2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #52528

Mapping Social and Environmental Dimensions in AI Ethics Education: A Syllabi Analysis

Kerrie Danielle Hooper, Florida International University

Kerrie Hooper is an Engineering and Computing Education Ph.D. candidate at Florida International University. She holds a Master's degree in Data Science and Artificial Intelligence from Florida International University and a Bachelor of Science in Computer Science from the University of Guyana. Prior to graduate studies, she worked in industry as a Data Analyst and Systems Administrator and has taught courses in programming, databases, and web development. Her research spans AI ethics education, responsible AI, sustainability, and technology policy, particularly how educational discourse shapes ethical perspectives and identities in emerging technology fields. She is broadly interested in interdisciplinary engineering education and advancing socially and ecologically responsible approaches to AI.

Mapping Social and Environmental Dimensions in AI Ethics Education

Introduction

Artificial intelligence (AI) is no longer a new concept in this modern era. AI is being integrated into many software applications, business models, government sectors, and even academia, making it a critical field in terms of its development, application, and broader implications [1]. As a result of the widespread adoption of AI applications, there has been an increase in the need and urgency for developing AI education [2]. Given the critical nature of the AI field, the ethics of AI have been developing similarly, particularly with the rapid adoption of generative AI (GenAI) and agentic AI, which has led to widening concerns about both social and environmental implications. These implications of AI span a wide range of areas, including human health, job security, intellectual property rights, bias and fairness, transparency, accountability, explainability, trustworthiness, environmental sustainability, energy consumption, and even climate change. However, despite this wide range of concerns, there is no clear understanding of how these topics/concerns are presented and structured within formal educational settings [3].

This study focuses on AI ethics curricula and aims to examine what is taught at the university level and how content is structured across AI ethics courses. Understanding how AI ethics is framed in higher education is meaningful because the students of today are tomorrow's AI practitioners, policymakers, and decision-makers [4]. The ways in which social and environmental ethics priorities are discussed or overlooked in educational discourse can shape how AI systems are ultimately designed, deployed, and governed. By analyzing AI ethics syllabi, this study seeks to shed light on which impacts are emphasized and where important ethical dimensions may remain underrepresented.

Literature Review

Social and Environmental Impacts as Key Categories

In literature, the implications of AI are discussed in terms of its social and environmental implications. Ojong [5], discussed a tri-level impact, which includes social, economic, and environmental. However, in this study, the economic impact is considered a social impact since the economic impact affects social groups in different ways. For instance, Gwagwa et al. [6], noted that economic exclusion caused by AI can reinforce forms of neocolonialism in various countries that were historically colonized. In cases where the focus is specifically on economic impact, it may be beneficial to address the concepts separately outside of discussions of who it affects or who bears the cost. Other studies, like [7] - [8], discussed the dual-edged nature of AI, which refers to leveraging AI for social and environmental sustainability while simultaneously noting the associated challenges or risks that come with the over-reliance on AI to solve sustainability challenges. Nosirov [9] examined those concepts from an environmental and social governance (ESG) standpoint and proposed metrics to quantify environmental and social effects, while positing that the path to sustainability is complex. Therefore, social and environmental implications of AI are known terms used in literature to generally categorize the implications of AI.

AI Ethics Course Contents

There are several ways in which AI ethics course contents were examined and dissected in the literature. Garret et al. [10] examined AI ethics course contents and proposed that ethics can be integrated into a single AI ethics course or a technical course. Garret et al. [10] also focused only on the social implications of AI, which can be attributed to the fact that the study was conducted prior to the GenAI boom, which began in 2022. Concerns for the environmental implications of AI started to become more apparent as the demand for power needed for data centers increased from 2,688 MW at the end of 2022 to 5,341 MW at the end of 2023 [11] – [12]. Overall, Garret et al. [10] were able to differentiate how AI ethics topics differed across technical AI courses and non-technical ones. Jared et al. [13] examined topics and placed them into five categories while cross-comparing them across syllabi from different continents, disciplines, and how they fit within Bloom's taxonomy. Other studies focused on piloting a new AI ethics course, where they laid out various topics and approaches to teaching the course [14] – [15]. One of the studies that stood out was [15], where the authors categorized AI ethics syllabi topics into four categories with perpendicular axes, see Figure 1. The four categories can potentially serve as a general way to classify syllabi regardless of the discipline. Additionally, this model positions AI ethics curriculum across two intersecting outer axes: Theory vs Practice and Philosophy vs Technology, which capture variation in the emphasis of course topics.

