http://dx.doi.org/10.21622/ACE.2026.06.1.1906

INTEGRATING ETHICS IN ARTIFICIAL INTELLIGENCE SYLLABUS IN UNIVERSITY AND CO-TEACHING WITH INDUSTRY EXPERTS: PERCEPTIONS AND EXPERIENCES OF STUDENTS FROM KENYA

Fredrick Mzee Awuor

Kisii University, School of Information Science and Technology, Kenya

Email: {fawuor@kisiiuniversity.ac.ke }

Received 20 January 2026 - Accepted 14 March 2026 - Published 14 May 2026

ABSTRACT

The common model adopted by most departments of Computer Science (CS) in universities in Kenya for teaching CS courses is the traditional faculty-led and content-driven approach that denies the learners the opportunity of learning from peers or industry experts or problem-based projects. At Kisii University in Kenya, we are trying a different model: the Faculty of CS teamed up with the faculties from the departments of psychology and law, and industry experts to redesign the syllabus for artificial intelligence (AI) course then delivered it using a collaborative co-teaching approach that integrated various methodologies including flipped classrooms, project-based learning, problem-based learning, and traditional lecture-based approaches. This paper reports on the perceptions and experiences of students on the proposed collaborative teaching framework for the ethical AI undergraduate level 300 course in the CS curriculum and how insights from these experiences can be transferred to other undergraduate CS courses. Our aim is to advance our understanding of how to promote student-centered learning, encourage collaboration between academia and industry, and across disciplines in the university for CS curriculum delivery, and advance dialogue on the integration of ethics in the AI syllabus.

Keywords: *Ethics, artificial intelligence (AI), computer science, co-teaching, collaborative teaching, pedagogical framework.*

1. INTRODUCTION

Among the revolutionary technologies of the 21st century, artificial intelligence (AI) has taken a lead role in driving every operation of economies by changing how we work, commute, communicate, and access services, and even our social lives, including how we connect and build relationships [1]. For developing countries, AI has positioned itself as an accelerant for economic transformation as it is creating new and scaling existing opportunities in all sectors of their economies [2]. For example, learners in rural and remote areas are now able to access education through learner support services [3], farmers can access markets faster through AI-based initiatives [4, 5], improve their farm production through precision farming [5], and access low-cost financial services on their mobile [6], among others.

However, these AI systems inherit and even amplify our own inherent biases as they are built on models and trained on datasets that reflect our prejudices, in addition to the cognitive and algorithmic bias that developers unconsciously introduce into these systems. As a result, we are witnessing a myriad of unethical AI applications - from applications that discriminate access to essential services such as credit facilities based on race or gender or sexual orientation or one's residence [7] to digital art applications such as LAION-5B [8] that are trained on datasets of personal images illegally harvested from social media[1] site. In the US, at least 9 out of 10 organizations reported having had instances when their AI systems were unethical, and 65% of executives in the same period reported being aware of instances when their AI systems were discriminatory and biased[2]. While these statistics are from the US, it is obvious that Kenya, just like most countries around the world, will

[1] https://laion.ai/blog/laion-5b/
[2] https://enterprisersproject.com/article/2020/10/artificial-intelligence-ai-ethics-14-statistics

soon grapple with these issues, given that organizations in Kenya are quickly adopting AI to improve their efficiency and productivity, and to stifle competition. For instance, Worldcoin [9] was barred from collecting iris-biometric data, sensitive personal data, from Kenyan in exchange for financial incentives[3]. Another example is the "Huduma Namba" project [10], an initiative rolled out by the Government of Kenya to centralize a national digital identity database that was later halted by the courts in 2021[4]. In both cases, the court faulted data protection and privacy concerns [10].

Collectively, these illustrations demonstrate that the extent to which we can exploit AI technology in solving our societal problems and challenges is a direct function of how committed we are as a humanity in mainstreaming and integrating ethics, responsibility, and human-centric computing in AI. To achieve this, we also need to develop architectures that designate users (humans) with the responsibility for AI. Users should be able to exercise control, authority, and accountability across the entire AI lifecycle, including its design, development, implementation, deployment, and adoption. This will enable holding AI system owners and engineers accountable when their AI systems are not aligned with our ethics, morals, and socialization as a society [2]. This demonstrate why Kenya, similar to other African countries, may need to urgently develop a generation of scientists, engineers and technologists who are human centric AI enthusiasts – AI experts who are not only technologically adept but who are also ethically, morally and socially grounded and able to think and reason ethically, and champion for ethical AI technologies in the society [9].

It follows, therefore, that AI engineers and researchers need to understand ethical concerns of the AI systems that they develop right at the conception of the project, and then leverage ethical and responsible AI frameworks and guidelines to inform their decision on whether to implement or drop the project. This is an essential skill towards achieving ethical AI that needs to be embedded in CS.[5] curriculum whose graduates end up as AI practitioners and experts in the industry and academia. Evidently, this could ensure that the need for the creation of ethical and responsible AI systems is internalised in the mindset of the future AI practitioners and experts at the earliest possible time, thus assisting in creating a community of ethical AI activists and champions. The premise is that CS students are the future drivers of AI innovations and hence they need to be empowered to formulate, design, develop and deploy AI applications that are ethical. As such, they need to be equipped and motivated to promote ethical AI initiatives in their AI projects and to champion it as "a must have" and not just an afterthought or a "good to have". In other words, if society expects organizations to create ethical AI, there is no doubt that this desire will remain unattainable until organizations that create these AI systems are able to identify and onboard skilled ethical AI engineers and scientists. Conversely, where would these organizations recruit ethical AI engineers and scientists from unless universities train them? This demands that the universities redesign their AI course syllabi to integrate and train students on ethical and responsible AI as required by the job market. Evidently, such a training should develop capacity among the students to implement ethical AI in their projects, and to advocate for ethical AI systems at their future workplaces. This argument is similar to the assertion by [9, 11] that CS curriculum can only prepare its graduates with technical skills, as well as ethical and socio-cultural awareness, only if these non-technical skills are intentionally integrated into the curriculum. This ensures that graduates develop the competence to understand the potential legal, social, and ethical implications and consequences of their designs, alongside technical competence. According to [12], the CS curriculum needs to empower students to develop a strong understanding of the complexities of power and underrepresentation in computing and how AI can exacerbate these imbalances in the context of power and governance, ethics, equity, race, and even social responsibility.

