Responsible Use Cards for Students Evaluating AI Technologies

Charlotte Dungan, Mark Cuban Foundation

Secondary Author: Harshvardhan Upadhyaya

Charlotte Dungan is the Chief Learning Officer of the Mark Cuban Foundation, which runs AI Bootcamps for high school students, teachers, and in partnership with informal educators. Charlotte is a leading expert in AI education, having worked with teams at TeachAI, Code.org, AI for Teachers, CS is Elementary, aiEDU, ISTE, MIT, and AI4K12. After 10 years working as a computer programmer, she earned a Master's Degree in Education from Harvard University. She has helped several states revise CS standards and add AI to their curricula. Charlotte supports expanding CS education in rural areas and underserved populations. She's also a lifelong Girl Scout and a first-generation college student.

Responsible Use Cards for Students Evaluating AI Technologies

Introduction, Problem Statement, and Need

Rapid AI integration has transformed K-12 education

The pace of AI adoption in K-12 schools has accelerated dramatically since 2023. In 2025, Gallup-Walton Family Foundation surveys show 60% of teachers have used AI at least once, with 32% using AI weekly and saving an average of 5.9 hours per week on lesson preparation (Walton Family Foundation).

District-level adoption has doubled rapidly, with 48% of districts providing AI training to teachers by Fall 2024, up from just 23% the previous year (Rand). Earlier projections have suggested that 74% of districts will have offered AI training to their educators by Fall 2025 (RAND). However, significant equity gaps persist, with 67% of low-poverty districts providing AI training compared to only 39% of high-poverty districts (rand) highlighting disparities that mirror broader digital inequities in CS education.

Student exposure to AI has grown correspondingly. The percentage of high school students who report using GenAI tools for schoolwork is growing, increasing to 84% in May 2025 (CollegeBoard). Yet despite this widespread exposure, only 35% of school districts provide AI training for students, and 19% of students reported receiving guidance from teachers on how to use AI for schoolwork (RAND). By 2026, only 30 states have established AI education frameworks or guidance documents and only 19% of teachers work in schools with formal AI policies (The 74 Million), leaving the majority to navigate ethical considerations without institutional guidance.

Computer Science education has evolved with increasing ethical focus

The last decade has witnessed a fundamental transformation in K-12 computer science education, with ethics emerging as a central concern. The CSTA K-12 Computer Science Standards, first published in 2011 and revised in 2017, are currently undergoing a comprehensive revision, scheduled for completion in the summer of 2026 (Csta) The Reimagining CS Pathways project, a collaboration between CSTA, IACE, ACM, Code.org, College Board, CSforALL, and ECEP Alliance, has identified Social Impacts and Ethics as the highest priority for the revised standards, being elevated to a "Pillar" to stress intentional integration throughout all CS learning (Csteachers).

State-level policy adoption has accelerated significantly. By 2024, 43 states had adopted K-12 CS standards (up from just 6 in 2017), 44 states established CS teacher certification, and 11 states implemented CS graduation requirements (CSTA). The percentage of U.S. public high schools offering foundational CS courses grew 60% in 2023-24 (CSTA), with 94% of Public High Schools (in 2023-24) offering foundational computer science courses (CSTA). Additionally, California's passing of AB 2876 in September 2024, mandating AI literacy integration into curriculum frameworks, signals growing recognition that AI ethics must be systematically addressed.(CODE.org)

Educator perspectives and needs

Recent CSTA surveys provide critical insight into educator perspectives on AI ethics education. According to a Spring 2024 survey conducted with TeachAI found that 80% of CS teachers believe core concepts in CS education should be updated to emphasize AI (CSTA) and 85% believe foundational CS experiences should include learning about AI (CSTA). Notably, 70% of CS teachers already teach AI in their classrooms, despite the lack of explicit AI standards (CSTA), which clearly demonstrates the recognition of AI's importance by educators.

However, confidence gaps emerge sharply in the data. Only 42% of teachers feel equipped to teach about AI (CSTA), while 88% indicate needing additional resources and professional learning (CSTA). With a growing emphasis on the need for ethical perspectives in AI education, The CSTA and AI4K12 partnership's 2025 publication "AI Learning Priorities for All K-12 Students" identified five content categories organized by grade band, with "Ethical AI System Design and Programming" and "Societal Impacts of AI" as essential learning priorities (CSTA). Additionally, a 2025 Digital Education study by Tripathi et al., identified "Ethical AI Engagement" as the most complex developmental area for teachers, noting they "grapple with responsible use principles while recognizing the need for structured guidance" (Frontiers).