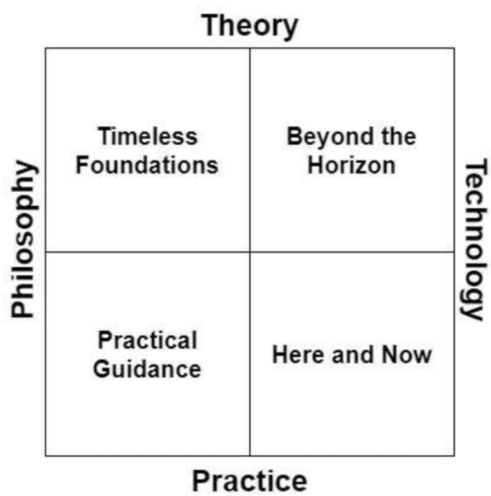

Figure 1: Categorization of study topics in AI ethics based on two perpendicular axes (Tuovinen & Rohunen, [15])

Social and Environmental Sustainability in AI Ethics Course Contents

Studies have examined a wide range of topics in engineering and computing ethics and technical courses. Fiesler et al. [16] found that some computing ethics courses focused specifically on the large-scale social or environmental impact, which included environmental concepts like green computing and geoengineering. Diaz et al. [17], through active learning techniques, embedded ethics in computer science courses, comprehensively noting computer scientists' responsibility towards society and the environment. Hooper & Lunn [18] examined AI ethics, investigated AI ethics syllabi pre-and-post ChatGPT era, noting changes in ethics topics, and called for syllabi to include global impact across society with a consideration for sustainability. Therefore, recent research suggests that efforts are being made to integrate social and environmental responsibility in engineering and computing education. However, there are not many studies that specifically discuss the social and environmental implications of AI ethics course contents. This study seeks to examine this particular niche by applying concepts from literature to explore recent AI ethics syllabi to map out the representation of social and environmental concepts in syllabi, with the research question of: How are social and environmental implications represented in recent AI ethics course syllabi?

Methods

Research Design

This study utilizes a conceptual content analysis method to examine how social and environmental concepts are presented within university-level AI ethics syllabi. Conceptual content analysis is one of the two general types of content analysis methods (conceptual analysis & relational analysis), which examines the occurrence of selected terms in qualitative data and is typically what people refer to as content analysis [19]. Syllabi were selected as the primary unit of analysis because they serve as formal instructional artifacts that communicate course priorities and ethical framing for students. This method allows for a systematic examination of both the presence and structuring of ethical concepts in AI ethics education. To guide the analysis, this study draws on the framework proposed by [15]. Their framework provides an established lens for examining how ethical considerations are framed and emphasized specifically in the AI ethics context. Using this framework, the study specifically investigates how social and environmental concepts are represented within AI ethics syllabi.

Data Collection

AI ethics syllabi were collected through Google searches of the first ten pages of results of each of the keywords “AI ethics course,” “AI ethics syllabus,” and “AI ethics curriculum.” The search and data collection process was conducted during the summer of 2025. The search focused on publicly available syllabi from 2023 to 2025, intentionally selected to reflect the period following the rapid expansion of GenAI technologies in late 2022. This timeframe was chosen to remain attentive to when environmental considerations related to AI may have begun to more visibly emerge in educational discourse. A total of 24 AI ethics syllabi were identified and included in the dataset. Inclusion criteria required that syllabi be obtained from accredited universities and contain sufficient information, like the course year, department, and institution, to enable analysis. Syllabi were selected within the defined study timeframe and across both STEM and non-STEM disciplines. Syllabi from massive open online courses (MOOCs), professional training programs, or courses outside the specified temporal frame were excluded. Additionally, syllabi lacking essential course details or incomplete documentation were not included. This approach ensured consistency, comparability, and relevance of the dataset for examining AI ethics syllabi discourse.

