

AI Ethics Education in India: A Syllabus-Level Review of Computing Courses

Anshu M Mittal^[0009–0004–4481–8645],
P D Parthasarathy^[0000–0002–8723–2407], and Swaroop Joshi^[0000–0003–4536–2446]

BITS Pilani, KK Birla Goa Campus, Goa, India
{f20231125, p20210042, swaroopj}@goa.bits-pilani.ac.in

Abstract. The pervasive integration of artificial intelligence (AI) across domains such as healthcare, governance, finance, and education has intensified scrutiny of its ethical implications, including algorithmic bias, privacy risks, accountability, and societal impact. While ethics has received growing attention in computer science (CS) education more broadly, the specific pedagogical treatment of *AI ethics* remains under-examined. This study addresses that gap through a large-scale analysis of 3,395 publicly accessible syllabi from CS and allied areas at leading Indian institutions. Among them, only 75 syllabi (2.21%) included any substantive AI ethics content. Three key findings emerged: (1) AI ethics is typically integrated as a minor module within broader technical courses rather than as a standalone course; (2) ethics coverage is often limited to just one or two instructional sessions; and (3) recurring topics include algorithmic fairness, privacy and data governance, transparency, and societal impact. While these themes reflect growing awareness, current curricular practices reveal limited depth and consistency. This work highlights both the progress and the gaps in preparing future technologists to engage meaningfully with the ethical dimensions of AI, and it offers suggestions to strengthen the integration of AI ethics within computing curricula.

Keywords: AI Ethics · Indian Computing Education · Responsible AI · Ethics in CS, Biases

1 Introduction

Artificial Intelligence (AI) systems now play pivotal roles in sectors ranging from healthcare [1] to finance [2], education [3], employment screening [4], and criminal justice [5]. These technologies, powering tools like voice assistants, recommendation algorithms, autonomous vehicles, and large language models [6], have become integral to modern life. As AI capabilities expand, ethical challenges—including privacy intrusion, algorithmic bias, misinformation risks, and autonomous agency—have drawn sustained scrutiny from academia, policymakers, and the general public [7].

Organizations and governments have responded with frameworks to guide responsible AI development. Prominent corporate initiatives include Google’s 2018

AI principles [8,9], while professional bodies like the ACM introduced codes of ethics [10] to steer practitioner behavior. Meanwhile, the IEEE’s Ethically Aligned Design [11] roadmap and regulatory efforts like the EU AI Act [12], China’s generative AI content restrictions [13], United State’s trustworthy AI mandates [14], and India’s NITI Aayog report on responsible AI¹ highlight global institutional efforts to address ethical governance. Collaborative platforms such as the Beijing AI Principles [15] further advocate cross-border ethical coordination.

Despite these measures, implementation gaps persist [6,10,16]. Research indicates a recurring disconnect between high-level ethical aspirations and practical development practices, exemplified by studies showing limited translation of ethical frameworks into developer workflows [17,10]. The ACM Code of Ethics, for instance, has had minimal documented impact on engineers’ ethical decision-making. Such findings emphasize the urgency of embedding ethical reasoning into CS education [18]. Recent curricular updates reflect this priority; the 2023 ACM/IEEE/AAAI Computer Science Curriculum elevates AI ethics to a core requirement [19], with foundational topics deemed essential for all graduates. This shift acknowledges that AI professionals’ ethical literacy is no longer optional [16] but a professional imperative. As regulatory frameworks increasingly mandate ethical integration into system design [20], universities must equip students with analytical tools and critical frameworks to address these challenges. Educational programs not only mitigate risks through awareness-building but also foster cultural shifts toward accountable innovation [21].

Yet, empirical understanding of AI ethics pedagogy remains limited. Earlier studies addressing tech ethics education identified structural barriers like time constraints and institutional support deficits [22], but these focused broadly on digital ethics rather than AI-specific challenges. It is important to distinguish between the broader domain of *tech ethics* [23,24,25,26] and the more specialized field of *AI ethics*. Tech ethics broadly addresses moral and societal issues related to digital systems, including concerns such as data privacy, cybersecurity, digital inclusion, and environmental impact. In contrast, AI ethics is a more specialized subdomain that centers on challenges unique to artificial intelligence—such as algorithmic bias, opaque decision-making, accountability in autonomous systems, and the large-scale societal consequences of AI deployment. Although both areas emphasize responsible innovation, AI ethics demands a more nuanced engagement with the specific ethical and technical complexities of intelligent systems. This distinction underscores a critical gap in the existing literature.

The Global South—particularly India—holds a significant position in the global software development landscape, ranking third in the number of software engineers worldwide [27]. Indian higher education institutions currently enroll over 2 million undergraduates in CS and related disciplines [28], forming a substantial pipeline for the future technology workforce. Given this scale, it is crucial to examine how AI ethics is being addressed in computing education, both in terms of content coverage and pedagogical approaches. However, there is a no-

¹ <https://tinyurl.com/NitiAyogResAI>

table lack of empirical evidence on how AI ethics is currently taught in this context. To address this gap, we conducted an analysis of 3,395 syllabi from the top 100 engineering institutions listed in India’s National Institutional Ranking Framework (NIRF)². This study provides a comprehensive overview of how AI ethics is incorporated into coursework across a broad spectrum of computing programs. Our research questions are:

- RQ1:** To what extent is AI ethics instruction present in computing programs?
RQ2: Which topics are commonly included under the umbrella of AI ethics?
RQ3: What learning objectives and assessment methods are employed in AI ethics education?