A significant gap exists between abstract AI ethics principles and practical K-12 implementation

Multiple systematic reviews confirm a substantial theory-practice gap in AI ethics education. A 2025 systematic review analyzing 68 peer-reviewed publications found that "AI ethics education remains significantly underprioritized in classroom practice, despite the global push for AI literacy in K-12 curricula" (ScienceDirect). This might be in large part due to the documented challenges teachers face in translating ethical frameworks to classroom practice. A systematic literature review by Gouseti et al., identified that students were “aware of the ethical issues and dangers associated with AI, but did not necessarily recognise or articulate these as ethical problems" (Taylor & Francis Online). Additionally, teachers reported feeling "disempowered" when curricula are developed by outside agencies without their input as co-designers (Taylor & Francis Online).

The consequences of this gap extend beyond missed learning opportunities. Akgun and Greenhow (2022) identified the biggest risks of AI integration in K-12 as, (1) perpetuating existing systemic bias and discrimination, (2) perpetuating unfairness for disadvantaged and marginalized students, and (3) amplifying racism, sexism, xenophobia, and other injustice (nihPubMed Central). Without practical tools for addressing ethics, these risks remain inadequately addressed despite growing AI integration.

Theoretical Framework and Literature Review

1. UNESCO's Recommendation on the Ethics of Artificial Intelligence

The UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted unanimously by all **193 Member States** on November 23, 2021, represents a global normative standard on AI ethics.

The Recommendation comprises of **four core values** (see figure 1.a.) – (1) Human Rights and Human Dignity (AI systems must respect, protect, and promote human rights throughout their lifecycle), (2) Living in Peaceful, Just, and Interconnected Societies (AI should contribute to peaceful societies respecting cultural diversity), (3) Ensuring Diversity and Inclusiveness (AI development must include diverse perspectives and ensure benefits reach all people), and (4) Environment and Ecosystem Flourishing (AI systems should support environmental sustainability).

These four values are further broken into **ten core principles** (see figure 1.b.) – (1) Proportionality and Do No Harm, (2) Safety and Security, (3) Right to Privacy and Data Protection, (4) Multi-stakeholder and Adaptive Governance, (5) Responsibility and Accountability, (6) Transparency and Explainability, (7) Human Oversight and Determination, (8) Sustainability, (9) Awareness and Literacy, and (10) Fairness and Non-Discrimination.

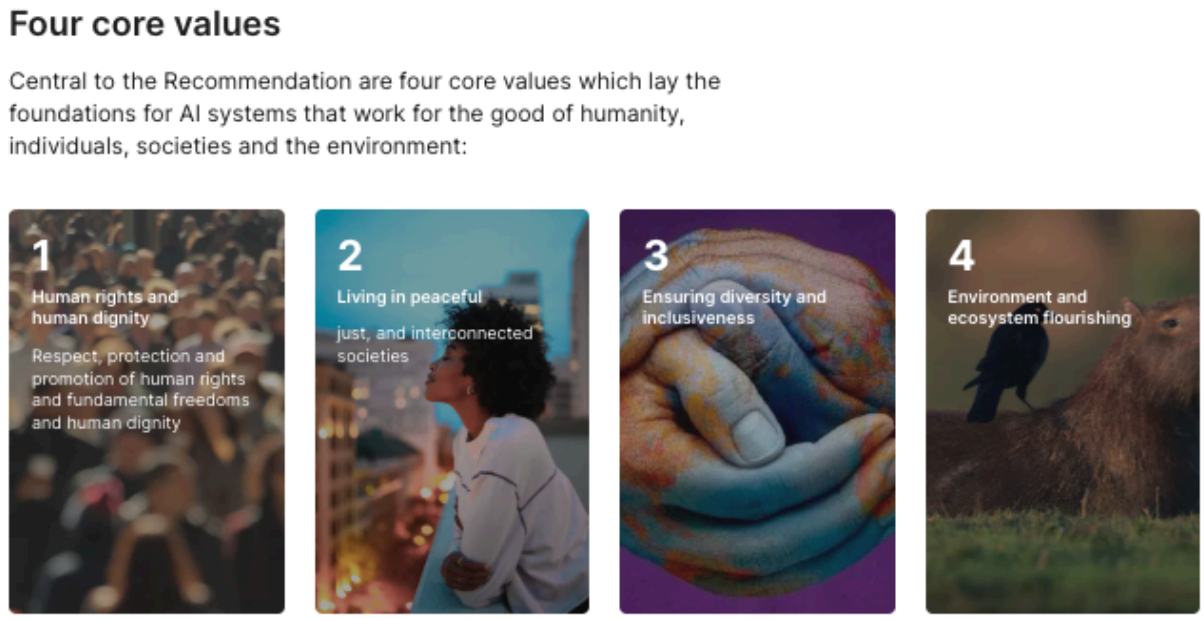

Figure 1.a.