For each syllabus, various elements were extracted, such as course year, department, course topics, learning objectives, and course description. These data were first organized in an Excel spreadsheet to support systematic comparison and record-keeping. Each syllabus was then converted into an individual PDF file and imported into ATLAS.ti to facilitate qualitative coding and analysis.

Meta-Data Description

ID	Year	Department Type	Level	Institution Status	Public Institution?
S3	2023	Non-STEM	Undergrad	R1	Yes
S8	2023	STEM	Grad	R1	Yes
S9	2023	Non-STEM	Undergrad	R1	No
S40	2023	Non-STEM	Undergrad	R1	Yes
S41	2023	Non-STEM	Undergrad	R1	No
S42	2023	Non-STEM	Undergrad	R1	Yes
S45	2023	STEM	Undergrad	R1	No
S55	2023	STEM	Grad	R1	Yes
S57	2023	Non-STEM	Undergrad	R1	Yes
S31	2024	Non-STEM	Undergrad	R1	No
S32	2024	Non-STEM	Undergrad	R1	Yes
S33	2024	STEM	Grad	R1	Yes
S36	2024	STEM	Grad	R1	Yes
S37	2024	STEM	Undergrad	R1	No
S39	2024	STEM	Grad	R1	Yes
S44	2024	Non-STEM	Undergrad	R1	Yes
S58	2024	STEM	Undergrad	R1	Yes
S60	2024	Non-STEM	Undergrad	R2	No
S38	2025	STEM	Undergrad	R1	No
S43	2025	STEM	Grad	R1	Yes
S48	2025	Non-STEM	Undergrad	R1	No
S49	2025	Non-STEM	Undergrad	R1	No
S50	2025	Non-STEM	Undergrad	R2	Yes
S59	2025	Non-STEM	Undergrad	R1	Yes

Table 1: Metadata description of AI ethics syllabi.

The primary data for analysis included **course topics, course descriptions, and learning objectives**, which served as the main sources for identifying and classifying AI ethics syllabi topics. These components formed the basis for all coding procedures.

Table 1 demonstrates that the sample includes both STEM and non-STEM courses, with a greater representation from non-STEM disciplines (58.3%, n = 14) compared to STEM fields (41.7%, n = 10). Most courses were offered at the undergraduate level (66.7%, n = 16), while graduate-level courses accounted for 33.3% (n = 8). Institutional context was primarily research-intensive, with the majority of courses drawn from R1 institutions (91.7%, n = 22) and a smaller proportion from R2 institutions (8.3%, n = 2). The public university syllabi represented 58.3% (n = 14) of the dataset, while the private university syllabi were 41.7% (n = 10). This dataset enables systematic examination of disciplinary and institutional differences in how AI ethics is framed across higher education.

Data Analysis

Next, the data analysis proceeded in two phases. In the first phase, syllabi were coded deductively using the four categories outlined by [15]. This phase focused on identifying how ethical concepts were framed within existing syllabi. In the second phase, coding focused specifically on the presence of social and environmental concepts, highlighting how these dimensions of AI ethics were emphasized across syllabi. Coding started at the course topics section of each syllabus, as this section most explicitly signals the instructional priorities. Then the learning objectives and course descriptions were coded to identify additional ethical themes or concepts that were not explicitly listed in the topics section. This helped ensure that content was captured comprehensively across different syllabus components.

To support research quality, one researcher conducted the primary coding of all syllabi. Rather than engaging in full co-analysis or independent double coding, an additional scholar with doctoral-level expertise reviewed and confirmed the coding decisions. This approach aligns with the reliability procedures described by Creswell & Creswell [20], who note that qualitative researchers may enhance trustworthiness through code cross-checking and intercoder agreement without requiring full independent coding. In this study, discrepancies were discussed and resolved through iterative refinement of the coding framework. Agreement between the primary coding and the confirmation process reached at least 70%, which is considered acceptable in qualitative research. This process helped to validate the category application, to strengthen the interpretive consistency, and to reduce potential bias in the analysis.

