

Ethical Literacy in Computing Education under the Pressure of Generative AI: A Mixed-Methods Study of Students and Educators in Brazil

Felippe Carvalho e Kipman  [University of Brasilia (UnB), Department of Computer Science | felipekipman@gmail.com]

João Rossi  [University of Brasilia (UnB), Department of Computer Science | joaorossi-borba@gmail.com]

Edna Dias Canedo  [Federal University of Pernambuco (UFPE), Computer Science Center (CIn) | ednacanedo@cin.ufpe.br]

Context: The rapid expansion of Artificial Intelligence (AI), particularly Generative AI, has intensified ethical challenges in computing education, exposing tensions between technical training and ethical responsibility. While ethical principles are widely endorsed in international AI governance frameworks, their effective integration into computing curricula remains uneven, especially in Global South contexts. **Goal:** This study investigates the ethical literacy of undergraduate computing students and examines how educators perceive their role, institutional constraints, and pedagogical challenges in fostering ethics education under the pressure of Generative AI. **Method:** We adopted a sequential explanatory mixed-methods design. First, a survey with 56 undergraduate computing students at the University of Brasília (UnB) examined familiarity with 21 AI ethical principles, perceived relevance, and curricular exposure. Second, semi-structured interviews with 31 computing educators explored how ethics integration has evolved, the impact of Generative AI on teaching practices, perceived student ethical maturity at graduation, and structural barriers to effective ethics education. Quantitative data were analyzed descriptively, and qualitative data were examined through iterative coding procedures. **Results:** Students demonstrated strong agreement regarding the importance of ethical principles in computing and reported high self-perceived preparedness to assess social impacts. However, a clear gap emerged between perceived importance and formal curricular exposure. Educators, in turn, identified Generative AI as an operational accelerator that intensified assessment reconfiguration, academic integrity concerns, and accountability tensions. Interviews also revealed a dominant perception of low ethical maturity at the transition to professional practice, alongside systemic barriers including curricular overload, limited faculty preparation, and institutional governance gaps. **Conclusions:** The findings suggest the existence of a structural ethical literacy gap in computing education, reinforced by curricular fragmentation and insufficient institutional support, and amplified by the rapid integration of Generative AI. Addressing this gap requires coordinated curricular reform, faculty development, and institutional governance strategies that integrate ethical reasoning as a core technical competency rather than a peripheral concern. This competency is particularly important for future software engineers who must evaluate AI-generated artifacts, assume accountability for AI-supported engineering decisions, and negotiate socio-technical trade-offs throughout the software lifecycle.

Keywords: Ethical literacy, Generative AI, AI ethics, computing education, mixed-methods study, curriculum integration, Software Engineering education

1 Introduction

In recent years, the field of Artificial Intelligence (AI) has undergone a profound transformation, marked by its rapid expansion across sectors such as healthcare, finance, public safety, and education (Palmquist et al., 2026; Yang et al., 2025). This explosion of applications has sparked widespread concern about ethical issues ranging from fairness, transparency, and privacy to accountability, explainability, and sustainability (Corrêa et al., 2023; Jobin et al., 2019; Jantunen et al., 2024). Despite broad consensus among scholars, practitioners, and international organizations on the relevance of ethical principles in AI, their practical implementation remains challenging (Morley et al., 2023). Well-publicized cases involving racial bias in facial recognition and misuse of personal data illustrate the real-world implications of ethical shortcomings in AI systems (Al-Kaswan and

Izadi, 2023).

The emergence of Generative AI (GenAI), particularly large language models such as ChatGPT, has further intensified ethical debates within computing education. Beyond traditional concerns related to algorithmic bias and data governance, GenAI introduces new dilemmas involving academic integrity, authorship, data provenance, automation of cognitive labor, and shifts in socio-technical dynamics within higher education institutions (Mak et al., 2026; Kılınç, 2024). These developments have not only reshaped how students learn but also how educators design assessments and integrate ethical reflection into their courses.

This scenario reveals a deeper structural issue: the ethical literacy gap among computing professionals. While developers and engineers are increasingly responsible for decisions with profound societal consequences, many have received limited, if any, training in ethics and responsible innovation.

This lack of preparation creates an urgent demand for ethics education embedded in computing curricula, enabling future professionals to critically reflect on the broader impacts of the systems they build.

Recent literature highlights persistent challenges in integrating ethics into computing programs. Approaches such as Embedded EthiCS advocate for distributing ethical reflection across technical courses rather than isolating it in standalone modules (Grosz et al., 2019; Pasricha, 2023a). Other initiatives propose competency-based data ethics modules and structured curricular reforms to ensure observable ethical learning outcomes (Kiesler et al., 2024). Nevertheless, studies consistently report barriers such as limited faculty preparation, curricular overload, lack of institutional incentives, and absence of consensus on core ethical competencies (Brown et al., 2024a; Smith et al., 2023).

Although ethics in computing has been debated in academic contexts for over five decades (Nielsen, 1972), the urgency has intensified with the global adoption of AI systems. The proliferation of digital technologies and the growing complexity of algorithmic decision-making make it imperative to revisit the role of ethics in undergraduate education (Kumar et al., 2024).

Recent studies have emphasized the need to strengthen ethics teaching in computer science programs (Grosz et al., 2019; Fiesler et al., 2020; Goetze, 2023), yet research on this topic remains scarce, especially in Global South contexts such as Brazil (De La Vega Hernández et al., 2023). Many Brazilian institutions still lack structured approaches to ethics instruction, and where present, such content is often treated as peripheral or purely theoretical (Carvalho et al., 2021).

Moreover, little empirical research has simultaneously examined both students' and educators' perspectives on AI ethics education in Latin American computing programs. Understanding this dual perspective is important to identify potential misalignments between curricular intentions, pedagogical practices, and students' perceived learning outcomes, particularly in the rapidly evolving context of Generative AI.

This study aims to address this gap by analyzing students' perceptions of AI ethical principles within the undergraduate computing programs at the University of Brasília (UnB). Specifically, we investigate (i) students' familiarity with and perceived importance of 21 ethical principles commonly cited in AI ethics frameworks; and (ii) whether students report having received formal education on these topics during their degree. These discussions are particularly relevant for Software Engineering education, where ethical literacy is emerging as an essential professional competency. As Generative AI is increasingly incorporated into software engineering activities, including requirements elicitation, code generation, testing, documentation, and architectural decision-making, future software engineers must be able not only to use AI tools effectively but also to critically evaluate their ethical implications, accountability boundaries, and socio-technical consequences throughout the software lifecycle.

To complement these quantitative findings, we conducted semi-structured interviews with computing educators to explore how they perceive the integration of ethics into their courses, the impact of Generative AI on teaching practices,

and the main barriers to effective ethics education. This mixed-methods design enables the study to examine the ethical literacy gap from both students' and educators' perspectives through complementary sources of evidence.

This article is an extended version of our earlier study (Kipman et al., 2025). While the initial version focused exclusively on students' perceptions of AI ethical principles, the present study substantially expands both scope and analytical depth by incorporating a second empirical phase with computing educators, refining the conceptual framing of ethical literacy, and explicitly examining the pedagogical and institutional challenges introduced by Generative AI. These extensions enable a broader interpretation of ethical literacy not only as individual awareness, but also as a phenomenon shaped by curricular, pedagogical, and institutional conditions.

To this end, we conducted a quantitative survey with 56 students enrolled in UnB's computing programs and semi-structured interviews with 31 computing educators. The survey instrument included Likert-scale items regarding students' familiarity with and perceived relevance of 21 widely recognized AI ethical principles, as well as their reported exposure to ethics-related coursework. The interview protocol explored how educators have integrated ethics over time, how Generative AI has reshaped pedagogical and assessment practices, how they perceive students' ethical maturity at graduation, and what structural barriers limit more effective ethics education.

Our findings provide evidence of a structural ethical literacy gap within the investigated institutional context. Although most students recognized the importance of ethical principles and report high self-perceived preparedness to assess the social impact of their work, a significant portion report limited formal curricular exposure to structured ethics instruction. From the educators' perspective, Generative AI has functioned primarily as an operational accelerator, intensifying assessment reconfiguration, academic integrity concerns, and accountability tensions, while simultaneously exposing institutional fragilities such as curricular fragmentation, insufficient faculty preparation, and governance deficiencies.

By integrating students' and educators' perspectives, this study contributes empirical evidence from a Latin American context to the ongoing debate on AI ethics education. Rather than treating the Brazilian context, mixed-methods design, educator inclusion, and Generative AI lens as isolated contributions, we position these elements as complementary dimensions supporting the study's central contribution: the identification of a structural ethical literacy gap emerging from tensions among students' perceived preparedness, educators' assessments of ethical maturity, and institutional conditions shaping ethics education. The mixed-methods design and dual perspective provide the empirical basis for understanding this phenomenon and for examining how Generative AI intensifies existing curricular and institutional challenges.

Rather than framing ethics education merely as a matter of adding content to curricula, our findings suggest a broader systemic challenge involving ethical literacy consolidation, institutional alignment, and pedagogical adaptation under the

rapid transformation driven by Generative AI.

The remainder of this article is organized as follows: Section 2 presents the background and related work on AI governance and ethics education in computing; Section 3 details the mixed-methods research design, including the student survey and educator interview protocol; Section 4 reports the quantitative findings from Study 1; Section 5 presents the qualitative results from Study 2; Section 6 integrates findings from both phases and discusses implications for ethical literacy, curriculum design, and institutional governance under the pressure of Generative AI; Section 7 outlines threats to validity; and Section 8 concludes with contributions and directions for future research.

2 Background and Related Work

2.1 AI Governance

The rapid proliferation of AI technologies has sparked a corresponding surge in academic, regulatory, and societal interest in understanding and managing their ethical implications (Corrêa et al., 2023). This rise has been accompanied by a growing number of ethical controversies involving AI systems ranging from racial bias in facial recognition to opaque decision-making in algorithmic platforms (Goetze, 2023; Fiesler et al., 2020; Ferreira and Vardi, 2021). In response, numerous AI ethics guidelines have been published by governments, academic consortia, and standard-setting bodies (Jobin et al., 2019; Floridi, 2023; Corrêa et al., 2023; Whittlestone et al., 2019; de Cerqueira et al., 2021, 2022). While these documents often converge on high-level principles such as fairness, transparency, and accountability, their practical implementation remains a significant challenge.

As pointed out by Morley et al. (2023) and de Cerqueira et al. (2022), many of these guidelines are perceived by practitioners as abstract, aspirational, and difficult to translate into technical requirements or governance processes. This gap between principle and practice contributes to what scholars have termed the AI ethics crisis. Moreover, there is still no universal agreement on the core ethical principles that should guide AI development and deployment (Jobin et al., 2019; Floridi, 2023), further fueling regulatory fragmentation and uncertainty. Although the European Union has taken a pioneering step with its 2023 AI Act¹, most countries remain on the periphery of international governance discussions (Corrêa et al., 2023).

Recent discussions also emphasize that governance challenges are increasingly shaped by the rapid integration of Generative AI into professional and educational environments. Unlike traditional predictive systems, large language models and generative tools introduce governance complexities related to authorship, epistemic reliability, automation of cognitive labor, and blurred accountability boundaries (Mak et al., 2026). These dynamics require not only regulatory responses but also pedagogical adaptation, as universities become frontline spaces where ethical tensions surrounding AI are experienced and negotiated.

In this context, the systematic literature review by Batool et al. (2025) offers valuable insights by mapping 28 AI governance frameworks across four key dimensions: who governs, what is governed, when, and how. Their analysis shows that few existing approaches effectively integrate all these dimensions, and many lack mechanisms for stakeholder inclusion or clear accountability structures. Notably, the review highlights that ethical principles are often cited but remain poorly operationalized in existing models. This reinforces the need for educational and organizational strategies that go beyond formal compliance, fostering a deeper ethical culture in AI development.

Another critical barrier lies in the limited ethical literacy among technical professionals. As highlighted in multiple studies (Morley et al., 2023; Porto et al., 2025), software engineers, data scientists, and AI developers often possess only a superficial understanding of ethical principles and their practical implications. This not only undermines responsible AI development but also leads to uncertainty about who holds responsibility for harmful outcomes, sometimes shifting ethical burdens to end-users (Morley et al., 2023; Pant et al., 2024). These findings emphasize the importance of embedding ethics education in computing curricula as a foundational step toward effective and accountable AI governance.

2.2 Ethics Education in Computing

Although discussions on the importance of teaching ethics in computer science and AI date back to the 1970s (Nielsen, 1972), there has been a noticeable surge in publications on the topic in recent years (Brown et al., 2024b), paralleling trends in AI governance. This renewed interest has led to meaningful advancements in how AI ethics is approached in educational contexts and the identification of several promising pedagogical directions.