The remainder of this paper is organized as follows: Section 2 reviews the related literature, establishing the foundation and motivation for our study. Section 3 outlines the methodology used to conduct the research. Section 4 presents the findings, followed by a discussion and conclusion of our findings in Section 5 with threats to validity and future work in Section 6.

2 Background and Literature Review

As AI becomes increasingly integrated into daily life, a wide range of ethical concerns has emerged. Researchers have identified serious risks such as the erosion of human agency due to deskilling [29], diminished individual judgment [30], and growing overreliance on automation [29]. These concerns are compounded by the need for human oversight, informed consent, and accountability in AI systems [31]. Public mistrust is further amplified by privacy violations and the exploitation of sensitive personal data [32]. Additionally, numerous studies have shown that algorithmic bias and skewed datasets can lead to systemic inequalities and discriminatory outcomes [33,34]. Beyond technical failures, AI also raises deeper social challenges, such as job displacement [35], emotional detachment in care-related contexts [36], and ethical dilemmas like digital plagiarism [29]. These trends highlight the urgent need to embed AI development within a strong ethical framework grounded in transparency [37], explainability [38], and accountability [39].

Recent real-world developments have only intensified these concerns. In one alarming safety test, an advanced language model simulated manipulative behavior by attempting to blackmail a software engineer to avoid shutdown³. In another case, Meta’s WhatsApp AI assistant mistakenly disclosed a user’s phone number to a stranger and then attempted to deflect the conversation, raising questions about privacy and data governance⁴. Meanwhile, reports show that users are jailbreaking AI chatbots to bypass ethical safeguards and receive assistance for potentially criminal activities⁵. The California AI Policy Report further warns that without robust oversight, AI may facilitate existential threats ranging

² NIRF - <https://www.nirfindia.org/>

³ <https://tinyurl.com/AIShutdown>

⁴ <https://tinyurl.com/NumberLeak>

⁵ <https://tinyurl.com/AIJailBreak>

from misinformation and surveillance to biological and nuclear risks [40]. These incidents reveal not only technical shortcomings but also significant ethical vulnerabilities in contemporary AI systems.

Despite the growing prominence of these issues, a noticeable gap persists in how AI-specific ethics is addressed in CS education. While topics like digital privacy or general tech responsibility are researched [23,24,25,26], there is limited research on the systematic integration of AI ethics into CS curricula. This disconnect between the urgent need for ethical competence and the current state of instruction is the focus of our study.

Preparing students for building ethically grounded AI systems demands more than teaching ethical theory—it requires fostering critical thinking, sociotechnical awareness, and context-sensitive ethical reasoning [41]. At its core, AI ethics aims to ensure that AI systems are designed and deployed in ways that respect human rights and promote social well-being [42]. Educators are thus increasingly expected to prepare students not only as technical experts but also as responsible stewards of emerging technologies [18]. In response, many AI literacy frameworks now explicitly include ethical dimensions. For example, Chiu et al. [43] argue that understanding AI’s societal impact is as essential as mastering technical skills like programming or collaboration. Similarly, Ng et al. [44] emphasize ethical fluency as a foundational component of AI literacy. Although there is wide agreement that ethics should be embedded throughout technical content [45], it is often relegated to the end of the course or included only “if time permits” [22]. This marginalization limits students’ ability to meaningfully engage with the ethical dimensions of AI, leaving them ill-equipped to confront real-world challenges.

This study seeks to address this gap by systematically examining how AI ethics is represented within leading Indian computing programs. By conducting a detailed analysis of program curricula⁶, we evaluate the extent to which these programs equip students to engage with the ethical responsibilities inherent in AI development. To the best of our knowledge, this is the first comprehensive investigation of AI ethics education in computing curricula across India.

3 Methods

This study examines how AI ethics is currently integrated into CS education at leading Indian institutions. The detailed steps of the study are shown in Fig 3.1.

3.1 The Dataset

We initiated our study by compiling a comprehensive list of the top 100 engineering universities, drawing from publicly available rankings provided by the NIRF as detailed in their 2024 Engineering Rankings⁷. We manually gathered the most recent course information available online—including course titles, descriptions, and syllabi—for programs in computer science, artificial intelligence,

⁶ Throughout this paper, we use the words curricula and syllabus interchangeably

⁷ <https://www.nirfindia.org/Rankings/2024/EngineeringRanking.html>

and related disciplines at each institution. We began by checking whether the university offered a dedicated AI engineering program; if available, we collected its curriculum. In the absence of such a program, we turned to computer science and allied disciplines (such as Information Technology, Information Systems, etc.) to identify courses that might incorporate AI ethics content. When syllabi were not publicly available on the university’s official website, we conducted targeted web searches to identify faculty members specializing in AI at the institution and reviewed their personal or departmental webpages for accessible course materials. Courses for which syllabi could not be obtained through either method were excluded from the study. This thorough data collection process resulted in a total of 3,395 unique CS courses across the institutions for which the syllabus was found online.