According to [2], an ethical AI system can only be formulated, designed, or developed, and advocated by a team that is intentional on delivering an ethical AI system. And to achieve this intentionality, the team must have this concept intertwined in the fabric of their training and professional development, which can only be achieved from the very onset of their career development in AI. This means that every AI course in the CS curriculum should integrate

[3] High Court of Kenya. (2025). Republic v. Tools for Humanity Corporation (US) & 8 others; Katiba Institute & 4 others (Judicial Review Application No. E119 of 2023), [2025] KEHC 5629 (KLR). Kenya Law. https://new.kenyalaw.org

[4] High Court of Kenya. (2021). Nubian Rights Forum & 2 others v. Attorney General & 6 others [2021] KEHC 282 (KLR). Kenya Law. https://new.kenyalaw.org

[5] The Computer Science (CS) is used as an umbrella term referring to all disciplines of computing sciences such information technology, computer engineering, computer science, software engineering, information systems among others.

ethical and responsible computing principles. In a typical CS curriculum [11], AI courses are taught separately from the general ethics ones, which mostly cover everything about professional ethics and legal issues in CS, with no specifics on the connection between ethics and CS principles, skills, and knowledge that the students cover in the CS curriculum. These courses on ethics are argued to appear as afterthoughts in the CS curriculum, and hence, students hardly give them much attention [13, 14]. A typical AI syllabus is, therefore, designed to scope and cover the core technical issues that the students need to learn to meet the learning expectations of the syllabus and assumes that ethics in AI is a peripheral non-issue which students will grasp from the professional ethics unit taught separately. As expected, this creates a disconnection between AI as a course and ethics in CS, as learners miss out on how to integrate ethical reasoning and thinking in AI projects [11, 15, 16]. Then it is impressive that to inculcate ethical AI in CS graduates, we must rethink and redesign the AI syllabus to integrate ethics, and teach it using a student-centered pedagogical framework that provides students with practical and real-world demonstrations.

1.1. CONTRIBUTION

In this paper, we describe an initiative proposed by the Department of CS at Kisii University to bridge the aforementioned gap by integrating ethical reasoning and thinking into the AI course syllabus in the CS curriculum and delivering it through a co-teaching approach with the Faculty of Law, the Faculty of Psychology, and ethical AI industry experts. Specifically, the paper reports on students' perception and experience on the proposed collaborative teaching of ethical AI course among undergraduate students of CS and how learning from this experience can be transferred to other undergraduate CS courses. Our aim is to advance our understanding of how to promote student-centered learning, encourage collaboration between academia and industry, and across disciplines in the university for CS curriculum delivery, and advance dialogue on the integration of ethics in the AI syllabus. Therefore, this paper interrogates two issues: (i) how can ethics be integrated in an AI syllabus and collaboratively taught by experts from the faculties of CS, law, and psychology, and ethical AI industry professionals; and (ii) what is the students' experience and perception in integrating ethics in an AI syllabus? The learning outcome indicates that the students appreciated that ethics in AI cannot be an afterthought after completing an AI project, but must be integral in AI system formulation, design, and implementation. Students reported that the learning sessions were engaging and interesting, and satisfactorily provided individual attention and learning support.

2. RELATED WORK

The emergence of AI technology is significantly changing how we live, work, interact, communicate, seek education, or health services, just to mention a few. As such, AI is becoming a core technology in our daily lives, a function that makes AI a technology that we cannot wish away. As the old aphorism goes, *"fire is a good servant but a bad master"*, so is AI technology, which has been shown to cause unintended but harmful consequences when its design, development, and deployment do not align with our ethics, morals, values, and culture as a society. Taxonomy developed by [17] classifies these unintended AI harm, risks and vulnerability as: adversarial threats; privacy and personal data protection violation; mis/disinformation, AI biases and discrimination; system safety and reliability failures; socioeconomic exploitation and inequality; carbon emission, environmental and ecological misuse; loss autonomy and weaponization; loss accountability, and human interaction and psychological harm.

Unfortunately, most of these harms [17, 18] only become evident after the AI system has been deployed to the market and not during the development. At this point, the developer's response becomes reactive, aimed at addressing these negative effects as activists and users demand accountability. In this case, ethical consideration is an afterthought being treated post-development of the AI system. Ideally, ethical thinking and reasoning should be integrated into the lifecycle of AI system development to enable early detection and treatment of potential harms. Achieving this, AI system owners and creators must develop

http://dx.doi.org/10.21622/ACE.2026.06.1.1906

a balance between maximizing return on investment and mitigating AI harm, and between pursuing efficiency, and ensuring inclusivity and representativeness in the datasets used, among others [19]. Nonetheless, at the center of this dilemma, what is fundamental is the ability to identify, reason through, and solve ethical problems in AI-based systems - a critical skill that CS departments must deliberately cultivate in CS graduates. According to [1], integrating ethical consideration in the development of an AI system must be sensitive to our cultural values, beliefs, and ethical principles. That is, at the core of an ethical AI system is social and cultural sensitivity and awareness to the diverse user population that it intends to serve.

For clarity, we define ethics in AI, similar to [14, 20], as concept and framework of social responsibility and social justice principled on social values, beliefs and moral guidelines that informs how AI systems can be formulated, designed and developed for societal good. Therefore, the social responsibility and framing of justice in the context of AI relates to aligning the development of AI systems to the ethical principles - fairness and non-discrimination, accountability, transparency, explainability, privacy and data protection, non-maleficence, inclusion and accessibility [21]. As such, ethics in computing education is intended to inculcate ethical and informed decision-making, and to leverage this reasoning in the creation of AI systems [15]. So, talking of ethical computing or ethical AI, the attention is on skills and practices that embed ethical thinking and reasoning into the formulation, creation, testing, deployment, and use of AI technology [16]. Given this discussion, we will use the ethical AI and responsible AI synonymously.

In integrating ethics in AI syllabus, the goal is to ensure that the AI syllabus meets and delivers the bare minimum of ethical (responsible) AI as per the above definition framework - with expectation that this contributes to building a community of graduates of CS who will palliate AI system harm and biases in the society by advocating for it through systems that they develop [20, 22, 23]. In addressing this need, universities have developed standalone computing ethics courses in the CS curriculum, while some have integrated ethics modules and discussions into pre-existing curricula [13]. For a while now, this has been a discussion in the community of practice of computer education on the role curriculum could play towards the attainment of an ethical AI system in society, and how this can be achieved. Such initiatives are exhaustively discussed and evaluated in [20] and [16].