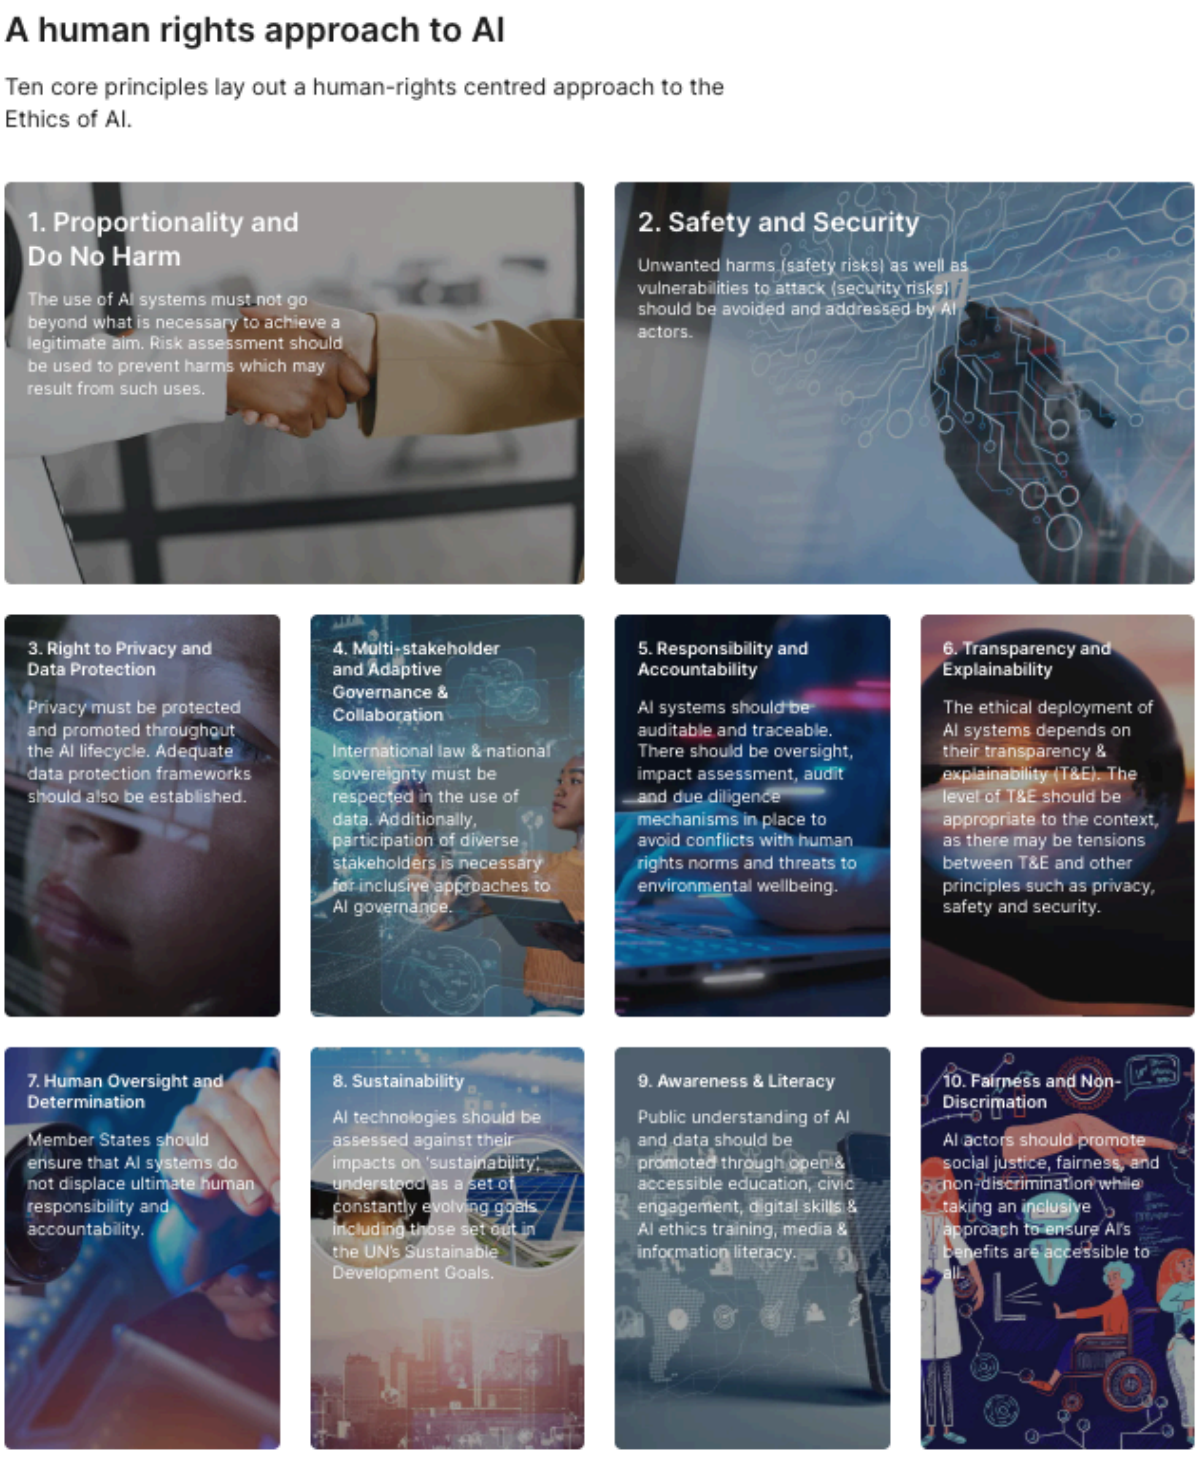

Figure 1.b.

2. The European Commission Ethics Guidelines for Trustworthy Artificial Intelligence

The European Commission's High-Level Expert Group (AI HLEG) published the Ethics Guidelines for Trustworthy Artificial Intelligence in April 2019, establishing that Trustworthy AI must be **(1) lawful, (2) ethical, and (3) robust**. The framework (see figure 2.) identifies **four ethical principles** – (1) Respect for Human Autonomy, (2) Prevention of Harm, (3) Fairness, and (4) Explicability. These are then translated into **seven key requirements** – (1) Human Agency and Oversight, (2) Technical Robustness and Safety, (3) Privacy and Data Governance, (4) Transparency, (5) Diversity, Non-Discrimination, and Fairness, (6) Societal and Environmental Well-being, and (7) Accountability. The EU AI Act (2024) explicitly references these principles in Recital 27, demonstrating the foundational role of the framework in global AI regulation and policy development.