More recent studies have begun to analyze the ethical implications of Generative AI specifically within computing education. For example, Mak et al. (2026) proposed an Ethical and Societal Impacts (ESI) framework to guide institutional decision-making in the adoption of GenAI tools. Their findings highlight tensions between academic integrity, equity of access, pedagogical redesign, and professional preparation. Similarly, Kılınç (2024) introduces a structured framework to ethically integrate GenAI into assessment practices, emphasizing responsible usage rather than prohibition. These studies indicate a shift from merely teaching AI ethics to rethinking how AI itself transforms ethical education.

While multidisciplinary teaching approaches are attractive due to their simplicity, they are often only a starting point. Since most computer science instructors lack formal training in ethics, it becomes challenging to integrate ethical content across entire curricula (Brown et al., 2024b). Effective interdisciplinary approaches typically require time-intensive collaboration among professionals from various fields, an investment that institutions may be reluctant to make (Goetze, 2023). As a result, when ethics is taught through standalone multidisciplinary courses, it often feels abstract and disconnected from the technical realities where ethical decisions must be made (Miller, 1988). This can lead students to perceive ethics as separate from the functional requirements of

¹<https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

system development (Goetze, 2023).

Furthermore, just as there is no universally agreed-upon set of ethical principles for AI, there are also few well-documented learning objectives specific to AI ethics (Aler Tubella et al., 2024). This lack of clarity hampers both the evaluation of student learning and the assessment of pedagogical methods (Brown et al., 2024b). Brown et al. also note the existence of multiple ethical frameworks, each potentially better suited to different disciplinary contexts, making it necessary for curricula to clearly articulate how ethics content aligns with course objectives and disciplinary scope (Brown et al., 2024b).

Beyond student perceptions, empirical research has also examined instructors' perspectives and institutional barriers to integrating ethics. Studies report recurring challenges such as limited faculty training, lack of curricular time, absence of consensus on core competencies, and insufficient institutional incentives (Smith et al., 2023; Brown et al., 2024a). These findings suggest that ethics integration depends not only on curricular design but also on structural and organizational support. However, there remains limited empirical evidence from Latin American institutions regarding how educators perceive these barriers, especially in the context of Generative AI.

Tran and Fiesler (2024) investigated how computing students perceive the integration of ethics into group projects, which are often used in higher education as simulations of real-world work environments. Through focus groups with 29 students, the authors found that, although ethics is recognized as important, it often plays a marginal role in project-based courses, particularly in software engineering classes. Many students reported that they did not view their projects as real enough to justify ethical concerns, partly due to the absence of real users or external stakeholders. Additionally, students expressed difficulties in applying ethical concepts due to a lack of clarity, time, and institutional support. In response, the authors proposed practical strategies for educators: incorporating ethical considerations early in the project timeline, using case studies, encouraging user testing, and designing rubrics that include ethical criteria. The study also highlights the potential of forming ethical consultants in academic settings to support students in their design decisions. The authors concluded that ethics should be treated as a functional and realistic component of software development, fostering students' preparedness for the ethical challenges they will face in the professional world.

Pasricha (2023b) also highlighted the growing urgency of embedding ethics into computing curricula due to the proliferation of AI, IoT, and embedded systems with significant societal impacts. The authors discussed a case at Colorado State University where ethics modules were incorporated into both undergraduate and graduate computer engineering courses. Unlike traditional standalone ethics courses, the approach infused ethical reflection throughout the curriculum, using real-world case studies and ethical theories such as utilitarianism, Kantianism, and virtue ethics. The study also emphasized that merely referencing professional codes of ethics such as the ACM Code of Ethics (Association for Computing Machinery, 2018) and the IEEE Code of Ethics (Institute of Electrical and Electronics Engineers, 2020) is insufficient; in-

stead, students benefit more from interactive pedagogy that encourages critical reasoning and ethical debate. The findings reinforced the value of curricular integration of ethics not as peripheral, but as a core element of technical education, supporting observations that students value ethical training but often lack formal exposure within computing programs.

2.3 Ethical Principles for AI

To evaluate students' familiarity with and perceptions of ethical issues in AI systems, this study adopts a comprehensive set of 21 ethical principles. These principles were derived from a synthesis of international guidelines, academic frameworks, and best practices identified in recent literature, including works by (Khan et al., 2023), (Vakkuri et al., 2022), and (de Cerqueira et al., 2022), among others (Khan et al., 2021; Porto et al., 2025). This selection was informed by an extensive review of over 80 global AI ethics initiatives and reflects the growing consensus around key values that should guide the responsible design, development, and deployment of intelligent systems.

The principles cover both technical and socio-ethical dimensions, including fairness, accountability, transparency, privacy, non-maleficence, and sustainability. While many of these values are widely cited, their interpretation and implementation can vary across contexts and stakeholders. For this reason, the questionnaire used in this study (Section 3) presented each principle with a concise, accessible definition to support consistent understanding by participants. These definitions were adapted from existing literature to reflect both normative goals and practical implications.

Table 1 provides the full list of the 21 ethical principles evaluated in this study, along with their corresponding definitions. These concepts served as the foundation for the self-assessment questions (Q6 in Table 2) and perceived relevance ratings (Q7 in Table 2) used to investigate student perceptions.

Although the 21 ethical principles adopted in this study are grounded in consolidated AI ethics frameworks and prior empirical studies (Khan et al., 2023; Vakkuri et al., 2022; Khan et al., 2021; Porto et al., 2025; de Cerqueira et al., 2022), some concepts may present partial semantic or conceptual overlap, such as explainability and interpretability or accountability and responsibility. These distinctions were intentionally preserved to maintain alignment with the terminology and granularity commonly adopted in the AI ethics literature. To reduce potential ambiguity during the student survey, participants were presented with the name of each principle together with its corresponding definition. Nevertheless, we acknowledge that conceptual proximity among certain principles may still have influenced students' perceptions and self-reported familiarity levels, particularly when interpreting related ethical concepts.

These principles have been widely cited in both theoretical and applied contexts, but there remains a lack of concrete tools, metrics, and educational strategies to operationalize them (Porto et al., 2025; de Cerqueira et al., 2022). Furthermore, even though organizations are beginning to adopt ethical AI frameworks, many practitioners struggle to apply these principles in day-to-day development due to vague def-

Table 1. Ethical Principles for AI Systems Used in This Study (Khan et al., 2023; Vakkuri et al., 2022; Khan et al., 2021; Porto et al., 2025; de Cerqueira et al., 2022)

ID	Principle	Detailed Description
P1	Transparency	The AI system should provide clear information about its decision-making process, data sources, and objectives to allow users and stakeholders to understand how decisions are made.
P2	Privacy	Users' personal data must be protected throughout the entire system lifecycle, including data collection, storage, processing, and sharing.
P3	Accountability (Justice Protection)	Assign responsibilities to developers and organizations for the outcomes of AI systems, ensuring mechanisms to investigate and respond to harmful consequences.
P4	Fairness	AI systems should treat individuals and groups equitably, avoiding biases and discrimination related to race, gender, age, or other factors.
P5	Autonomy	Ensure that AI systems support and enhance human autonomy rather than undermining users' capacity for decision-making.
P6	Explainability	Provide meaningful explanations for AI decisions that can be understood by different stakeholders, including non-experts.
P7	Justice	Promote social justice and inclusion, ensuring AI systems contribute to the reduction of inequality and discrimination.
P8	Non-maleficence	Prevent harmful applications and uses of AI that could cause psychological, physical, social, or economic damage.
P9	Human Dignity	Respect human values and rights, ensuring that AI systems do not compromise the intrinsic worth of individuals.
P10	Beneficence	AI systems should actively promote human well-being, safety, and social progress.
P11	Responsibility (Design Intention)	AI systems must be intentionally designed to empower humanity and serve public interest with ethical foresight.
P12	Safety	Ensure that AI systems operate reliably, securely, and predictably under expected conditions to avoid unintended harm.
P13	Data Security	Protect data integrity and security from unauthorized access, leaks, or manipulation.
P14	Sustainability	AI systems should consider environmental, economic, and social sustainability, minimizing ecological impacts and promoting long-term viability.
P15	Freedom	Respect and promote human freedoms, including civil liberties and the right to free expression and thought.
P16	Solidarity	Promote shared benefits and responsibilities, ensuring that both risks and rewards of AI systems are fairly distributed.
P17	Prosperity	Support collective prosperity, promoting economic and social development through the ethical use of AI.
P18	Effectiveness	Ensure that AI systems are technically sound and socially appropriate to effectively address the intended problems.
P19	Accuracy	Ensure that AI outputs are precise, consistent, and based on accurate and up-to-date information.
P20	Predictability	Design AI systems in ways that their behaviors and long-term impacts are foreseeable and understandable.
P21	Interpretability	Provide tools or interfaces that help users interpret how an AI system arrived at a particular outcome, beyond technical explainability.

initions and limited institutional support (Morley et al., 2023; Pant et al., 2024).

Although these principles are widely recognized, recent literature argues that their educational operationalization remains underdeveloped (Kiesler et al., 2024). Competency-based approaches have been proposed to translate abstract ethical principles into observable learning outcomes, yet empirical validation of such approaches is still limited, particularly in AI-focused curricula.

In this study, ethical literacy is conceptually defined and analytically operationalized as a multidimensional construct rather than measured as a single latent variable. Specifically, we interpret ethical literacy through four complementary di-

mensions investigated across the two empirical phases: (i) familiarity with ethical principles, (ii) perceived relevance of these principles, (iii) curricular exposure to ethics education, and (iv) ethical maturity, as perceived by computing educators. Together, these dimensions provide the analytical basis for interpreting the structural ethical literacy gap discussed throughout this paper.

2.4 Ethics Education in Computing Curricula

The inclusion of ethics in computing curricula has been advocated for over five decades (Nielsen, 1972), yet only in recent years has it received sustained attention due to ethical failures

in AI deployments (Fiesler et al., 2020; Goetze, 2023; Ferreira and Vardi, 2021). Educators and professional organizations now emphasize the need to integrate ethics across technical training, not as isolated lectures or standalone courses, but as a core component of software and systems design (Grosz et al., 2019).

Despite these advances, challenges persist. Most instructors lack formal training in ethics (Brown et al., 2024b), and institutions are often reluctant to allocate resources for interdisciplinary collaboration required to teach ethical reasoning alongside technical content. Additionally, students often perceive ethics courses as detached from real engineering work, which undermines engagement and knowledge retention (Goetze, 2023; Miller, 1988).

Another key issue is the absence of standard learning outcomes and evaluation methods tailored to AI ethics. As highlighted by Aler Tubella et al. (2024), clearly defined objectives are important for measuring pedagogical effectiveness and ensuring that ethics education results in actual behavioral change.

A recent study by Brown et al. (2024c) highlighted a persistent misalignment between students' perceptions of learning outcomes and the actual curricular goals related to ethics and social responsibility in computing programs. Although many computing programs aim to instill ethical reasoning, students often struggle to identify when and how they are learning these skills. The authors emphasize that ethics content is frequently implicit, leaving students unaware of its presence in technical courses. This aligns with previous findings that isolated or decontextualized ethics instruction can limit student engagement and comprehension. The study advocates for explicit and integrated pedagogical approaches, reinforcing the importance of making ethical principles visible and contextualized within real-world technical scenarios. These findings support our argument that curricula should include structured, transparent, and applied instruction on ethical principles, especially in light of the low exposure to formal ethics education reported in our survey.

In the Brazilian context, studies such as Carvalho et al. (2021) have shown that a significant number of computing courses still either omit ethics or address it superficially. This educational gap becomes particularly concerning given the growing demand for AI systems that align with democratic values and human rights.

Recent research has further highlighted that the integration of ethics into computing curricula is being reshaped by the emergence of Generative AI, which introduces new educational demands related to ethical reasoning, AI literacy, and changing pedagogical practices (da Rocha Braz et al., 2026; Yang et al., 2025). Unlike previous waves of AI technologies, GenAI tools directly affect classroom practices, assessment design, authorship norms, and students' epistemic autonomy (Mak et al., 2026). This shift places new responsibilities on educators, who must not only teach ethical principles but also critically address how AI tools transform learning environments.

Moreover, studies examining educators' perspectives reveal persistent structural barriers to effective ethics integration, including limited training, curricular overload, lack of institutional support, and uncertainty about how to opera-

tionalize abstract principles into technical contexts (Smith et al., 2023). These findings suggest that successful ethics education depends not only on curricular content but also on institutional culture and faculty development.

Despite this growing body of international literature, there is limited empirical evidence examining how computing educators in Latin American contexts perceive their role in AI ethics education, particularly in the rapidly evolving context of Generative AI. Addressing this gap is important for understanding whether curricular reforms align with local pedagogical realities and institutional constraints.

