# Computing Students' Perceptions of Ethical Principles in AI and Curricular Coverage

Felippe Kipman
University of Brasília (UnB)
Computer Science Department, P.O. Box 4466
Brasília–DF, Brazil
fkipman@hotmail.com

Gilmar dos Santos Marques
University of Brasília (UnB)
Computer Science Department, P.O. Box 4466
Brasília–DF, Brazil
gilmar.marx@gmail.com

Edna Dias Canedo
University of Brasília (UnB)
Computer Science Department, P.O. Box 4466
Brasília–DF, Brazil
ednacanedo@unb.br

## ABSTRACT

**Context**: As artificial intelligence (AI) systems become increasingly integrated into society, ensuring that computing professionals are equipped with ethical competencies is a growing concern. While many international guidelines emphasize ethical principles in AI development, their integration into computing curricula remains inconsistent, especially in the Brazilian context. **Goal**: This study investigates computing students' perceptions of ethical principles in AI, their self-assessed preparedness to apply them, and the presence of ethics-focused disciplines in undergraduate curricula. **Method**: We conducted a survey with 56 undergraduate computing students at the University of Brasília (UnB). The questionnaire included Likert-scale items on students' familiarity with 21 AI ethical principles, their perceived importance, and whether these topics were covered during their studies. Descriptive statistics were used to analyze the results, and we discuss them in light of related literature. **Results**: Students overwhelmingly perceive ethical principles as important and report high agreement with their relevance to professional practice. However, a significant gap exists between perceived importance and curricular exposure: although most students agree on the need to learn ethical principles during their studies, a substantial number report not having had a dedicated course on the topic. Some principles, such as Privacy, Justice, and Transparency, were rated both highly important and well-understood, while others like Prosperity and Solidarity showed lower familiarity. **Conclusions**: The findings highlight the urgent need to integrate dedicated ethics courses into computing curricula. Ethical literacy is important not only for technical competency but also for responsible decision-making in AI development. Future work includes extending the study to other institutions and exploring pedagogical strategies for teaching AI ethics in computing programs.

## KEYWORDS

ethics education, embedded ethics, ethics course, higher education, interdisciplinary teaching and learning

## 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. This explosion of applications has sparked widespread concern about ethical issues—ranging from fairness, transparency, and privacy to accountability, explainability, and sustainability [7, 16, 17]. Despite broad consensus among scholars, practitioners, and international organizations on the relevance of ethical principles in AI, their practical implementation remains challenging [22]. 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 [1].

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.

Although ethics in computing has been debated in academic contexts for over five decades [23], 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 [20].

Recent studies have emphasized the need to strengthen ethics teaching in computer science programs [12, 14, 15], yet research on this topic remains scarce, especially in Global South contexts such as Brazil [10]. Many Brazilian institutions still lack structured approaches to ethics instruction, and where present, such content is often treated as peripheral or purely theoretical [6].

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.

To this end, we conducted a quantitative survey with 56 students enrolled in UnB's computing programs. The questionnaire included Likert-scale items regarding students' agreement with, and relevance attributed to, each principle. The results reveal that while students overwhelmingly recognize the importance of ethical principles in their future professional activities, a considerable number report not having studied these topics in depth during their coursework.

These findings highlight the need for more robust curricular strategies to ensure future professionals are equipped to design AI systems responsibly.

The remainder of this article is organized as follows: Section 2 presents related work on ethics education in computing; Section 3 details the methodology; Section 4 presents the survey results; Section 5 discusses implications for computing curricula and professional practice; Section 6 outlines threats to validity; and Section 7 concludes with directions for future research.

# 2 Background and Related Work

## 2.1 AI Governance

The rapid proliferation of artificial intelligence (AI) technologies has sparked a corresponding surge in academic, regulatory, and societal interest in understanding and managing their ethical implications [7]. 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 [11, 12, 14]. In response, numerous AI ethics guidelines have been published by governments, academic consortia, and standard-setting bodies [7–9, 13, 17, 30]. 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. [22] and Cerqueira et al. [8], 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 [13, 17], further fueling regulatory fragmentation and uncertainty. Although the European Union has taken a pioneering step with its 2023 AI Act[1], most countries remain on the periphery of international governance discussions [7].

