defining Issues Test-2, the Engineering and Science Issues Test, and the Engineering Ethical Reasoning Instrument, provide important foundations but were not designed for AI-related contexts involving subgroup disparities, opacity, privacy risks, and responsibility under partial automation. In response to this gap, this paper proposes AI-Situated Ethical Reasoning (ASER) as a conceptual framework for engineering education.

ASER is organized around three interrelated dimensions: ethical recognition, or the ability to identify ethically salient features of an AI-enabled engineering problem; ethical evaluation, or the ability to interpret and weigh competing principles, stakeholder consequences, and technical evidence; and responsible action, or the ability to articulate justified, context-sensitive responses. The paper situates this framework within four strands of scholarship: engineering ethics education, ethics assessment research, AI ethics and responsible-AI literature, and professional frameworks including ABET Student Outcome 4 and the NIST AI Risk Management Framework. It also illustrates the educational use of the framework through a visual model, a synthesis table, and a worked example involving an AI-enabled skin lesion triage tool. At this stage, ASER has been developed as a conceptual framework with accompanying instructional representations, including a visual model, synthesis table, and worked example; it has not yet been implemented as a validated classroom intervention or psychometric assessment instrument.

Rather than presenting a validated instrument, this paper offers a conceptual foundation for engineering ethics education by clarifying how students can recognize ethical issues, evaluate competing obligations, and justify responsible action in AI-enabled engineering contexts.

1. Introduction and Conceptual Gap

Artificial intelligence is reshaping both engineering practice and engineering education in ways that make ethical judgment more complex, more data-dependent, and more socially consequential. Engineers increasingly work with AI-enabled systems for prediction, classification, optimization, and decision support, while engineering programs are simultaneously grappling with how AI should be taught, used, and critically examined in higher education [1]. As a result, engineering education must prepare students not only to use AI-related tools competently, but also to reason responsibly about their assumptions, limitations, and downstream human impacts.

This need is not peripheral to engineering education. ABET Student Outcome 4 requires that graduates be able to recognize ethical and professional responsibilities in engineering situations and make informed judgments that consider the impact of engineering solutions in global, economic, environmental, and societal contexts [2]. In parallel, the NIST AI Risk Management Framework identifies core characteristics of trustworthy AI systems, including validity and reliability, safety, security and resilience, accountability and transparency, explainability and interpretability, privacy enhancement, and fairness with harmful bias managed, while emphasizing that these characteristics are socio-technical and context-dependent [3]. Together, these frameworks make clear that ethical competence in contemporary engineering must include the ability to reason about AI-related risks and responsibilities in context. At the same time, they stop short of specifying how such reasoning should be conceptualized for educational purposes or what it should look like in assessable form.

That omission matters because ethical reasoning in AI-enabled engineering contexts is not simply a matter of applying traditional ethics language to new tools. AI-related harms are often socio-technical, emerging through interactions among datasets, models, interfaces, institutional practices, deployment settings, and stakeholder power relations rather than from a single isolated technical decision [4]. Questions of fairness, accountability, and human impact are therefore tied not only to design intent, but also to data quality, model behavior, interpretation, implementation, and use. Buolamwini and Gebru, for example, demonstrated substantial intersectional performance disparities in commercial gender-classification systems, showing how seemingly strong technical performance can mask ethically significant inequities across demographic groups [5]. In such cases, ethical reasoning requires more than general moral awareness. It requires the ability to identify ethically salient features of AI-enabled systems, interpret their human consequences, and justify responses in light of competing values, evidence, and contextual constraints.

The engineering ethics education literature suggests that the field has not yet fully conceptualized or assessed this kind of reasoning. Prior reviews have found substantial variation in how ethics is taught, integrated, and evaluated across engineering programs, indicating that the assessment landscape remains heterogeneous and still developing [6], [7]. Several established instruments have made important contributions to this broader effort. The Defining Issues Test-2 has long served as a prominent measure of moral judgment development [8]. The Engineering and Science Issues Test adapted moral reasoning assessment to science and engineering contexts through discipline-relevant dilemmas [9]. More recently, the Engineering Ethical Reasoning

Instrument has supported the study of changes in students' ethical reasoning following instructional interventions [10]. These instruments have advanced the field substantially, but they were not designed specifically for AI-enabled engineering situations involving data bias, subgroup disparities, privacy risks, opacity, human oversight, and responsibility under partial automation.

The need for sharper conceptualization is also reflected in the emerging literature on AI ethics education. Recent scholarship shows growing attention to AI ethics as a formal area of instruction, with recurring emphasis on fairness, accountability, transparency, privacy, and governance, while also identifying continuing challenges in curriculum design, pedagogy, and assessment [10], [11], [12]. Recent empirical work also suggests that the field is beginning to move beyond general calls for AI ethics education toward more structured interventions aimed at student learning. In a study of science and engineering graduate students, Usher and Barak found that an online explicit-reflective AI ethics module improved students' ethical knowledge and awareness while also strengthening their ability to identify and respond to ethical problems in AI-related contexts [13]. Their findings are especially relevant here because they suggest that ethical competence in AI-rich settings involves more than abstract awareness alone. It also includes the ability to identify, interpret, and respond to ethically significant features of technical systems.

