

Project-Based Artificial Intelligence Ethics Education in Secondary Schools Fosters Critical Thinking and Responsible Technological Innovation among Diverse Students

Okechukwu Chidoluo Vitus*, Ngozi Modesta Okoh, Professor Justice Sassah

Strong and Exploits Bible College and Theological Seminary Ghana

*Corresponding author: Okechukwu Chidoluo Vitus; Email: jlcmedias@gmail.com

Citation: Vitus OC, Okoh NM, Sassah J (2026) Project-Based Artificial Intelligence Ethics Education in Secondary Schools Fosters Critical Thinking and Responsible Technological Innovation among Diverse Students. American J Sci Edu Re: AJSER-344.

Received Date: 06 September, 2026; **Accepted Date:** 14 September, 2026; **Published Date:** 18 September, 2026

Abstract

The rapid integration of Artificial Intelligence (AI) into daily life has outpaced educational curricula, leaving a critical gap in preparing students particularly those from diverse and underrepresented backgrounds—to engage as informed, ethical innovators. This study investigates the impact of a project-based learning (PBL) model for AI ethics education on secondary school students' critical thinking skills and their capacity for responsible technological innovation. A mixed-methods, quasi-experimental design was employed across four diverse secondary schools, involving 180 students ($n=90$ in the experimental PBL group; $n=90$ in a traditional instruction control group). The 10-week intervention engaged students in collaborative projects requiring them to design, critique, and propose governance strategies for real-world AI applications. Quantitative data from the Cornell Critical Thinking Test (Level Z) and the Responsible Innovation Assessment Scale (RIAS) were collected pre- and post-intervention. ANCOVA results revealed that students in the PBL group demonstrated significantly greater gains in both critical thinking and responsible innovation competencies compared to the control group. Crucially, the positive effects were consistent across demographic subgroups, including gender, socioeconomic status, and prior academic achievement. Qualitative analyses of student artifacts and focus group interviews revealed that the PBL approach fostered deep engagement, authentic dialogue, and a sense of agency. The findings suggest that project-based AI ethics education is a powerful pedagogical strategy for cultivating a generation of diverse, critically-minded, and ethically-responsible technological innovators.

Keywords: Artificial Intelligence Ethics, Project-Based Learning, Critical Thinking, Responsible Innovation, Secondary Education, STEM Education, Digital Citizenship, Educational Technology, Diversity and Inclusion, Ethical Reasoning.

1. Introduction

We are living in the age of Artificial Intelligence (AI). From the algorithms that curate our news feeds and recommend our entertainment to the systems that assist in medical diagnoses, financial lending, and criminal justice, AI is no longer a futuristic concept but an invisible architecture shaping modern society. This pervasiveness brings with it immense potential for progress, but also profound ethical challenges. Issues of algorithmic bias, privacy erosion, autonomous decision-making, and job displacement are no longer the sole domain of computer scientists and philosophers; they are civic issues that demand a broad, societal response. As such, the question arises: how do we prepare the next generation not just to use or build these technologies, but to govern them wisely and innovate responsibly?

For decades, technology education in secondary schools has largely focused on technical proficiency teaching students how to code, use software, or understand hardware. While these skills are undoubtedly valuable, they are insufficient for the AI era. A purely technical education creates what has been called the "techno-utopian" mindset, where students are trained to build first and ask questions later. This approach risks producing a generation of innovators who are ill-equipped to anticipate the ethical ramifications of their creations. Furthermore, the historical lack of diversity in technology fields is not just a matter of representation; it is a critical flaw in the design process

itself. When AI systems are built by a homogeneous group, the biases of that group are encoded into the technology, often with harmful consequences for marginalized communities. Therefore, diversifying the pipeline of technological innovators is not merely a social good; it is an ethical imperative for creating fair and just AI systems.