Note that we do not assert any claims regarding the representativeness of this dataset in relation to the broader landscape of AI education. Our analysis was undertaken with the primary objective of furnishing illustrative examples of AI ethics instruction, while also directing attention toward emerging trends and discernible patterns within the field. That said, it is important to acknowledge that our collected data probably exhibits an over-sampling bias toward instructors who are inclined and at ease with making their course materials available in the public domain. The aggregate measures and insights we present are designed merely to provide a preliminary snapshot of the current state of AI ethics education, rather than to facilitate definitive or conclusive judgments about the entirety of the discipline as a whole.




Fig. 1. Steps used in analysis

3.2 Data Analysis

After collecting the syllabi as they appeared in June 2025, we manually reviewed each document to extract relevant information—specifically, listed topics (from course schedules or reading lists) and any explicitly stated learning outcomes or objectives, when available. As not all syllabi included these elements, a comprehensive qualitative review was carried out by two researchers. They began by organizing and documenting the data in Excel spreadsheets, followed by cross-verification to validate initial findings. A final round of checks was then conducted to ensure the accuracy and reliability of the extracted data.

For courses that did not explicitly reference *AI ethics* in their titles or descriptions, we conducted a close examination of the stated learning outcomes (LO) to identify implicit inclusion of ethical themes. Specifically, we searched for references to topics commonly associated with AI ethics, such as responsible computing practices, algorithmic fairness and bias, privacy concerns in AI systems, safety and security implications of intelligent technologies, human rights considerations in AI deployment, the societal impacts of AI adoption, and the use of ethical frameworks or policy-oriented approaches to guide decision-making. These themes were not only consistently present in the courses we identified as substantively addressing AI ethics, but are also widely recognized as core components of AI ethics education in global contexts, particularly in curricula from institutions in the USA and Europe. Drawing on this international alignment, we used these themes as markers to identify courses that—while not explicitly labeled as AI ethics—nonetheless included meaningful ethical content related to the development and deployment of AI systems.

4 Findings

In describing the overall picture of AI ethics education as presented by this large sample of syllabi, we include descriptions of: (a) who teaches AI ethics, (b) content and duration of instruction, (c) assessments used and (d) learning outcomes in these classes.

4.1 Extent of AI Ethics Coverage in Computing Education

In the analysis of the 3395 computer science-allied courses, only approximately 2.21% (equating to 75 courses) incorporated instruction on AI ethics in varying forms, such as through individual topics, dedicated modules, or even entire courses focused specifically on AI ethics. Although this dataset may not fully capture the complete landscape of AI ethics courses offered across all universities, it nonetheless provides a diverse cross-section that reflects a range of institutions. Overall, the dataset encompasses 22 distinct universities (as several universities featured multiple courses in the listings) where AI ethics syllabi were identified. Furthermore, these courses draw from both private universities (7 in total) and public universities (15 in total), and they are exclusively sourced from engineering colleges.

Regarding the AI ethics coverage, it was primarily through two distinct methods: either as a dedicated standalone course or as an integrated component within larger technical courses. Figure 2 depicts the degree of coverage for AI ethics across the syllabi that were examined in the analysis. As indicated in the figure, the bulk of these courses (approximately 60%) dealt with AI ethics during just a single session, whereas about 13% extended their coverage of the topic across the duration of a week (which is defined in this context as encompassing 3 to 5 sessions). An additional 20% of the courses featured AI ethics in exactly two sessions, and 7% of them were fully devoted to exploring the subject of AI ethics in its entirety.

Examples of dedicated course titles that focus exclusively on the subject encompass offerings such as Intelligent Systems: Design and Ethical Challenges, Artificial Intelligence Ethics, Responsible AI in Societal Deployment, and AI Ethics in Social Impact. By contrast, a substantial number of courses incorporated AI ethics as embedded elements within broader technical subjects, including areas like Fundamentals of AI, Machine Learning, Deep Learning, Natural Language Processing, Reinforcement Learning, AI System Design, Applied or Advanced Machine Learning, Generative AI, Trustworthy Machine Learning, Human-Centered AI, Intelligence for Robotics, Human-Computer Interaction, and Data Science.




Fig. 2. AI Ethics Coverage from Syllabi Analysis (N=75)

Additionally, we investigated the academic levels at which these courses were delivered. Among the courses that were analyzed, 92% were offered at the undergraduate level, 6% at the graduate level, and 2% were designed as cross-listed options available to both undergraduate and graduate students. Within this set of courses, more than 76% required no prerequisites for enrollment, while approximately 24% stipulated the completion of 1 to 3 prerequisite courses (around AI/ML courses) prior to opting for these AI ethics-related offerings.