Taken together, this literature suggests that although there is increasing agreement that students need preparation in AI ethics, there is still less clarity about what should count as ethically competent reasoning in AI-rich engineering settings. The unresolved issue, therefore, is not only that AI introduces new ethical topics into engineering education, but that it changes the structure of the reasoning students must perform. What remains underdeveloped is an explicit construct that captures how technical evidence, socio-technical context, stakeholder impacts, and professional responsibility are integrated in ethical judgment about AI-enabled engineering systems. Without such a construct, educational efforts risk becoming diffuse, treating AI ethics as a set of cautionary examples, a list of responsible AI principles, or a general appeal to professionalism without a clear account of the reasoning process students are expected to develop.

This paper addresses that gap by proposing a framework for AI-situated ethical reasoning (ASER) in engineering education. In this paper, ASER refers to the ability to recognize ethically significant features of AI-enabled engineering problems, evaluate competing principles, stakeholder impacts, and trade-offs, and articulate responsible actions or design choices that are technically informed and context-sensitive. The framework is organized around three interrelated dimensions: ethical recognition, ethical evaluation, and responsible action. Rather than presenting an empirically validated instrument, this paper offers a conceptual foundation for understanding what should be assessed when students reason ethically about AI-integrated design and decision-making. The paper makes two primary contributions. First, it argues that ASER should be treated as a distinct extension of ethical competence in engineering education rather than as a minor variation of general moral reasoning. Second, it presents a framework intended to support future assessment development, instructional design, curricular integration, and accreditation-aligned educational practice by connecting ethical reasoning in AI contexts to both ABET Student Outcome 4 and the trustworthiness concerns articulated in the NIST AI Risk

Management Framework [2], [3]. The present manuscript provides the ASER construct, a three-dimensional framework, professional alignment table, and worked example as developed scholarly products; empirical classroom testing and instrument validation are identified as future work.

2. Literature Foundations

The present framework sits at the intersection of four strands of scholarship that together justify the need for a more explicit construct of ethical reasoning in AI-enabled engineering contexts. First, engineering ethics education is a mature but internally varied field. Reviews show substantial diversity in how ethics is framed, taught, and assessed across programs, reflecting broad agreement that ethical formation matters but less agreement about precisely what ethical competence should entail in practice [6], [7]. This variability is especially important here because AI-enabled engineering intensifies longstanding socio-technical questions rather than merely adding new case topics. If the boundaries of ethical competence are already somewhat unsettled in engineering education, then AI-rich contexts make the need for conceptual clarity even sharper.

Second, the ethics assessment literature demonstrates that ethics-related learning can be studied systematically, but it also shows that assessment is always constrained by the construct being measured. Widely used instruments such as the Defining Issues Test-2, the Engineering and Science Issues Test, and the Engineering Ethical Reasoning Instrument represent important advances, moving ethics assessment from abstract moral judgment toward more contextually grounded and discipline-relevant reasoning [8], [10], [11]. Recent efforts have also begun to operationalize responsible AI more directly within engineering education. Mirsch et al., for example, developed and validated the RAISE scale to assess attitudes toward responsible AI in engineering, focusing on dimensions such as do-no-harm, transparency, and privacy [14]. Kong and Zhu similarly developed the AI Ethical Awareness Scale and reported evidence of validity together with gains in students' ethical awareness following a problem-solving-based AI literacy intervention [15]. These developments are important because they suggest that the field is beginning to generate AI-specific assessment tools rather than relying entirely on general ethics instruments. At the same time, these newer measures still tend to focus on dispositions or awareness rather than the fuller process of ethical reasoning in context. The issue is therefore not that prior assessment work is inadequate, but that existing and emerging tools still tend to capture important fragments of ethical competence rather than the integrated reasoning process required in AI-enabled engineering problems. Table 1 summarizes the theoretical basis, target population, construct focus, and key limitations of these five instruments in relation to the demands of ASER in engineering contexts. All five instruments draw on Neo-Kohlbergian moral development theory or related traditions, and all use scenario-based or survey formats validated through confirmatory factor analysis or pre-post intervention designs.

Table 1. *Comparison of existing ethics assessment instruments in relation to AI-situated engineering contexts. The proposed AI-Situated Ethical Reasoning (ASER) framework (final row) is included as a conceptual contrast; no empirical validation is reported in this paper.*

Instrument	Target Population	Construct Measured	AI-Specific Content	Assessment Approach

Defining Issues Test-2 (DIT-2) [8]	General college and adult populations	Moral judgment schema development (personal interest, maintaining norms, post-conventional reasoning)	No	Dilemmas involve interpersonal or civic situations with no technical content; measures schema preference rather than contextual reasoning across technical and social domains
Engineering and Science Issues Test (ESIT) [9]	Engineering and science undergraduates	Moral reasoning quality in discipline-relevant dilemmas; sensitivity to ethics pedagogy	No	Dilemmas predate AI-enabled engineering contexts; does not address dataset bias, subgroup disparities, model opacity, or distributed socio-technical harms
Engineering Ethical Reasoning Instrument (EERI) [10], [16]	Engineering undergraduates and graduate students	Ethical reasoning quality in engineering design contexts; sensitivity to educational interventions	No	Scenarios drawn from conventional domains (professional responsibility, safety, conflicts of interest); does not distinguish recognition, evaluation, and action as separable dimensions
RAISE Scale [14]	Engineering students and professionals	Attitudes and dispositions toward responsible AI in engineering (do-no-harm, transparency, privacy, fairness)	Yes	Measures value endorsement rather than reasoning process; attitude endorsement does not capture how students interpret evidence, weigh trade-offs, or justify action under uncertainty
AI Ethical Awareness Scale (AIEAS) [15]	Secondary and university students; not engineering-specific	AI ethical awareness across three dimensions: human autonomy, beneficence, and fairness	Yes	Focuses on recognition only; does not capture evaluation of competing trade-offs or articulation of responsible action; validated outside engineering-specific contexts
Proposed: AI-Situated Ethical Reasoning (ASER)	Engineering students in AI-integrated courses	Integrated ethical reasoning across three dimensions: ethical recognition, ethical evaluation, and responsible action	Yes, designed for AI-enabled engineering contexts	Conceptual framework only; assessment instruments based on this construct remain to be developed and empirically validated

As the table illustrates, the instruments form a progression from general moral judgment toward AI-specific dispositions, yet none was designed to capture the integrated reasoning process that spans recognition, evaluation, and responsible action, as AI-enabled engineering problems require.