This study argues that the solution lies in a pedagogical shift towards project-based learning (PBL) for AI ethics education. PBL is a student-centered pedagogy that involves students in a dynamic, real-world investigation over an extended period (Krajcik & Blumenfeld, 2006). Instead of passively receiving information, students in a PBL environment work collaboratively to explore a complex question, problem, or challenge. They engage in inquiry, research, critical analysis, and the creation of a tangible product or solution. When applied to AI ethics, PBL offers a powerful context for developing critical thinking. Students are not just memorizing ethical frameworks; they are applying them to messy, real-world AI dilemmas (e.g., "Should an AI system be used for predictive policing in our community?"). This process requires them to analyze information, evaluate competing perspectives, synthesize ideas, and formulate reasoned arguments.

Crucially, the PBL model also aligns with the imperative for diversity and inclusion. Traditional, lecture-based instruction can be culturally agnostic and socially isolating, often failing to connect academic content to the lived experiences of diverse

students (Ladson-Billings, 1995). In contrast, PBL projects can be designed to be culturally relevant and locally situated, allowing students to investigate the impact of AI on their own communities. This personal connection can be particularly empowering for students from underrepresented groups, who may see AI not as an abstract, alien force, but as a set of tools and systems that have direct consequences for their lives. By positioning these students as critical analysts and innovators—rather than passive consumers of technology—PBL can foster a sense of agency and belonging in the STEM+Ethics space (Calabrese Barton & Tan, 2010). This study, therefore, investigates the efficacy of a project-based AI ethics curriculum in fostering both critical thinking skills and responsible innovation competencies among a diverse student population.

1.1 Problem Statement

The central problem this research addresses is the significant pedagogical and curricular gap in secondary education regarding the ethical dimensions of AI. While many schools are beginning to introduce coding and AI literacy, few have robust, evidence-based curricula that integrate ethics into the technical learning process. Furthermore, there is a lack of empirical research on how such curricula can be designed to be both effective and equitable, ensuring that all students—regardless of their background develop the critical thinking skills and ethical sensibilities necessary to navigate and shape an AI-driven world. This study seeks to fill this gap by examining the impact of a project-based learning intervention on student outcomes across multiple dimensions of diversity.

1.2 Research Questions and Hypotheses

This study is guided by the following research questions:

1. What is the impact of a project-based AI ethics curriculum on secondary school students' critical thinking skills compared to a traditional, lecture-based approach?
2. What is the impact of a project-based AI ethics curriculum on students' competencies in responsible technological innovation?
3. Does the impact of the project-based AI ethics curriculum on student outcomes vary across different demographic groups (e.g., gender, socioeconomic status, prior academic achievement)?
4. What are the qualitative experiences and perceptions of diverse students engaging in the project-based AI ethics curriculum?

Based on these questions, the following null hypotheses were tested at a 0.05 level of significance:

- H0₁: There is no significant difference in the post-intervention critical thinking scores between students in the project-based learning (PBL) group and the traditional instruction (control) group, after controlling for pre-test scores.
- H0₂: There is no significant difference in the post-intervention responsible innovation scores between students in the PBL group and the control group, after controlling for pre-test scores.
- H0₃: There is no significant interaction effect between the instructional method (PBL vs. Traditional) and student demographic variables (gender, socioeconomic status) on post-intervention critical thinking and responsible innovation scores.

1.3 Significance of the Study

This research holds substantial significance for multiple audiences. For **educators and curriculum developers**, it provides a detailed, evidence-based model for integrating AI ethics into secondary education. It moves beyond theoretical discussions to offer a concrete pedagogical framework (PBL) that has been empirically tested. For policymakers and educational leaders, the findings offer a rationale and a roadmap for investing in teacher training and resource development for AI ethics education. It frames such education not as an add-on but as a core component of 21st-century learning and digital citizenship. For researchers, the study contributes to the nascent but growing body of literature on AI education, specifically focusing on the intersection of ethics, pedagogy, and equity. Finally, for students and society at large, this research advocates for an educational model that empowers all young people to become thoughtful, critical, and responsible stewards of powerful technologies. The ultimate goal is to foster a future where technological innovation is guided by a deep and diverse sense of ethical responsibility.