To further enhance transparency and replicability, a detailed codebook documenting coding definitions, decision rules, and representative examples was developed (See Table 2). This documentation supports the auditability of the analytic process and allows future researchers to apply or adapt the framework in similar contexts.

Codebook

Code ID	**Code Name**	**Definition**	**Theoretical Reference**	**Signal Words / Indicators**	**Example (Applies)**	**Example (Does Not Apply)**
SI	Social Implication	The social effects, consequences, and changes arising from the use of artificial intelligence, including impacts on human interactions,	UNESCO. [21] *Societal Implications Of AI: Key Insights From Emerging Signals.*	Signal words include AI ethics principle that involves humans with AI such privacy, security, transparency, explainability, agency, etc. This also includes problems like that affect people like bias, automation, labor, health,	Topics addressing algorithmic bias in hiring or the social impact of automation on the workforce.	Technical discussions of model accuracy without reference to human or social impact.

		relationships, and social structures.		social equity, the economy, etc.		
EI	Environmental Implications	The environmental effects and consequences of artificial intelligence, including its resource use, energy demands, and its influence on how humans interact with and manage the natural environment.	UNEP. [22] *Artificial Intelligence (AI) end-to-end: The Environmental Impact of the Full AI Lifecycle*	Signal words include anything related to ecological consequences like climate, energy or power, footprints, consumptions, resource gathering, water usage, sustainability, etc.	Discussion of AI energy consumption or carbon emissions from data centers.	Ethical theory discussions unrelated to environmental or ecological concerns.
TF	Timeless Foundations	Established concepts, theories and traditions in ethics;	Tuovinen & Rohunen [15]	Signal words include definitions or introductions to AI, Ethics, ethical theories like utilitarian or deontology, and any topics that are foundational to AI ethics regardless of eras.	Topics on utilitarianism or Kantian ethics in the context of AI.	Course modules focused only on specific AI applications without theoretical grounding
PG	Practical Guidance	Applied ethics principles and guidelines relevant to AI	Tuovinen & Rohunen [15]	Signal words include any ethical principle or guidelines that are positive e.g fairness, equity, security, ethical considerations etc	Discussion of ethical AI design principles or responsible AI checklists.	Speculative discussions of future AI risks without actionable recommendations.
HN	Here and Now	AI applications of the present and the near future and the ethical issues associated with them	Tuovinen & Rohunen [15]	Signal words include negative implications of AI e.g. algorithmic discrimination, algorithmic bias; Signal words also include well a practical applications e.g predictive policing, GenAI writing, ethical implications or challenges of AI	Case studies on GenAI in education or predictive policing, or economic or equity issues.	Purely futuristic or hypothetical AI scenarios.
BH	Beyond the Horizons	Future potential of AI and the ethical implications of hypothetical scenarios.	Tuovinen & Rohunen [15]	Signal words include futuristic concepts like the singularity and AGI.	Discussions of AGI governance or AI singularity scenarios or transcendence.	Present-day applications such as facial recognition without speculative framing.

Table 2: Analytical Code book used for analysis: combining both AI ethics syllabi topic classification and social and environmental implications of AI.

Results

AI ethics topics identified in the syllabi were classified into four categories adapted from [15]: Timeless Foundations (TF), Practical Guidance (PG), Here and Now (HN), and Beyond the Horizon (BH). The following results describe the extent to which the themes of social and environmental implications were present within each of the four categories. To assess the prominence of these themes, prominence was operationalized as the frequency of coded topics across syllabi. The unit of analysis was the individual topic listed in the syllabi discourse of course descriptions, topics, and learning objectives. Each coded topic was treated as a unique instance and counted once per syllabus to avoid overrepresentation of any repeated topic or concepts within a single course. Using this approach allowed for normalization across syllabi of varying length and level of detail, to ensure that no single course would disproportionately influence the results.