According to [24], CS curriculum in most universities in Africa hardly integrates ethical computing, a reason we can attribute to the low ethical AI activism and championing from the AI professionals in Africa. This implies that while the job market would ordinarily expect AI graduates to have these ethical reasoning and thinking, and ethical dilemma-solving skills by the time they enter the job market, universities have not been responsive enough to redesign their syllabi to meet this need. Potentially, this has also created a skillset mismatch between the market demand and that possessed by CS graduates, which arguably contributes to an increase in youth unemployability in Africa. In this regard, authors such as [24, 25] call for urgent interrogation of how Faculty teach, how students learn, and what baselines of ethics are covered in AI and related courses, such as machine learning, data sciences, among others, in view of developing necessary polices and frameworks to support integration of ethical computing in syllabi of these courses.

Given the foregoing, the CS syllabus must be aligned accordingly to equip its graduates with an understanding that just like any other technology, AI is neither neutral nor does it exist in a vacuum. It contains within it inherent biases that reflect the perspective, preference, norm, culture, and perception of its creators and the data from which it is created. Specifically, such biases may arise from unrepresentative or incomplete training data, or due to the way the data was collected and sampled, so called data bias, or bias introduced by the design or assumptions from the AI algorithm itself, so called algorithmic bias, or bias reflecting the assumptions or perspectives of developers or dominant cultural norms, so called developer or cultural bias among other forms of biases [26]. In this regard, [26] draws our attention to the reality that technologies are hardly flawless and that it would be a mirage to purport that we can mitigate all ethical concerns in an AI system. Nonetheless, we have an obligation to optimize the societal benefits of AI while reducing its potential harm. This is in addition to developing these AI systems in an open manner, to the extent of sharing the model and data used to create them.

While it is evident from [13, 15] that most CS departments in the global north understand the value of integrating ethics in the CS curriculum, a couple of factors have been listed to derail this initiative. These issues include: Faculty not being in control of course curriculum and therefore does not have flexibility to adjust a course syllabus as and when they need; Faculty feeling that ethics is for other departments and does not have a place in CS curriculum; feeling that CS course content does not relate to ethics; feeling that technical concepts and skills in computing are of more importance than the ethical skills among others. To surmount these barriers, [13] proposes a collaborative approach built on inter-departmental and inter-disciplinary collaboration between the Faculty of CS and other faculties to co-teach some courses in the CS curriculum. Evidently, this can be enriched further by involving guest lectures from the ethical AI experts and professionals from the industry. They could deliver workshops and seminars on ethics and AI and develop learning resources, including case studies and reflections on practical ethical AI dilemmas that CS faculty may grapple with. In Section 3, we build on this discussion to propose a pedagogical framework to support CS faculty to integrate ethics in the AI syllabus and co-teach it with their colleagues from the faculties of psychology and law, and ethical AI industry professionals and experts. The framework is a multi-pedagogical strategy that incorporates the traditional content-based approaches, such as lecture methods, with emerging active learning approaches, such as problem-based learning, flipped classroom, and project-based learning. This builds on the Embedded EthiCS approach in [19] that limited the collaborative teaching to the faculties of CS and philosophy only. Different from the Embedded EthiCS approach, this paper proposes joint teaching with the ethical AI industry practitioners, which has the advantage of providing the students with real-world case studies of ethical AI dilemmas and reflections on how they handle them, and lessons learnt.

3. MATERIALS AND METHODS

This study adopted a design-based research methodology to enable us to iteratively co-design, co-develop, and co-evaluate interdisciplinary teaching materials and pedagogical frameworks collaboratively with the Faculty from departments of law, CS, and psychology, and ethical AI industry experts. This formed a team of eighteen experts – nine from the Faculty of CS, and two each from the faculties of law and psychology, and five ethical AI experts from the industry. The AI syllabus was redesigned to incorporate principles of ethics, ethical reasoning and thinking, and strategies for solving ethical dilemma problems. This resulted in an ethical AI (EAI) syllabus that was used to teach 130 third-year undergraduate students in the department of CS in the first semester of the 2024/2025 academic year at Kisii University. The team also co-developed a pedagogical framework for implementing this syllabus in addition to course materials such as course notes and tutorials, class and lab assignments, workshop and discussion tasks, guidelines for implementing the EAI class mini-project, and an end of semester exams.

During student onboarding, informed consent was obtained, and a pre-test was administered to assess students' understanding of ethical AI. Students then carried out the ethical AI mini-project in groups of five, under the supervision of a CS faculty member and an expert from either the faculties of law or psychology, or from industry.

Both qualitative and quantitative data were collected through students' feedback using surveys, AI mini-project implementation progress monitor, final presentation evaluation reports, and performance in exams and assignments. We also carried out focus group discussions with students, Faculty, and EAI industry practitioners, and collected reflection write-ups from the participants. Additionally, we also analysed teaching materials and the pedagogical framework developed. Qualitative data were analysed using content analysis to identify patterns related to learning experience and pedagogical effectiveness, while quantitative data were analysed using descriptive statistics to examine trends in student perceptions and learning experience.

http://dx.doi.org/10.21622/ACE.2026.06.1.1906

3.1. FROM AI SYLLABUS TO ETHICAL AI SYLLABUS

The AI course is a level 300 course taught to CS students aimed at introducing fundamental concepts, techniques, and applications of AI, covering the legacy and state-of-the-art AI technologies. At the end of the course, the students are expected to demonstrate theoretical understanding and practical ability to implement AI to solve a typical societal problem. Both the course objectives and content were redesigned to integrate the foundation and principles of ethics, and ethical reasoning and thinking to enable students to identify instances of ethical dilemma and to solve them, besides mastering the technical AI competencies. The structures for the AI course (before integration with ethics) and the revised AI course (after integrating with ethics, so-called ethical AI) are presented in Table 1.