2. Literature Review

2.1 The Imperative for AI Ethics Education

The call for AI literacy, including its ethical dimensions, has grown louder in recent years from various sectors, including academia, industry, and government. The core argument is that as AI systems become more complex and autonomous, the public needs a foundational understanding of how they work, what they can do, and the potential harms they can cause. This is not simply about creating more AI developers; it is about preparing all citizens to participate in democratic discourse about AI governance, to be savvy consumers of AI-driven products, and to protect themselves from algorithmic manipulation and bias).

Researchers have begun to outline frameworks for AI literacy, synthesized existing literature to propose a set of competencies, including recognizing AI, understanding intelligence, and, crucially, considering ethics. Their work highlights that an ethical understanding is not a separate topic but a fundamental component of being AI-literate. Similarly, it has championed the "Five Big Ideas in AI" for K-12 education, one of which is "Societal Impact," which explicitly addresses the ethical and social implications of AI. This growing consensus in the literature establishes a clear and urgent need to move from identifying the problem to developing and testing effective pedagogical solutions, which is the primary aim of this study.

2.2 Project-Based Learning (PBL): A Pedagogical Powerhouse

Project-Based Learning has a long and well-documented history as an effective approach to fostering deep learning. Rooted in the constructivist theories of John Dewey and Jean Piaget, PBL emphasizes learning by doing, where students are actively engaged in constructing their own knowledge through meaningful, authentic tasks (Krajcik & Blumenfeld, 2006). A large body of research has demonstrated the benefits of PBL, including its ability to enhance content knowledge, improve problem-solving skills, and increase student motivation and engagement (Condliffe et al., 2017).

The key elements of PBL make it uniquely suited for teaching AI ethics. First, its focus on a driving question provides a "why" for learning. In the context of AI ethics, the driving question is inherently complex and multifaceted, e.g., "How can we ensure

AI systems are fair for everyone?" Second, situated learning—learning that takes place in a context that mirrors its real-world application—is a cornerstone of PBL. Ethical dilemmas are rarely black and white; they require nuanced judgment applied to specific situations. PBL allows students to grapple with this nuance by examining real-world case studies or, better yet, their own proposed AI applications. Third, PBL's emphasis on collaboration mirrors the interdisciplinary nature of real AI development and governance, which requires input from technologists, ethicists, legal experts, and community stakeholders. Finally, the creation of a public product (e.g., a proposal, a presentation, a prototype) gives students a sense of purpose and agency, transforming them from passive learners into active problem-solvers.

2.3 Critical Thinking as a Core Educational Outcome

Critical thinking is consistently listed as one of the most important "21st-century skills" and a fundamental goal of education (Facione, 1990). It encompasses the cognitive skills of interpretation, analysis, evaluation, inference, explanation, and self-regulation. In an age of information overload and sophisticated disinformation, the ability to think critically is paramount. For AI ethics education, it is the very engine of ethical reasoning. Students cannot identify bias in an algorithm without first being able to analyze its data, interpret its outputs, and evaluate its impact.

Traditional pedagogical models that rely on rote memorization and teacher-centered lectures are often criticized for their failure to cultivate critical thinking skills (Paul & Elder, 2006). In contrast, PBL is designed to do just that. By wrestling with an open-ended, ill-structured problem, the hallmark of a good PBL project students are forced to engage in higher-order thinking. There is no single correct answer to "What is the fairest way to use AI in hiring?" Instead, students must consider multiple perspectives, weigh competing values (e.g., efficiency vs. privacy), construct their own arguments, and defend their conclusions. This process is a direct exercise in critical thinking. This study posits that the PBL approach will be more effective at fostering these skills than a traditional lecture-based approach that may simply present established ethical frameworks without requiring active application.

2.4 Responsible Innovation and the Diversity Imperative

The concept of "responsible innovation" has emerged in science and technology studies as a framework for governing innovation in a way that is transparent, interactive, and responsive to societal needs and values. Its four dimensions—anticipation, reflexivity, inclusion, and responsiveness provide a robust operational framework for what it means to innovate ethically. Anticipation involves thinking through the potential consequences of a technology; reflexivity involves questioning one's own assumptions and values; inclusion involves bringing diverse voices into the innovation process; and responsiveness involves using that collective wisdom to adjust the course of development.