Category	Environmental Implications n (%)	Social Implications n (%)	Total Codes
TF (Timeless Foundations)	0 (0.0%)	8 (14.3%)	56
PG (Practical Guidance)	1 (1.1%)	70 (75.3%)	93
HN (Here and Now)	8 (6.3%)	112 (88.2%)	127
BH (Beyond the Horizon)	0 (0.0%)	4 (44.4%)	9

Table 3: Distribution of Social and Environmental Codes by Analytical Category

Environmental Implications Across Classification Buckets

Environmental implications were sparsely represented across all four categories (Table 3). Within TF, 0% of environmental codes were identified. This indicates that foundational AI ethics concepts rarely incorporate environmental considerations. In PG, environmental implications were minimally present, comprising approximately 1% of total coded content. While this category emphasizes applied ethics and professional responsibility, environmental sustainability appears largely absent from discussions of AI practice and governance. The HN category contained the highest proportion of environmental implications, though these remained limited, accounting for approximately 6% of total codes. Environmental topics in this category were typically mentioned within broader discussions of societal impact rather than framed as a primary ethical concern. Finally, BH showed virtually 0% engagement with environmental implications. Despite BH being future-oriented, this category included little attention to long-term environmental or ecological consequences of AI systems. Overall, environmental implications remain marginal across all AI ethics classification categories, even in categories oriented toward societal impact and future considerations.

Social Implications Across Classification Buckets

In contrast, social implications were prominently represented across all four categories (Table 3). Within TF, social implications accounted for approximately 14% of total codes, suggesting some attention to social concerns at the foundational level, though these were not dominant. The PG category showed a substantial emphasis on social implications, which constituted approximately 75% of the total codes. This reflects a strong focus on issues such as fairness, accountability, professional responsibility, and ethical decision-making. The highest amount of social implication codes appeared within the HN category, where social implication codes accounted for approximately 88% of total content. This category emphasized immediate societal concerns, including bias, job security, governance, and the social consequences of AI deployment. Within BH, social implications accounted for approximately 44% of total codes. Although this category had relatively few overall codes, social considerations were still present, indicating some engagement with future societal trajectories and systemic impacts.

Comparative Patterns

Taken together, the results reveal a clear imbalance in how AI ethics curricula address social versus environmental concerns. Social implications are consistently foregrounded across all classification buckets, particularly within PG and HN. Environmental implications, by contrast, remain underemphasized, appearing infrequently even in categories that emphasize ethical reflection, societal impact, and future-oriented thinking. These patterns suggest that recent AI ethics education at the university level remains strongly anthropocentric, with limited integration of environmental sustainability as a core ethical concern.

Limitations

This study has several limitations that should be considered when interpreting the findings. First, the sample consists of 24 AI ethics syllabi identified through publicly available Google searches, which may not capture the full diversity of AI ethics curricula globally or those that are not publicly available. Second, the analysis relied only on the syllabi topics as an estimation rather than actual classroom activities or observation, meaning that the presence or absence of social and environmental concepts reflects only what is formally stated in syllabi, not necessarily what is actually emphasized in classroom practice or activities.

Discussion

The results indicate a clear imbalance in how social and environmental implications of AI are represented across AI ethics curricula. Social implications are consistently emphasized across categories, particularly within the Here and Now and Practical Guidance buckets, suggesting that AI ethics education prioritizes immediate societal concerns such as fairness, bias, accountability, labor, and governance. This aligns with existing literature that positions AI ethics as primarily human-centered and policy-driven, often framed around present-day risks and regulatory responses. Environmental implications, in contrast, are sparsely represented overall and more present to the Here and Now category, with minimal presence in Timeless Foundations and near absence in Beyond the Horizon. This pattern suggests that environmental concerns such as energy consumption, resource extraction, and climate impacts are treated as secondary or emerging issues rather than foundational ethical considerations in AI ethics. The limited integration of environmental topics supports prior critiques that AI ethics discourse tends to foreground social harms while sidelining ecological dimensions, despite growing evidence of AI's environmental footprint [2], mainly since the GenAI boom. Overall, these findings suggest that while AI ethics education has made significant progress in addressing social concerns but remains limited in its treatment and representation of environmental implications.

To support translation of these findings into practice, a coverage checklist aligned with the four analytical categories identified in this study was developed (see list below). This artifact provides educators with a structured tool to assess and redesign AI ethics curricula, particularly to strengthen the integration of environmental sustainability alongside social and human-centered concerns when teaching topics that span across the categories of Tuovinen & Rohunen's [15] categorization.