Table 1: Course structures for AI and Ethical AI

Week	AI course outline	Ethics component integrated	Ethical AI course outline
1	Introduction to AI: definition, scope, history, and evolution of AI; strong AI vs weak AI; applications and limitations of AI	Benefits and harms of AI	Introduction to AI: definition, scope, history, and evolution of AI; strong AI vs weak AI; applications and limitations of AI; benefits and harms of AI
2	Intelligent agents: rational agents and environments; performance measures and rationality	Balancing utility functions vs human values	Intelligent agents: rational agents and environments; performance measures and rationality; balancing utility functions of AI vs human values
3 - 4	Problem solving and search: problem formulation; informed and uninformed search strategies; complexity analysis; heuristics and optimization	Optimization for whom -balancing efficiency vs fairness trade-offs; explore hidden assumptions in heuristics	Search and optimization: uninformed and informed search; heuristics and optimization; optimization for whom - efficiency vs fairness trade-offs; hidden assumptions in heuristics
5	Adversarial search and games: game playing and minimax algorithm; alpha-beta pruning; applications in strategic decision-making	Moral reasoning in adversarial contexts and responsibility in strategic system design – case of AI in conflict, surveillance, and warfare	Adversarial search: minimax algorithm; alpha-beta pruning; competitive AI vs cooperative outcomes; AI in conflict, surveillance, and warfare
6	Knowledge representation: logic-based representations; ontologies and semantic networks	Whose knowledge is represented? Cultural bias in ontologies and exclusion through formalization	Knowledge representation: logic-based representations; ontologies and semantic networks; cultural bias in ontologies and exclusion through formalization; inclusive knowledge modelling
7	Reasoning and inference: logical and rule-based inference; forward and backward chaining	Determinism vs human judgment; automation bias; when should humans override AI decisions?	Reasoning and inference: logical and rule-based inference; forward and backward chaining; determinism vs human judgment; human-in-the-loop reasoning
8	Uncertainty and probabilistic models: Bayesian reasoning; reasoning under uncertainty	Risk distribution across populations; thresholds and moral responsibility	Uncertainty and probabilistic models: Bayesian reasoning; reasoning under uncertainty; risk and harm anticipation
9	Machine learning: learning paradigms (supervised, unsupervised, reinforcement learning); model evaluation	Data bias and representation; consent and data provenance; historical bias in training data	Machine learning: learning paradigms (supervised, unsupervised, reinforcement learning); model evaluation; bias identification; ethical data practices
10	Machine learning algorithms: classification, regression, and clustering	Fairness metrics and trade-offs; comparative ethical reasoning	Machine learning algorithms: classification; regression; clustering; fairness metrics and trade-offs; comparative ethical reasoning; algorithmic biases
11	Neural networks and deep learning: neural architectures and model complexity	Transparency reasoning; opacity, explainability, and trust in black-box systems	Neural networks and deep learning: neural architectures and model complexity; opacity and explainability; transparency reasoning
12	AI mini-project: students formulate a basic problem and design and develop an AI solution for it.	End-to-end ethical evaluation of an AI solution.	AI mini-project: students formulate a basic problem and design and develop an AI solution for it; end-to-end integration of ethics in the AI solution

In Table 1, the first column is the ordering of course content in the syllabus that covers 12 weeks of learning. The second column presents the course outline for the original AI course as offered by the department. The third column captures the ethics component that the faculties of CS, psychology, and law, and the ethical AI industry experts collectively identified for integration into the respective AI topics per week. The last column presents the Ethical AI course, that is, the AI course integrating ethics developed by the team from the brainstorming sessions – as shown in Figures 1 and 2. The continuous assessments – assignments, quizzes, group work, lab work, AI application mini-project, and exams all contributed to the final score for the course.

Figure 1: Team brainstorming session redesigning the AI curriculum for the level 300 CS curriculum to Ethical AI by integrating ethics, including ethical reasoning and thinking

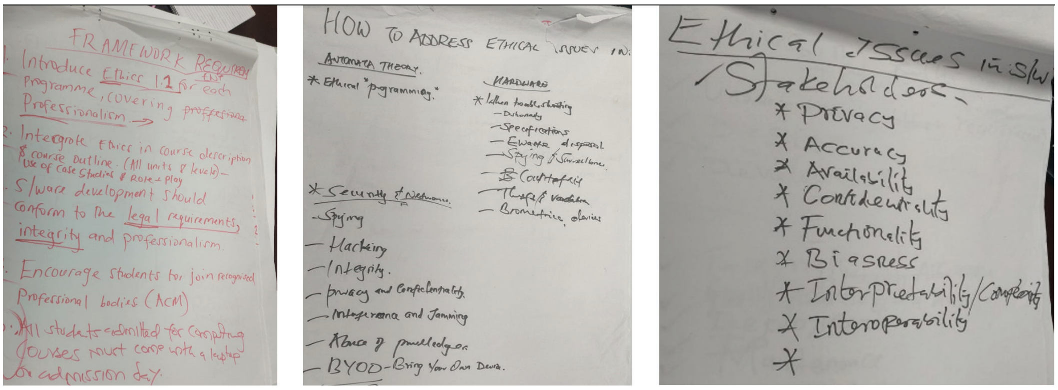

Figure 2: Sample of brainstorming output from the team

The proposed Ethical AI course aimed at integrating core concepts of ethics such as accountability, fairness, non-discrimination and bias, privacy and data protection, transparency and explainability, human in-the-loop/human oversight, AI harm prevention, inclusivity, and social impact into the AI syllabus. The premise here was that upon undertaking the redesigned Ethical AI syllabus, the students would be able to develop both technical AI competencies and skills, ethical responsibility, and professional judgement to align human values in AI systems that they create now and in the future.

3.2. FRAMEWORK FOR COLLABORATIVE TEACHING OF ETHICAL AI SYLLABUS

The team held multiple brainstorming sessions (see Figure 3) to formulate and design a co-teaching approach to structure and inform how the Faculty of CS will collaborate with their colleagues from the faculties of law and psychology, and the ethical AI industry experts and professionals in teaching the course.

http://dx.doi.org/10.21622/ACE.2026.06.1.1906

Figure 3: Experts from the faculties of CS, law, and psychology, and the ethical AI industry professionals, are brainstorming on the EAI collaborative teaching approach

The team proposed that the ethical AI course should adopt a student-centered approach that fosters active learning, enabling students to construct knowledge through discussions, problem-solving, mini-projects, and lab practicals, while guidance and support are provided by Faculty and industry experts. The approach was also designed to promote collaborative learning, including group work, peer-to-peer interactions, and co-teaching by Faculty across disciplines and industry professionals. To operationalize this, the team identified four essential components of the framework: problem-based learning, flipped classroom, case study and scenario-based learning, and project-based learning. These strategies were applied across lecture sessions, workshops, seminars, lab work, and project presentations, as summarized in Table 2.