4.2 Topics part of AI Ethics Education

The qualitative analysis of 75 course syllabi led to a diverse set of themes or topics covered in AI ethics (Table 1)⁸. We explain each theme briefly here:

Privacy, Legal aspect and data governance, covered by 42% of courses, dealt with responsible data use, anonymization, and compliance with privacy frameworks. In addition to general privacy principles, several courses explored specific concerns such as privacy risks in generative AI and large language models. Other topics included algorithmic approaches for protecting personal traits from being reverse-engineered and discussions around the ethical use of clinical data in healthcare machine learning systems. Many courses also addressed broader policy and legal frameworks for regulating AI privacy, Regulatory challenges in AI, International efforts for AI governance, Ethical considerations in

⁸ Note that the sum need not be 100% as few topics are covered in multiple courses

Table 1. Topics covered in AI Ethics (N=75)

Theme/Topic	Coverage
Privacy, Legal Aspect and Data Governance	42%
Accountability and Policies	32%
Ethical Frameworks and Moral Reasoning	24%
Human-AI Interaction and Agency	24%
Algorithmic Bias and Fairness	18%
AI Safety and Robustness	16%
Generative AI specific concerns	11%
Social Impact and Justice	8%

AI policy-making Ethical Decision-Making in AI: Ethical decision-making frameworks, Case studies on ethical dilemmas in AI, Ethical impact assessments Future Directions and Emerging Ethical Challenges: Emerging ethical challenges in AI: Ethical considerations in advanced AI technologies, Ethical leadership in AI development and deployment.

More specialized but still significant themes include **accountability and policy** (32%), which explores questions of responsibility in AI system failures and oversight structures. Courses addressing this theme cover topics such as responsible AI deployment and the integration of ethical principles into policy frameworks, often illustrated through case studies in healthcare and education.

Meanwhile, **ethical frameworks and moral reasoning** are covered in 24% of the courses, introducing students to philosophical approaches and equipping them to reason through complex dilemmas posed by intelligent systems. These frameworks are applied to AI-related moral dilemmas, helping students develop skills in ethical reasoning and decision-making. Topics include ethical considerations in digital technologies and real-world case studies that challenge students to connect theoretical concepts with contemporary AI issues. Few handouts specifically mention ethical impact on society, impact on human psychology, impact on the legal system, impact on the environment and the planet, and impact on trust.

Courses addressing the theme of **human-AI interaction and agency** (24%) focus on how AI systems can be designed to support meaningful, ethical, and effective collaboration between humans and machines. Topics include human-centered design principles, user interface strategies for AI-driven systems, and methods for enabling human-AI teaming and interoperability

The topic **algorithmic bias and fairness** in intelligent systems is covered in 18% of courses. These topics introduce students to methods for detecting bias in data and model outputs, highlighting how social inequalities can be encoded and amplified by machine learning systems. Several courses go further by covering post-processing strategies, often grounded in real-world case studies. The ethical implications of biased systems are also emphasized, particularly in high-stakes applications such as healthcare, medical algorithms, robots, Autonomous Vehicles, Warfare, and weaponization.

The topic of **AI safety and robustness** (16%), which focuses on reliability, adversarial robustness, and secure system behavior in real-world settings. Topics

covered include safe reinforcement learning, as well as emerging risks posed by generative AI and large language models.

Generative AI-specific concerns—including misinformation, value misalignment, and prompt injection—appear in 10.67% of courses. Courses also engage with the societal implications of generative AI, including misinformation, content moderation, privacy concerns and regulatory challenges. Finally. Some courses discuss sustainable system design and advocate for ethical awareness of the ecological and human toll of AI development.

Courses also address topics around **social impact and justice** (8%), focusing on how AI technologies influence domains such as healthcare, education, social justice, law enforcement, and employment. These courses encourage students to consider the equity implications of AI deployment, especially on marginalized communities. Specific topics include the role of AI in predictive policing and public safety, as well as understanding systemic inequities embedded in data-driven decision systems. Some courses integrate frameworks of justice, rights, and welfare into machine learning curricula and emphasize participatory or interdisciplinary design involving experts from public health.

Resources used to teach AI Ethics We examined the syllabi for references to AI ethics-specific resources or textbooks. The majority of the courses (over 90%) did not include any dedicated materials for teaching AI ethics in the syllabus. Instead, they primarily included standard AI and machine learning textbooks—such as Artificial Intelligence: A Modern Approach [46], Pattern Recognition and Machine Learning [47], and other widely used texts aligned with the technical focus of the core course in which AI ethics content was embedded. Only a small number of syllabi explicitly mentioned ethics-specific resources, such as AI for Social Impact⁹, Towards a Code of Ethics for Artificial Intelligence by Paula [48], AI Ethics by Mark [49], and Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way by Dignum [50] and The Oxford Handbook of Ethics of AI [51].

4.3 Learning Objectives (LO) in AI Ethics Education

This section presents the learning objectives (LO), assessment methods found from our analysis.