In this context, the systematic literature review by Batool et al. [3] 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 [22, 26], 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 [22, 24]. 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 [23], there has been a noticeable surge in publications on the topic in recent years [4], 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.

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 [4]. Effective interdisciplinary approaches typically require time-intensive collaboration among professionals from various fields—an investment that institutions may be reluctant to make [14]. 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 [21]. This can lead students to perceive ethics as separate from the functional requirements of system development [14].

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 [2]. This lack of clarity hampers both the evaluation of student learning and the assessment of pedagogical methods [4]. 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 [4].

Tran and Fiesler [28] 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 et al. [25] 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

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

the curriculum, using real-world case studies and ethical theories such as utilitarianism, Kantianism, and virtue ethics. The study also emphasized that merely referencing codes of ethics like those from the ACM and IEEE is insufficient; instead, 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 with 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. [18], Vakkuri et al. [29], and Cerqueira et al. [8], among others [19, 26]. 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) and perceived relevance ratings (Q7) used to investigate student perceptions.

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 [8, 26]. 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 definitions and limited institutional support [22, 24].

## 2.4 Ethics Education in Computing Curricula

The inclusion of ethics in computing curricula has been advocated for over five decades [23], yet only in recent years has it received sustained attention due to ethical failures in AI deployments [11, 12, 14]. 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 [15].

Despite these advances, challenges persist. Most instructors lack formal training in ethics [4], 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 [14, 21].

Another key issue is the absence of standard learning outcomes and evaluation methods tailored to AI ethics. As highlighted by Aler Tubella et al. [2], 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. [5] 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. [6] 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.

This paper contributes to this debate by presenting a survey conducted at the University of Brasília (UnB), analyzing students' exposure to and perceptions of key ethical principles in AI, and identifying curriculum gaps that need to be addressed to foster more ethically aware computing professionals.

## 3 Study Settings

This study investigates students' perceptions of ethics in computing education and their self-assessment regarding ethics-related learning in undergraduate Computer Science programs at the University of Brasília (UnB). Our aim is to understand how students evaluate the presence, quality, and relevance of ethical aspects in their training.

To guide the investigation, we formulated the following research questions:

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.

**Table 1: Ethical Principles for AI Systems Used in This Study [8, 18, 19, 26, 29]**

| 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. |

To address RQ.1 and RQ.2, we designed and conducted a survey targeting undergraduate students enrolled in computing programs. The goal was to explore their perceptions regarding the relevance of ethical principles in artificial intelligence and their preparedness to engage with such principles in academic and professional contexts. In addition, students were asked to evaluate the importance of learning about ethics in computing education and to indicate whether they had taken any formal course that addressed these topics. The following subsections describe the study's target population, the development and refinement of the survey instrument through a pilot test, the dissemination strategy employed to reach participants, and the procedures used for data collection.

## 3.1 Target Audience

The target audience comprises 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 are directly engaged with computing curricula and represent the future professionals of the field. Their feedback provides valuable insights into how ethics is integrated and perceived in computing education.

## 3.2 Survey Design

The survey instrument was collaboratively designed by the authors and consisted of 9 closed-ended questions, as shown in Table 2, and organized into four thematic sections:

- **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.

## 3.3 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 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.

## 3.4 Survey Dissemination 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 data analysis presented in this study.

# 4 Results

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 3, all respondents (100%) accepted the informed consent. 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 3).

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.

## 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 1 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%)

**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: Profile of survey respondents (Q1–Q3, $n$ = 56)**

| 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 |

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.