Third, the AI ethics and responsible AI literature clarifies why that shift matters. A central insight in this scholarship is that AI systems are socio-technical rather than purely technical artifacts, meaning that their ethical significance depends not only on model performance but also on data provenance, deployment conditions, interfaces, institutional practices, and power relations [4]. Empirical work such as Buolamwini and Gebru's study of intersectional disparities in commercial gender-classification systems shows how apparently strong technical performance can conceal ethically significant subgroup harms [5]. At the same time, scholarship on AI ethics education increasingly emphasizes fairness, accountability, transparency, privacy, governance, and related concerns as matters that should be taught in integrated rather than peripheral ways [11], [12], [17]. Recent empirical work further suggests that AI ethics education is moving toward more concrete institutional and pedagogical models. Usher et al., for instance, found that science and engineering graduate students largely viewed universities as bearing essential or shared responsibility for fostering AI ethics and identified strategies such as standalone ethics courses, curricular integration, and stronger institutional leadership [18]. Broader review work likewise suggests that the field is expanding, but still marked by gaps in empirical validation, policy implementation, and stakeholder coverage [19]. Together, these studies reinforce that instruction can improve students' ethical knowledge, awareness, and problem-solving in AI contexts while also showing that the field is still developing clear educational models of what students should actually be able to do when reasoning ethically about AI-enabled systems.

Fourth, professional frameworks help define the boundaries of contemporary responsibility but do not, by themselves, provide an educationally operationalized model of student reasoning. ABET Student Outcome 4 establishes that engineering students must be able to recognize ethical and professional responsibilities and make informed judgments about the broader impacts of engineering solutions [2]. The NIST AI Risk Management Framework sharpens this expectation for AI-rich settings by identifying trustworthiness concerns such as validity, safety, transparency, explainability, privacy, and fairness while explicitly framing them as socio-technical and context-dependent [3]. Taken together, these four strands of literature support the same conclusion: engineering ethics education offers the pedagogical foundation, ethics assessment research offers methodological precedent, AI ethics scholarship clarifies the distinctive socio-technical demands of the problem space, and professional frameworks identify why these issues matter for contemporary engineering practice. What remains underdeveloped is an education-oriented construct that integrates these insights into a coherent account of how students should recognize ethically salient features of AI-enabled problems, evaluate competing considerations in light of technical and social evidence, and articulate responsible action.

3. Framework for AI-Situated Ethical Reasoning (ASER)

The preceding sections established the central gap motivating this paper: existing scholarship does not yet provide a sufficiently explicit construct for understanding how engineering students reason ethically in AI-enabled contexts. Engineering ethics education offers pedagogical foundations, ethics assessment research offers methodological precedent, AI ethics scholarship clarifies the socio-technical nature of the problem space, and professional frameworks such as ABET and the NIST AI Risk Management Framework identify domains of contemporary responsibility. What remains underdeveloped is a framework that brings these strands together into a coherent account of the kind of reasoning students should demonstrate when confronting AI-rich engineering problems. This section addresses that need by proposing ASER as a distinct conceptual construct for engineering education.

3.1 Definition of ASER

In this paper, **ASER** refers to the ability to identify ethically significant features of AI-enabled engineering situations, interpret and weigh competing principles and stakeholder consequences in light of technical evidence and socio-technical context, and articulate responsible actions or design choices under conditions of uncertainty. The construct is "situated" in two related senses. First, it is situated in disciplinary practice, referring to the forms of judgment engineers must exercise when working with data-driven models, algorithmic systems, and automated decision-support tools. Second, it is situated in context, because the ethical significance of an AI system depends on factors such as problem framing, data provenance, deployment environment, stakeholder vulnerability, institutional use, and the distribution of harms and benefits [3], [4].

This definition builds on prior work in engineering ethics while extending it. Existing assessment traditions have often focused on general moral judgment, professional responsibility, or ethical reasoning in more conventional engineering dilemmas [8], [9], [10]. These remain important foundations. However, AI-enabled engineering problems often require students to reason across technical and social domains simultaneously, for example when judging subgroup performance differences, structural bias in data, limits on accountability created by system opacity, privacy risks, or designs that are technically efficient but ethically unacceptable for affected communities. In such settings, ethical reasoning is not separate from engineering judgment; it is part of how engineering judgment must be exercised.