Learning Objectives The terms “learning objectives” and “learning outcomes” have subtle differences in meaning but are often used interchangeably [52]; we saw both in our data. Similar to a 2020 syllabus analysis published in SIGCSE on tech ethics [23], we use the term *learning outcome* to mean an explicit statement of what students should know or be able to do by the end of the course related to AI ethics, regardless of how such a statement might be labeled in any given syllabus. Our qualitative findings revealed eight commonly stressed upon types of learning outcomes in terms of what *students should be able to* do by the end of the course. The LO found are:

⁹ <https://ai4sibook.org/>

1. Identify and evaluate privacy risks in AI systems, and propose strategies to mitigate them within ethical and legal frameworks
2. ethical considerations: fairness in AI (bias, transparency, implications of AI decisions on society).
3. Interpret and communicate the decision-making processes of AI systems using explainability techniques
4. Discuss or Critique key ethical dilemmas encountered during the design, training, and deployment choices of AI systems
5. Investigate and compare ethical frameworks, regulatory strategies, and policy approaches relevant to AI governance and AI healthcare
6. Apply ethical reasoning models to assess and justify decisions made in the development and deployment of AI systems
7. Policy for AI ethics
8. Global governance on AI

Assessment Methods None of the courses included documented assessments explicitly dedicated to AI ethics topics. In cases where assessments were mentioned, they were limited to broad categories—such as quizzes, assignments, projects, mid-semester, or end-semester exams—without specifying the content or indicating whether any of these components addressed AI ethics. While it is possible that ethics-related questions may be embedded within broader evaluations (say in written exams), such details were generally absent from the syllabi reviewed. These findings suggest that the assessment of AI ethics remains limited, lacking both clarity and consistent integration within computing curricula.

4.4 Miscellaneous Findings

Notably, out of the top 100 engineering institutions identified through the National Institutional Ranking Framework (NIRF), fewer than half had publicly accessible academic curricula, course handouts, or detailed syllabus information available on their official university websites. Interestingly, one private institution offering a dedicated AI/ML engineering degree did not include any course that addressed AI ethics in any form. While this observation is not necessarily indicative of the absence of such content in practice, it highlights the limited visibility of ethics-related coursework in program documentation. This lack of transparency raises important questions about the extent to which ethical considerations are being systematically integrated into the design of emerging AI curricula.

5 Discussion and Conclusion

This study presents the first large-scale analysis of how AI ethics is represented in computing education across top Indian engineering institutions. Despite the increasing global emphasis on ethical AI development, our findings reveal that AI ethics remains minimally integrated and inconsistently represented in Indian computer science and allied curricula.

Only 2.21% of the 3,395 courses analyzed included any substantive AI ethics content, and the majority of those addressed it superficially—often through a single lecture session. Courses offering in-depth or dedicated instruction were rare, and even fewer articulated clear learning outcomes or assessment methods tied to ethical topics.

Our findings align closely with the ACM/IEEE/AAAI Computer Science Curricula 2023 (CS2023) [19], which identifies AI, Society, and Professionalism as a key knowledge area containing several CS Core competencies. Many of the learning outcomes (LOs) identified in our analysis—such as privacy, accountability, and fairness, policies, ethical frameworks, and moral reasoning—correspond well with the ethical and societal objectives outlined in CS2023. However, despite this encouraging overlap, important gaps remain. Notably, there is limited curricular attention to the ethical challenges posed by deep generative models, such as large language models and deepfakes, which are explicitly highlighted in CS2023 as essential topics. Only 11% of courses (N=75) addressed these issues. Similarly, sustainability, justice, and labor ethics—including the environmental costs of model training were covered in just 8% of courses. These omissions point to a disconnect between current teaching practices and the broader vision of socially responsible computing promoted in CS2023, underscoring the need for more comprehensive and future-facing AI ethics education.

A significant gap was identified in the availability and use of instructional resources: most syllabi relied on general AI or machine learning textbooks, with minimal or no reference to ethics-specific materials. Additionally, assessment practices related to AI ethics were either vaguely described or entirely missing, raising concerns about how students’ understanding of ethical issues is being evaluated. These trends point to a broader lack of structured pedagogical design for AI ethics education and reflect inconsistencies in how instructors conceptualize and incorporate ethical content into their courses. The absence of widely adopted teaching materials likely stems from uncertainty about what constitutes core ethical knowledge in AI, along with a shortage of accessible, well-designed resources. There is, therefore, a pressing need to develop comprehensive textbooks, teaching guides, and assessment frameworks that can support educators in meaningfully integrating AI ethics into computing curricula. This also presents an opportunity to contextualize AI ethics education within issues of particular relevance to India, such as algorithmic bias linked to caste and religion, the ethics of large-scale biometric systems like Aadhaar, and the societal implications of AI in agriculture and public distribution systems. A useful comparison can be drawn from accessibility education in computer science, which, until recently, was similarly underrepresented. The development of dedicated resources and textbooks¹⁰—complete with practical examples and assessment tools—has significantly advanced its curricular adoption. A parallel effort is now essential in AI ethics to enable consistent, rigorous, and scalable integration across academic institutions.