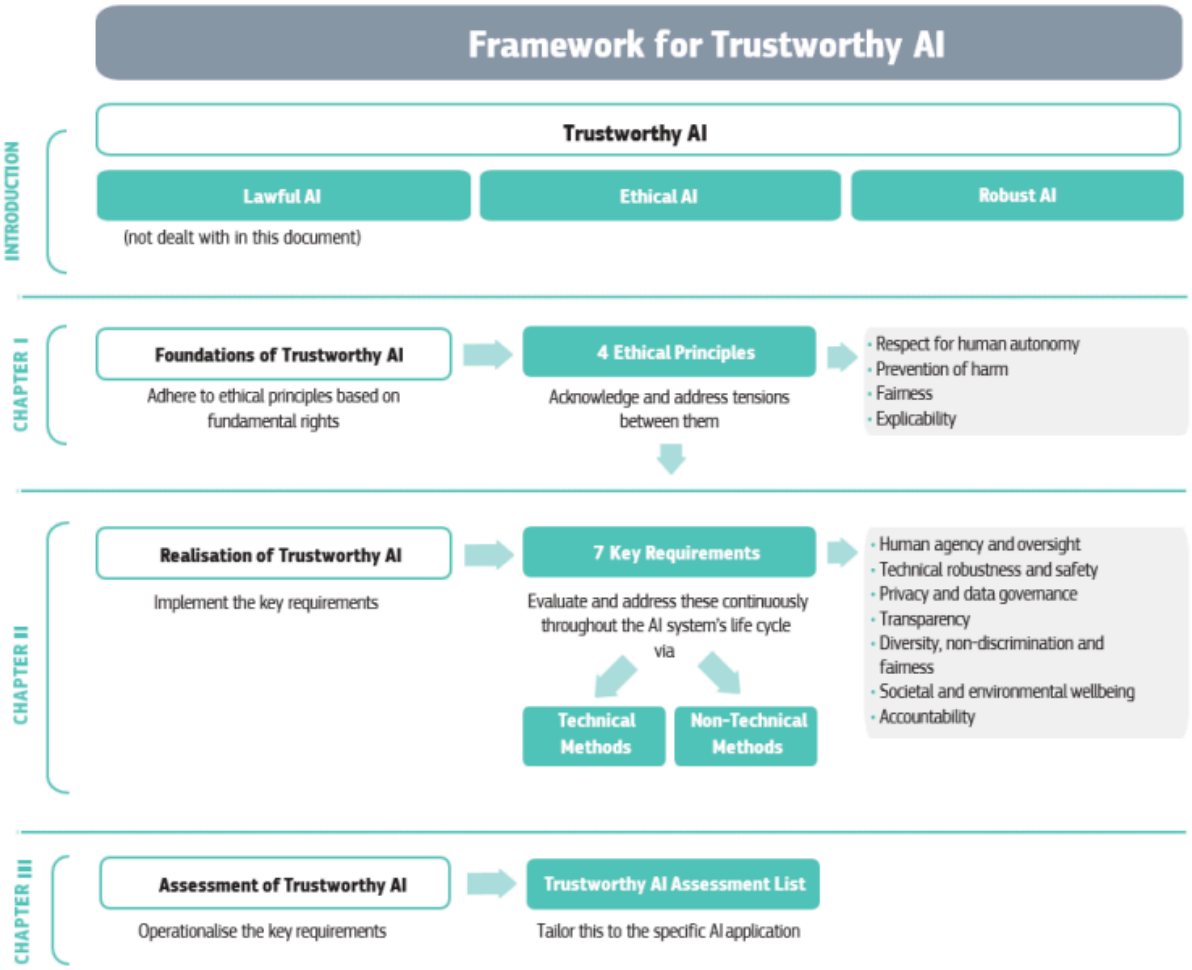

Figure 2.

3. AI4People—An Ethical Framework for a Good AI Society

The AI4People framework, represents one of the most influential contributions to contemporary AI ethics, with over 2,500 citations (on Springer) and direct influence on EU AI Act

development (AI4People/About). Rather than generating new ethical principles, AI4People synthesized existing frameworks from multiple stakeholder organizations to identify **five core principles** (see figure 3.) – (1) beneficence (promoting well-being and human dignity), (2) non-maleficence (ensuring privacy, security, and capability caution), (3) autonomy (preserving human decision-making power), (4) justice (promoting prosperity and solidarity), and (5) explicability (ensuring AI decisions can be understood and challenged). In addition to the framework, AI4People offers 20 specific recommendations for assessing, developing, incentivizing, and supporting ethical AI development in a systematic attempt to move "from principles to proposed policies, best practices, and concrete recommendations." (AI4People Framework)