This paper contributes to this debate by presenting a mixed-methods study conducted at the University of Brasília (UnB). Building upon the initial survey of students' exposure to and perceptions of key AI ethical principles, this extended version also incorporates semi-structured interviews with computing educators. By combining students' and instructors' perspectives, the study identifies curricular gaps, institutional barriers, and pedagogical tensions associated with integrating ethics into computing programs, particularly under the transformative pressures introduced by Generative AI.

3 Study Settings

This study adopts a sequential explanatory mixed-methods design (Storey et al., 2025) to investigate ethics education in undergraduate computing programs at the University of Brasília (UnB). The design combines quantitative data from a student survey with qualitative data from semi-structured interviews with computing educators.

The sequential approach was chosen to first capture students' perceptions regarding their exposure to and understanding of AI ethical principles (RQ.1 and RQ.2), and subsequently explore instructors' perspectives to contextualize and explain curricular practices, institutional barriers, and emerging challenges related to Generative AI (RQ.3). This design aligns with methodological guidance for mixed-methods research in computing and software engineering education (Storey et al., 2025), which emphasizes the integration of complementary data sources to provide a more comprehensive understanding of educational phenomena.

The integration of findings occurs at the interpretation stage, where quantitative patterns identified in the student survey are examined alongside qualitative insights from educators to identify convergences, tensions, and explanatory factors.

This integrative interpretation operationalizes ethical literacy as a multidimensional construct supported by complementary evidence regarding familiarity with ethical principles, perceived relevance, curricular exposure, and ethical maturity.

Figure 1 operationalizes this interpretation by summarizing how the quantitative and qualitative dimensions investigated in this study jointly support the interpretation of ethical literacy, the identification of the structural ethical literacy gap, and its implications for AI-assisted Software Engineering.

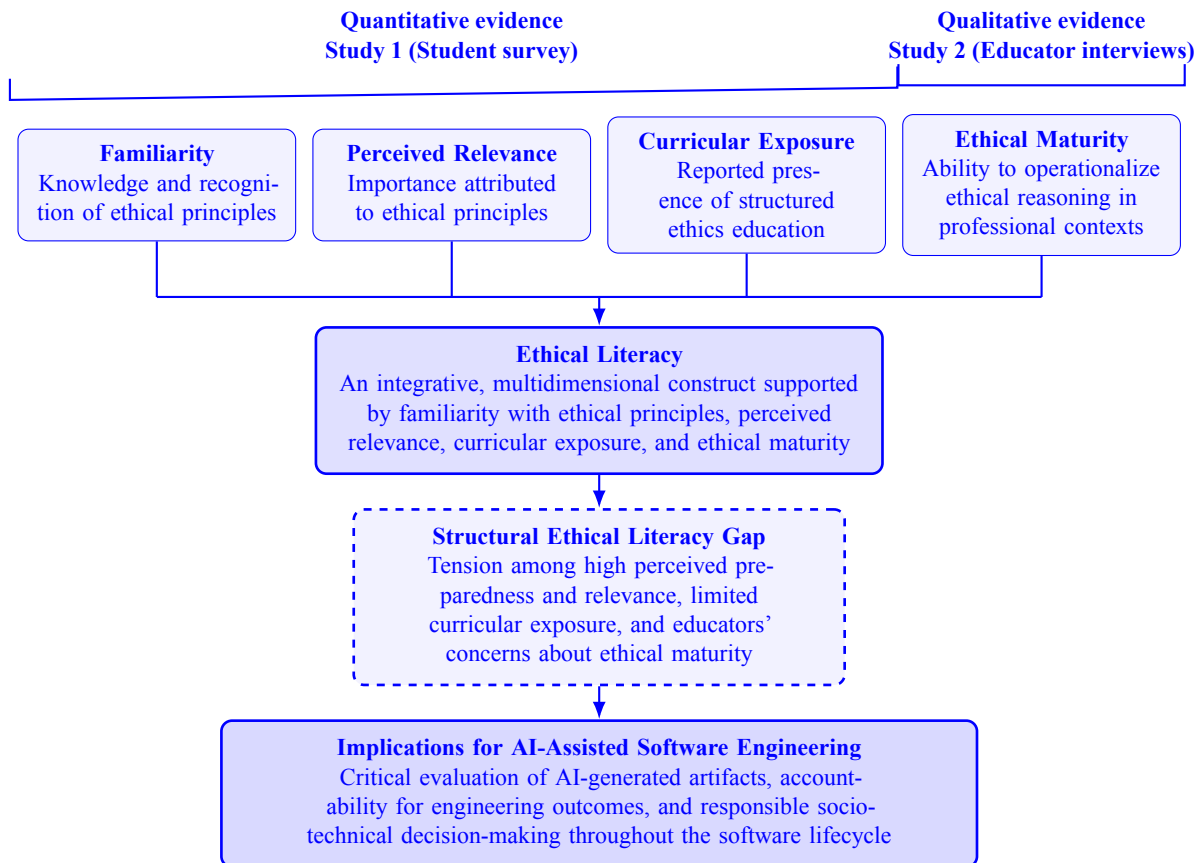


Figure 1. Conceptual model of ethical literacy as a multidimensional construct integrating quantitative and qualitative evidence. The ethical literacy gap emerges from tensions among students' familiarity, perceived relevance, curricular exposure, and educators' assessments of ethical maturity, with implications for responsible AI-assisted Software Engineering.

To guide the investigation, we formulated the following research questions (RQ):

RQ.1: How do students assess their learning about ethics during the undergraduate program?

This question explores students' self-perception regarding the exposure to and assimilation of ethical content throughout their academic journey.

RQ.2: What are students' perceptions of the role and relevance of ethics in the computing field?

This question aims to understand how students perceive the importance, application, and impact of ethical principles in computing.

RQ.3: How do educators perceive their role, the main challenges, and institutional responsibilities in fostering ethics education in computing, particularly in the context of Generative AI?

This question aims to understand how educators interpret their responsibilities, the institutional constraints they face, and the challenges of integrating ethics in relation to Generative AI into computing curricula.

To address RQ.1 and RQ.2, we designed and conducted a quantitative survey targeting undergraduate students enrolled in computing programs. The instrument explored students' familiarity with AI ethical principles, their perceived importance, and their exposure to formal ethics instruction

during their degree. Detailed results from this survey were initially reported in Kipman et al. (2025) and are here revisited and analytically extended within a broader mixed-methods framework.

To address RQ.3, we conducted semi-structured interviews with computing educators affiliated with the same programs. The interviews aimed to capture instructors' perceptions regarding curricular integration of ethics, pedagogical challenges, institutional support, and the implications of Generative AI for teaching practices. By integrating quantitative student data with qualitative instructor insights, the study seeks to identify convergences and tensions between perceived learning outcomes and institutional realities.

The following subsections describe the study population, instrument design, pilot validation procedures, data collection strategies, and analytical methods employed for both phases of the study.

3.1 Target Audience

The study involved two complementary participant groups. The first group comprised undergraduate students enrolled in Computing-related programs at the University of Brasília (UnB), including Computer Science, Software Engineering, Computer Engineering, Network Engineering, and the Teaching Degree in Computing. These students represent the primary beneficiaries of curricular decisions regarding ethics integration.

The second group consisted of computing educators affiliated with the same programs. These instructors are directly involved in curricular design, course delivery, and pedagogical decision-making, particularly regarding the integration of ethics and emerging AI technologies into technical courses.

By including both students and educators, the study captures perspectives from those who experience ethics education and those who are responsible for implementing it, enabling a more comprehensive understanding of curricular dynamics.

3.2 Instruments and Data Collection

Two complementary instruments were developed to address the research questions: a structured survey targeting undergraduate students (RQ.1 and RQ.2) and a semi-structured interview protocol targeting computing educators (RQ.3). Both instruments were collaboratively designed by the authors and aligned with the study's sequential explanatory mixed-methods design.

3.2.1 Student Survey Instrument

The survey instrument consisted of 9 closed-ended questions organized into four thematic sections, as presented in Table 2. The instrument was structured as follows:

- **Section 1: Terms and Conditions** – Presented an informed consent statement. The first question asked: *“Do you agree with the terms and conditions to participate in this research?”*
- **Section 2: Demographic Profile** – Included questions about the student's course and current semester.
- **Section 3: Self-Assessment on Ethics Learning** – Included questions such as whether students felt prepared to assess the social impacts of their decisions, and if they were exposed to ethical discussions during their education.
- **Section 4: Perception of Ethics in Computing** – Focused on students' familiarity with ethical principles, the perceived relevance of these principles, and whether ethics was covered during their coursework.

3.2.2 Educator Interview Protocol

To address RQ.3, a semi-structured interview protocol was developed to capture instructors' perceptions regarding ethics integration in computing curricula, institutional responsibilities, pedagogical challenges, and the impact of Generative AI on teaching practices.

The interview included demographic questions and five open-ended questions directly aligned with RQ.3, as presented in Table 3. The protocol was designed to elicit reflective and experience-based responses, enabling a deeper understanding of curricular dynamics and institutional constraints that may not be visible through student self-report data alone.

3.3 Pilot Study

A pilot study was conducted with three undergraduate students enrolled in Computing programs to evaluate the clarity, relevance, and overall comprehensibility of the survey questionnaire. Based on their feedback, minor refinements were made to improve the wording and structure of several questions. The average completion time was approximately 10 minutes. It is important to note that the responses collected during the pilot phase were not included in the final data analysis. After finalizing the instrument, the survey was made available online for a period of two months to maximize participation and ensure broad dissemination.

A pilot test was also conducted for the educator interview protocol. Two computing professors participated in this preliminary phase to assess the clarity, sequencing, and depth of the interview questions. Based on their feedback, minor adjustments were made to improve question phrasing and ensure alignment with the study's objectives. The pilot interviews confirmed the feasibility of the protocol and helped refine prompts to encourage more reflective responses. As with the student pilot, data from this phase were not included in the final qualitative analysis. The full interviews conducted in Study 2 lasted between 25 and 30 minutes on average.

3.4 Participant Recruitment and Data Collection

The final version of the survey was developed using the Google Forms platform, chosen for its ease of use and accessibility. The link to the questionnaire was disseminated through multiple communication channels frequently used by students in Computing programs at the University of Brasília (UnB), including WhatsApp groups, Instagram profiles managed by student associations, and internal academic networks. The dissemination strategy aimed to reach undergraduate students across various semesters and specializations within the university's Computing-related programs.

Participation in the study was entirely voluntary, and no incentives were offered. To ensure ethical compliance, an informed consent form was presented at the beginning of the questionnaire. Only respondents who explicitly agreed to the terms of participation were permitted to access and complete the survey. The informed consent included details about the study's objectives, anonymity assurances, data usage, and contact information for further clarification.

The survey remained open for a period of two months, providing ample time for engagement from a diverse group of students. In total, 56 valid responses were collected, after excluding incomplete or invalid entries. These responses formed the basis of the quantitative data analysis presented in this study.

For the qualitative phase (RQ.3), computing educators affiliated with the same undergraduate programs were invited to participate in semi-structured interviews. Invitations were sent via institutional email and reinforced through direct contact via WhatsApp, using publicly available professional contact information.

Participation was voluntary, and no incentives were offered. Of the 49 faculty members invited, 31 agreed to par-

Table 2. Survey Structure and Corresponding Research Questions

ID	Survey Question	Related RQ
Section 1 – Terms and Conditions		
Q1	Do you agree with the terms and conditions to participate in this research?	–
Section 2 – Demographic Profile		
Q2	What is your degree program?	–
Q3	Which semester are you currently enrolled in?	–
Section 3 – Self-Assessment on Ethics Learning		
Q4	I feel capable and prepared to assess the social impact of my decisions as a computing professional.	RQ1
Q5	Throughout my education, I was frequently exposed to ethical discussions and questions.	RQ1
Q6	I am familiar with ethical principles and know how to apply them in practice.	RQ1
Section 4 – Perception of Ethics in Computing		
Q7	Please rate the relevance of each ethical principle according to your understanding and experience.	RQ2
Q8	Do you think it is important to learn about these ethical principles during your undergraduate studies?	RQ2
Q9	Did your course include a subject that addressed these ethical principles?	RQ2

Table 3. Semi-structured interviews applied to computing educators: questions and related research questions

ID	Section	Question (English)	Related RQ
IQ1	Identification	What is your email address?	–
IQ2	Consent	Do you agree with all terms and conditions to participate in this study?	–
IQ3	Demographics	Which academic unit are you affiliated with?	–
IQ4	Demographics	Which degree program are you affiliated with?	–
IQ5	Demographics	For how many years have you taught within the broader Computing field?	–
IQ6	Ethics teaching over time	How has the way you incorporate ethics into your courses changed over the years, and more specifically, how has it been influenced by the emergence of Generative AI?	RQ3
IQ7	Role of educators	In your view, what is the role of educators in shaping ethical dialogue in the field, given the limited consensus on which concerns matter most—and even less agreement on how they should be taught?	RQ3
IQ8	Institutional and regulatory responsiveness	Given how quickly AI advances have been occurring, how can institutions and legislators keep pace with the potential harms that new technologies may cause to society?	RQ3
IQ9	Students' ethical maturity	How would you assess students' ethical maturity and their understanding of the severity and breadth of the consequences of their actions as they transition from university to professional practice?	RQ3
IQ10	Barriers	In your view, what are the main barriers that currently limit more effective ethics education in computing programs?	RQ3

ticipate in the interview phase. Prior to each interview, participants received an informed consent statement outlining the study's objectives, confidentiality procedures, voluntary participation, and data usage policies.