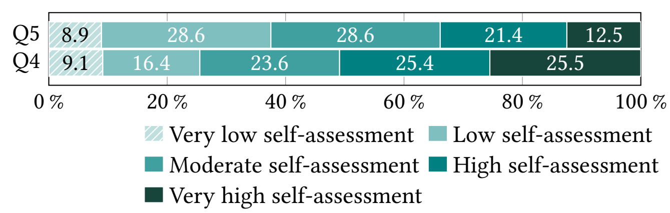


**Figure 1: Distribution of responses to Q4 and Q5 based on self-assessed exposure and preparedness**

Table 4 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 well-understood principles 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 4.

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 4.

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 discussed 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 5 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: Privacy (75.0%), Safety (75.0%), Data Security (75.0%), and Accountability (70.2%). These results suggest that students are particularly concerned with principles that safeguard user rights and prevent harm, reflecting current concerns about data misuse and algorithmic risks.

Principles such as Human Dignity, Justice, and Non-maleficence also received high levels of agreement (above 60% in Strongly agree), reinforcing a focus on fairness and social responsibility. In contrast, principles like Solidarity, Prosperity, and Sustainability showed relatively lower levels of Strongly agree, with 39.3%, 44.6%, and 48.2% respectively (Table 5). While still seen as important (with most responses still falling under Agree), this may indicate a gap in how these principles are contextualized or emphasized in computing curricula.

The proportion of neutral responses was minimal across all principles (mostly under 10%), 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 5 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 2, 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 2, only 14.3% strongly agreed and 26.8% agreed that they had taken a subject covering ethical principles. A considerable portion of 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.

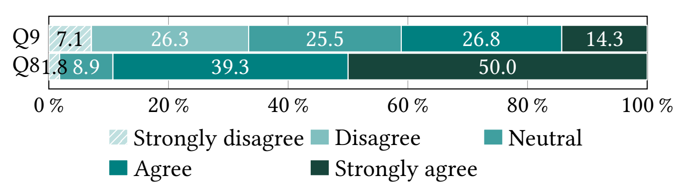


**Figure 2: Distribution of responses to Q8 and Q9 on the perceived importance of learning ethical principles (Q8) and the extent to which such content was addressed in their curriculum (Q9)**

**RQ.2 Summary**:Students demonstrated strong agreement on the relevance of ethical principles in AI systems, particularly those tied to user protection and harm prevention—such as Privacy, Safety, Data Security, and Accountability. Conversely, principles like Solidarity, Prosperity, and Sustainability were perceived as less central, possibly due to their abstract nature or lower curricular emphasis. Most of 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 limits students' ability to critically engage with the ethical dimensions of computing.

**Table 4: 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% |

**Table 5: 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% | 0.0% | 23.2% | 64.3% |
| Privacy | 0.0% | 0.0% | 0.0% | 23.2% | 73.2% |
| Accountability | 0.0% | 3.6% | 0.0% | 25.0% | 64.3% |
| Fairness | 1.8% | 5.4% | 0.0% | 14.3% | 75.0% |
| Autonomy | 1.8% | 3.6% | 0.0% | 32.1% | 44.6% |
| Explainability | 1.8% | 0.0% | 0.0% | 35.7% | 55.4% |
| Justice | 3.6% | 1.8% | 0.0% | 25.0% | 62.5% |
| Non-maleficence | 1.8% | 0.0% | 0.0% | 21.4% | 71.4% |
| Human dignity | 1.8% | 0.0% | 0.0% | 19.6% | 73.2% |
| Beneficence | 7.1% | 7.1% | 0.0% | 37.5% | 35.7% |
| Responsibility | 1.8% | 5.4% | 0.0% | 23.2% | 64.3% |
| Safety | 0.0% | 1.8% | 0.0% | 21.4% | 73.2% |
| Data security | 0.0% | 0.0% | 0.0% | 19.6% | 78.6% |
| Sustainability | 5.4% | 0.0% | 0.0% | 28.6% | 60.7% |
| Freedom | 0.0% | 0.0% | 0.0% | 17.9% | 71.4% |
| Solidarity | 1.8% | 3.6% | 0.0% | 28.6% | 51.8% |
| Prosperity | 7.1% | 1.8% | 0.0% | 21.4% | 48.2% |
| Effectiveness | 3.6% | 1.8% | 0.0% | 35.7% | 53.6% |
| Accuracy | 7.1% | 1.8% | 0.0% | 21.4% | 58.9% |
| Predictability | 3.6% | 1.8% | 0.0% | 30.4% | 48.2% |
| Interpretability | 1.8% | 3.6% | 0.0% | 26.8% | 55.4% |