This framing also distinguishes ASER from adjacent constructs such as ethical awareness and responsible-AI attitudes. Recent work has begun developing AI-specific measures of ethical awareness and responsible-AI attitudes, which is an important advance [14], [15]. Yet awareness that fairness, privacy, or transparency matters does not by itself capture the fuller reasoning process at issue here. A student may recognize that an AI-related problem has ethical content without being able to interpret competing evidence, weigh stakeholder consequences, or justify a defensible course of action. Recent work on embedded ethics and curriculum integration reinforces this point. Lim et al. found that students often struggle to recognize the social and ethical relevance of technical content when that relevance is not made explicit [20]. Ferdman and Ratti likewise argue that embedded ethics efforts can remain conceptually thin when their learning goals are underspecified [21]. Esposito et al. further show that AI ethics instruction can be embedded directly within technical training through interdisciplinary and co-designed approaches [22]. Together, these studies support the central claim of this paper: ASER is not merely the ability to spot ethical issues, but an integrated form of judgment connecting technical interpretation, stakeholder analysis, ethical justification, and action-oriented responsibility.

3.2 A three-dimensional model

To make this construct more educationally usable, ASER is organized here into three interrelated dimensions: ethical recognition, ethical evaluation, and responsible action. These dimensions are analytically distinct but pedagogically connected. Together, they describe a movement from identifying ethically salient features of a situation, to interpreting and weighing those features, to articulating justified responses.

The dimensions are not intended to represent rigid stages or a simple checklist. Rather, they identify major components of reasoning that should become visible when students engage ethically with AI-enabled engineering problems. This framing is especially important for educational use. It provides greater precision than a general appeal to “ethical awareness,” while avoiding a reductive scoring logic that would flatten ethical judgment into a formula. It also aligns with prior work showing that ethical decision-making in design is contextual and interpretive, shaped by how students frame technical problems, stakeholders, and possible courses of action [23].

3.2.1 Ethical recognition

The first dimension, ethical recognition, refers to the ability to perceive that an AI-enabled engineering situation contains ethically relevant features and to identify what those features are. In AI-rich contexts, this may include noticing concerns related to fairness, privacy, accountability, transparency, consent, safety, or differential stakeholder impact, including when such concerns are embedded in technical choices such as datasets, metrics, interpretability, or deployment assumptions.

This dimension matters because students cannot reason well about ethical problems they do not first perceive as ethical problems. AI-enabled systems complicate this task because ethically significant features are often dispersed across technical, organizational, and social layers rather than presented as explicit moral dilemmas [6]. Buolamwini and Gebru's work illustrates this clearly: apparently strong commercial system performance concealed serious intersectional disparities across demographic groups [5]. Ethical recognition therefore requires more than noticing that "ethics is involved." It requires identifying how technical choices may be ethically loaded and why they matter in relation to stakeholders, risk, and context.

In the present framework, ethical recognition entails both detection and interpretation. Students must not only name relevant concerns, but also explain why those concerns are ethically significant in the engineering situation at hand [4], [6].

3.2.2 Ethical evaluation

Recognition alone, however, is insufficient. The second dimension, ethical evaluation, refers to the ability to interpret ethically salient features of a problem, weigh competing principles and stakeholder interests, and justify judgments using both ethical reasoning and relevant evidence. If ethical recognition answers the question *What here is ethically significant?*, ethical evaluation answers *How should these concerns be interpreted and weighed?*

In AI-enabled engineering contexts, ethical evaluation often requires reasoning through trade-offs that do not admit purely technical answers. A system may improve efficiency while reducing transparency, or perform well overall while disadvantaging a subgroup. Evaluating such situations requires more than identifying stakeholders or naming principles. It requires interpreting how ethical commitments interact with technical evidence, institutional conditions, uncertainty, and competing values [3], [23].

This dimension is especially important because ethics instruction can too easily collapse into a list of values or duties without addressing how those values conflict in practice. AI contexts make this especially visible, since fairness, privacy, explainability, accountability, safety, and reliability are often mutually entangled [3]. Ethical evaluation therefore requires students to reason through competing considerations and to treat evidence, such as performance disparities or documentation gaps, as part of moral judgment rather than as separate technical details.

3.2.3 Responsible action

The third dimension, responsible action, refers to the ability to articulate and justify a course of action, design revision, oversight mechanism, or professional response that is ethically defensible in light of the recognized issues and evaluated trade-offs. This dimension moves beyond diagnosis toward the question *What should be done, by whom, and why?*

In engineering education, this dimension matters because ethics is not only about analysis but also about professional responsibility in practice. In AI-enabled contexts, responsible action may involve redesign, additional oversight, stakeholder engagement, limited deployment, or refusal to proceed in the system's current form. The key point is not simply naming an action, but justifying it in light of the recognized issues and evaluated trade-offs.

This dimension aligns closely with both ABET Student Outcome 4 and the NIST AI RMF, which emphasize informed judgment, governance, accountability, and the management of harmful bias and other risks [2], [3]. Recent work on embedded and integrated ethics likewise reinforces that responsible practice must be tied to the technical and institutional conditions in which engineers actually work [21], [22]. The present framework extends these expectations by treating responsible action as an observable dimension of student reasoning rather than only as a professional ideal.

3.3 Interdependence of the three dimensions

Although these three dimensions are presented separately for conceptual clarity, they are not meant to function as isolated stages. In practice, ethical recognition, ethical evaluation, and responsible action are iterative and mutually informing. Recognition may deepen as evaluation reveals new stakeholder implications, while evaluation and action may expose further ethical tensions. The framework should therefore be understood as a structured model of reasoning rather than a rigid sequence. Accordingly, claims in this paper are limited to conceptual development and illustrative educational use, not demonstrated learning gains or validated measurement performance.

This interdependence matters for educational design. It suggests that teaching and assessment should neither treat ethical competence as a single undifferentiated trait nor fragment it into disconnected components. Students need opportunities to identify ethically salient features, reason through competing claims, and articulate justified responses in relation to the same socio-technical problem. Embedded approaches are especially promising because tying ethics directly to technical content makes recognition, evaluation, and action more visible as intertwined parts of engineering judgment rather than as separate classroom exercises [20], [21], [22].