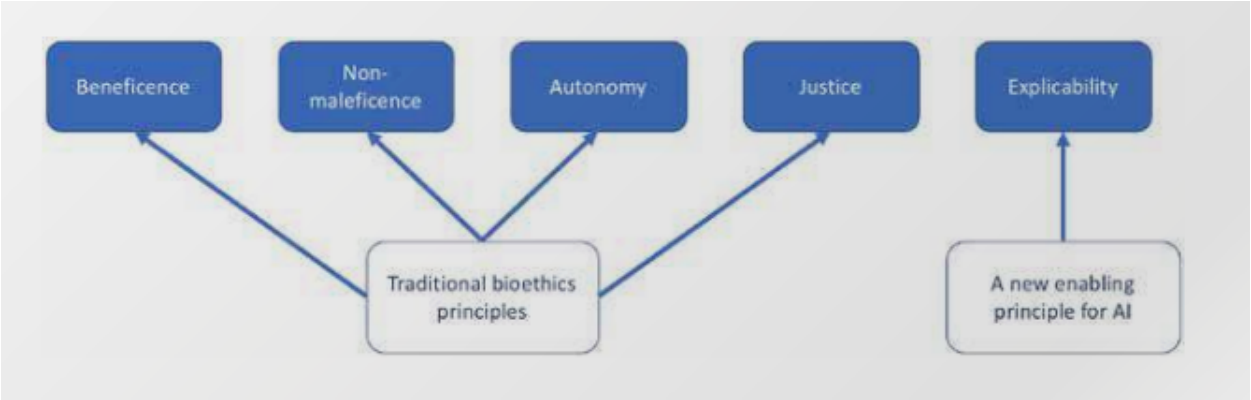

Figure 3.

4. **IBM's Five Pillars Framework**

IBM Research developed the Five Pillars of Trustworthy AI in an attempt to operationalize ethics in organizational settings, emerging from work initiated in 2015-2016 when Francesca Rossi joined IBM Research to focus on AI ethics (World Economic Forum). The framework comprises of **five core pillars** (see figure 4.) – (1) **Explainability** (AI systems should explain how and why they arrived at conclusions), (2) **Fairness** (properly calibrated AI can assist humans in making fairer choices, countering human biases), (3) **Robustness** (AI systems must be defended from adversarial attacks and handle exceptional conditions), (4) **Transparency** (users must be able to see how services work, evaluate functionality, and comprehend limitations), and (5) **Privacy** (systems must prioritize and safeguard data rights, collecting only minimum necessary data). (IBM Foundations of Trustworthy AI)