Interviews were scheduled individually and conducted remotely. All interviews were recorded with participants' authorization and later transcribed for analysis. These 31 interviews constitute the qualitative dataset used to address RQ.3.

3.5 Data Analysis

3.5.1 Quantitative analysis (Study 1)

Survey responses were analyzed using descriptive statistics. We computed absolute and relative frequencies for all items and summarized Likert-scale distributions (1–5) for Q4–Q9. For the ethical principles assessed in Q6 and Q7, we exam-

ined response distributions across scale categories to characterize students' self-assessed knowledge and perceived relevance.

3.5.2 Qualitative analysis (Study 2)

Interview data were analyzed through an iterative coding process inspired by Grounded Theory procedures (Hoda, 2025), focusing on categorization rather than theory building. First, we conducted line-by-line coding to identify concepts and meaning units expressed by educators. Next, focused coding was applied to iteratively group semantically related codes into broader analytical categories representing recurrent perceptions, barriers, and responsibilities regarding ethics education in computing and the implications of Generative AI. Throughout the process, the codebook was refined iteratively to improve consistency and resolve overlaps. Category sta-

bilization was considered achieved when subsequent interviews no longer contributed substantively new analytical distinctions relevant to the research questions.

An illustrative example of the qualitative coding process is presented based on participant P#4's response to IQ6 in Figure 2. Initially, the initial coding process generated codes such as progressive curricular formalization of ethics, emphasis on academic integrity, use of professional codes (ACM and SBC), and increased need for ethical clarity following the emergence of Generative AI.

Through focused coding, the line-by-line codes extracted from the raw excerpt were iteratively grouped into broader analytical categories. Figure 2 illustrates this process, showing the progression from the original participant response, to line-by-line coding, and finally to the higher-level categories generated through abstraction. The categories shown in Figure 2 correspond only to the concepts emerging from participant P#4 and illustrate how raw excerpts contributed to category development. The complete coding process across all interviews resulted in additional categories presented in Table 7.

Illustrative participant excerpts supporting the qualitative categories, together with anonymized participant identifiers (e.g., P1–P31), are provided in the supplementary material and codebook to improve transparency and traceability of the qualitative interpretation process.

This process allowed us to systematically categorize instructors' perspectives without aiming to construct a grounded theory, but rather to organize recurring patterns and explanatory dimensions emerging from the interviews.

3.5.3 Integration of quantitative and qualitative findings

Following the sequential explanatory mixed-methods design, integration occurred at the interpretation stage. Quantitative patterns from the student survey were used as analytic anchors to guide qualitative interpretation, and educators' accounts were examined to explain, contextualize, and qualify those patterns. We compared convergences and divergences across data sources, highlighting tensions between students' perceived preparedness and reported curricular exposure, as well as structural and pedagogical factors that may explain such discrepancies.

For conceptual clarity, this study distinguishes related constructs frequently discussed throughout the analysis. Familiarity refers to participants' awareness and recognition of ethical principles; preparedness concerns self-perceived capability to address ethical issues; relevance refers to the perceived importance attributed to ethical principles; ethical maturity reflects the ability to operationalize ethical reasoning in realistic professional contexts; and ethical literacy is interpreted as an overarching construct integrating ethical knowledge, reflection, and applied reasoning.

4 Results Study 1

This section presents the results of the study, organized to address the research questions defined in Section 3. We begin

by outlining the demographic profile of the participants to provide context for interpreting the findings. Next, we present the results related to **RQ.1**, which explores their exposure to ethical principles and curricular experiences. Finally, we address **RQ.2**, focusing on how students perceive the importance and applicability of ethical principles in their future professional practice.

4.1 Demographic Profile

A total of 56 students completed the survey and consented to participate by agreeing to the terms and conditions presented at the beginning of the questionnaire (Q1). As shown in Table 4, all respondents (100%) accepted the informed consent. One participant did not provide responses to the demographic questions on degree program and current semester; therefore, the distributions reported for Q2 and Q3 are based on 55 responses. Regarding their degree programs (Q2), the majority of respondents were enrolled in the Teaching Degree in Computing program (38.2%), followed by Computer Engineering (23.6%), and Network Engineering (21.8%). Fewer responses were recorded from students in Computer Science (9.1%) and Software Engineering (7.3%).

In terms of academic progression (Q3), participants came from a wide range of semesters. Most were in the early stages of their programs, with 32.7% in the first semester and 20.0% in the third semester. Intermediate and advanced semesters were also represented, including the seventh (10.9%), fifth (9.1%), and sixth semesters (7.3%). A smaller portion of respondents were in later semesters, including the eleventh or beyond (5.5%), and other scattered stages such as the second, fourth, eighth, ninth, and tenth semesters (Table 4).

These results indicate that the sample includes students from all phases of the undergraduate journey, which supports a diverse set of perspectives on ethics education across different levels of exposure to the curriculum.

Table 4. Profile of survey respondents (Q1: $n = 56$; Q2–Q3: $n = 55$)

Consent to Participate	#	%
Yes	56	100.0
Degree Program	#	%
Teaching Degree in Computing	21	38.2
Computer Engineering	13	23.6
Network Engineering	12	21.8
Computer Science	5	9.1
Software Engineering	4	7.3
Current Semester	#	%
1st semester	18	32.7
3rd semester	11	20.0
7th semester	6	10.9
5th semester	5	9.1
6th semester	4	7.3
11th or beyond	3	5.5
2nd semester	2	3.6
4th semester	2	3.6
9th semester	2	3.6
8th semester	1	1.8
10th semester	1	1.8

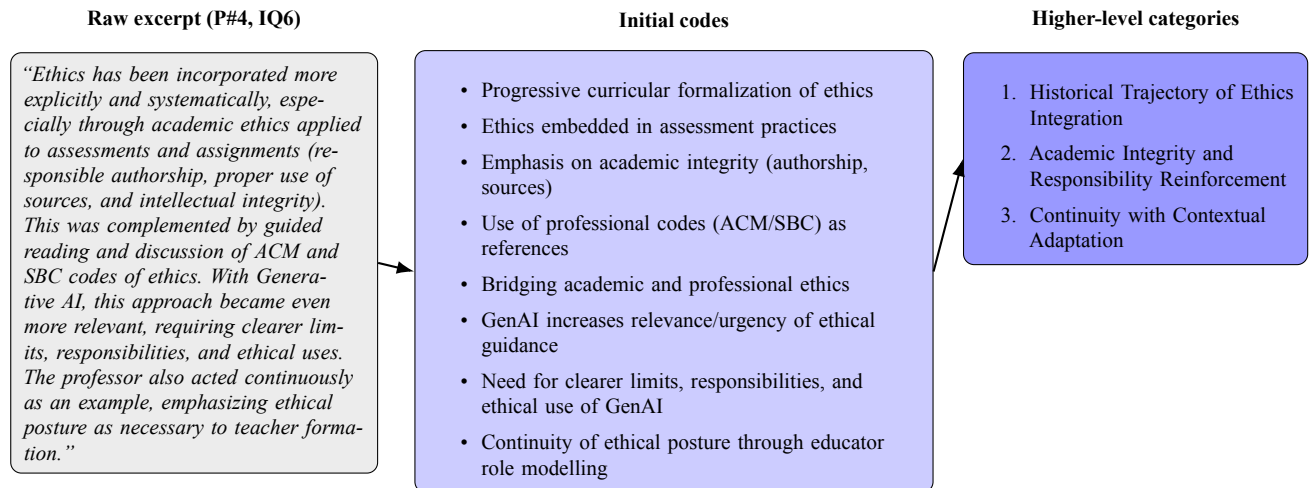


Figure 2. Illustrative example of the qualitative coding procedure showing the progression from a raw participant excerpt to initial codes and higher-level categories (Participant P#4, IQ6).

4.2 RQ.1. How do students assess their learning about ethics during the undergraduate program?

This subsection explores students' perceptions of their own learning and preparedness regarding ethical issues during their undergraduate program. First, we analyze their self-assessed preparedness to deal with ethical dilemmas and their exposure to ethics-related topics throughout the curriculum (Q4 and Q5). Then, we examine the depth of their self-reported knowledge across a range of ethical principles (Q6). Together, these findings shed light on potential gaps between students' confidence in ethical reasoning and the actual presence and consistency of ethics instruction in the computing curriculum.

Figure 3 presents the distribution of responses to Q4 and Q5, based on a five-point self-assessment scale ranging from "Very low" to "Very high". Q4 asked students to assess their own preparedness to evaluate the social impact of their decisions as computing professionals. The responses indicate a strong perception of readiness, with more than half of the participants selecting the highest categories: 25.5% rated their preparedness as "Very high" and another 25.4% as "High". Only 9.1% of students reported a "Very low" sense of preparedness, while intermediate categories "Moderate" (23.6%) and "Low" (16.4%) also gathered substantial responses. These results suggest that the majority of students feel at least moderately confident in their ability to reflect on the ethical and social implications of their work.

Q5 investigated students' perceived exposure to ethical discussions throughout their academic training. In contrast to Q4, the responses were more evenly distributed and somewhat lower in intensity. The most common responses were "Moderate" (28.6%) and "Low" (28.6%), suggesting that many students recognize a limited or occasional engagement with ethics-related content. Only 12.5% indicated "Very high" exposure, while 8.9% reported "Very low" exposure. These findings point to a potential gap: while students feel relatively well-prepared to evaluate ethical impacts (Q4), their formal or frequent exposure to structured ethical discussions (Q5) appears to be lacking.

Together, these results may reflect a disconnect between students' perceived ethical maturity and the depth or consistency of ethics education within the curriculum. While self-confidence regarding ethical responsibility is high, the educational processes underpinning that confidence may not be equally robust or widespread.




Figure 3. Distribution of responses to Q4 and Q5 based on self-assessed exposure and preparedness

Table 5 presents the distribution of students' self-assessed knowledge and ability to apply various ethical principles in practice (Q6). The responses were collected using a 5-point Likert scale, ranging from "Strongly disagree" to "Strongly agree". Overall, students reported high levels of perceived knowledge for the majority of the principles. Justice, Fairness, and Privacy stand out as the most internalized principles: Justice received the highest rate of Strongly agree responses (73.2%). Fairness followed closely (67.9%), with Privacy at 64.3%.

Other principles with comparatively higher levels of self-reported familiarity include Accuracy (60.7%), Data Security (60.7%), Beneficence (59.0%), and Human Dignity (58.9%). On the other hand, Autonomy (35.7% Strongly agree) and Solidarity (44.6%) were perceived as less familiar to students. These principles also had the highest proportions of neutral or disagreeing responses: Autonomy had 26.8% neutral responses and 14.2% in disagreement (7.1% Disagree and 7.1% Strongly disagree), as shown in Table 5.

Solidarity had 30.3% of respondents either neutral or disagreeing. Similarly, Prosperity, Interpretability, and Predictability showed lower familiarity levels, with more than 25% of participants selecting Neutral or below, as shown in Table 5.

This distribution suggests that students are more confident in their understanding of well-established principles such as Privacy, Justice, and Safety, while more abstract or less dis-

Table 5. Self-reported knowledge and practical understanding of ethical principles in AI systems (Q6)

Ethical Principle	Strongly disagree	Disagree	Neither agree or disagree	Agree	Strongly agree
Transparency	3.6%	8.9%	16.1%	37.5%	33.9%
Privacy	0.0%	1.8%	7.1%	26.8%	64.3%
Accountability	3.6%	5.4%	8.9%	33.9%	48.2%
Fairness	0.0%	1.8%	5.4%	25.0%	67.9%
Autonomy	7.1%	7.1%	26.8%	23.2%	35.7%
Explainability	1.8%	8.9%	10.7%	25.0%	53.6%
Justice	0.0%	1.8%	3.6%	21.4%	73.2%
Non-maleficence	1.8%	5.4%	10.7%	30.4%	51.8%
Human dignity	1.8%	5.4%	8.9%	25.0%	58.9%
Beneficence	0.0%	1.8%	7.1%	32.1%	59.0%
Responsibility	1.8%	5.4%	7.1%	26.8%	58.9%
Safety	0.0%	3.6%	8.9%	32.1%	55.4%
Data security	0.0%	3.6%	8.9%	26.8%	60.7%
Sustainability	1.8%	1.8%	14.3%	28.6%	53.6%
Freedom	0.0%	3.6%	10.7%	30.4%	55.4%
Solidarity	1.8%	7.1%	21.4%	25.0%	44.6%
Prosperity	1.8%	5.4%	21.4%	26.8%	44.6%
Effectiveness	0.0%	3.6%	14.3%	26.8%	55.4%
Accuracy	0.0%	3.6%	8.9%	26.8%	60.7%
Predictability	1.8%	7.1%	14.3%	30.4%	46.4%
Interpretability	1.8%	5.4%	17.9%	26.8%	48.2%

cussed concepts like Autonomy and Solidarity may require greater emphasis in the curriculum.