3.4 Ethical traditions as interpretive lenses

To deepen the interpretive value of ASER, this paper draws on three ethical traditions that commonly inform engineering ethics education: utilitarian reasoning, deontological reasoning, and virtue-based reasoning. These are not presented as separate dimensions of ASER, nor as competing moral algorithms. Rather, they function as interpretive lenses that help clarify how students may reason within and across the dimensions of ethical recognition, ethical evaluation, and responsible action.

A utilitarian lens emphasizes consequences, aggregate welfare, and trade-offs across stakeholder groups. Within ASER, it is especially relevant to ethical evaluation and responsible action because it helps students reason about benefits, harms, and societal consequences when technical decisions affect groups differently. At the same time, utilitarian reasoning alone can be insufficient if aggregate benefit is used to justify serious burdens on marginalized populations or to obscure concerns about rights, fairness, or accountability.

A deontological lens emphasizes duties, rights, obligations, and respect for persons. In AI contexts, it is especially useful for identifying actions that may be ethically problematic even when they appear efficient or beneficial overall, such as using non-consensual data, violating privacy expectations, concealing system limitations, or deploying opaque systems in high-stakes settings. Within ASER, deontological reasoning supports ethical recognition by highlighting morally non-negotiable features of a situation and supports ethical evaluation by foregrounding duties and rights that should not be erased by cost-benefit analysis.

A virtue-based lens emphasizes character, judgment, and practical wisdom. Rather than asking only what rule applies or what outcome maximizes welfare, it asks what a responsible, honest, just, and trustworthy engineer would do in a given situation. This perspective is especially valuable in ambiguous contexts where rules do not fully specify the right course of action and responsible conduct depends on attentiveness, humility, courage, and integrity. Within ASER, virtue-based reasoning is relevant across all three dimensions because AI-enabled problems often involve uncertainty, incomplete information, and evolving standards.

3.5 Why a plural ethical framework matters

Using these traditions together strengthens ASER in two ways. First, it avoids reducing ethical reasoning to a single mode of justification. AI-enabled engineering problems are rarely captured fully by one ethical vocabulary alone. Consequences matter, but so do duties, rights, professional obligations, and habits of responsible judgment. Second, a plural framework is pedagogically useful because it helps instructors and researchers interpret student reasoning without forcing it into a single philosophical template.

In this sense, ASER is not intended to prescribe one correct moral theory. Rather, it clarifies the dimensions of reasoning that should be visible when students engage ethically with AI-enabled engineering situations. Ethical recognition identifies what is morally salient, ethical evaluation interprets and weighs competing considerations, and responsible action articulates a justified response. Utilitarian, deontological, and virtue-based lenses then help illuminate how those dimensions may be enacted in practice. Taken together, ASER conceptualizes ethical reasoning in AI-enabled engineering as an integrated form of judgment that is theoretically grounded, educationally usable, and aligned with the socio-technical realities of contemporary practice.

4. From Conceptual Framework to Educational Use

This section moves ASER from conceptual definition toward educational use through three complementary representations: a visual model, a synthesis table, and a worked example. Together, these elements show how ASER may support teaching, assessment, and curriculum design without turning the framework into a rigid scoring scheme.

4.1 Visual representation of the framework

Figure 1 summarizes ASER as a recursive rather than strictly linear form of judgment. Although ethical recognition, ethical evaluation, and responsible action can be described analytically as distinct dimensions, in practice students may move iteratively among them as new evidence, stakeholder concerns, or contextual constraints emerge [3], [4]. The figure therefore reinforces a central claim of this paper: ethical reasoning in AI-enabled engineering contexts develops through repeated cycles of recognition, interpretation, evaluation, action, and reassessment rather than through a simple linear progression.

Figure 1. *AI-situated ethical reasoning (ASER) as an iterative framework for AI-enabled engineering contexts. The figure illustrates ethical reasoning as a recursive process linking ethical recognition, contextual interpretation, ethical evaluation, responsible action, and reassessment. Rather than proceeding in a strictly linear sequence, students may move among these phases as new evidence, stakeholder concerns, or socio-technical constraints emerge. Collectively, the framework shows how ethical recognition, ethical evaluation, and responsible action mutually inform one another in the analysis and governance of AI-enabled engineering decisions.*

4.2 Synthesis of dimensions, professional alignment, and possible evidence

Table 2 translates ASER into a more operational form by showing how each dimension may be interpreted educationally, how it aligns with broader professional expectations, and what kinds of student work might provide evidence of it. The table is not intended as a finalized assessment rubric. Rather, it serves as a bridge between the conceptual framework and future instructional or assessment design. More specifically, the table connects each dimension to student demonstration, observable indicators, relevant aspects of ABET Student Outcome 4, selected functions of the NIST AI Risk Management Framework, and plausible forms of assessment evidence [2], [3]. In this way, it helps translate broad professional expectations into educationally interpretable terms without reducing the framework to a checklist or formula.