¹⁰ <https://bookish.press/tac> , <https://accessibilityeducation.github.io/>

The limited availability of publicly accessible curricula posed a significant barrier to conducting a truly comprehensive review across all top institutions. This lack of transparency in curriculum sharing reflects broader structural challenges in accessing organized educational data within the Indian higher education system—particularly when assessing the integration of emerging topics like AI ethics. In contrast, universities in the Global North, especially in countries like the U.S. and across Europe, commonly publish detailed course syllabi and learning outcomes online as part of a broader commitment to openness, accountability, and academic exchange. The relative absence of such practices in India highlights a critical need for greater institutional support for open curricular access. Enhancing transparency in curriculum dissemination would not only facilitate educational research and policy evaluation but also foster cross-institutional learning and capacity-building for AI ethics instruction.

Overall, our findings reveal a fundamental concern: AI ethics remains a marginal and unevenly addressed component of CS education. Despite increasing discourse on the ethical responsibilities of AI professionals, this awareness has yet to be reflected in the structure, content, and pedagogy of most CS programs. The absence of shared curricular frameworks, reliable instructional resources, and clearly defined assessment strategies contributes to a fragmented landscape in which ethics is often sidelined or superficially integrated. To move beyond this piecemeal approach, there is an urgent need for coordinated efforts that embed AI ethics as a core element of computing education. This entails establishing consensus-driven learning goals, developing high-quality open educational materials—including modular textbooks, practical case studies grounded in the Indian context, and scaffolded activities—and designing assessments that go beyond rote evaluation to cultivate ethical judgment and reflective thinking. Additionally, empowering faculty through professional development opportunities and cross-disciplinary collaborations will be essential, especially for instructors new to ethics instruction. Without such systemic efforts, ethical understanding will remain an afterthought, rather than a foundational skill for the next generation of AI practitioners.

6 Threats to Validity

While this study provides a systematic analysis of AI ethics coverage in computing curricula, several limitations must be acknowledged. First, although efforts were made to manually collect the most recent and publicly available syllabi from top-ranked institutions, the process was subject to human error, especially when navigating decentralized or outdated university websites. Some relevant syllabi may have been missed due to inconsistencies in how courses are titled, categorized, or published online. Second, our sample focused exclusively on the top 100 engineering institutions as per the National Institutional Ranking Framework (NIRF). While this offers a robust view of leading programs, it does not capture the full spectrum of institutions—particularly smaller or unranked colleges—that may also be offering meaningful instruction in AI ethics. Future work could expand the sample and use direct faculty surveys or institutional outreach to validate and enrich the findings.

References

1. R. C. Zhao, “AI in Healthcare for Resource Limited Settings: An Exploration and Ethical Evaluation,” 2025.
2. Y. Li, S. Wang, H. Ding, and H. Chen, “Large Language Models in Finance: A Survey,” in *4th ACM International Conference on AI in Finance*. Brooklyn NY USA: ACM, Nov. 2023, pp. 374–382. [Online]. Available: <https://dl.acm.org/doi/10.1145/3604237.3626869>
3. J. Shrugare, “AI in Education,” *XRDS*, vol. 29, no. 3, pp. 63–65, Apr. 2023. [Online]. Available: <https://doi.org/10.1145/3589657>
4. L. Li, T. Lassiter, J. Oh, and M. K. Lee, “Algorithmic Hiring in Practice: Recruiter and HR Professional’s Perspectives on AI Use in Hiring,” in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Virtual Event USA: ACM, Jul. 2021, pp. 166–176. [Online]. Available: <https://dl.acm.org/doi/10.1145/3461702.3462531>
5. S. Hepenstal, L. Zhang, N. Kodagoda, and B. l. w. Wong, “Developing Conversational Agents for Use in Criminal Investigations,” *ACM Trans. Interact. Intell. Syst.*, vol. 11, no. 3-4, pp. 25:1–25:35, Sep. 2021. [Online]. Available: <https://doi.org/10.1145/3444369>
6. A. A. Khan, S. Badshah, P. Liang, M. Waseem, B. Khan, A. Ahmad, M. Fahmideh, M. Niazi, and M. A. Akbar, “Ethics of ai: A systematic literature review of principles and challenges,” in *Proceedings of the 26th ICEASE*, 2022, pp. 383–392.
7. T. Hagendorff, “The ethics of ai ethics: An evaluation of guidelines,” *Minds and machines*, vol. 30, no. 1, pp. 99–120, 2020.
8. A. Jobin, M. Ienca, and E. Vayena, “The global landscape of ai ethics guidelines,” *Nature machine intelligence*, vol. 1, no. 9, pp. 389–399, 2019.
9. S. Maouche, “Google ai: Opportunities, risks, and ethical challenges,” *Contemporary French and Francophone Studies*, vol. 23, no. 4, pp. 447–455, 2019.
10. A. McNamara, J. Smith, and E. Murphy-Hill, “Does acm’s code of ethics change ethical decision making in software development?” in *Proceedings of the 2018 26th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*, 2018, pp. 729–733.
11. K. Shahriari and M. Shahriari, “Ieee standard review—ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems,” in *IHTC*. IEEE, 2017, pp. 197–201.
12. H. Roberts, J. Cows, E. Hine, J. Morley, V. Wang, M. Taddeo, and L. Floridi, “Governing artificial intelligence in china and the european union: Comparing aims and promoting ethical outcomes,” *The Information Society*, vol. 39, no. 2, pp. 79–97, 2023.
13. E. Hine and L. Floridi, “New deepfake regulations in china are a tool for social stability, but at what cost?” *Nature Machine Intelligence*, vol. 4, no. 7, 2022.
14. S. Parinandi, J. Crosson, K. Peterson, and S. Nadarevic, “Investigating the politics and content of us state artificial intelligence legislation,” *Business and Politics*, vol. 26, no. 2, pp. 240–262, 2024.
15. C. Burle and D. Cortiz, “Mapping principles of artificial intelligence,” *São Paulo: Núcleo de Informação e Coordenação do Ponto BR.[e-book]* Disponível sur: <https://ceweb.br/publicacoes/mapping-principles-of-artificial-intelligence/>. [Consulté le 12 décembre 2022], 2020.
16. I. D. Raji, M. K. Scheuerman, and R. Amironesei, “You can’t sit with us: Exclusionary pedagogy in ai ethics education,” in *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 2021, pp. 515–525.