# 5 Discussion

## 5.1 The Curricular Relevance of Ethical Principles in Computing Education

Our findings reveal a strong consensus among students about the importance of learning ethical principles in computing education: over 85% of respondents agreed or strongly agreed with the relevance of these topics (Q8, see Figure 2). Despite this consensus, only 41.1% confirmed having taken a course that explicitly addressed these principles (Q9), pointing to a significant disconnection between students' expectations and curricular practices. This result reinforces previous findings by Silva et al. [27] and Carvalho et al. [6], who highlighted the fragmented and often superficial treatment of ethics in Brazilian computing programs.

This gap is particularly concerning given the increasing presence of AI and data-driven systems in socially sensitive domains. Ethical principles such as transparency, fairness, privacy, and accountability are frequently referenced in international guidelines [13, 17] and have been recognized as central to responsible AI development [18]. However, the effective teaching of these principles demands more than abstract discussion—it requires pedagogical strategies that integrate ethics into real-world problem-solving contexts, particularly within software engineering, AI, and data science curricula.

Khan et al. [18] found that without practical, contextualized education, ethical principles often remain symbolic rather than actionable. Our study echoes this concern: while students reported relatively high familiarity with many principles (Q6), a subset—such as autonomy, solidarity, and prosperity—remained less understood or internalized (see Table 4). These results suggest that curriculum reform should not only include ethics content but also diversify the treatment of principles beyond the most popular ones (e.g., privacy and fairness), ensuring that less intuitive values are equally addressed.

## 5.2 Impacts on Students' Professional Performance

The findings also suggest a tension between students' perceived ethical maturity and their actual curricular exposure. As shown in responses to Q4 and Q5 (Figure 1), students often rated themselves as highly or moderately prepared to assess the social impact of their work. However, their limited exposure to structured discussions on ethics suggests that such preparedness may be more aspirational than reflective of actual training. This mismatch has been discussed by Morley et al. [22], who noted that many professionals feel responsible for ethical considerations but lack guidance or formal education on how to navigate them.

These self-perceptions, though optimistic, may reflect an overreliance on personal intuition rather than evidence-based ethical literacy. As Tran et al. [28] emphasized, many students struggle to apply ethical concepts when faced with real or simulated design decisions, particularly in the absence of structured reflection, institutional support, or practical tools. Our study corroborates this: although students express confidence, their exposure to ethical reasoning activities appears insufficient, reinforcing the need for more embedded pedagogical approaches.

To address this gap, several authors recommend embedding ethics education throughout the curriculum rather than limiting it to isolated courses [14, 15]. Activities such as ethical design scenarios, reflective essays, and interdisciplinary collaboration have shown promise in enhancing student engagement and ethical reasoning. Our results align with these recommendations and emphasize the need to transform ethical literacy into a practical competency—something that can only be achieved if ethical principles are taught not just as theoretical constructs, but as tools for everyday decision-making.

Ultimately, preparing computing students to design responsible technologies requires a dual investment: in curricular reform and in the training of educators. As highlighted by Brown et al. [4], many computing instructors lack formal background in ethics, and institutions often under-invest in interdisciplinary collaboration. A robust ethics education strategy must therefore include faculty development, updated learning outcomes, and evaluation mechanisms that recognize the complexity of ethical practice in the digital age.

# 6 Threats to Validity

Following the framework proposed by Wohlin et al. [31], we discuss below the main threats to the validity of this study and the strategies adopted to mitigate them.

**Conclusion Validity** concerns the accuracy of statistical inferences. One potential threat in this study lies in the relatively small sample size (n=56), which limits the statistical power for more robust inferential analyses. To mitigate this risk, we focused on descriptive statistics and avoided broad generalizations unsupported by the data. The representativeness of the sample was considered carefully when interpreting results.