RQ.1 Summary: Students report a high level of self-perceived preparedness to assess the social impact of their actions as computing professionals. However, their exposure to ethical discussions during their undergraduate studies is more limited and uneven. There is a clear discrepancy between students' confidence and the depth of curricular engagement with ethics. Most students are familiar with well-established principles such as Justice, Fairness, and Privacy. Conversely, abstract principles like Autonomy, Solidarity, and Prosperity were less understood and rated lower in self-assessed knowledge.

4.3 RQ.2: What are students' perceptions of the role and relevance of ethics in the computing field?

This subsection examines students' views regarding the significance of ethical principles in AI and computing (Q7), as well as their perspectives on the inclusion of such content in undergraduate computing curricula (Q8 and Q9). The results provide insight into which ethical principles students deem most important, and whether they feel their academic training has adequately addressed these topics. Combined, these findings help identify curricular strengths and gaps in ethics education.

Table 6 presents the distribution of responses regarding the perceived relevance of 21 ethical principles in AI systems, according to students' self-assessed understanding and experience (Q7). The majority of students rated almost all principles as highly or very highly relevant, indicating a strong perceived importance of ethical considerations in computing.

The top-rated principles in terms of Strongly agree were Data Security (80.4%), Fairness (78.5%), Privacy (75.0%), Human dignity (75.0%), and Safety (73.2%). These results

suggest that students prioritize principles directly associated with user protection, harm prevention, and technical robustness, reflecting contemporary concerns about data misuse and algorithmic risks.

Principles such as Accountability, Non-maleficence, Freedom, and Transparency also received high levels of agreement (above 69% in Strongly agree), reinforcing an emphasis on fairness and social responsibility. In contrast, Beneficence, Prosperity, and Predictability showed comparatively lower levels of Strongly agree responses (35.7%, 48.3%, and 48.3%, respectively) as shown in Table 6. Although these principles were still predominantly rated as Agree or Strongly agree, the distribution suggests comparatively lower salience or curricular emphasis.

The proportion of neutral responses was minimal across all principles (mostly under 20%), and disagreement levels were negligible suggesting students do not reject the importance of any of the listed ethical values but do perceive them with varying degrees of significance. Overall, the data from Table 6 indicate that students clearly value ethical concerns in AI and computing systems, with greater emphasis on technical robustness, privacy, and fairness-oriented principles.

Regarding the perceived importance of ethical principles in computing education (Q8), the vast majority of participants expressed a positive attitude. As illustrated in Figure 4, 50.0% of respondents strongly agreed and 39.3% agreed that learning about ethical principles during their undergraduate studies is important. Only 8.9% indicated a neutral stance, and minimal disagreement was observed (1.8% strongly disagreed and no respondents selected disagree). These results indicate a clear consensus among students about the relevance of ethics in computing curricula, reinforcing the argument for formally integrating such content across courses.

In contrast, the responses to Q9 (During your course, did you have a subject that addressed these ethical principles?) revealed a mismatch between perceived importance and actual curricular experience. As shown in Figure 4, only 14.3% strongly agreed and 26.8% agreed that they had taken a subject covering ethical principles. A considerable portion of

Table 6. Perceived relevance of ethical principles in AI (Q7)

Ethical Principle	Strongly disagree	Disagree	Neither agree or disagree	Agree	Strongly agree
Transparency	0.0%	1.8%	1.8%	26.8%	69.6%
Privacy	0.0%	1.8%	0.0%	23.2%	75.0%
Accountability	0.0%	3.6%	0.0%	25.0%	71.4%
Fairness	1.8%	5.4%	0.0%	14.3%	78.5%
Autonomy	1.8%	3.6%	0.0%	32.1%	62.5%
Explainability	1.8%	0.0%	0.0%	35.7%	62.5%
Justice	3.6%	1.8%	0.0%	25.0%	62.5%
Non-maleficence	1.8%	0.0%	0.0%	26.8%	71.4%
Human dignity	1.8%	3.6%	0.0%	19.6%	75.0%
Beneficence	7.1%	19.7%	0.0%	37.5%	35.7%
Responsibility	1.8%	5.4%	0.0%	28.5%	64.3%
Safety	0.0%	3.6%	0.0%	23.2%	73.2%
Data security	0.0%	0.0%	0.0%	19.6%	80.4%
Sustainability	5.4%	0.0%	0.0%	33.9%	60.7%
Freedom	0.0%	0.0%	10.7%	17.9%	71.4%
Solidarity	1.8%	3.6%	0.0%	28.6%	66.0%
Prosperity	7.1%	5.4%	17.8%	21.4%	48.3%
Effectiveness	3.6%	1.8%	0.0%	35.6%	59.0%
Accuracy	7.1%	1.8%	10.7%	21.4%	59.0%
Predictability	3.6%	1.8%	15.9%	30.4%	48.3%
Interpretability	7.1%	3.6%	7.1%	26.8%	55.4%

students reported a neutral response (25.5%) or expressed disagreement (26.3% disagree and 7.1% strongly disagree). This distribution suggests that although students highly value ethics education, most of them did not perceive sufficient formal exposure to these topics during their academic training.




Figure 4. Distribution of responses to Q8 and Q9 on the perceived importance of learning ethical principles (Q8) and the extent to which such content was addressed in their curriculum (Q9)

RQ.2 Summary: Students demonstrated strong agreement regarding the relevance of ethical principles in AI systems, particularly those associated with user protection, fairness, and harm prevention such as Data Security, Fairness, Privacy, Human Dignity, and Safety. These results suggest that students prioritize principles directly linked to safeguarding individuals and ensuring responsible system behavior. Conversely, principles such as Beneficence, Prosperity, and Predictability received comparatively lower levels of strong endorsement, which may indicate either lower conceptual familiarity or reduced curricular emphasis. Most students consider learning about ethical principles important during their undergraduate program. However, a significant gap remains between this perceived importance and their actual curricular exposure to ethics-related content. This disconnect highlights a curricular shortfall that may limit students' ability to critically engage with the broader ethical dimensions of computing.

5 Results Study 2

5.1 Demographic Profile

A total of 31 computing educators participated in the interview phase. All participants provided an email address (IQ1) and accepted the informed consent terms (IQ2), resulting in a 100% consent rate.

Regarding academic affiliation (IQ3), most participants were linked to the Department of Computer Science (CIC) (58.1%), followed by the Faculty of Science and Technology in Engineering (FCTE) (29%) and the Department of Electrical Engineering (ENE) (12.9%).

Concerning the degree programs with which educators were affiliated (IQ4), the largest shares reported Bachelor's in Computer Science (40%) and Software Engineering (30%), followed by Network Engineering (13.3%) and Mechatronics Engineering (6.7%). Smaller proportions were affiliated with the Teaching Degree in Computing (6.7%) and Computer Engineering (3.3%).

In terms of teaching experience in the broader Computing field (IQ5), the sample was composed predominantly of senior educators: 83.9% reported teaching for 10 years or more. The remaining participants reported 5–10 years (9.7%) and 2–5 years (6.5%) of experience. No participants reported 1 year or less of teaching experience.

5.2 RQ.3: How do educators perceive their role, the main challenges, and institutional responsibilities in fostering ethics education in computing, particularly in the context of Generative AI?

To address RQ.3, we conducted an iterative qualitative coding process based on 31 semi-structured interviews with computing educators. The analysis considered both temporal evolution and contemporary perceptions regarding educators' roles, challenges, and institutional responsibilities in ethics education within the context of Generative AI.

5.2.1 IQ6: Evolution of Ethics Integration

The coding process for IQ6 revealed five interrelated dimensions that describe how ethics incorporation has evolved over time and how Generative AI has reshaped pedagogical practices. Table 7 presents the emergent categories and sub-categories identified in the analysis.

The most recurrent pattern refers to *Pedagogical and Assessment Reconfiguration* (n=8). Educators reported that the most concrete changes did not occur at the level of abstract ethical principles, but rather in evaluation strategies and classroom practices. The emergence of Generative AI prompted assessment reform, increased adoption of oral examinations, individualized verification of knowledge, explicit AI-use policies, and process-oriented evaluation models. These adaptations were frequently described as necessary to ensure authentic learning and prevent misuse of generative tools.

Closely related, the category *Academic Integrity and Responsibility Reinforcement* (n=7) captures concerns regarding authorship, plagiarism amplification, automation without understanding, and the need to reinforce student accountability. Several educators described a shift from prohibition-based approaches to models centered on transparency, disclosure, and formative guidance regarding ethical AI use. In this sense, Generative AI intensified existing integrity challenges rather than creating entirely new ethical problems.

At the same time, a significant group of participants (n=6) framed the transformation as one of *Continuity with Contextual Adaptation*. These educators emphasized that ethical concerns have long been part of their teaching practices and that Generative AI represents an additional technological layer requiring regulation and critical interpretation, but not a fundamental rupture in ethical stance.

The category *Historical Trajectory of Ethics Integration* (n=6) reflects educators' narratives about how ethics evolved from being treated in a normative or punctual manner toward more explicit, transversal, and systematically integrated approaches. Some participants described periods in which ethics lost curricular prominence and later regained relevance, particularly in response to emerging technologies.

Finally, a smaller but analytically relevant group (n=4) explicitly described Generative AI as a *Catalyst for Ethical Intensification*. These responses highlighted that AI expanded ethical complexity, exposed insufficiencies in prior approaches, and revealed gaps in AI ethical literacy among both students and faculty. As one educator stated: "*With generative AI, there was a shift in level: ethics stopped being just a topic and became a condition of technical quality*" (P#18).

Generative AI appears less as a disruptive ethical rupture and more as an operational accelerator that intensified existing concerns, reconfigured assessment practices, and made accountability tensions more visible within computing education.

Whereas IQ6 captured temporal and pedagogical transformations triggered or intensified by Generative AI, IQ7 shifts the analytical focus to educators' normative self-positioning, perceived responsibilities, and institutional challenges in ethics education.

5.2.2 IQ7 – Educators' perceived roles and responsibilities

Table 8 presents the distribution of participants across the identified macro-categories (IQ7). The frequencies indicate how many participants explicitly articulated ideas aligned with each theme. As interviews often contained overlapping dimensions, individual responses could contribute to more than one category.

The most salient theme concerned the central ethical responsibility of educators (8 participants), emphasizing the role of the professor as a primary ethical driver and moral reference in professional formation. This perception was reflected in one participant's statement: "*If it does not come from the faculty, it will not come from anywhere*" (P#2). This was followed by curricular structuring and integration of ethics (7 participants), reflecting concerns about whether ethics should be transversal, embedded, or structured as dedicated coursework. The ethical mediation of AI integration (6 participants) emerged as a distinct contemporary challenge, highlighting the need for AI literacy and responsible technological adoption. Dialogical and reflective ethics facilitation (6 participants) emphasized pedagogical mediation and structured debate, while institutional governance and structural gaps (4 participants) pointed to perceived deficiencies in formal guidance and coordinated institutional response.

5.2.3 IQ8: Institutional and regulatory strategies to keep pace with rapid AI advances

Question IQ8 investigated how educators believe institutions and legislators can keep pace with the rapid evolution of AI in order to anticipate and mitigate potential societal harms. The qualitative coding of the 31 responses revealed five macro-categories (Table 9), reflecting a spectrum of governance strategies that combine regulatory design, institutional capacity, stakeholder engagement, and public-oriented measures.

The most frequent theme, *Adaptive and Risk-Based Regulatory Approaches* (9 responses), emphasized that regulation should prioritize the *use* of AI rather than attempting to control technological development itself. As one educator stated: "*There is no way to regulate the development of technological innovation... The law should address the ethical and legal issues of technology use, not its development per se*" (P#1). Educators also highlighted the need for preventive and adaptive regulation, including risk-based frameworks, temporary or principle-based mechanisms, accelerated decision-making processes, and stronger enforcement to address misuse (e.g., disinformation). Several participants also stressed institutional liability and accountability as central levers for mitigating harms under uncertainty.