Table 2: Learning approaches adopted

Learning approach	How to use	When to use	Who to use
Problem-based learning	To equip students with skills to integrate technical competence with ethical thinking and ethical reasoning. For instance, students are asked to analyse problems and identify technical and ethical issues in them, and to develop solutions.	During lectures, lab work, workshops, and seminar sessions	Faculty of CS, Law, Psychology, and Ethical AI industry experts
Flipped classroom	Students were provided with recorded lectures, reading assignments, and tutorials prior to class – these were then discussed in class.	During lectures and lab work	Faculty of CS
Case study and scenario-based learning	Students provided with practical, real-world AI ethical dilemma cases and scenarios - situations and experiences – to analyse and apply their knowledge to and come up with an ethically aligned and informed decision.	During the workshop and seminar, lectures, and lab work sessions	EAI industry experts, Faculty of Law and Psychology
Project-based learning	In groups of five, students develop an AI mini-project to solve a practical problem, considering ethical AI reasoning and thinking, and assessments.	During project consultation and presentation	Faculty of CS, law, psychology, and industry experts

By leveraging these four learning strategies, the teaching of ethical AI transitions from a traditional content-focused, faculty-led model to student-centered active learning. This approach encourages students to discover and construct knowledge through inquiry and reasoning, and to apply it in addressing real-world problems. It fosters a holistic learning environment that equips learners with both technical and ethical knowledge of AI, while providing opportunities to develop AI systems that are socially, culturally, and morally responsible. The proposed pedagogical framework, referred to as the CF2P framework, is presented in Figure 4.

To implement the ethical AI syllabus using the CF2P framework, the course was structured over twelve modules, comprising three lecture hours per week, two hours of lab sessions, one hour of weekly project consultation and reporting, two six-hour workshops, one three-

http://dx.doi.org/10.21622/ACE.2026.06.1.1906

hour seminar, and a dedicated day for the mini-project final presentation, totalling 103 contact hours. Lecture and lab sessions were facilitated by computer science faculty, while workshops and seminars were led by ethical AI industry experts and Faculty from law and psychology. Mini-projects were jointly supervised and evaluated by Faculty from CS, law, and psychology, alongside industry experts.

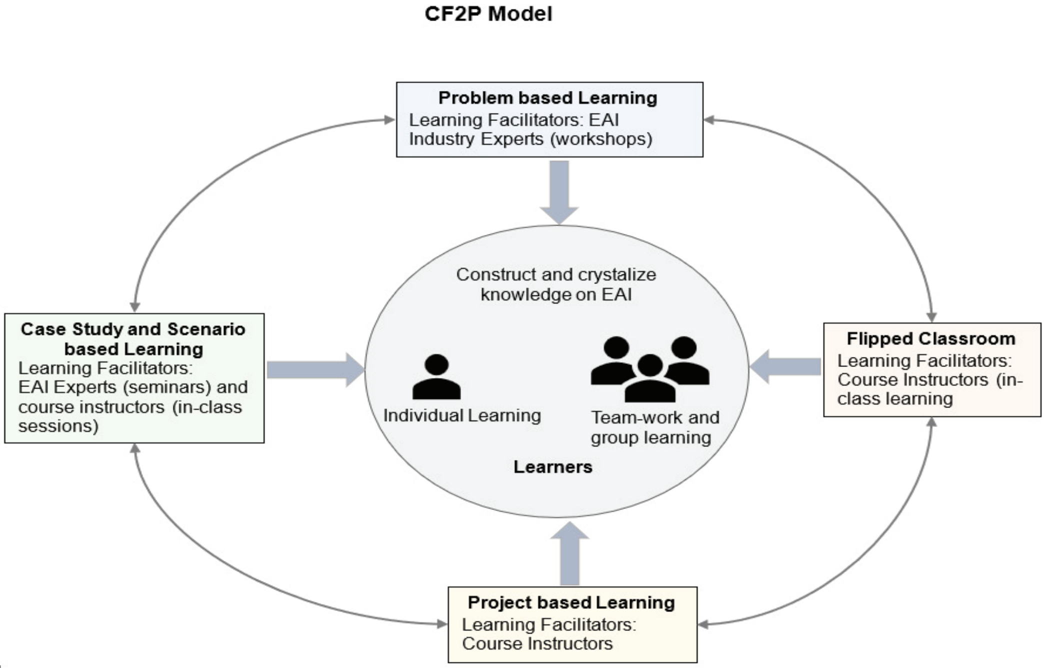

Figure 4: CF2P Pedagogical Framework

The proposed CF2P pedagogical framework engages CS faculty, Faculty from the departments of psychology and law, and ethical AI industry experts to co-deliver the syllabus, equipping learners with both technical competencies in AI and a deep understanding of ethics and moral responsibilities as AI practitioners. Collaboration with multidisciplinary Faculty and industry experts provided CS instructors with the necessary support to integrate ethics effectively into the course. This partnership also allowed Faculty to learn from peers in academia and industry, incorporating real-world case studies and documented experiences on ethical AI to enrich their broader teaching.

4. RESULTS AND DISCUSSION

The aim of this study was to redesign the Level 300 undergraduate AI course syllabus to incorporate principles of ethics, ethical reasoning, and critical thinking, and to implement it through a co-teaching approach involving industry experts and multidisciplinary Faculty from CS, law, and psychology. A summary of this implementation is presented in Table 3.