**Internal Validity** refers to the degree to which causal conclusions can be drawn. Given the self-reported nature of the survey, there is a risk of social desirability bias, whereby participants may have provided answers they perceived as more socially acceptable—particularly regarding the importance of ethical principles. To address this, participants were assured of the anonymity of their responses, and it was emphasized that there were no right or wrong answers.

**Construct Validity** relates to how accurately the survey measures the concepts it intends to capture. The survey items were based on well-established ethical AI principles from sources such as UNESCO and the European Commission, and reviewed by experts in computing ethics. Nevertheless, some technical concepts (e.g., explainability or non-maleficence) may not have been uniformly understood by all respondents. To mitigate this, brief conceptual explanations accompanied each principle.

**External Validity** refers to the generalizability of findings. A primary threat in this study is the use of a convenience sample, composed primarily of students from a single Brazilian public university. This limits the generalizability of results to all computing programs in Brazil. To mitigate this, the findings were interpreted in light of national and international literature to identify broader trends and patterns.

# 7 Conclusion and Future Work

This study explored computing students' perceptions of ethical principles in artificial intelligence and their presence in undergraduate curricula. The results reveal a strong consensus on the importance of learning about ethical principles during undergraduate education, but a significant portion of respondents reported not having attended any dedicated course or discipline that explicitly addressed these topics. This disconnect between perceived importance and curricular coverage suggests a gap in current computing curricula.

Moreover, the analysis of Q6 and Q7 demonstrated that students feel more familiar with certain principles—such as privacy, fairness, and transparency—than others, such as solidarity or prosperity. Similarly, when evaluating the perceived relevance of the principles, those with closer ties to technical and legal concerns were more frequently rated as highly important, while principles associated with social and collective impacts were less emphasized. These results align with previous findings in the literature, reinforcing the need to integrate ethical education more systematically and comprehensively into computing programs.

Future work involves expanding the study to include students and professionals from multiple institutions and regions, in order to improve generalizability and to capture curricular variations across programs. Additionally, we plan to conduct qualitative interviews to better understand students' experiences and expectations regarding ethics in computing. Finally, the results of this study can inform the design of pedagogical strategies, learning materials, and policies aimed at reinforcing the role of ethical principles in computing education and professional development.

## ACKNOWLEDGMENTS

This study was financed in part by the Project No. 514/2023 – Call Nº 10/2023 – FAPDF Learning Program, a strategic development initiative in the macro areas of Agro Learning, Bio Learning, Gov Learning, and Tech Learning; and Conselho Nacional de Desenvolvimento Científico e Tecnológico CNPq (Grant 300883/2025-0).

## DECLARATION OF GENERATIVE AI IN SCIENTIFIC WRITING

We employed generative AI in the writing process, especially aiming to improve the text quality. We asked for help to improve the text readability and to find suitable synonyms for words based on the context.

## REFERENCES

[1] Ali Al-Kaswan and Maliheh Izadi. 2023. The (ab)use of Open Source Code to Train Large Language Models. doi:10.48550/arXiv.2302.13681 arXiv:2302.13681 version: 2.

[2] Andrea Aler Tubella, Marçal Mora-Cantallops, and Juan Carlos Nieves. 2024. How to teach responsible AI in Higher Education: challenges and opportunities. *Ethics and Information Technology* 26, 1 (March 2024), 3. doi:10.1007/s10676-023-09733-7

[3] Amna Batool, Didar Zowghi, and Muneera Bano. 2025. AI governance: a systematic literature review. *AI Ethics* 5, 3 (2025), 3265–3279. doi:10.1007/S43681-024-00653-W

[4] Noelle Brown, Benjamin Xie, Ella Sarder, Casey Fiesler, and Eliane S. Wiese. 2024. Teaching Ethics in Computing: A Systematic Literature Review of ACM Computer Science Education Publications. *ACM Transactions on Computing Education* 24, 1 (March 2024), 1–36. doi:10.1145/3634685