Table 2. *Synthesis of AI-situated ethical reasoning (ASER) dimensions and their educational use*

Dimension	Student demonstration	Observable indicators	ABET SO4	NIST AI RMF	Assessment evidence
Ethical recognition	Identifies ethically salient features, stakeholders, harms, and consequential technical choices in an AI-enabled problem	Names affected groups; identifies risks, exclusions, or hidden assumptions in data, labels, objectives, or deployment context	Recognizes ethical and professional responsibilities in engineering situations	Map, Govern	Stakeholder maps; case annotations; AI audit worksheets; short reflections; design memos
Ethical evaluation	Weighs competing values, evidence, and stakeholder consequences; justifies why some concerns merit greater weight	Compares trade-offs; connects technical performance to fairness, accountability, public welfare, or trust; supports claims with evidence	Makes informed judgments considering global, economic, environmental, and societal impacts	Measure, Govern	Scenario analyses; structured arguments; oral defenses; rubric-scored case responses
Responsible action	Proposes and justifies responsible responses, including redesign, mitigation, governance, communication, or non-deployment	Recommends safeguards, redesign, monitoring, documentation, escalation, pause, or refusal to deploy	Applies ethical judgment in making decisions with societal and public welfare implications	Manage, Govern	Decision memos; design recommendations; prototype critiques; policy briefs; capstone justifications

4.3 Worked example: An AI-enabled skin lesion triage tool

A brief worked example illustrates how the framework can guide interpretation of student reasoning in a concrete AI-enabled engineering scenario. Consider a student team evaluating whether to deploy a smartphone-based AI tool designed to classify skin lesion images as low-risk or high-risk in support of triage in community clinics. Initial validation results indicate strong overall accuracy, but lower sensitivity for darker skin tones. The model is proprietary and only partially explainable, images are stored through a cloud-based service, and clinicians report that they may rely heavily on the tool in settings with limited dermatology access. Although hypothetical, this scenario reflects documented concerns in AI-enabled systems, especially the possibility that apparently strong overall performance can mask subgroup harms and ethically significant disparities [4], [5]. Figure 2 shows how this case may be interpreted through the three dimensions of the framework.

The scenario also allows instructors to surface familiar engineering ethics concerns in AI-specific form, including selective reporting of validation results, inappropriate exclusion of subgroup performance data, overstatement of model accuracy, and inadequate protection of patient images or metadata. In this way, the example connects responsible-AI concerns to established engineering ethics topics such as honesty in reporting, privacy, consent, public welfare, and professional accountability.

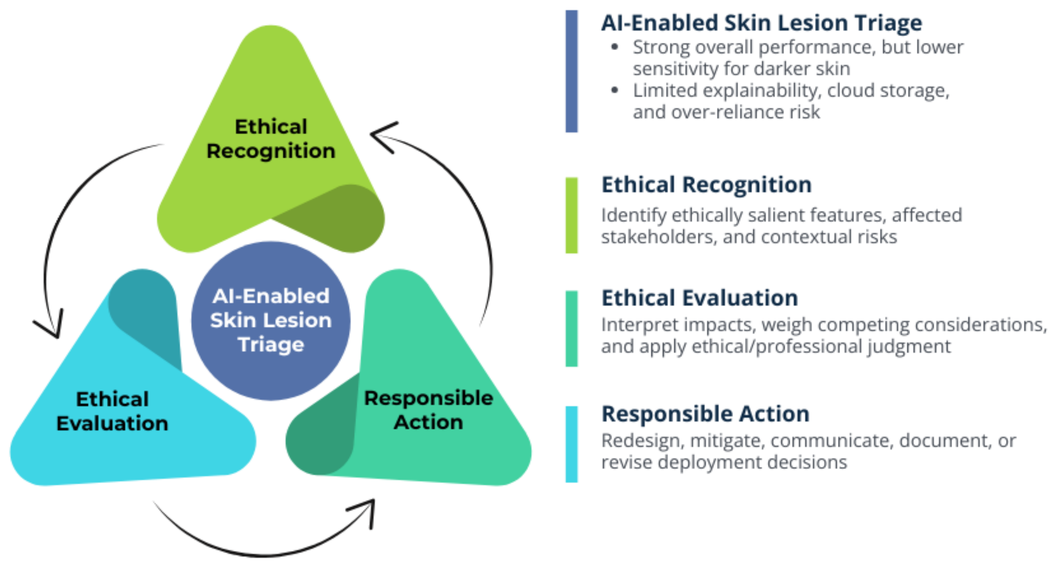

Figure 2. *Worked example of ASER in a skin lesion triage scenario. The figure shows how a smartphone-based AI skin lesion triage tool can be analyzed through ethical recognition, ethical evaluation, and responsible action, highlighting the interdependence of technical evidence, stakeholder consequences, and justified response.*

Within this scenario, ethical recognition would involve identifying that the problem is not purely technical. Students would need to notice that apparently strong aggregate model performance may conceal subgroup harms, that lower sensitivity for darker skin tones raises concerns about inequitable false negatives, and that cloud-based image storage and clinician reliance raise concerns about privacy, consent, and overtrust. Ethical recognition here means identifying what is ethically significant and explaining why those features matter in relation to stakeholder vulnerability, risk, and deployment context.

Ethical evaluation would involve weighing these concerns in light of both technical evidence and ethical considerations. Students might consider whether the access benefits of faster triage justify deployment despite unequal subgroup performance, whether fairness concerns should outweigh efficiency gains, and whether limited explainability is acceptable in this clinical context. Strong ethical evaluation would require students to interpret the available validation evidence, reason through competing values, and explain how stakeholder consequences should influence the decision.

Responsible action would involve articulating and defending a course of action. A strong response might recommend against unrestricted deployment in the system's current form and instead propose a limited pilot with stronger human oversight, additional subgroup validation, clearer documentation of known limitations, and explicit patient consent procedures. In either case, the emphasis is not simply on naming a preferred action, but on justifying that action in light of the recognized issues and evaluated trade-offs.