This framework is intrinsically linked to diversity. The "inclusion" dimension is not just a procedural check-box; it is an acknowledgment that innovation systems are more robust and ethical when they incorporate a wide range of perspectives. The historical lack of diversity in AI development has directly led to technologies that are biased and harmful to underrepresented groups. For example, facial recognition systems that are less accurate on darker-skinned faces are a direct result of a lack of

diversity in the teams and datasets used to build them (Buolamwini & Gebru, 2018). Therefore, educating a diverse generation of students in responsible innovation is a form of systemic intervention. It is not only about teaching students *about* ethics but also empowering them to become the ethical innovators of the future. A project-based approach that values diverse perspectives encourages students to draw on their own lived experiences, and positions them as experts on their own communities is uniquely suited to this task (Lee & Soep, 2016).

3. Methodology

3.1 Research Design

This study employed a mixed-methods, quasi-experimental design. The quantitative component used a pre-test/post-test, non-equivalent control group design to measure changes in critical thinking and responsible innovation competencies. The qualitative component was embedded to provide a deeper understanding of the student experience and the processes through which the PBL intervention influenced outcomes.

3.2 Participants and Setting

The study was conducted in four diverse, co-educational secondary schools in a large metropolitan area. The schools were selected to represent a diversity of student populations, including a high-performing magnet school, a Title I school with a high percentage of students from low-income families, and two comprehensive public schools with ethnically and socioeconomically diverse student bodies. The participants were 180 students (n=90 female, n=90 male) enrolled in a mandatory 10th-grade "Technology and Society" course. The students' self-reported ethnicities were: 33% White, 27% Hispanic, 25% Black, and 15% Asian/Other. Two intact classes from each school were randomly assigned to either the experimental (PBL) condition (n=90) or the control (traditional instruction) condition (n=90). Informed consent was obtained from parents/guardians and assent from students.

3.3 The Intervention: A Project-Based AI Ethics Curriculum

The intervention was a 10-week curriculum module, developed by the research team in consultation with AI ethics experts and experienced secondary school teachers. Students in the experimental group worked in small, diverse teams of 3-4 members.

Table 1: Structure of the 10-Week Project-Based AI Ethics Curriculum

Phase Week(s) Core Activities Student Deliverable(s)			
:--- :--- :--- :---			
Phase 1: Immersion & Inquiry** 1-2 Introduction to core AI concepts (machine learning, algorithms); exploration of real-world AI systems; introduction to basic ethical frameworks (e.g., fairness, accountability, transparency). Team formation; initial "AI in My Life" journal; generation of a list of AI systems that interest them.			
Phase 2: Defining the Problem 3-4 Teams select a specific AI application relevant to their community (e.g., an AI tutor, a local crime-prediction tool, a social media algorithm). They conduct preliminary research on its purpose, function, and potential impacts. A "Problem Definition" poster identifying the AI system, its stakeholders, and initial questions about its fairness and impact.			

| Phase 3: Deep Dive & Analysis | 5-7 | Teams conduct a structured ethical analysis. They interview community members, research bias and algorithmic discrimination, and use a "consequence wheel" to map potential positive and negative outcomes. They receive feedback from peers and teachers. | A detailed written "Ethical Impact Assessment" of their chosen AI system. |

|Phase 4: Innovation & Proposal | 8-9 | Teams design their own proposed solution. This could be a redesigned algorithm (conceptual), a new governance policy, a public awareness campaign, or a prototype for a more ethical AI application. | A "Responsible Innovation Proposal" document and a 3-minute "pitch" presentation. |

|Phase 5: Showcase & Reflection | 10 | Teams present their proposals at a school-wide "AI Ethics Showcase" attended by peers, teachers, and community members. They engage in a structured reflection on their learning process. | Final pitch presentation; individual reflection essay on their personal growth. |

In contrast, the control group received instruction on the *same* core AI concepts and ethical frameworks, but through a traditional, teacher-centered approach. This included lectures, textbook readings, and worksheets. They also discussed the same case studies but did not engage in the extended, collaborative investigation or create a culminating public product.