<u>1. Timeless Foundations (Theory - Philosophy)</u>

- ☐ Definition of: AI, ML, Data, and Ethics
- ☐ Ethical theories (e.g., utilitarianism, deontology, virtue ethics)
- ☐ Ethical dilemmas (e.g., Trolley Problem, Chinese Room)
- ☐ Historical perspectives on technology and ethics
- ☐ Concepts of human-centeredness and agency
- ☐ Critical and philosophical analysis of AI and Data

<u>2. Practical Guidance (Practice - Philosophy)</u>

- ☐ Fairness, accountability, transparency, and explainability
- ☐ Diversity, Justice, Rights and Responsibility
- ☐ Responsible and trustworthy AI frameworks
- ☐ Governance, standards, and professional ethics
- ☐ Ethical decision-making and design trade-offs
- ☐ Value-sensitive or human-centered design
- ☐ Risk and impact assessment
- ☐ Stakeholder engagement and participatory approaches

<u>3. Here and Now (Practice - Technology)</u>

- ☐ Algorithmic bias and discrimination
- ☐ Privacy, surveillance, and data governance
- ☐ Labor, automation, and economic inequality
- ☐ Generative AI and large language models
- ☐ Predictive policing and criminal justice
- ☐ Misinformation, deepfakes, and information integrity
- ☐ AI in healthcare, education, and public services
- ☐ Platform economies and algorithmic management
- ☐ Ethical risks in real-world deployment

4. Beyond the Horizon (Theory - Technology)

- ☐ Artificial general intelligence (AGI) and superintelligence
- ☐ The singularity and existential risk
- ☐ Human enhancement and human-AI integration
- ☐ Long-term governance and global coordination
- ☐ Speculative and future-oriented ethical scenarios

5. Environmental Implications (Theory, Philosophy, Practice, and Technology)

- ☐ Energy and carbon impacts of AI systems
- ☐ Sustainable AI and green computing
- ☐ Resource extraction and hardware supply chains
- ☐ Water and infrastructure impacts
- ☐ Environmental justice and global inequality
- ☐ AI for climate change and adaptation

Conclusion

This study examined how social and environmental implications of AI are presented in university-level AI ethics curricula using an AI ethics course topics classification framework. The findings show that AI ethics education strongly emphasizes social concerns, particularly those situated in immediate and practical contexts, while environmental implications remain limited and unevenly integrated. When environmental topics do appear, they are largely framed as present-day issues rather than foundational ethical principles or even long-term considerations. This imbalance suggests that current AI ethics curricula continue to prioritize human-centered and policy-driven narratives, with less attention to AI's environmental and sustainability implications. As AI systems, in particular GenAI, continue to expand in scale and application with many data centers being developed, integrating environmental considerations across dimensions of AI ethics education will be critical for preparing students to engage with the full spectrum of ethical challenges posed by AI.

References

[1] Russell, S. J., & Norvig, P. (2021). Artificial Intelligence: A Modern Approach (4th ed.). Pearson.
https://www.amazon.com/Artificial-Intelligence-A-Modern-Approach/dp/0134610997

[2] Uffelen, van, Lauwaert, L., Coeckelbergh, M., & Kudina, O. (2024). *Towards an Environmental Ethics of Artificial Intelligence*. ArXiv.org. https://arxiv.org/abs/2501.10390

[3] Wiese, L. J., Patil, I., Schiff, D. S., & Magana, A. J. (2025). AI Ethics Education: A Systematic Literature Review. Computers and Education Artificial Intelligence, 8, 100405–100405. https://doi.org/10.1016/j.caeai.2025.100405

[4] Hooper, K., & Fletcher, T. (2022, August). An analysis of engineering and computing students' attitudes to AI and ethics. In *2022 ASEE Annual Conference & Exposition*.

[5] Ojong, N. (2025). Interrogating the Economic, Environmental, and Social Impact of Artificial Intelligence and Big Data in Sustainable Entrepreneurship. Business Strategy and the Environment.