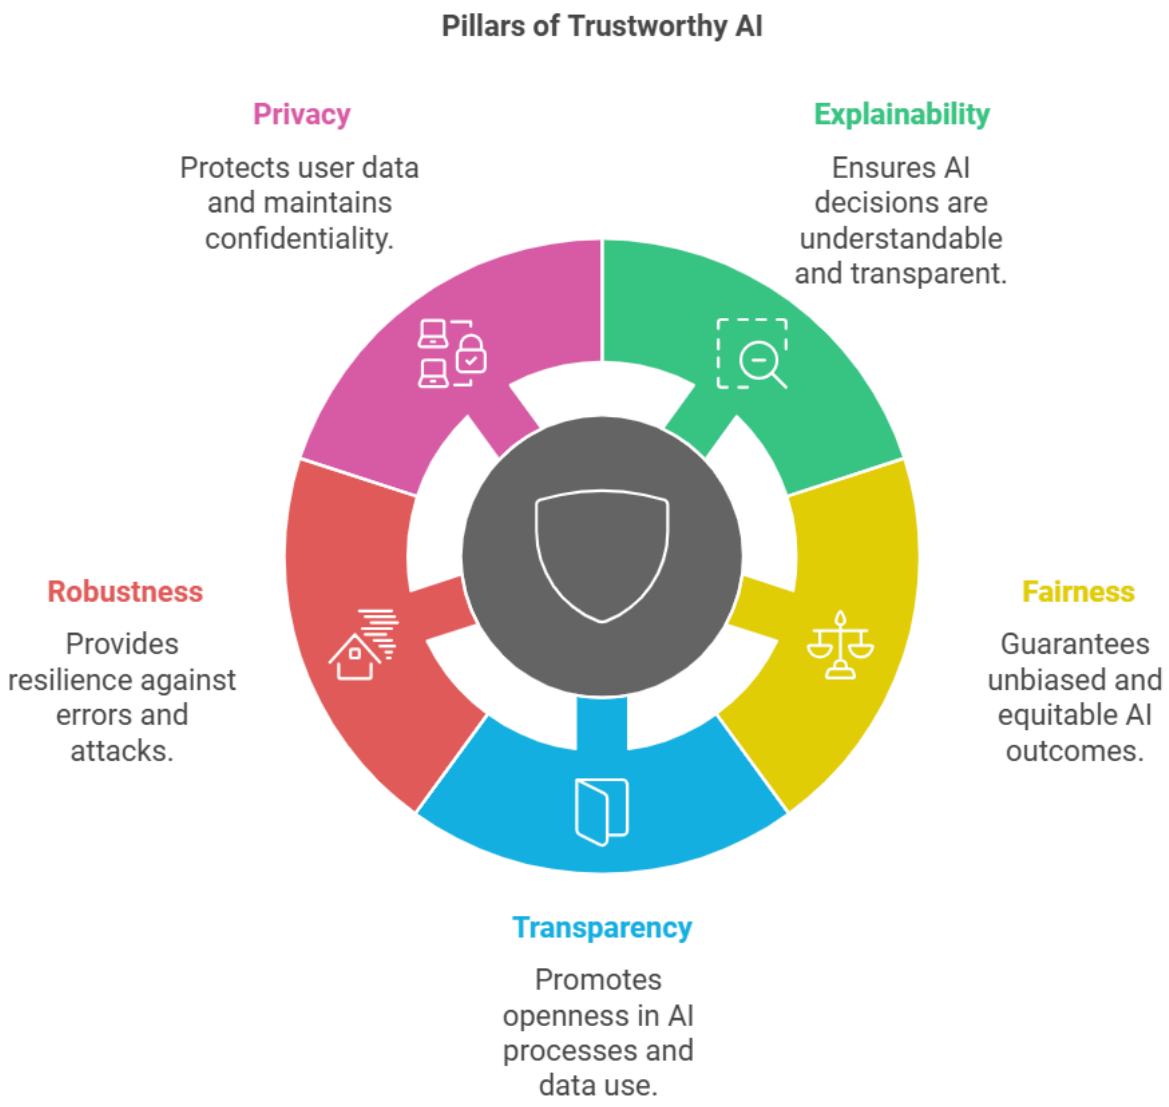

Figure 4.

Model Cards

Model Cards were introduced by Mitchell et al. (2019) in "Model Cards for Model Reporting" at the ACM Conference on Fairness, Accountability, and Transparency (2019). The paper describes Model Cards as "short documents accompanying trained machine learning models that provide benchmarked evaluation in a variety of conditions," similar to a nutrition label for AI models. Model Cards include relevant information including intended use cases, usage minimization in unsuitable contexts, evaluation metrics, increased transparency, and to support the "responsible democratization of machine learning" (Mitchell et al. 2019).

The Model Card structure comprises **nine sections** (see figure 5.) – (1) Model Details (basic information including type, architecture, license), (2) Intended Use (primary intended uses and out-of-scope applications), (3) Factors (relevant demographic, cultural, or phenotypic groups), (4) Metrics (performance measures and decision thresholds), (5) Evaluation Data, (6) Training Data, (7) Quantitative Analyses (disaggregated performance results), (8) Ethical Considerations, and (9) Caveats and Recommendations. The paper has received over 2,000 citations (on Semantic Scholar) and although the naming convention across organizations differ, it has influenced industry adoption at organizations like NVIDIA, Google, and Hugging Face.