Table 3: Team co-teaching ethical AI course

SN	Description	Department	Number
1.	Internal Faculty	Computing Sciences	7
2.	External Faculty	Faculty of Law Faculty of Psychology	2
3.	External Faculty (Ethical AI industry experts)	Data Protection and Privacy Expert, Cybersecurity and AI Security Expert, AI Governance and Policy Expert, Human-Centered AI Expert, AI Product Manager, AI Ethics Lead	8
4	Students enrolled	Faculty of Computing Science	135

4.1. LEARNING ETHICAL AI

At the start of the course, 55.6% of students reported being unfamiliar with ethical AI, and 84.3% indicated they had no prior experience working on or developing ethically-centered AI applications. Despite this, 96% of students agreed that AI systems should be ethical, yet 84.3% reported not knowing how to implement ethical considerations in their projects. Notably, 4% of students did not perceive a need for ethical AI systems. These findings suggest that the existing computer science curriculum primarily emphasizes technical computing skills, with limited focus on equipping students with the knowledge and capacity to be responsible and accountable for the societal impacts of the systems they develop. For some students, the responsibility for ethical computing appears externalized, considered an afterthought rather than an integral part of system design, consistent with observations reported in [11, 22, 27].

At the end of the course, 98% of students reported being familiar with ethics in AI systems, an increase from 44.4% at the start. However, only 55% indicated they could effectively integrate and apply ethical principles in a practical AI project, despite the course incorporating project-based and problem-based learning with mentorship in labs and workshops. Analysis of submitted projects revealed that approximately 45% did not adequately demonstrate understanding of key ethical AI principles, including accountability, explainability, transparency, data privacy and security, human agency, safety, fairness, and bias avoidance, all of which were covered in the course. This outcome can be attributed to the course assessment structure: class assignments (20%), continuous assessments (10%), a mini-class project (10%), and end-of-semester exams (70%). Consequently, students gave limited attention to the project component. Furthermore, as projects were completed in groups of five, it was observed that the majority of groups relied on one or two members with stronger programming skills, rather than collaborating equally on the project.

4.2. COLLABORATIVE TEACHING OF ETHICAL AI

At least nineteen out of twenty students, that is, 95.5%, reported being satisfied with the amount of individual attention and support provided during the learning sessions, indicating that the teaching approach ensured that all students participated and got involved in the learning. The students reported that the learning sessions were engaging and interesting, and that there were different and new learning activities, including guest lectures from industry experts, and the Faculty from the departments of law and psychology, and workshops and seminars that made learning fun. They also reported learning new soft skills such as teamwork, pitching and communication of science, human-centered AI system design and development, and presentation.

Their response also showed that case studies and scenario-based learning, and problem-based learning (30% and 28% respectively) were the most preferred learning approaches in the CS2P framework. They considered that these two approaches provided more engagement among students and between them and the facilitators. Other learning approaches that were reported to provide active participation were lectures and demonstrations (20.1%), followed by project-based learning (13%). The low preference for project-based learning could be attributed to low skill level in computer system project work, which is usually taught in year three, and as such, these students were learning computer system project work in a different course and were required to implement the skills in the ethical AI course. The flipped classroom was the least preferred learning approach by students (8.9%). This could be attributed to the need for pre-class personal engagement with course materials demanded in the flipped classroom approach that students found demanding. Notably, these different approaches being leveraged in this framework complemented each other, as 97% of the students appreciated that the teaching method used significantly enhanced their understanding and mastery of the course content. When asked about the likelihood of recommending this teaching method to their lecturers in other courses, 90% of the students indicated that they are likely to recommend this learning pedagogical framework. While students did not show a preference for any of the core components of the pedagogy if implemented as a standalone, we notice that combining them builds a holistic approach to learning and is preferred by the students. This is because the CS2P pedagogical framework leverages the strength of each individual learning framework to deliver a student-centric learning experience. To account for the 10% who did not find this approach of teaching to be effective, we argue that being a new approach, there is a need for mindset change among students and capacity building to promote acceptance and shift the mindset from the traditional lecture-based approach that they are used to. A summary of the activities carried out in this study and the results is presented in Table 4 below.

Table 4: Summary of the activities and results

Description	Results	Learnings	Challenges/Obstacles
AI course syllabus	An ethical AI course was developed by redesigning the AI course to integrate ethics.	The ethical AI syllabus is too broad to be covered in one semester – it can be divided into two levels, covering two semesters, to enable mastery of the content and practical learning.	Faculty of psychology and law needed to have some understanding of core concepts of AI to enable them to develop valuable feedback and insights on the integration of ethics in the AI syllabus.
Pedagogical framework for collaborative teaching	The student-centered co-teaching approach, the CF2P model, that integrates case studies, flipped classroom, project-based learning, and problem-based learning, was co-developed to co-teach the ethical AI syllabus.	The approach implements active learning and focuses the learning activity on the student. Preliminary results are promising, but need further investigations to see if this model can supplement the traditional lecture-based method of teaching in other computing courses, too.	Need to build the capacity of the Faculty on active learning teaching approaches. Need for mindset change from students and Faculty – for instance, students found the approach a little too involving, as they are used to the lecture approach that needed minimal involvement from them.
Learning materials, including recorded lessons, case studies, problems, etc	Course learning materials were developed, including assignments, case studies, video and audio recordings, group discussion tasks, tutorials, and lecture notes that students accessed from the university's learning management system.	Need for retooling and re-skilling of Faculty on instructional design and development of interactive content using the emerging authoring tools.	Most students could not afford Internet connectivity while off campus. Students relied on the university's internet connection for their learning.
AI class mini-projects integrating ethical AI concepts learnt	26 projects developed by the students, 55% of which integrated ethics – ethical reasoning and thinking in solving the problem using AI.	Ethical reasoning and thinking need to be integrated in all CS courses and class mini-projects.	Need for capacity building to CS faculty and students on the trade-off between the protection of intellectual property rights of students' innovations (projects) versus project mentorship by the faculty.

4.3. POTENTIAL GAINS

In collaborative teaching with the ethical AI industry experts and the faculties of law and psychology, the CS students reported having forged an inter-departmental and multidisciplinary network from these faculties and the industry, in addition to appreciating how ethical computing cuts across multiple disciplines. Students learnt so many other skills besides the AI and ethics skills – including marketing, communicating science in their project, communication skills, teamwork, pitching among others.

The diversity in the teaching approach and diversity in profiles of the Faculty and industry experts who delivered the course enabled the students to appreciate why ethics, morals, values, and culture are key in AI systems and need to develop AI systems with multidisciplinary experts. For instance, 93% of the students reported that the teaching method provided them with opportunities for collaboration and interaction with other learners in a manner they had never experienced before, and that they now appreciate the need for teamwork in developing computing applications even with peers from non-computing disciplines.