[5] Noelle Brown, Benjamin Xie, Ella Sarder, Casey Fiesler, and Eliane Stampfer Wiese. 2024. Teaching Ethics in Computing: A Systematic Literature Review of ACM Computer Science Education Publications. *ACM Trans. Comput. Educ.* 24, 1 (2024), 6:1–6:36. doi:10.1145/3634685

[6] Luiz Paulo Carvalho, Jonice Oliveira, and Flávia Maria Santoro. 2021. A presenÇa de conteÚdos sobre Ética computacional na literacia em computaÇÃo institucional brasileira. In *ABCiber Virtual Meeting Proceedings (Online). ABCiber Virtual Meeting*. SBC, Encontro Virtual ABCiber 2021, 1–22.

[7] Nicholas Kluge Corrêa, Camila Galvão, James William Santos, Carolina Del Pino, Edson Pontes Pinto, Camila Barbosa, Diogo Massmann, Rodrigo Mambrini, Luiza Galvão, Edmund Terem, and Nythamar De Oliveira. 2023. Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns* 4, 10 (Oct. 2023), 100857. doi:10.1016/j.patter.2023.100857

[8] José Antonio Siqueira de Cerqueira, Anayran Pinheiro De Azevedo, Heloise Acco Tives Leão, and Edna Dias Canedo. 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*. ScholarSpace, http://hdl.handle.net/10125/80015, 1–10.

[9] José Antonio Siqueira de Cerqueira, Heloise Acco Tives Leão, and Edna Dias Canedo. 2021. Ethical Guidelines and Principles in the Context of Artificial Intelligence. In *SBSI 2021: XVII Brazilian Symposium on Information Systems, Uberlândia, Brazil, June 7 - 10, 2021*, Rafael D. Araújo, Fabiano A. Dorça, Renata Mendes de Araujo, Sean W. M. Siqueira, and Awdren L. Fontão (Eds.). ACM, https://doi.org/10.1145/3466933.3466969, 36:1–36:8. doi:10.1145/3466933.3466969

[10] Iván Manuel De La Vega Hernández, Angel Serrano Urdaneta, and Elias Carayannis. 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 (Feb. 2023), 1699–1729. doi:10.1007/s10462-022-10206-4

[11] Rodrigo Ferreira and Moshe Y. Vardi. 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*. ACM, Virtual Event USA, 1041–1047. doi:10.1145/3408877.3432449

[12] Casey Fiesler, Natalie Garrett, and Nathan Beard. 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*. ACM, Portland OR USA, 289–295. doi:10.1145/3328778.3366825

[13] Luciano Floridi. 2023. *The ethics of artificial intelligence: principles, challenges, and opportunities.* Oxford University Press, Oxford ; New York, NY. OCLC: on1376727516.

[14] Trystan S. Goetze. 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*. ACM, Toronto ON Canada, 645–651. doi:10.1145/3545945.3569792

[15] Barbara J. Grosz, David Gray Grant, Kate Vredenburgh, Jeff Behrends, Lily Hu, Alison Simmons, and Jim Waldo. 2019. Embedded EthiCS: integrating ethics across CS education. *Commun. ACM* 62, 8 (July 2019), 54–61. doi:10.1145/3330794

[16] Marianna Jantunen, Richard Meyes, Veronika Kurchyna, Tobias Meisen, Pekka Abrahamsson, and Rahul Mohanani. 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*. ACM, https://doi.org/10.1145/3643690.3648238, 24–31. doi:10.1145/3643690.3648238

[17] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (Sept. 2019), 389–399. doi:10.1038/s42256-019-0088-2

[18] Arif Ali Khan, Muhammad Azeem Akbar, Mahdi Fahmideh, Peng Liang, Muhammad Waseem, Aakash Ahmad, Mahmood Niazi, and Pekka Abrahamsson. 2023. AI Ethics: An Empirical Study on the Views of Practitioners and Lawmakers. *IEEE Trans. Comput. Soc. Syst.* 10, 6 (2023), 2971–2984. doi:10.1109/TCSS.2023.3251729