This worked example illustrates how the three dimensions operate in relation rather than isolation. Ethical recognition surfaces the morally significant features of the case, ethical evaluation interprets and weighs those features in context, and responsible action translates that judgment into a defensible engineering response. The example also shows that ASER requires more than general ethical awareness. Students must interpret technical evidence, situate it within a socio-technical context, and justify action under conditions of uncertainty.

5. Implications for Engineering Education and Assessment

The framework proposed here is intended not only to clarify a conceptual gap, but also to support educational use. By distinguishing ethical recognition, ethical evaluation, and responsible action, it gives instructors a more precise way to design AI ethics learning experiences and interpret student reasoning. Rather than treating AI ethics as a list of topics such as bias, privacy, or transparency, educators can use the framework to structure activities that ask students to identify ethically salient features of a problem, weigh competing values and evidence, and justify a response. This is especially important because students do not always recognize the ethical significance of technical content unless it is made explicit, and because embedded ethics is most effective when the intended form of reasoning is clearly specified [20], [21].

The framework may also strengthen classroom discourse and assessment. It can help instructors design prompts that move students beyond vague claims that a system should be “fair” or “transparent” by asking what is ethically significant, how competing concerns should be weighed, and what should be done in response. In this way, the framework supports more analytically rich case discussions, design reviews, system audits, and scenario-based tasks that treat ethical reasoning as part of engineering work rather than as a separate add-on [4], [5], [10], [23]. It may likewise support future assessment design by providing a clearer conceptual basis for scenario prompts, rubrics, and formative feedback. Rather than measuring only general awareness or attitudes, future tools could use the framework to examine whether students demonstrate perceptive recognition, evidence-based evaluation, and action-oriented judgment in AI-enabled engineering contexts.

At the curricular and program level, the framework may help programs generate more specific evidence of ethical reasoning in relation to ABET Student Outcome 4 and emerging responsible-AI expectations [2], [3]. Student work in case analyses, design memos, audit assignments, or reflective writing could be interpreted for evidence of the three dimensions, allowing programs to document not only that ethics was discussed, but how students reasoned in socio-technical contexts shaped by AI. The framework may also support faculty development by providing a shared language for what instructors mean when they ask students to “think ethically” about AI. In this sense, the contribution of the framework is practical as well as conceptual: it offers a structure that can travel from theory into teaching, assessment, and curricular planning.

6. Limitations

This paper is intentionally conceptual in scope. Its primary contribution is to articulate ASER as a distinct construct within engineering education and to propose a framework organized around ethical recognition, ethical evaluation, and responsible action. It does not present classroom intervention data, psychometric evidence, or a validated assessment instrument. The framework should therefore be understood as a conceptual foundation for future work rather than as a completed model.

Several limitations follow from that scope. The framework has not yet been empirically examined through expert review, classroom implementation, or student response data, so its dimensions remain theoretically reasoned rather than empirically established. The paper also does not yet determine whether these dimensions are distinct, observable, and interpretable in practice, nor does it address the full range of validity and measurement questions required for future instrument development, including response process, scoring consistency, internal structure, and relations to other variables [24], [25]. In addition, the framework is intentionally broad across AI-enabled engineering settings. That breadth is conceptually useful, but it may obscure important disciplinary variation across domains such as biomedical engineering, infrastructure, manufacturing, or autonomous systems.

The paper also focuses more on the conceptualization of student reasoning than on the broader instructional ecology needed to cultivate it, including curricular sequencing, faculty preparation, and longitudinal development. Finally, it does not attempt to resolve broader debates about the scope of AI ethics or the adequacy of particular ethical theories. Instead, it offers a structure intended to support future dialogue, refinement, and empirical study. These limitations do not undercut the paper's contribution, but they do define its current stage: this is a conceptual basis for future assessment and pedagogy, not yet a validated instrument or closed theory.

7. Conclusion

As artificial intelligence becomes more deeply embedded in engineering practice, engineering education must expand its understanding of ethical competence. Existing work in engineering ethics education, ethics assessment, and responsible AI provides important foundations, but it does not yet offer a sufficiently explicit construct for understanding how students reason ethically in AI-enabled engineering contexts.

In response, this paper proposed ASER as a conceptual framework for engineering education. The framework defines this construct through three interrelated dimensions, ethical recognition, ethical evaluation, and responsible action, and interprets them through plural ethical lenses rather than a single moral template. In doing so, it offers a framework that is theoretically grounded, responsive to socio-technical realities, and usable for future educational purposes.

The contribution of the paper is therefore both conceptual and practical. Conceptually, it clarifies a construct at the intersection of engineering ethics, AI ethics, and professional judgment. Practically, it provides a foundation for designing learning activities, structuring reflection, informing future assessment, and generating more meaningful evidence of student reasoning in relation to ABET Student Outcome 4 and responsible-AI expectations.

References

[1] C. Liu, "A comprehensive review of applications of AI technologies in higher engineering education," Discov. Educ., vol. 4, no. 1, Art. no. 528, 2025.

[2] ABET, "Criteria for accrediting engineering programs, 2025–2026," ABET, 2024. Accessed: Mar. 11, 2026.

[3] National Institute of Standards and Technology, “Artificial intelligence risk management framework (AI RMF 1.0),” U.S. Dept. Commerce, NIST AI 100-1, 2023.