Additionally, majority of the students (87.6%) indicated that the ethical AI course made a great impact in their future career paths by motivating them to explore unexpected career paths such as ethical AI, AI safety and trust among others.

4.4. LESSONS LEARNED FROM INTEGRATING ETHICS IN AI

The baseline results reveal a significant gap in students' prior exposure to ethical AI. More than half of the students (55.6%) initially reported unfamiliarity with ethical AI, while 84.3% had no experience developing ethically centred AI applications. Despite this lack of exposure, a large majority (96%) acknowledged the necessity of ethical AI systems, though 84.3% did not know how to operationalize ethical principles in practice. This gap highlights a structural limitation in conventional CS curricula, which tend to prioritize technical competencies while giving limited attention to the societal and ethical implications of computing technologies. The perception among a small minority (4%) that ethical responsibility lies elsewhere further reinforces the notion that ethics has traditionally been treated as peripheral rather than integral to AI system development.

Following the Ethical AI course, the results demonstrated a substantial improvement in students' conceptual awareness of ethical AI. Familiarity increased to 98%, suggesting that the curriculum successfully introduced and contextualized ethical principles in AI development. However, the translation of conceptual understanding into practice remained limited, as only 55% of students demonstrated the ability to integrate ethical principles into AI projects. The analysis of submitted projects supports this finding, with approximately 45% failing to adequately demonstrate key ethical AI principles such as fairness, accountability, transparency, data privacy, and bias mitigation. This outcome appears to be influenced by assessment design, where the project component constituted only 10% of the overall grade and was completed in groups, often resulting in uneven participation.

The perception of students on the collaborative and interdisciplinary teaching approach was highly positive. The involvement of experts from the Faculty of CS, law, psychology, and industry enhanced engagement and broadened students' understanding of ethical AI as a multidisciplinary domain. High satisfaction rates (95.5%) and strong appreciation for interactive pedagogies - particularly case study/scenario-based and problem-based learning - indicate that such approaches effectively support deeper engagement with ethical considerations in AI.

The findings indicate that integrating ethics into AI education through collaborative and multidisciplinary teaching can enhance students' awareness and appreciation of ethical AI. However, the results also show that translating this understanding into practice remains a challenge. Strengthening assessment structures, particularly those that emphasize practical application, may help ensure that students are able to effectively integrate ethical principles into the design and development of AI systems.

While the results appear to provide insights on the integration of ethics in AI and collaborative teaching with interdisciplinary faculty and industry professionals, we appreciate that the study is limited to the context of Kisii University, with findings based on self-reported experiences and perceptions from students who registered for the course, which may be subject to response bias therein, affecting the generalizability of the findings. Subsequently, the responses may be influenced by response bias, which may limit the extent to which the findings may be generalized. Future studies could address this limitation by examining similar interventions across multiple institutions, courses, and contexts. Such broader investigations would allow for comparative analysis and provide stronger evidence regarding the applicability of the results in different educational settings.

5. CONCLUSION AND FUTURE WORK

While society expects organizations to develop and deploy ethical AI systems, this goal is unlikely to be achieved without a pipeline of ethically trained AI graduates from universities. Consequently, universities must redesign their AI curricula to integrate ethics to producing graduates capable not only of formulating, designing, and deploying ethical AI systems but also advocating for responsible AI practices. To address this objective, this paper examines two questions: (1) how ethics can be integrated into the AI syllabus and collaboratively taught by Faculty from computer science, law, psychology, and industry experts, and (2) what would be the experiences and perceptions of students enrolled in an ethics-integrated AI course.

To this end, a team of Faculty from CS, law, and psychology, together with ethical AI industry experts, co-developed an ethical AI syllabus by integrating ethics into the Level 300 undergraduate AI course and implemented it for *135 CS students* using a novel *student-centered, collaborative* teaching framework. The aim was to cultivate AI graduates who understand that ethics is not an afterthought, but a core and integral component of AI system formulation, design, and implementation.

Preliminary results indicate that the proposed CS2P pedagogical framework can enhance student engagement and mastery of ethical AI concepts. According to student self-reports, the framework positions learners at the center of the educational process, with Faculty serving as facilitators rather than primary drivers, as in traditional lecture-based approaches. Most students (95.5%) found this approach satisfying, and 90% stated that they would recommend it for other computer science courses. By the end of the course, familiarity with ethics in AI increased from 44% to 98%.

These findings demonstrate that ethical principles, including reasoning and critical thinking, can be embedded within the core AI syllabus. Additionally, these show that learning can be enriched through collaborative teaching with the faculty of law and psychology, and ethical AI industry professionals providing students with a multi-faceted understanding of ethical AI. Moreover, students gained access to real-world case studies, networking opportunities, and potential collaborations with industry experts and Faculty across disciplines.

Looking ahead, it is essential to build the capacity of CS faculty to integrate ethical considerations into the core curriculum and to collaborate effectively with experts from other disciplines and industry in co-teaching. Achieving this requires a mindset shift among Faculty and motivation to embrace student-centered pedagogical approaches.

ACKNOWLEDGEMENT

This work was supported by the Mozilla Foundation through the Responsible Computing Challenge grant award to Kisii University. Our thanks to the faculty, industry experts, and students who assisted in the implementation of this intervention, with special thanks to Benard Maake, Ruth Chweya, Jane Maina, and Teresa Abuya for assisting in conducting the study.

All data and artifacts used in this study can be made available upon request.

REFERENCES

[1] E. Brynjolfsson *et al.*, "Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence," *digitaleconomy.stanford.eduE Brynjolfsson, B Chandar, R ChenStanford Digital Economy Lab. Published August, 2025•digitaleconomy.stanford.edu*, 2025.

[2] D. O. Eke, S. S. Chintu, and K. Wakunuma, "Towards Shaping the Future of Responsible AI in Africa," 2023, pp. 169–193. doi: 10.1007/978-3-031-08215-3_8.

[3] A. U. Karimy, J. Rasuli, and P. C. Reddy, "AI in Online Educational Access: Breaking Barriers for Female Learners," in *2024 21st International Conference on Information Technology Based Higher Education and Training (ITHET)*, IEEE, Nov. 2024, pp. 1–9. doi: 10.1109/ITHET61869.2024.10837681.