Equally prominent, *Dialogical and Multi-Stakeholder Governance* (9 responses) framed institutional readiness as dependent on sustained dialogue and collaboration among legislators, academia, professionals, and society. Participants advocated debate-based engagement, legislator researcher collaboration, research-driven monitoring, and the university's role as an intermediary that translates evidence into public communication and policy relevant insights. This perspective emphasized that effective governance requires con-

Table 7. Emergent categories from IQ6: Evolution of ethics integration and influence of Generative AI

Category	Sub-category	#
Pedagogical and Assessment Reconfiguration	Shift from theoretical discussion to practical ethical integration	8
	Assessment reform to prevent AI misuse	
	Oral assessment to verify authorship	
	Formal regulation of AI usage	
	Process-oriented assessment under AI influence	
	Individual verification of knowledge	
	Reactive adjustment to AI misuse	
Academic Integrity and Responsibility Reinforcement	Process monitoring and ethical collaboration under AI influence	7
	Learning vs. automation tension	
	Shift from prohibition to transparency and accountability	
	Explicit student accountability	
	Authorship challenges under AI	
	Plagiarism amplification due to AI	
	AI-induced academic integrity risk	
Continuity with Contextual Adaptation	Formative guidance for ethical AI use	6
	Ethics as complementary to technical focus + AI-driven critical verification	
	Ethical stance independent of AI influence	
	Intensification without rupture	
	Stable ethical posture despite AI	
	Long-standing ethical commitment	
	AI as contextual tool, not rupture	
Historical Trajectory of Ethics Integration	Progressive formalization of ethics + responsibility clarification under AI	6
	Curricular decline and reintegration of ethics	
	Dedicated ethical discussion sessions	
	Curricular incorporation via updated materials	
	Professional ethical framework integration	
	Critical ethical perspective expansion	
Generative AI as Catalyst for Ethical Intensification	AI as turning point for ethical awareness	4
	Insufficient AI ethical literacy	
	Expansion of ethical complexity due to AI	

Table 8. Macro-categories identified in IQ7 and number of participants ($n = 31$)

Macro-Category	#
Educator-Centered Ethical Responsibility	8
Curricular Structuring and Integration of Ethics	7
Ethical Mediation of AI Integration	6
Dialogical and Reflective Ethics Facilitation	6
Institutional Governance and Structural Gaps	4

textualized understanding of real-world impacts and continuous, collaborative revision of norms and practices.

The category *Institutional Capacity Building and Technical Expertise* (6 responses) focused on strengthening the State's and institutions' technical ability to evaluate and supervise AI systems. Educators recommended creating specialized units or departments, investing in technical competence, requiring transparency and auditability, and establishing oversight bodies such as ethics committees or multidisciplinary review groups. Some participants also explicitly deferred regulatory guidance to domain specialists, acknowledging limits of personal expertise.

In parallel, *Ethical Awareness, Education, and Societal Responsibility* (6 responses) underscored that formal regulation alone is insufficient without broader ethical culture and public preparedness. Responses highlighted continuous ethical awareness, normative clarity communicated through campaigns and multiple channels, responsiveness to societal demands, and long-term educational strengthening (including improvements in basic education) as foundational conditions for reducing harms and enabling responsible adoption.

Table 9. Macro-categories identified in IQ8 and number of participants ($n = 31$)

Macro-Category	#
Adaptive and Risk-Based Regulatory Approaches	9
Dialogical and Multi-Stakeholder Governance	9
Institutional Capacity Building and Technical Expertise	6
Ethical Awareness, Education, and Societal Responsibility	6
Uncertainty / Non-position	2

Finally, a small subset of responses expressed *Uncertainty/Non-position* (2 responses), indicating either inability to propose strategies or lack of confidence to address regulatory questions. Although infrequent, these responses suggest that uncertainty and perceived knowledge gaps may also be present among educators when reflecting on governance and legislative responsibilities.

5.2.4 IQ9: Ethical Maturity at the Transition from University to Professional Practice

The analysis of the responses to IQ9 reveals a predominantly critical perception of students' ethical maturity at the moment of transition from university to professional practice. The most recurrent macro-category was *Low Ethical Maturity at Graduation* (n=13), indicating that many educators perceive students as insufficiently prepared to fully understand the gravity and breadth of the consequences of their technical decisions. Responses highlighted age-related immaturity, limited theoretical grounding in ethical reasoning, low awareness regarding misuse of AI tools (e.g., plagiarism), insufficient perception of accountability, and a general underestimation of the broader social and professional implications of their actions. As one educator observed: "*They relativize it because they are not yet legally accountable*" (P#19).

At the same time, the second most frequent category, *Developing but Incomplete Ethical Awareness* (n=7), suggests a more nuanced perspective (Table 10). Several educators observed that while students often demonstrate emerging ethical sensitivity, particularly regarding issues such as bias, privacy, and responsible AI use, they frequently struggle to operationalize that awareness into concrete technical procedures. Ethical understanding, in these cases, remains conceptual rather than procedural, and maturity appears heterogeneous across students.

Table 10. Summary of Macro-Categories Identified in IQ9

Macro-Category	#
Low Ethical Maturity at Graduation	13
Developing but Incomplete Ethical Awareness	7
Epistemic Uncertainty and Transitional Adaptation	4
Maturity as Experiential and Context-Dependent	4
Institutional and Curricular Responsibility for Ethical Formation	3

The category *Epistemic Uncertainty and Transitional Adaptation* (n=4) reflects the perception that ethical maturity is not exclusively a student limitation, but part of a broader societal learning process. Some participants emphasized that both students and professors are still adapting to rapid technological changes, especially those driven by generative AI. In this view, ethical formation occurs within a context of collective uncertainty and evolving norms.

Closely related, *Maturity as Experiential and Context-Dependent* (n=4) captures the understanding that ethical consolidation tends to occur through real-world responsibility. Several educators argued that students' maturity increases progressively across the academic trajectory and is significantly strengthened in professional environments where ac-

countability, documentation, negotiation of constraints, and exposure to real impacts become unavoidable. Ethical development, therefore, is seen as experiential rather than fully achieved at graduation.

Finally, although less frequent (n=3), *Institutional and Curricular Responsibility for Ethical Formation* indicates that low maturity is also attributed to insufficient structured debate, limited curricular integration of ethics, and lack of institutional reinforcement mechanisms. These responses suggest that the transition gap cannot be understood solely as a developmental issue, but also as a reflection of structural educational limitations.

The findings reveal a dominant perception of insufficient ethical maturity at graduation, tempered by recognition that ethical awareness is emerging but not yet fully consolidated. Ethical competence at the transition to professional practice is therefore perceived as heterogeneous, evolving, and strongly influenced by experiential learning and institutional support structures.

5.2.5 IQ10: Barriers to Effective Ethics Education in Computing Programs

The analysis of the 31 responses to IQ10 indicates that the barriers to a more effective ethics education in computing programs are predominantly structural and formative rather than merely individual or conceptual. Instead of attributing limitations solely to student disinterest or to the intrinsic complexity of ethical debates, educators emphasized systemic constraints related to faculty preparation, curriculum design, and institutional governance.

The most recurrent macro-category was *Lack of Faculty Preparation and Ethical Literacy* (n=8), as shown in Table 11. Participants highlighted insufficient specific training in ethics applied to computing, conceptual misunderstandings regarding what constitutes ethical reflection, limited exposure to sociotechnical issues during their own academic formation, and insecurity in mediating complex ethical discussions. Several responses also pointed to limited practical engagement with ethical applications in real-world contexts. As one participant emphasized: "*If the professor does not know, it becomes difficult to guide students*" (P#10). These findings suggest that the effectiveness of ethics education depends not only on curricular inclusion but also on the educator's ability to contextualize, translate, and operationalize ethical reasoning within technical domains.

Table 11. Summary of Macro-Categories Identified in IQ10

Macro-Category	#
Lack of Faculty Preparation and Ethical Literacy	8
Curricular Overload and Fragmentation	8
Institutional Gaps and Governance Deficiencies	7
Student Engagement and Maturity Challenges	3
Technological Acceleration and External Pressures	3
Technicist Culture and Humanistic Disconnection	2

With the same frequency (n=8), the category *Curricular Overload and Fragmentation* emerged as a central barrier. Educators reported that computing curricula remain predominantly technical, leaving little room for structured ethical

reflection. Ethics is often treated as a peripheral or isolated discipline, rather than as a transversal dimension integrated into engineering processes, system design, and development cycles. The expansion of technological domains within computing programs further intensifies competition for curricular space, limiting opportunities for interdisciplinary integration and sustained ethical engagement.

The third most frequent category was *Institutional Gaps and Governance Deficiencies* (n=7). Respondents noted the absence of formalized institutional guidelines, weak normative alignment between discourse and curricular practice, limited structured debate within departments, and a lack of accountability mechanisms to ensure ethical integration. Some also highlighted disconnections between legal frameworks (e.g., data protection regulations) and pedagogical implementation. These findings indicate that strengthening ethics education requires coordinated institutional policies and governance mechanisms, not solely individual faculty initiative.

Less frequent but still significant were *Student Engagement and Maturity Challenges* (n=3), referring to perceived student immaturity, low interest in ethical discussions, and difficulty in recognizing the long-term consequences of unethical technical decisions; and *Technological Acceleration and External Pressures* (n=3), which captured concerns about the rapid pace of technological change and the influence of corporate and market forces on academic priorities.

Finally, although less recurrent (n=2), the category *Technicist Culture and Humanistic Disconnection* reveals a deeper epistemological tension within computing education: the persistent separation between technical and humanistic domains. Some educators argued that overcoming this divide requires the humanization of computing education and a broader recognition that ethical reasoning is constitutive of technical practice rather than external to it. The findings suggest that the barriers to effective ethics education in computing are perceived as systemic and multidimensional, involving curricular structure, faculty formation, institutional governance, and broader cultural dynamics within the field.

Taken together, IQ9 and IQ10 suggest that the perceived limitations in students' ethical maturity are not merely developmental, but structurally connected to curricular fragmentation, insufficient faculty preparation, and institutional governance gaps identified in earlier responses.

RQ.3 Summary: The qualitative analysis of 31 computing educators revealed that Generative AI has functioned primarily as an operational accelerator rather than a disruptive ethical rupture. In IQ6, participants described pedagogical and assessment reconfiguration as the most prominent transformation, followed by reinforcement of academic integrity and contextual continuity in ethical stance. In IQ7, educators positioned themselves as central ethical agents while simultaneously acknowledging curricular and institutional gaps. In IQ8, governance strategies emphasized adaptive, risk-based regulation and multi-stakeholder dialogue. IQ9 exposed a dominant perception of low ethical maturity at graduation, although signs of emerging ethical awareness were identified. Finally, IQ10 highlighted structural barriers, particularly related to faculty preparation, curricular overload, and institutional governance deficiencies. Overall, the findings suggest that ethics education in computing is evolving, but remains constrained by systemic, curricular, and formative limitations that hinder full consolidation of ethical competence in the context of rapid AI transformation.

6 Discussion

The integrated interpretation proposed in Figure 1 provides the conceptual basis for discussing the empirical findings presented in this section. Rather than interpreting ethical literacy as a single observable competency, this study conceptualizes it as a multidimensional construct integrating ethical knowledge, curricular exposure, perceived relevance, and the ability to operationalize ethical reasoning within professional Software Engineering contexts Morley et al. (2023); Yang et al. (2025). Accordingly, the structural ethical literacy gap identified in this study should be understood not merely as limited ethical awareness, but as a broader difficulty in translating ethical principles into professional practice. This interpretation is consistent with recent studies arguing that the main challenge in AI ethics lies not only in defining ethical principles but also in operationalizing them through concrete organizational and educational practices Andrade and Canedo (2026). From this perspective, familiarity, perceived relevance, preparedness, and perceived ethical maturity represent complementary dimensions through which ethical literacy can be understood Grosz et al. (2019); Palmquist et al. (2026). Together, these complementary dimensions provide the analytical basis for interpreting ethical literacy throughout this study.

Figure 1 synthesizes this multidimensional perspective by illustrating how familiarity, perceived relevance, curricular exposure, and ethical maturity jointly support the conceptualization of ethical literacy and the identification of the structural ethical literacy gap. Rather than representing a theoretical framework proposed *a priori* or a validated causal model, the figure summarizes the analytical relationships emerging from the integration of the quantitative and qualitative findings.

6.1 The Ethical Literacy Gap: Perceived Preparedness versus Structural Exposure

The sequential explanatory integration of quantitative and qualitative findings reveals a central cross-phase tension that we conceptualize as a structural ethical literacy gap. Quantitative evidence indicated high levels of student self-perceived preparedness regarding ethical responsibility, whereas qualitative findings from educators revealed concerns regarding ethical maturity and students' ability to operationalize ethical reasoning in realistic professional situations. On the one hand, students report high levels of perceived preparedness to assess the social impact of their decisions (Q4) and demonstrate strong agreement regarding the importance of ethical principles in computing education (Q8). On the other hand, a significant portion report limited formal exposure to structured ethics instruction (Q5 and Q9).