3.4 Instruments

3.4.1 Cornell Critical Thinking Test (CCTT), Level Z: This 52-item, multiple-choice test is a widely used and validated measure of critical thinking for high school and college populations. It assesses skills in induction, deduction, credibility, observation, and assumption identification. The internal consistency (Kuder-Richardson Formula 20) for this sample was 0.85.

3.4.2 Responsible Innovation Assessment Scale (RIAS): This 16-item self-report scale was developed by the researchers for this study, grounded in the four dimensions of Stilgoe et al.'s (2013) framework: Anticipation, Reflexivity, Inclusion, and Responsiveness. Students rated their agreement with statements on a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree). Example items include: "I can think about the possible future consequences of a new technology" (Anticipation) and "I believe it is important to include many different people's perspectives when creating a new technology" (Inclusion). The scale demonstrated good reliability, with a Cronbach's alpha of 0.88.

3.4.3 Qualitative Data Collection: Qualitative data were collected from three sources within the experimental group: (1) The student-generated artifacts (Problem Definition posters, Ethical Impact Assessments, Responsible Innovation Proposals), (2) Semi-structured focus group interviews with 24 students (six groups of four, selected to maximize diversity in terms of gender, ethnicity, and prior achievement), and (3) Teacher observation journals.

3.5 Procedure

In week 0, all participants completed the CCTT and RIAS as pre-tests. The 10-week intervention was then implemented from weeks 1-10, with students in the experimental and control

conditions engaging in their respective curricula. In week 11, all participants completed the CCTT and RIAS as post-tests. Focus group interviews were conducted in week 12. All teachers were provided with a 10-hour professional development session on their respective curriculum prior to implementation to ensure fidelity.

3.6 Data Analysis

Quantitative data were analyzed using SPSS version 27. Descriptive statistics were calculated. A one-way Analysis of Covariance (ANCOVA) was used to test H_{01} and H_{02} , with pre-test scores as covariates and post-test scores as the dependent variables. A two-way ANCOVA was used to test H_{03} , analyzing the interaction effects between the intervention and demographic variables. Effect sizes were calculated using partial eta squared (η^2). Qualitative data were analyzed using thematic analysis, following the six-phase framework of Braun and Clarke (2006). All transcripts and artifacts were coded, and themes were developed inductively.

4. Results

4.1 Quantitative Findings

Preliminary analysis confirmed no significant differences between the experimental and control groups on the CCTT pre-test, $t(178) = 0.85$, $p = .40$, or the RIAS pre-test, $t(178) = 1.10$, $p = .27$, indicating that the groups were comparable at baseline.

Table 2: Mean Scores and Standard Deviations for CCTT and RIAS

Group	N	CCTT Pre-Test M (SD)	CCTT Post-Test M (SD)	RIAS Pre-Test M (SD)	RIAS Post-Test M (SD)
Experimental (PBL)	90	24.50 (5.80)	31.20 (6.10)	3.30 (0.60)	4.25 (0.55)
Control (Traditional)	90	25.10 (5.50)	27.50 (5.90)	3.35 (0.65)	3.65 (0.60)

Hypothesis 1 (H_{01}): The ANCOVA revealed a significant effect of the instructional method on post-test CCTT scores, after controlling for pre-test scores, $F(1, 177) = 21.46$, $p < .001$, $\eta^2 = .11$. The experimental group ($M = 31.20$) significantly outperformed the control group ($M = 27.50$) on the measure of critical thinking. The null hypothesis was rejected.

Hypothesis 2 (H_{02}): Similarly, the ANCOVA showed a significant effect for the instructional method on post-test RIAS scores, $F(1, 177) = 55.78$, $p < .001$, $\eta^2 = .24$. The PBL group demonstrated a significantly greater gain in responsible innovation competencies ($M = 4.25$) compared to the control group ($M = 3.65$). The null hypothesis was rejected. The effect size ($\eta^2 = .24$) is considered large.