[4] R. Vyas, V. Kesharwani, C. Rohra, and R. Sundrani, "AI-Driven Platform for Direct Market Access for Farmers," 2025, pp. 713–730. doi: 10.2991/978-94-6463-738-0_57.

[5] R. Ladhar *et al.*, "AI-based Market Intelligence Systems for Farmer Collectives: A Case Study from India," *ACM Journal on Computing and Sustainable Societies*, vol. 1, no. 1, pp. 1–32, Sep. 2023, doi: 10.1145/3609262.

[6] M. C. Parlasca, C. Johnen, and M. Qaim, "Use of mobile financial services among farmers in Africa: Insights from Kenya," *Glob. Food Sec.*, vol. 32, p. 100590, Mar. 2022, doi: 10.1016/j.gfs.2021.100590.

[7] R. H. N. Sitinjak, S. M. A. Purba, S. Majidah, and N. A. Azis, "Towards a Non-Discriminatory Artificial Intelligence in Healthcare: Ensuring Equal Access and Utilization for Rural and Underdeveloped Areas," *Jurnal HAM*, vol. 16, no. 3, pp. 177–196, Dec. 2025, doi: 10.30641/ham.2025.16.177-196.

[8] J. Taylor, W. Agnew, M. Sap, S. E. Fox, and H. Zhu, "The Algorithmic Gaze: An Audit and Ethnography of the LAION-Aesthetics Predictor Model," *arXiv preprint arXiv:2601.09896*, 2026.

[9] C. Mbogho, S. Azeka, and J. Edward-Gill, "RESPONSIBLE COMPUTING IN KENYAN UNDERGRADUATE CURRICULA: INSIGHTS FROM A YEAR-LONG MULTIVERSITY INITIATIVE," Nov. 2025, pp. 7519–7528. doi: 10.21125/iceri.2025.2090.

[10] N. Kiilu, "Indirect Discrimination: Huduma Namba (Digital Identification) and the Plight of the Nubian Community in Kenya," *Strathmore Law Review*, vol. 7, no. 1, pp. 17–47, Oct. 2022, doi: 10.52907/slr.v7i1.188.

[11] L. Cohen, H. Precel, H. Triedman, and K. Fisler, "A New Model for Weaving Responsible Computing Into Courses Across the CS Curriculum," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, New York, NY, USA: ACM, Mar. 2021, pp. 858–864. doi: 10.1145/3408877.3432456.

[12] J. J. Ryoo and T. Blunt, "'Show Them the Playbook That These Companies Are Using': Youth Voices about Why Computer Science Education Must Center Discussions of Power, Ethics, and Culturally Responsive Computing," *ACM Transactions on Computing Education*, vol. 24, no. 3, pp. 1–21, Sep. 2024, doi: 10.1145/3660645.

[13] 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*, New York, NY, USA: ACM, Mar. 2023, pp. 367–373. doi: 10.1145/3545945.3569855.

[14] G. Musumba, M. Kagiri, H. Mwangi, D. Wachepele, and D. Miano, "Embedded Responsible Computing: Mainstreaming Ethics into Computer Science Curriculum: A Case Study of Dedan Kimathi University of Technology (DeKUT)," in *2024 International Conference on Information and Communication Technology for Development for Africa (ICT4DA)*, IEEE, Nov. 2024, pp. 247–252. doi: 10.1109/ICT4DA62874.2024.10777260.

[15] S. Pasricha, "Ethics in Computing Education: Challenges and Experience with Embedded Ethics," in *Proceedings of the Great Lakes Symposium on VLSI 2023*, New York, NY, USA: ACM, Jun. 2023, pp. 653–658. doi: 10.1145/3583781.3590240.

[16] S. A. Doore, M. Trim, J. Streator, R. L. Blumenthal, A. , Rudra, and R. B. Schnabel, "Multiple Approaches for Teaching Responsible Computing," *arXiv preprint arXiv:2502.10856*, 2025.

[17] N. Seghid, F. Iqbal, K. Al-Room, and Á. MacDermott, "Emerging threats in AI: a detailed review of misuses and risks across modern AI technologies," *Frontiers in Communications and Networks*, vol. 6, Feb. 2026, doi: 10.3389/frcmn.2025.1727425.

[18] D. Acemoglu, "Harms of AI," Cambridge, MA, Sep. 2021. doi: 10.3386/w29247.

[19] B. J. Grosz *et al.*, "Embedded EthiCS: integrating ethics across CS education," *Commun. ACM*, vol. 62, no. 8, pp. 54–61, Jul. 2019, doi: 10.1145/3330794.

[20] N. Brown, B. Xie, E. Sarder, C. Fiesler, and E. S. Wiese, "Teaching Ethics in Computing: A Systematic Literature Review of ACM Computer Science Education Publications," *ACM Transactions on Computing Education*, vol. 24, no. 1, pp. 1–36, Mar. 2024, doi: 10.1145/3634685.

[21] 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*, New York, NY, USA: ACM, Mar. 2021, pp. 515–525. doi: 10.1145/3442188.3445914.

[22] N. Garrett, N. Beard, and C. Fiesler, "More Than 'If Time Allows,'" in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, New York, NY, USA: ACM, Feb. 2020, pp. 272–278. doi: 10.1145/3375627.3375868.

[23] J. Saltz *et al.*, "Integrating Ethics within Machine Learning Courses," *ACM Transactions on Computing Education*, vol. 19, no. 4, pp. 1–26, Dec. 2019, doi: 10.1145/3341164.

[24] I. T. Sanusi and F. G. Deriba, "What do we know about computing education in Africa? A systematic review of computing education research literature," *arXiv preprint arXiv:2406.11849*, 2024.

[25] CUE, "Kisii University: Initiative to Promote Ethical Use of AI among Students puts Varsity on the Map," Mar. 2024, *Varsity Digest 1*. [Online]. Available: https://fliphtml5.com/yvqhd/oaxr/

[26] B. Hofmann, "Biases in AI: acknowledging and addressing the inevitable ethical issues," *Front. Digit. Health*, vol. 7, Aug. 2025, doi: 10.3389/fdgth.2025.1614105.

[27] T. S. Goetze, "Integrating Ethics into Computer Science Education," in *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, New York, NY, USA: ACM, Mar. 2023, pp. 645–651. doi: 10.1145/3545945.3569792.