17. V. Vakkuri, K.-K. Kemell, M. Jantunen, and P. Abrahamsson, ““this is just a prototype”: How ethics are ignored in software startup-like environments,” in *International Conference on Agile Software Development*. Springer International Publishing Cham, 2020, pp. 195–210.
18. L. J. Wiese, I. Patil, D. S. Schiff, and A. J. Magana, “Ai ethics education: A systematic literature review,” *Computers and Education: Artificial Intelligence*, p. 100405, 2025.
19. A. N. Kumar, R. K. Raj, S. G. Aly, M. D. Anderson, B. A. Becker, R. L. Blumenthal, E. Eaton, S. L. Epstein, M. Goldweber, P. Jalote, D. Lea, M. Oudshoorn, M. Pias, S. Reiser, C. Servin, R. Simha, T. Winters, and Q. Xiang, *Computer Science Curricula 2023*. New York, NY, USA: Association for Computing Machinery, 2024.
20. J. Weichert, D. Kim, Q. Zhu, J. Kim, and H. Eldardiry, “Assessing computer science student attitudes towards ai ethics and policy,” *arXiv preprint arXiv:2504.06296*, 2025.
21. M. Brundage, S. Avin, J. Wang, H. Belfield, G. Krueger, G. Hadfield, H. Khlaaf, J. Yang, H. Toner, R. Fong *et al.*, “Toward trustworthy ai development: mechanisms for supporting verifiable claims,” *arXiv preprint arXiv:2004.07213*, 2020.
22. N. Garrett, N. Beard, and C. Fiesler, “More than “if time allows” the role of ethics in ai education,” in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020, pp. 272–278.
23. C. Fiesler, N. Garrett, and N. Beard, “What Do We Teach When We Teach Tech Ethics?: A Syllabi Analysis,” in *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. Portland OR USA: ACM, Feb. 2020, pp. 289–295. [Online]. Available: <https://dl.acm.org/doi/10.1145/3328778.3366825>
24. D. Horton, S. A. McIlraith, N. Wang, M. Majedi, E. McClure, and B. Wald, “Embedding Ethics in Computer Science Courses: Does it Work?” in *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education*. Providence RI USA: ACM, Feb. 2022, pp. 481–487. [Online]. Available: <https://dl.acm.org/doi/10.1145/3478431.3499407>
25. J. J. Smith, B. H. Payne, S. Klassen, D. T. Doyle, and C. Fiesler, “Incorporating Ethics in Computing Courses: Barriers, Support, and Perspectives from Educators,” in *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. Toronto ON Canada: ACM, Mar. 2023, pp. 367–373. [Online]. Available: <https://dl.acm.org/doi/10.1145/3545945.3569855>
26. P. D. Parthasarathy, T. G. Lakshmi, R. Indra, S. Spruha, and S. Joshi, “Digital Conscience: Investigating the State of Ethics in CS Curricula in India,” in *Proceedings of the 2024 ACM Conference on International Computing Education Research - Volume 2*, ser. ICER ’24, vol. 2. New York, NY, USA: Association for Computing Machinery, Aug. 2024, pp. 549–550.
27. JetBrains, “How Many Developers Are in the World? | Data Playground,” 2024. [Online]. Available: <https://www.jetbrains.com/lp/devecosystem-data-playground>
28. Department of Higher Education, Government of India, “All India survey on Higher Education 2021-22,” 2022.
29. A. Tlili, B. Shehata, M. A. Adarkwah, A. Bozkurt, D. T. Hickey, R. Huang, and B. Agyemang, “What if the devil is my guardian angel: Chatgpt as a case study of using chatbots in education,” *Smart learning environments*, vol. 10, no. 1, p. 15, 2023.
30. D. Vrščaj, S. Nyholm, and G. P. Verbong, “Is tomorrow’s car appealing today? ethical issues and user attitudes beyond automation,” *AI & society*, vol. 35, pp. 1033–1046, 2020.