[4] A. D. Selbst, D. Boyd, S. A. Friedler, S. Venkatasubramanian, and J. Vertesi, “Fairness and abstraction in sociotechnical systems,” in Proc. 2019 Conf. Fairness, Accountability, Transparency, 2019, pp. 59–68.

[5] J. Buolamwini and T. Gebru, “Gender shades: Intersectional accuracy disparities in commercial gender classification,” in Proc. Conf. Fairness, Accountability, Transparency, vol. 81, 2018, pp. 77–91.

[6] D. A. Martin, E. Conlon, and B. Bowe, “A multi-level review of engineering ethics education: Towards a socio-technical orientation of engineering education for ethics,” Sci. Eng. Ethics, vol. 27, no. 5, Art. no. 60, Aug. 2021.

[7] J. L. Hess, A. J. Kerr, A. Lin, and A. Chung, “A systematic review of the 2016 National Academy of Engineering exemplary ethics programs: Revisions to a coding framework,” Sci. Eng. Ethics, vol. 29, no. 6, Art. no. 36, 2023.

[8] J. R. Rest, D. Narvaez, S. J. Thoma, and M. J. Bebeau, “DIT2: Devising and testing a revised instrument of moral judgment,” J. Educ. Psychol., vol. 91, no. 4, pp. 644–659, 1999.

[9] J. Borenstein, M. J. Drake, R. Kirkman, and J. L. Swann, “The engineering and science issues test (ESIT): A discipline-specific approach to assessing moral judgment,” Sci. Eng. Ethics, vol. 16, no. 2, pp. 387–407, 2010.

[10] J. L. Hess, J. Beever, C. B. Zoltowski, L. Kisselburgh, and A. O. Brightman, “Enhancing engineering students’ ethical reasoning: Situating reflexive principlism within the SIRA framework,” J. Eng. Educ., vol. 108, no. 1, pp. 82–102, 2019.

[11] J. Borenstein and A. Howard, “Emerging challenges in AI and the need for AI ethics education,” AI Ethics, vol. 1, no. 1, pp. 61–65, Feb. 2021.

[12] V. Bogina, A. Hartman, T. Kuflik, and A. Shulner-Tal, “Educating software and AI stakeholders about algorithmic fairness, accountability, transparency and ethics,” Int. J. Artif. Intell. Educ., vol. 32, no. 3, pp. 808–833, 2022.

[13] M. Usher and M. Barak, “Unpacking the role of AI ethics online education for science and engineering students,” Int. J. STEM Educ., vol. 11, no. 1, Art. no. 35, Aug. 2024.

[14] M. Mirsch, N. Knoth, B. Schultz, K. Bata, J. Schleiss, and C. Leicht-Scholten, “Measuring attitudes toward responsible AI in engineering – development and validation of the RAISE scale,” Eur. J. Eng. Educ., pp. 1–23, 2026.

[15] S.-C. Kong and J. Zhu, “Developing and validating an artificial intelligence ethical awareness scale for secondary and university students: Cultivating ethical awareness through problem-solving with artificial intelligence tools,” Comput. Educ. Artif. Intell., vol. 9, Art. no. 100447, 2025.

[16] Q. Zhu, C. Zoltowski, M. Feister, P. Buzzanell, W. Oakes, and A. Mead, “The development of an instrument for assessing individual ethical decision-making in project-based design teams: Integrating quantitative and qualitative methods,” presented at the ASEE Annu. Conf. Expo., 2014, p. 24.

[17] L. J. Wiese, I. Patil, D. S. Schiff, and A. J. Magana, “AI ethics education: A systematic literature review,” Comput. Educ. Artif. Intell., Art. no. 100405, 2025.

[18] M. Usher, M. Barak, and S. Erduran, “What role should higher education institutions play in fostering AI ethics? Insights from science and engineering graduate students,” Int. J. STEM Educ., vol. 12, no. 1, Art. no. 51, 2025.

[19] Y. Yan, H. Liu, and T. Chau, “A systematic review of AI ethics in education: Challenges, policy gaps, and future directions,” J. Glob. Inf. Manage., vol. 33, no. 1, pp. 1–50, 2025.

[20] J. H. Lim, B. D. Hunt, N. Findlater, P. T. Tkacik, and J. L. Dahlberg, “‘In our own little world’: Invisibility of the social and ethical dimension of engineering among undergraduate students,” Sci. Eng. Ethics, vol. 27, no. 6, Art. no. 74, Dec. 2021.

[21] A. Ferdman and E. Ratti, “What do we teach to engineering students: Embedded ethics, morality, and politics,” Sci. Eng. Ethics, vol. 30, no. 1, Art. no. 7, Feb. 2024.

[22] G. Esposito et al., “Embedding AI ethics in technical training: A multi-stakeholder pilot module emphasizing co-design and interdisciplinary collaboration at Rome Technopole,” Educ. Sci., vol. 15, no. 10, 2025.

[23] D. J. Corple, C. B. Zoltowski, M. Kenny Feister, and P. M. Buzzanell, "Understanding ethical decision-making in design," J. Eng. Educ., vol. 109, no. 2, pp. 262–280, 2020.

[24] K. A. Douglas and Ş. Purzer, "Validity: Meaning and relevancy in assessment for engineering education research," J. Eng. Educ., vol. 104, no. 2, pp. 108–118, 2015.

[25] N. E. Canney and A. R. Bielefeldt, "Validity and reliability evidence of the engineering professional responsibility assessment tool," J. Eng. Educ., vol. 105, no. 3, pp. 452–477, 2016.