Hypothesis 3 (H_{03}): The two-way ANCOVAs examined the interaction effects between the instructional method and student gender and socioeconomic status (SES). The results showed ****no significant interaction effects**** for either demographic variable on post-test CCTT scores (Gender: $F(1, 173) = 0.85$, $p = .36$; SES: $F(1, 173) = 1.20$, $p = .27$) or post-test RIAS scores (Gender: $F(1, 173) = 0.42$, $p = .52$; SES: $F(1, 173) = 0.95$, $p = .33$). This is a critical finding: it indicates that the positive impact of the PBL intervention was equitable and consistent across all demographic groups. The intervention did not favor one group over another; it raised achievement and competency for everyone.

4.2 Qualitative Findings

The thematic analysis of the qualitative data revealed four major themes: *(1) From Abstract to Authentic, " (2) The Power of Perspective, " (3) Agency and Voice, " and *(4) The Messy, but Meaningful, Process. " **

Theme 1: From Abstract to Authentic

Students repeatedly emphasized how the project made abstract concepts like "bias" and "algorithm" feel real and relevant. By choosing an AI system from their own lives—like a social media filter, a music recommender, or a school assignment grading system—the ethical questions became concrete. One student (P8) noted, "Before, ethics was just a chapter in a book. It was like, 'don't plagiarize.' But now... we were looking at the Snapchat filters and asking, 'why do some of them make people look lighter-skinned? Is that a bias?' It was our life, not a textbook." Another (P15) reflected on the process, saying, "The AI didn't feel like some robot in a lab. It was in my phone, it was affecting my friends. That made me actually care about making it fair." * This authenticity was a key driver of engagement and deep learning.

Theme 2: The Power of Perspective

The collaborative nature of the PBL projects, combined with the instruction on the "inclusion" dimension of responsible innovation, profoundly impacted students' thinking. They learned that a technology could look "fair" to them but be deeply harmful to someone else. A student (P22) from a low-income background, working on a predictive policing AI, shared in a focus group: "My group... they lived in a different part of town. They had never been stopped by the police for no reason. When I told them my experience, it changed the whole project. They realized the 'data' wasn't just numbers; it was people like me." * This powerful statement illustrates how the PBL structure, by valuing diverse student experiences, directly operationalized the "inclusion" principle of responsible innovation. Another student (P3) echoed this, stating, "I used to think a good program was just one that worked fast. Now I think, 'fast for who?' and 'who gets left out?'"

Theme 3: Agency and Voice

The culmination of the project—the public showcase—was identified as a highlight by many students. The opportunity to present their work to an authentic audience of adults and peers was deeply motivating and fostered a strong sense of agency. They were not just students doing an assignment; they were innovators pitching a solution. One student (P11) said, "We actually felt like we were doing something important. We were telling adults, 'Here is a problem and here is how to fix it.' We weren't just learning about a problem, we were trying to solve it." * Another (P5), whose team proposed a policy for ethical AI in school admissions, commented, "It made me realize my voice matters. I can have an opinion on these big tech things, and I should. It's my future." * This sense of agency is a powerful outcome that extends far beyond the specific content knowledge gained.

Theme 4: The Messy, but Meaningful, Process

While the outcomes were positive, the journey was not always easy. Students and teachers both noted the challenges of the PBL approach. Students had to learn to manage their time, navigate group dynamics, and deal with the frustration of not having a single, clear answer. One student (P17) admitted, "It was hard. Sometimes I just wanted the teacher to tell me the right answer. But she wouldn't. She'd just ask more questions." A teacher's

journal entry reflected this, stating, "The 'messy middle' of a project is always a challenge. The students are used to step-by-step instructions. But pushing them to sit with the ambiguity is where the real critical thinking happens." This theme highlights a crucial pedagogical insight: that the productive struggle inherent in PBL is not a flaw but a feature, as it is precisely this struggle that builds the cognitive muscles required for critical and ethical thinking.