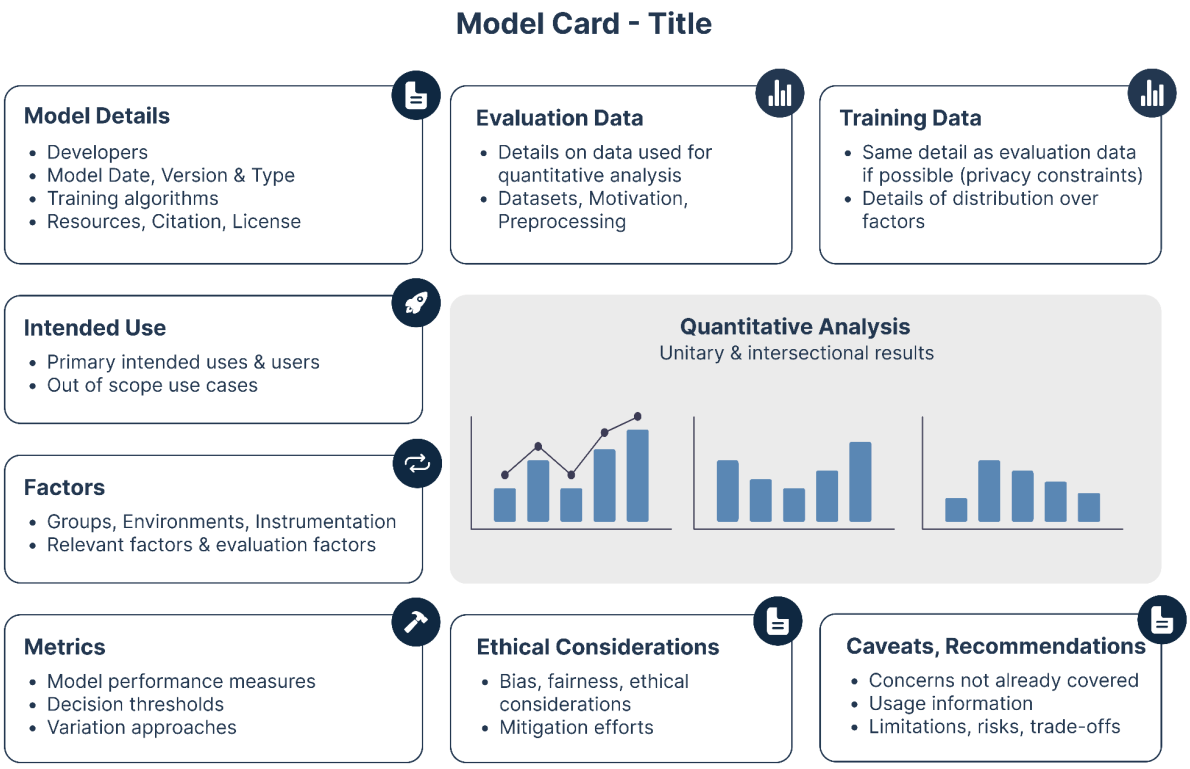

Figure 5.

Responsible Use Cards as Extensions of Model Cards

Multiple comprehensive frameworks have emerged to guide ethical AI development, including as discussed earlier UNESCO's global Recommendation on AI Ethics (2021), the EU High-Level Expert Group's Ethics Guidelines for Trustworthy AI (2019), and the highly influential AI4People framework (Floridi et al., 2018), which synthesized existing ethical principles and directly informed the development of the EU AI Act. These frameworks establish important foundational principles. However, these are largely policy-oriented frameworks, which while comprehensive in scope and influential in governance contexts, can present challenges for K-12 implementation due to their abstractness, technical complexities, lack of pedagogically practical guidance that students can meaningfully apply to their own AI development projects.

IBM's Five Pillars framework on the other hand, offers a more pedagogically accessible approach for educational contexts. By combining the concrete documentation structure of Model Cards with the systematic ethical evaluation categories of IBM's framework, Responsible Use Cards transform both approaches into a practical pedagogical tool. Students can engage with the familiar 'nutrition label' concept of Model Cards with the five pillar framework providing concrete, actionable categories that students can directly evaluate in their own AI implementations, while maintaining comprehensive coverage of essential ethical considerations. Additionally, the Five Pillars are connected to industry practices that students will encounter in future AI careers, bridging educational and professional contexts.