[19] Arif Ali Khan, Sher Badshah, Peng Liang, Bilal Khan, Muhammad Waseem, Mahmood Niazi, and Muhammad Azeem Akbar. 2021. Ethics of AI: A Systematic Literature Review of Principles and Challenges. arXiv:2109.07906 [cs.CY]

[20] Amruth N. Kumar, Rajendra K. Raj, Sherif G. Aly, Monica D. Anderson, Brett A. Becker, Richard L. Blumenthal, Eric Eaton, Susan L. Epstein, Michael Goldweber, Pankaj Jalote, Douglas Lea, Michael Oudshoorn, Marcelo Pias, Susan Reiser, Christian Servin, Rahul Simha, Titus Winters, and Qiao Xiang. 2024. *Computer Science Curricula 2023*. ACM, New York, NY, USA. doi:10.1145/3664191

[21] Keith Miller. 1988. Integrating Computer Ethics into the Computer Science Curriculum. *Computer Science Education* 1, 1 (Jan. 1988), 37–52. doi:10.1080/0899340880010104

[22] Jessica Morley, Libby Kinsey, Anat Elhalal, Francesca Garcia, Marta Ziosi, and Luciano Floridi. 2023. Operationalising AI ethics: barriers, enablers and next steps. *AI & SOCIETY* 38, 1 (Feb. 2023), 411–423. doi:10.1007/s00146-021-01308-8

[23] Norman R. Nielsen. 1972. Social responsibility and computer education. *SIGCSE Bull.* 4, 1 (March 1972), 90–96. doi:10.1145/873684.873708

[24] Aastha Pant, Rashina Hoda, Chakkrit Tantithamthavorn, and Burak Turhan. 2024. Ethics in AI through the practitioner's view: a grounded theory literature review. *Empirical Software Engineering* 29, 3 (May 2024), 67. doi:10.1007/s10664-024-10465-5

[25] Sudeep Pasricha. 2023. Ethics in Computing Education: Challenges and Experience with Embedded Ethics. In *Proceedings of the Great Lakes Symposium on VLSI 2023, GLSVLSI 2023, Knoxville, TN, USA, June 5-7, 2023*, Himanshu Thapliyal, Ronald F. DeMara, Inna Partin-Vaisband, and Srinivas Katkoori (Eds.). ACM, https://doi.org/10.1145/3583781.3590240, 653–658. doi:10.1145/3583781.3590240

[26] Daniel Porto, Renata Prado, Gilmar Marques, André Serrano, Fabio Mendonça, and Edna Canedo. 2025. Ethical Requirements in the Age of Artificial Intelligence: A Systematic Literature Review. In *Anais do XXI Simpósio Brasileiro de Sistemas de Informação* (Recife/PE). SBC, Porto Alegre, RS, Brasil, 663–672. doi:10.5753/sbsi.2025.246613

[27] João Silva, Paula Cruz, and Lenise Rodrigues. 2024. Proposta de Atividades para o Ensino de Ética em Computação. In *Anais do IV Simpósio Brasileiro de Educação em Computação* (Evento Online). SBC, Porto Alegre, RS, Brasil, 234–244. doi:10.5753/educomp.2024.237728

[28] Michelle Tran and Casey Fiesler. 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*. ACM, https://doi.org/10.1145/3632620.3671109, 517–526.

[29] Ville Vakkuri, Kai-Kristian Kemell, Joel Tolvanen, Marianna Jantunen, Erika Halme, and Pekka Abrahamsson. 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* (Gothenburg, Sweden) *(EASE '22)*. Association for Computing Machinery, https://doi.org/10.1145/3530019.3530030, 100–109.

[30] Jess Whittlestone, Rune Nyrup, Anna Alexandrova, and Stephen Cave. 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*. ACM, Honolulu HI USA, 195–200. doi:10.1145/3306618.3314289

[31] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in Software Engineering*. Springer, https://doi.org/10.1007/978-3-642-29044-2. doi:10.1007/978-3-642-29044-2