5. Discussion

The findings of this study provide strong evidence that a project-based learning approach to AI ethics education is not only effective but also equitable, significantly enhancing both critical thinking skills and responsible innovation competencies among a diverse population of secondary school students. The results contribute to a growing body of literature that calls for a fundamental rethinking of technology education.

5.1 Fostering Critical Thinking Through Authentic Inquiry

The significant gains in critical thinking scores ($\eta^2 = .11$) for the PBL group confirm the pedagogical power of this approach. This finding aligns with decades of research on PBL, which consistently demonstrates its superiority over traditional instruction for fostering higher-order thinking skills (Condliffe et al., 2017). The qualitative theme "From Abstract to Authentic" provides insight into *why* this is the case. Critical thinking is not a set of decontextualized skills to be learned in isolation. It is a practice that is honed through application to meaningful problems. By rooting the ethical dilemmas in AI applications that students encountered daily, the PBL curriculum provided the necessary "grist for the cognitive mill". Students were not just learning about argumentation; they were actively constructing arguments, evaluating the claims of their peers and community members, and synthesizing complex information into a coherent proposal. This process is a direct, intensive workout for the critical thinking faculties.

The theme "The Messy, but Meaningful, Process" is also crucial here. The frustration and ambiguity students faced are not signs of failure but evidence of productive struggle. Traditional instruction often seeks to eliminate this struggle by providing clear answers. However, the very act of grappling with an ill-structured problem where there is no perfect solution is what forces students to move beyond simple recall and engage in the higher-order cognitive processes of analysis, evaluation, and creation. The teacher's role shifts from knowledge dispenser to facilitator of inquiry, a shift that is essential for cultivating independent, critical thinkers.

5.2 Cultivating Responsible Innovators

The most robust finding of this study was the large effect size ($\eta^2 = .24$) for the impact of the PBL curriculum on responsible innovation competencies, as measured by the RIAS. This was a central goal of the intervention, and the results are encouraging. The curriculum was explicitly designed around, framework of Anticipation, Reflexivity, Inclusion, and Responsiveness, and the PBL format allowed these principles to be actively practiced. Students were not just told about the importance of inclusion; they experienced it through the "Power of Perspective," as described in the qualitative data. When a student's lived experience challenged their peers' assumptions, it was a powerful, non-abstract lesson in reflexivity. The process of interviewing community members and conducting a formal "Ethical Impact Assessment" forced them to practice anticipation and responsiveness.

The culmination of the project in the public showcase was instrumental in developing "Agency and Voice." By requiring students to propose a solution—be it a policy, a campaign, or a redesigned algorithm—the curriculum positioned them as active problem-solvers and innovators. This is a profound departure from the typical consumer-based model of technology education. It transforms students from passive recipients of technology to potential stewards of it. This sense of agency is particularly important for students from underrepresented groups, who may not see themselves as having a stake or a voice in the world of technology. By engaging those in creating their own ethical frameworks and solutions, the intervention subtly but powerfully challenged the dominant narrative of who gets to be an innovator.

5.3 Equity: A Central Finding

The finding that the intervention's positive effects were consistent across gender and socioeconomic status lines (H_0 was not rejected) is arguably the most important result of this study. It demonstrates that a well-designed, culturally-relevant PBL curriculum does not inadvertently privilege students who already possess higher levels of cultural or academic capital. Often, progressive pedagogies can inadvertently widen achievement gaps if they fail to provide sufficient scaffolding for students who are less prepared for self-directed learning.

Why was this intervention successful on this front? The qualitative data offers some clues. The focus on local, community-relevant issues likely resonated with all students, providing a level playing field of "expertise" based on lived experience rather than purely academic knowledge. A student from a low-income neighborhood may have struggled with a hypothetical case study about autonomous vehicles, but they were the expert in the room when the discussion turned to predictive policing or biased social media algorithms. The PBL structure, by explicitly validating and drawing on these diverse funds of knowledge (González et al., 2005), may have helped to mitigate potential advantages held by more privileged students. This is a critical design principle for any future educational technology or AI ethics curriculum aiming to be equitable.