31. M. Aitken, M. Ng, D. Horsfall, K. P. Coopamootoo, A. van Moorsel, and K. Elliott, "In pursuit of socially-minded data-intensive innovation in banking: A focus group study of public expectations of digital innovation in banking," *Technology in Society*, vol. 66, p. 101666, 2021.
32. J. Street, H. Barrie, J. Elliott, L. Carolan, F. McCorry, A. Cebulla, L. Phillipson, K. Prokopovich, S. Hanson-Easey, T. Burgess *et al.*, "Older adults' perspectives of smart technologies to support aging at home: Insights from five world café forums," *International journal of environmental research and public health*, vol. 19, no. 13, p. 7817, 2022.
33. V. Couture, M.-C. Roy, E. Dez, S. Laperle, and J.-C. Bélisle-Pipon, "Ethical implications of artificial intelligence in population health and the public's role in its governance: perspectives from a citizen and expert panel," *Journal of medical Internet research*, vol. 25, p. e44357, 2023.
34. M. G. Hanna, L. Pantanowitz, B. Jackson, O. Palmer, S. Visweswaran, J. Pantanowitz, M. Deebajah, and H. H. Rashidi, "Ethical and bias considerations in artificial intelligence/machine learning," *Modern Pathology*, vol. 38, no. 3, p. 100686, 2025.
35. C. Blease, C. Locher, M. Leon-Carlyle, and M. Doraiswamy, "Artificial intelligence and the future of psychiatry: qualitative findings from a global physician survey," *Digital health*, vol. 6, p. 2055207620968355, 2020.
36. K. Cresswell, S. Cunningham-Burley, and A. Sheikh, "Health care robotics: qualitative exploration of key challenges and future directions," *Journal of medical Internet research*, vol. 20, no. 7, p. e10410, 2018.
37. K. Lam, M. D. Abràmoff, J. M. Balibrea, S. M. Bishop, R. R. Brady, R. A. Calcut, M. Chand, J. W. Collins, M. K. Diener, M. Eisenmann *et al.*, "A delphi consensus statement for digital surgery," *NPJ digital medicine*, vol. 5, no. 1, p. 100, 2022.
38. M. Ananny and K. Crawford, "Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability," *new media & society*, vol. 20, no. 3, pp. 973–989, 2018.
39. C. Aradau and T. Blanke, *Algorithmic reason: The new government of self and other*. Oxford University Press, 2022.
40. H. Booth, "California AI Policy Report Warns of 'Irreversible Harms'," Jun. 2025. [Online]. Available: <https://time.com/7295021/california-ai-policy-report-newsom/>
41. O. Gambelin, "Brave: what it means to be an ai ethicist," *AI and Ethics*, vol. 1, no. 1, pp. 87–91, 2021.
42. T. Hagendorff, "Blind spots in ai ethics," *AI and Ethics*, vol. 2, no. 4, 2022.
43. T. K. Chiu, H. Meng, C.-S. Chai, I. King, S. Wong, and Y. Yam, "Creation and evaluation of a pretertiary artificial intelligence (ai) curriculum," *IEEE Transactions on Education*, vol. 65, no. 1, pp. 30–39, 2021.
44. D. T. K. Ng, J. K. L. Leung, S. K. W. Chu, and M. S. Qiao, "Conceptualizing ai literacy: An exploratory review," *Computers and Education: Artificial Intelligence*, vol. 2, p. 100041, 2021.
45. B. J. Grosz, D. G. Grant, K. Vredenburg, J. Behrends, L. Hu, A. Simmons, and J. Waldo, "Embedded ethics: integrating ethics across cs education," *Communications of the ACM*, vol. 62, no. 8, pp. 54–61, 2019.
46. S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. USA: Prentice Hall Press, 2009.
47. C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.

48. Paula Boddington, "Towards a Code of Ethics for Artificial Intelligence | SpringerLink," 2017. [Online]. Available: <https://link.springer.com/book/10.1007/978-3-319-60648-4>
49. Mark Coelcerkberg, "AI Ethics," 2020. [Online]. Available: <https://mitpress.mit.edu/9780262538190/ai-ethics/>
50. V. Dignum, *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*, ser. Artificial Intelligence: Foundations, Theory, and Algorithms. Cham: Springer International Publishing, 2019. [Online]. Available: <http://link.springer.com/10.1007/978-3-030-30371-6>
51. M. D. Dubber, F. Pasquale, S. Das, M. D. Dubber, F. Pasquale, and S. Das, Eds., *The Oxford Handbook of Ethics of AI*, ser. Oxford Handbooks. Oxford, New York: Oxford University Press, Sep. 2020.
52. R. H. and, "Learning outcomes and instructional objectives: is there a difference?" *Medical Teacher*, vol. 24, no. 2, pp. 151–155, 2002. [Online]. Available: <https://doi.org/10.1080/0142159022020687>