This discrepancy becomes more pronounced when contrasted with educators' perceptions. In IQ9, the dominant view among instructors was that students graduate with low or insufficient ethical maturity, particularly regarding the gravity and breadth of professional consequences. While some emerging ethical awareness was acknowledged, educators frequently emphasized students' difficulty in operationalizing ethical reasoning in concrete technical decisions.

These findings suggest that ethical confidence may not necessarily reflect consolidated ethical competence. The integration of both phases reveals a cross-phase explanatory pattern: students reported relatively high perceived preparedness and strong recognition of ethical principles, despite limited formal exposure to ethics instruction, whereas educators frequently perceived ethical maturity as insufficient and difficult to operationalize in realistic professional contexts. This interpretation is consistent with the multidimensional conceptualization proposed in Figure 1, where ethical literacy emerges from the convergence of complementary dimensions rather than from any single indicator.

This multidimensional interpretation also aligns with recent international discussions indicating that AI literacy education increasingly requires interdisciplinary competencies that combine technical understanding with ethical and socio-emotional dimensions Palmquist et al. (2026); Yang et al. (2025).

Although these findings suggest the presence of a structural ethical literacy gap, this interpretation should be understood within the specific institutional and educational context investigated in this study. The results are grounded in data collected from students and educators at a single Brazilian public university and may reflect contextual characteristics related to curricular organization, institutional culture, and participant distribution across computing programs and academic stages. Therefore, rather than proposing a universally generalizable condition, we interpret the identified tensions as an empirically grounded pattern observed within this setting that may inform broader discussions on ethics education and Generative AI in computing programs.

6.2 Generative AI as an Ethical Accelerator Rather than a Rupture

Educators' responses to IQ6 indicate that the emergence of Generative AI has not fundamentally altered their ethical stance, but rather intensified and operationalized pre-existing tensions. Many participants framed Generative AI as an accelerator of ethical concerns, particularly in relation to assessment practices, authorship verification, academic integrity, and accountability.

Rather than describing a conceptual rupture, instructors emphasized pedagogical reconfiguration: adoption of oral assessments, process-oriented evaluation, explicit AI-use policies, and closer monitoring of collaborative work. Generative AI exposed vulnerabilities in evaluation models that previously relied heavily on product-based outputs. In this sense, the technology functioned less as a new ethical domain and more as a stress test for existing curricular fragilities.

This interpretation aligns with the broader perception that many ethical challenges attributed to Generative AI, including plagiarism, superficial automation, and epistemic opacity are extensions of long-standing issues related to technological mediation and academic responsibility. However, the scale and accessibility of GenAI tools amplified these dynamics, making ethical governance and pedagogical adaptation more urgent.

6.3 Curricular Fragmentation and Institutional Governance Constraints

The barriers identified in IQ10 reinforce the structural nature of the ethical literacy gap. Faculty preparation deficits, curricular overload, and institutional governance gaps emerged as dominant constraints. Ethics is often treated as peripheral, isolated, or dependent on individual instructor initiative rather than embedded as a transversal dimension of technical formation.

This fragmentation generates a paradox: while institutions rhetorically endorse responsible AI development, curricular structures remain predominantly technicist, leaving limited space for sustained ethical reasoning. The expansion of emerging technical domains such as AI, data science, and cybersecurity further intensifies competition for curricular time, frequently displacing humanistic integration.

Moreover, educators reported limited institutional guidance, absence of coordinated ethical frameworks, and weak alignment between regulatory discourse and pedagogical implementation. These findings are consistent with recent studies showing that, although organizations increasingly recognize the importance of AI ethics, they continue to face challenges in translating ethical principles into concrete governance mechanisms and operational practices (Andrade and Canedo, 2026; Canedo et al., 2026). Accordingly, strengthening ethical literacy depends not only on curricular inclusion or individual competencies but also on institutional governance structures capable of formalizing expectations, supporting faculty development, and operationalizing ethical principles across educational and professional contexts.

6.4 Ethical Maturity at Transition: Experiential Development and Academic Responsibility

The perception of low ethical maturity at graduation raises important questions regarding the locus of ethical consolidation. Several educators described ethical maturity as experiential and context-dependent, strengthening primarily through real-world accountability and professional exposure. While experiential development is a natural component of professional growth, reliance on post-graduation maturation may leave transitional vulnerabilities.

If ethical competence develops progressively as ethical maturity evolves under professional pressure, academic environments risk underpreparing students for complex sociotechnical decision-making. Universities therefore hold a formative responsibility not merely to introduce ethical vocabulary, but to simulate conditions of responsibility, trade-off negotiation, and impact assessment within the curriculum itself.

The heterogeneity observed in students' maturity levels further indicates that ethical development is uneven and contingent upon individual engagement, course exposure, and faculty mediation. Strengthening structured ethical formation may reduce this variability and support more consistent professional readiness.

6.5 Implications for AI-Assisted Software Engineering Education and Practice

The findings point toward several implications for computing education under the pressure of Generative AI. First, ethical literacy should be conceptualized as a core technical competency rather than an auxiliary normative add-on. Embedding ethical reasoning into development cycles, design documentation, and assessment rubrics may facilitate operational integration.

Second, faculty development emerges as a central requirement. Without structured support, training opportunities, and interdisciplinary collaboration, instructors may feel insecure mediating complex ethical debates, particularly in rapidly evolving AI contexts.

Third, institutional governance mechanisms should articulate explicit guidelines regarding AI usage, assessment design, and accountability structures. Such alignment reduces reliance on isolated initiatives and promotes coherent ethical culture across departments.

From a Software Engineering perspective, the findings indicate that ethical literacy becomes a foundational competency throughout AI-assisted software development workflows, including requirements elicitation, software design, code generation, testing, debugging, documentation, and maintenance. In these activities, ethical literacy supports the critical evaluation of AI-generated artifacts, assessment of code and data provenance, negotiation of quality and socio-technical trade-offs, and accountability for AI-supported engineering outcomes. As Generative AI becomes increasingly embedded in everyday software development, ethical reasoning can no longer remain restricted to isolated classroom discussions; it must inform decisions throughout the software

lifecycle. Consequently, the structural ethical literacy gap identified in this study represents not only an educational challenge but also a potential risk to responsible Software Engineering practice. This observation also reinforces recent evidence that responsible AI depends on organizational governance structures capable of translating ethical principles into operational practices throughout the software lifecycle (Canedo et al., 2026).

Overall, the integration of students' and educators' perspectives demonstrates that ethics education in computing is evolving, yet remains constrained by structural, curricular, and governance limitations. Addressing the ethical literacy gap requires systemic coordination, pedagogical intentionality, and institutional commitment, particularly in contexts where Generative AI is rapidly reshaping both professional practice and educational environments.

Our findings also resonate with evidence reported in international contexts outside Latin America. Previous studies on AI literacy and ethics education, particularly in European and broader international educational settings, suggest that AI-related competencies should extend beyond technical knowledge and include ethical awareness, socio-emotional competencies, and critical reflection regarding the societal implications of AI systems Palmquist et al. (2026); Yang et al. (2025). Similar concerns emerged in our findings, especially regarding students' perceived preparedness, institutional constraints, and the need for interdisciplinary ethical formation. However, while prior studies mainly discuss broad educational frameworks and policy recommendations, our results provide empirical evidence from a Latin American context, where such discussions remain comparatively underexplored. This comparison suggests that although ethical and educational challenges surrounding AI appear globally shared, regional educational structures and institutional conditions may shape how these challenges are experienced and addressed.

7 Threats to Validity

Following the framework proposed by Wohlin et al. (2012), we discuss the main threats to the validity of this mixed-methods study and the strategies adopted to mitigate them.

Conclusion Validity concerns the accuracy of statistical inferences and the robustness of interpretive claims. In the quantitative phase, a potential threat lies in the relatively small student sample ($n=56$ students), which limits statistical power for inferential analysis. To mitigate this risk, we relied on descriptive statistics and avoided causal or general claims beyond the scope of the data. In the qualitative phase ($n=31$ educators), interpretive validity may be affected by researcher bias in coding and category consolidation. To address this, the coding process followed an iterative procedure, and emerging categories were discussed among the authors to improve consistency and reduce individual interpretive bias.

Internal Validity refers to the extent to which observed patterns can be attributed to the phenomena under investigation rather than external factors. In the survey phase, the self-reported nature of responses introduces the risk of so-

cial desirability bias, particularly regarding agreement with ethical principles. To mitigate this, anonymity was ensured and participants were informed that there were no right or wrong answers. In the interview phase, educators may have provided normatively framed responses aligned with professional expectations. The semi-structured format and the emphasis on reflective experiences were intended to encourage nuanced accounts rather than socially desirable positioning.

Construct Validity relates to how accurately the instruments capture the intended concepts. In the quantitative instrument, ethical principles were derived from established AI governance frameworks and accompanied by concise definitions to reduce conceptual ambiguity. Nevertheless, variation in students' interpretation of abstract concepts (e.g., autonomy, solidarity, interpretability) may have influenced responses. In the qualitative phase, the operationalization of categories such as "ethical maturity" or "institutional barriers" depends on interpretive coding decisions. To strengthen construct validity, coding was conducted systematically, and category refinement occurred iteratively to ensure conceptual coherence and alignment with participants' expressed meanings. Additionally, some ethical principles may present conceptual proximity (e.g., explainability versus interpretability and accountability versus responsibility), which may have affected participants' differentiation and interpretation despite the definitions provided.

External Validity concerns the generalizability of findings. Both phases relied on participants from a single Brazilian public university. Furthermore, the student sample presented a concentration of participants in earlier semesters and uneven representation across computing programs, which may have influenced perceptions regarding curricular exposure and ethical maturity. Cultural, regulatory, and curricular differences may produce distinct dynamics in other regions or countries. However, the integration of findings with national and international literature strengthens analytical generalization and situates the results within broader debates on AI ethics education. Future multi-institutional and cross-cultural studies would further enhance generalizability.

Finally, as this study integrates quantitative and qualitative findings in a sequential explanatory design, an additional validity consideration concerns the coherence of integration. While the two phases were conducted sequentially and analyzed independently, integration occurred at the interpretation stage. Although this approach allows complementary insights, alternative integration strategies (e.g., joint displays or triangulation matrices) might further strengthen cross-phase analytic rigor in future research.

8 Conclusion and Future Work

This study investigated ethical literacy in computing education under the pressure of Generative AI through a sequential explanatory mixed-methods design combining students' and educators' perspectives. By integrating quantitative evidence from undergraduate students with qualitative insights from computing educators, the research advances understanding of how ethical principles are perceived, operationalized, and institutionally supported within a Brazilian university con-

text.

The findings suggest patterns associated with a structural ethical literacy gap within the investigated educational context. Students overwhelmingly recognize the importance of ethical principles and report confidence in their ability to assess social impacts. However, many indicate limited formal curricular exposure to structured ethics instruction. From the educators' perspective, ethical maturity at graduation is often perceived as insufficient or heterogeneous, and Generative AI has intensified pedagogical tensions related to assessment, authorship, and accountability. Rather than constituting a conceptual rupture, Generative AI appears to function as an operational accelerator, exposing curricular fragmentation, faculty preparation deficits, and institutional governance gaps.

These results suggest that strengthening ethics education in computing requires systemic alignment rather than isolated curricular adjustments. Ethical literacy should be institutionalized as a core technical competency embedded across development processes, assessment strategies, and institutional policies, supported by faculty development, interdisciplinary collaboration, and explicit governance frameworks regarding AI use.

More broadly, this study contributes a multidimensional conceptual interpretation of ethical literacy that explains the tensions observed between students' perceptions, educators' assessments, and institutional conditions in AI-assisted computing education. By framing ethics education as a structural and institutional challenge amplified by Generative AI, the research extends discussions beyond content inclusion toward broader questions of curricular coherence and professional formation.

From a Software Engineering perspective, our findings reinforce that ethical literacy is becoming an essential competency for professionals working in AI-assisted development environments. As Generative AI increasingly supports activities such as requirements elicitation, software design, code generation, testing, documentation, software maintenance, and architectural decision-making, future software engineers must be prepared to critically evaluate AI-generated artifacts, assume accountability for engineering decisions, and negotiate socio-technical trade-offs throughout the software lifecycle. Accordingly, ethical literacy should be treated not as an auxiliary educational topic, but as a foundational competency for responsible AI-assisted Software Engineering.