[6] Gwagwa, A., Kazim, E., & Hilliard, A. (2022). The role of the African value of Ubuntu in global AI inclusion discourse: A normative ethics perspective. Patterns, 3(4).

[7] Turyasingura, B., Ayiga, N., Byamukama, W., Kayusi, F., Tumuhimbise, M., & Tumushabe, T. (2024). Application of artificial intelligence (AI) in environment and societal trends: Challenges and opportunities. Babylonian Journal of Machine Learning, 2024, 177-182.

[8] Mulita, R., Nasaraia, Z., Konarivska, O., Boiko, S., & Paniuk, T. (2024). Societal Impacts And Ethical Considerations Of AI In The Business. Business Management/Biznes Upravlenie, (4).

[9] Nosirov, I., Yormatov, I., Yuldasheva, N., & Avulchayeva, F. (2024, April). AI and Corporate Sustainability: Exploring the Environmental and Social Impacts of AI Integration. In 2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS) (Vol. 1, pp. 1-7). IEEE.

[10] Garrett, N., Beard, N., & Fiesler, C. (2020, February). More than" If Time Allows" the role of ethics in AI education. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (pp. 272-278).

[11] Bashir, N., Donti, P., Cuff, J., Sroka, S., Ilic, M., Sze, V., ... & Olivetti, E. (2024). The climate and sustainability implications of generative AI.

[12] Beets, K., & Hartnett, M. (2024). North America Data Center Report | H2 2023. Technical Report. JLL Research. https://www.us.jll.com/en/trends-and-insights/research/na-data-center-report.

[13] Javed, R. T., Nasir, O., Borit, M., Vanhée, L., Zea, E., Gupta, S., ... & Qadir, J. (2022). Get out of the BAG! Silos in AI ethics education: Unsupervised topic modeling analysis of global AI curricula. Journal of Artificial Intelligence Research, 73, 933-965.

[14] Alam, A. (2023). Developing a Curriculum for Ethical and Responsible AI: A University Course on Safety, Fairness, Privacy, and Ethics to Prepare Next Generation of AI Professionals. Lecture Notes on Data Engineering and Communications Technologies, 879–894. https://doi.org/10.1007/978-981-99-1767-9_64

[15] Tuovinen, L., & Rohunen, A. (2021). Teaching AI Ethics to Engineering Students: Reflections on Syllabus Design and Teaching Methods Long paper. https://ceur-ws.org/Vol-3069/FP_02.pdf

[16] Fiesler, C., Garrett, N., & Beard, N. (2020). What Do We Teach When We Teach Tech Ethics? Proceedings of the 51st ACM Technical Symposium on Computer Science Education. https://doi.org/10.1145/3328778.3366825

[17] Díaz, P., Onorati, T., & Aedo, I. (2024). Using Active Learning and Critical Thinking to Identify and Apply Ethical Values in Engineering Education. 2024 IEEE Global Engineering Education Conference (EDUCON), 1–10. https://doi.org/10.1109/educon60312.2024.10578610

[18] Hooper, K., & Lunn, S. J. (2024). Values in Education: Exploration of Artificial Intelligence Ethics Syllabi Using Natural Language Processing Analyses. 2024 IEEE Frontiers in Education Conference (FIE), 1–8. https://doi.org/10.1109/fie61694.2024.10893167

[19] Columbia University. (2025). *Content Analysis*. Www.publichealth.columbia.edu. https://www.publichealth.columbia.edu/research/population-health-methods/content-analysis

[20] Creswell, J. W., & Creswell, J. D. (2017). *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications.

[21] UNESCO. (2025). *SOCIETAL IMPLICATIONS OF AI: KEY INSIGHTS FROM EMERGING SIGNALS*. https://www.unesco.org/sites/default/files/medias/fichiers/2025/12/Societal%20implications%20of%20AI%20-%20Insight%20from%20Signals.pdf

[22] UNEP. (2024). *Artificial Intelligence (AI) end-to-end: The Environmental Impact of the Full AI Lifecycle Needs to be Comprehensively Assessed - Issue Note*. https://doi.org/10.59117/20.500.11822/46288