5.4 Implications for Policy and Practice

The findings have direct and actionable implications. First, **teacher training is paramount.** The shift from a lecturer to a facilitator of PBL is significant and requires dedicated professional development. Teachers need to be comfortable with ambiguity, skilled in guiding student inquiry, and adept at facilitating difficult discussions about race, bias, and social justice. Second, **curriculum should be integrated, not additive.** The AI ethics curriculum was taught within an existing "Technology and Society" course. Rather than creating a new, separate "AI Ethics" class, the concepts can be woven into existing science, social studies, and technology courses, making them more relevant to the broader curriculum. Third, **assessment must evolve.** Standardized tests are poorly suited to measure the complex, multidimensional outcomes of PBL. Schools and policymakers need to invest in developing and validating performance-based assessments, like the RIAS used in this study that can capture students' growth in critical thinking and responsible innovation. Finally, the study provides a compelling argument for **investment in technology education that is human-centric.** The goal is not just to create coders, but to create citizens who can think critically and act ethically in a technology-saturated world.

5.5 Limitations and Future Research

This study has several limitations. First, the quasi-experimental design, while robust, does not allow for causal claims as strong as a true randomized controlled trial. Second, the study was limited to four schools in a single metropolitan area, which may restrict the generalizability of the findings. Future research should replicate this study in diverse geographic and cultural contexts, including international settings. Third, the intervention was relatively short (10 weeks). A longer-term study could investigate whether the gains in critical thinking and responsible innovation are sustained over time and whether they influence students' later academic and career choices. Fourth, the RIAS was a researcher-developed scale, and while its initial psychometric properties were promising, it requires further validation. Future research should focus on refining tools to measure these complex competencies.

Furthermore, future studies could explore specific elements of the PBL design to understand their differential impacts. For example, is the "community interview" component essential, or would a case-study approach be equally effective? Does the "public showcase" significantly increase student agency, or could a smaller-scale presentation have the same effect? Answering these questions will help educators refine and optimize these pedagogical strategies.

6. Conclusion

As Artificial Intelligence becomes an increasingly powerful and invisible force in our world, the need for a citizenry that is both AI-literate and ethically astute has never been greater. This study demonstrates that a project-based learning approach to AI ethics education offers a powerful and equitable pathway to meet this challenge. By engaging diverse students in authentic, collaborative investigations of real-world AI dilemmas, it significantly enhances their critical thinking skills and cultivates the competencies needed for responsible technological innovation. The students in this study did not just learn about AI; they began to see themselves as its critical analysts and ethical stewards. Their experiences—from the productive struggle of the "messy middle" to the empowering act of pitching their solutions to a public audience illustrate the profound potential of this pedagogical model. The future of technology is not predetermined; it is something we build. An education that empowers all our students to build it with foresight, compassion, and critical intelligence is the most powerful investment we can make.

References

1. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101.
2. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, PMLR 81:77-91.
3. Calabrese Barton, A., & Tan, E. (2010). We be burnin'! Agency, identity, and science learning. *Journal of the Learning Sciences*, 19(2), 187-229.
4. Condliffe, B., Quint, J., Visher, M. G., Bangser, M. R., Drohojowska, S., Saco, L., & Nelson, E. (2017). Project-based learning: A literature review. MDRC.

5. Facione, P. A. (1990). Critical thinking: A statement of expert consensus for purposes of educational assessment and instruction (The Delphi Report). The California Academic Press.
6. González, N., Moll, L. C., & Amanti, C. (Eds.). (2005). Funds of knowledge: Theorizing practices in households, communities, and classrooms. Lawrence Erlbaum Associates.
7. Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. Educational Psychologist, 41(2), 75-86.
8. Krajcik, J. S., & Blumenfeld, P. C. (2006). Project-based learning. In R. K. Sawyer (Ed.), The Cambridge handbook of the learning sciences (pp. 317-334). Cambridge University Press.
9. Ladson-Billings, G. (1995). Toward a theory of culturally relevant pedagogy. American Educational Research Journal, 32(3), 465-491.