Future work includes expanding the study to multiple institutions and regions to enhance analytical generalization and examine contextual variations. Longitudinal studies could further explore how ethical maturity develops across academic trajectories and early professional experiences. Additionally, future research may investigate the effectiveness of specific pedagogical interventions such as embedded ethics modules, competency-based frameworks, and AI-integrated assessment redesign in strengthening ethical literacy in computing programs.

Data Availability

All materials produced during this research are publicly available to ensure transparency and reproducibility. The dataset includes: (i) the full student survey instrument; (ii) anonymized responses from the 56 undergraduate participants; (iii) the semi-structured interview protocol applied to computing educators; (iv) the anonymized qualitative interview dataset; and (v) the complete qualitative codebook developed during the iterative coding process, including categories, sub-categories, and coding descriptions.

These materials are available on Zenodo at <https://zenodo.org/records/18828956>.

Acknowledgements

This study was financed in part by the National Council for Scientific and Technological Development (CNPq), Grants No. 300883/2025-0 and No. 406266/2025-5.

References

- Al-Kaswan, A. and Izadi, M. (2023). The (ab)use of Open Source Code to Train Large Language Models. arXiv:2302.13681 version: 2.
- Aler Tubella, A., Mora-Cantalops, M., and Nieves, J. C. (2024). How to teach responsible AI in Higher Education: challenges and opportunities. *Ethics and Information Technology*, 26(1):3.
- Andrade, R. B. and Canedo, E. D. (2026). The gap between ethical discourse and organizational practice in the responsible development, deploy, and use of AI systems: An exploratory study. In *Proceedings of the 19th International Conference on Cooperative and Human Aspects of Software Engineering, CHASE 2026, Rio de Janeiro, Brazil, April 12-18, 2026*, pages 110–114, <https://doi.org/10.1145/3794860.3794900>. ACM.
- Association for Computing Machinery (2018). Acm code of ethics and professional conduct.
- Batool, A., Zowghi, D., and Bano, M. (2025). AI governance: a systematic literature review. *AI Ethics*, 5(3):3265–3279.
- Brown, N., Xie, B., Sarder, E., Fiesler, C., and Wiese, E. S. (2024a). Teaching ethics in computing: A systematic literature review of acm computer science education publications. *ACM Trans. Comput. Educ.*, 24(1).
- Brown, N., Xie, B., Sarder, E., Fiesler, C., and Wiese, E. S. (2024b). Teaching Ethics in Computing: A Systematic Literature Review of ACM Computer Science Education Publications. *ACM Transactions on Computing Education*, 24(1):1–36.
- Brown, N., Xie, B., Sarder, E., Fiesler, C., and Wiese, E. S. (2024c). Teaching ethics in computing: A systematic literature review of ACM computer science education publications. *ACM Trans. Comput. Educ.*, 24(1):6:1–6:36.
- Canedo, E. D., Mendes, F. F., Andrade, R., de Cerqueira, J. S., and Abrahamsson, P. (2026). Operationalizing AI ethics in the public sector: A cross-context replication in brazil. In *Proceedings of the 19th International Conference on Cooperative and Human Aspects of Software Engineering, CHASE 2026, Rio de Janeiro, Brazil, April 12-18, 2026*, pages 126–137, <https://doi.org/10.1145/3794860.3794909>. ACM.
- Carvalho, L. P., Oliveira, J., and Santoro, F. M. (2021). A presença de conteúdos sobre ética computacional na literacia em computação institucional brasileira. In *ABCiber Virtual Meeting Proceedings (Online)*. *ABCiber Virtual Meeting*, pages 1–22. SBC.
- Corrêa, N. K., Galvão, C., Santos, J. W., Del Pino, C., Pinto, E. P., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., and De Oliveira, N. (2023). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns*, 4(10):100857.
- da Rocha Braz, A. C., Rossi, J., and Canedo, E. D. (2026). Bridging data protection and AI ethics: A two-study empirical examination of LGPD principles and ethical AI in brazil. *J. Interact. Syst.*, 17(1):257–273.
- de Cerqueira, J. A. S., Azevedo, A. P. D., Leão, H. A. T., and Canedo, E. D. (2022). Guide for artificial intelligence ethical requirements elicitation - RE4AI ethical guide. In *55th Hawaii International Conference on System Sciences, HICSS 2022, Virtual Event / Maui, Hawaii, USA, January 4-7, 2022*, pages 1–10, <http://hdl.handle.net/10125/80015>. ScholarSpace.
- de Cerqueira, J. A. S., Leão, H. A. T., and Canedo, E. D. (2021). Ethical guidelines and principles in the context of artificial intelligence. In Araújo, R. D., Dorça, F. A., de Araujo, R. M., Siqueira, S. W. M., and Fontão, A. L., editors, *SBSI 2021: XVII Brazilian Symposium on Information Systems, Uberlândia, Brazil, June 7 - 10, 2021*, pages 36:1–36:8, <https://doi.org/10.1145/3466933.3466969>. ACM.
- De La Vega Hernández, I. M., Urdaneta, A. S., and Carayannis, E. (2023). Global bibliometric mapping of the frontier of knowledge in the field of artificial intelligence for the period 1990–2019. *Artificial Intelligence Review*, 56(2):1699–1729.
- Ferreira, R. and Vardi, M. Y. (2021). Deep Tech Ethics: An Approach to Teaching Social Justice in Computer Science. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, pages 1041–1047, Virtual Event USA. ACM.
- Fiesler, C., Garrett, N., and Beard, N. (2020). What Do We Teach When We Teach Tech Ethics?: A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, pages 289–295, Portland OR USA. ACM.
- Floridi, L. (2023). *The ethics of artificial intelligence: principles, challenges, and opportunities*. Oxford University Press, Oxford ; New York, NY. OCLC: on1376727516.
- Goetze, T. S. (2023). Integrating Ethics into Computer Science Education: Multi-, Inter-, and Transdisciplinary Approaches. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, pages 645–651, Toronto ON Canada. ACM.
- Grosz, B. J., Grant, D. G., Vredenburg, K., Behrends, J., Hu, L., Simmons, A., and Waldo, J. (2019). Embedded EthiCS: integrating ethics across CS education. *Commun.*

- ACM, 62(8):54–61.
- Hoda, R. (2025). Qualitative research with socio-technical grounded theory a practical guide to qualitative data analysis and theory development in the digital world. *Innovations*, pages 1–379.
- Institute of Electrical and Electronics Engineers (2020). Ieee code of ethics, ieee-cs/acm joint task force on software engineering ethics and professional practices.
- Jantunen, M., Meyes, R., Kurchyna, V., Meisen, T., Abrahamsson, P., and Mohanani, R. (2024). Researchers' concerns on artificial intelligence ethics: Results from a scenario-based survey. In *Proceedings of the 7th ACM/IEEE International Workshop on Software-intensive Business, IWSiB 2024, Lisbon, Portugal, 16 April 2024*, pages 24–31, <https://doi.org/10.1145/3643690.3648238>. ACM.
- Jobin, A., Ienca, M., and Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399.
- Khan, A. A., Akbar, M. A., Fahmideh, M., Liang, P., Waseem, M., Ahmad, A., Niazi, M., and Abrahamsson, P. (2023). AI ethics: An empirical study on the views of practitioners and lawmakers. *IEEE Trans. Comput. Soc. Syst.*, 10(6):2971–2984.
- Khan, A. A., Badshah, S., Liang, P., Khan, B., Waseem, M., Niazi, M., and Akbar, M. A. (2021). Ethics of ai: A systematic literature review of principles and challenges.
- Kiesler, N., Opel, S., and Thorbrügge, C. (2024). With great power comes great responsibility - integrating data ethics into computing education. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. I, ITiCSE 2024*, page 471–477. Association for Computing Machinery, New York, NY, USA.
- Kılınç, S. (2024). Comprehensive ai assessment framework: Enhancing educational evaluation with ethical ai integration. *Journal of Educational Technology and Online Learning*, 7(4-ICETOL 2024 Special Issue):521–540.
- Kipman, F., dos Santos Marques, G., and Canedo, E. D. (2025). Computing students' perceptions of ethical principles in AI and curricular coverage. In Santos, G., Reinehr, S. S., de Farias Júnior, I., Gadelha, B., Barcellos, M., Freire, S., de França, B. B. N., Canedo, E. D., Oran, A. C., Matsubara, P., and Parizi, R., editors, *Proceedings of the 24th Brazilian Symposium on Software Quality, SBQS 2025, São José dos Campos, SP, Brazil, November 4-7, 2025*, pages 431–441. SBC.
- Kumar, A. N., Raj, R. K., Aly, S. G., Anderson, M. D., Becker, B. A., Blumenthal, R. L., Eaton, E., Epstein, S. L., Goldweber, M., Jalote, P., Lea, D., Oudshoorn, M., Pias, M., Reiser, S., Servin, C., Simha, R., Winters, T., and Xiang, Q. (2024). *Computer Science Curricula 2023*. ACM, New York, NY, USA.
- Mak, J., Nakatumba-Nabende, J., Clear, T., Clear, A., Abluwi, I., Andrei, O., Angeli, L., MacNeil, S., Oyelere, S. S., Hale Rattigan, M., Sheard, J., and Zhu, T. (2026). Navigating the ethical and societal impacts of generative ai in higher computing education. In *Proceedings of the 2025 Working Group Reports on Innovation and Technology in Computer Science Education, ITiCSE-WGR 2025*, page 313–359. Association for Computing Machinery, New York, NY, USA.
- Miller, K. (1988). Integrating Computer Ethics into the Computer Science Curriculum. *Computer Science Education*, 1(1):37–52.
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., and Floridi, L. (2023). Operationalising AI ethics: barriers, enablers and next steps. *AI & SOCIETY*, 38(1):411–423.
- Nielsen, N. R. (1972). Social responsibility and computer education. *SIGCSE Bull.*, 4(1):90–96.
- Palmquist, A., Sigurdardottir, H. D. I., and Myhre, H. (2026). Exploring interfaces and implications for integrating social-emotional competencies into ai literacy for education: a narrative review. *Journal of Computers in Education*, 13(1):127–163.
- Pant, A., Hoda, R., Tantithamthavorn, C., and Turhan, B. (2024). Ethics in AI through the practitioner's view: a grounded theory literature review. *Empirical Software Engineering*, 29(3):67.
- Pasricha, S. (2023a). Ethics in computing education: Challenges and experience with embedded ethics. In *Proceedings of the Great Lakes Symposium on VLSI 2023, GLSVLSI '23*, page 653–658, New York, NY, USA. Association for Computing Machinery.
- Pasricha, S. (2023b). Ethics in computing education: Challenges and experience with embedded ethics. In Thapliyal, H., DeMara, R. F., Partin-Vaisband, I., and Katkooori, S., editors, *Proceedings of the Great Lakes Symposium on VLSI 2023, GLSVLSI 2023, Knoxville, TN, USA, June 5-7, 2023*, pages 653–658, <https://doi.org/10.1145/3583781.3590240>. ACM.
- Porto, D., Prado, R., Marques, G., Serrano, A., Mendonça, F., and Canedo, E. (2025). Ethical requirements in the age of artificial intelligence: A systematic literature review. In *Anais do XXI Simpósio Brasileiro de Sistemas de Informação*, pages 663–672, Porto Alegre, RS, Brasil. SBC.
- Smith, J. J., Payne, B. H., Klassen, S., Doyle, D. T., and Fiesler, C. (2023). Incorporating ethics in computing courses: Barriers, support, and perspectives from educators. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. I, SIGCSE 2023*, page 367–373, New York, NY, USA. Association for Computing Machinery.
- Storey, M.-A., Hoda, R., Maciel Paz Milani, A., and Baldassarre, M. T. (2025). Guiding principles for mixed methods research in software engineering. *Empirical Software Engineering*, 30(5):138.
- Tran, M. and Fiesler, C. (2024). "it's not exactly meant to be realistic": Student perspectives on the role of ethics in computing group projects. In *Proceedings of the 2024 ACM Conference on International Computing Education Research-Volume 1*, pages 517–526, <https://doi.org/10.1145/3632620.3671109>.
- Vakkuri, V., Kemell, K.-K., Tolvanen, J., Jantunen, M., Halme, E., and Abrahamsson, P. (2022). How do software companies deal with artificial intelligence ethics? a gap analysis. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering, EASE '22*, page 100–

- 109, <https://doi.org/10.1145/3530019.3530030>. Association for Computing Machinery.
- Whittlestone, J., Nyrup, R., Alexandrova, A., and Cave, S. (2019). The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 195–200, Honolulu HI USA. ACM.
- Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B., and Wesslén, A. (2012). *Experimentation in Software Engineering*. Springer.
- Yang, Y., Zhang, Y., Sun, D., He, W., and Wei, Y. (2025). Navigating the landscape of ai literacy education: insights from a decade of research (2014–2024). *Humanities and Social Sciences Communications*, 12(1):1–12.