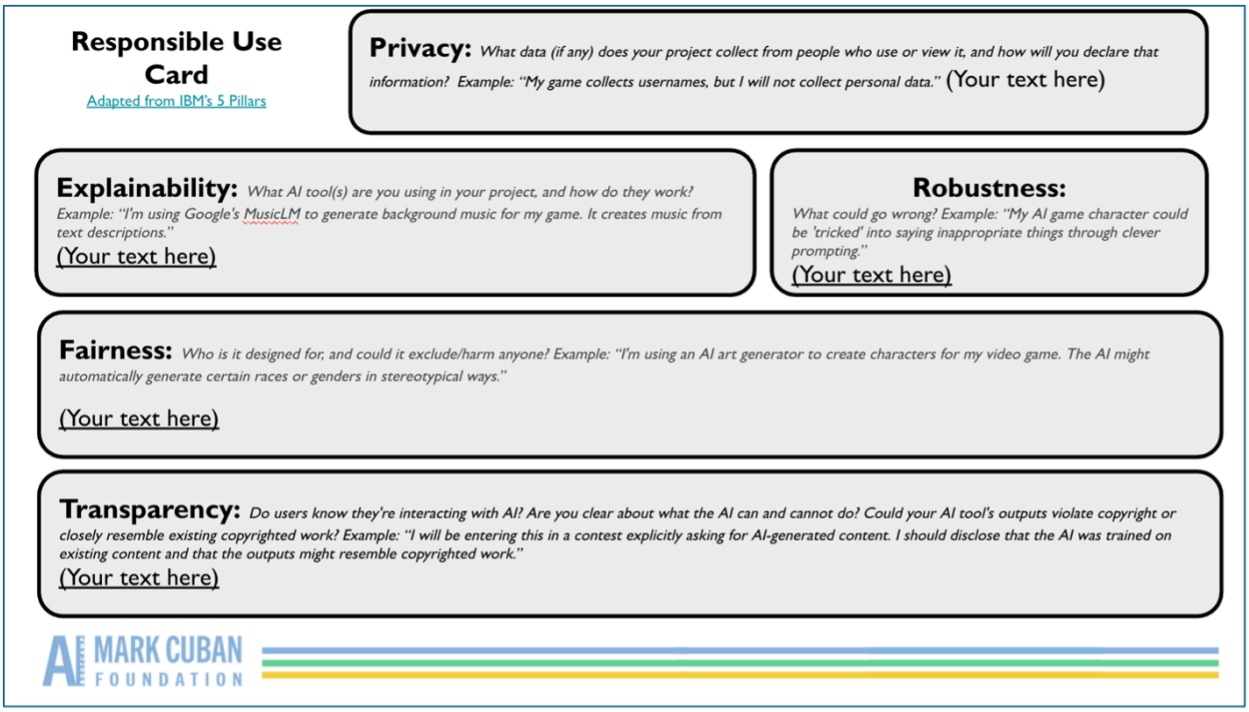

Figure 6.

While Model Cards traditionally serve as post-hoc documentation, Responsible Use Cards are designed to function as a proactive design framework. By introducing the Five Pillars during the ideation phase, students can use the card to anticipate ethical friction points and adjust their project's technical architecture—such as data collection methods or interface transparency—before implementation begins.

Compared to existing K-12 AI ethics resources, such as the MIT's DAILY curricula or common case-study-based discussions, Responsible Use Cards offer a more artifact-centric approach. While case studies focus on evaluating external technologies, this framework requires students to apply ethical rigor to their own code and design decisions, bridging the gap between theoretical awareness and technical accountability.

Setting and Implementation Design

Setting Context : Mark Cuban Foundation AI Bootcamps

The Responsible Use Cards framework was developed and implemented through the Mark Cuban Foundation's AI Student Bootcamps, a comprehensive three-day program designed to introduce high school students to AI technologies and their practical applications. These bootcamps serve students in grades 9-12 nationwide, with 2025 programs spanning 20 locations across the United States. The program explicitly welcomes students from low and underserved communities, with an ethics integrated curriculum designed to accommodate diverse technical backgrounds while maintaining rigorous learning outcomes.

The bootcamp structure follows an intensive format across three consecutive Saturdays, with each day building foundational AI concepts while progressing towards a launchpad project – foundations of a capstone project that students can refine and pursue on their own. Beginning with AI fundamentals including Large Language Models, prompting techniques, and machine learning concepts, then advancing through data analysis, ethics education, and image generation before introducing entrepreneurship and design thinking methodologies. Students engage with AI through **six specialized career track modules** – (1) Healthcare, (2) Sports Science, (3) Arts & Entertainment, (4) Business & Entrepreneurship, (5) Computer Science, and (6) Ecology & Food Security. This track-based approach allows students to explore AI applications within their areas of personal interest, increasing meaningful engagement and real-world relevance.

Activity Structure and Flow

The Responsible Use Cards activity is embedded within a comprehensive "Model Cards and Ethical Analysis" lesson that spans 30 minutes and serves as a required ethics evaluation for all launchpad projects. The activity is not a terminal assessment; rather, it is positioned as an **iterative design tool**. Students are encouraged to draft their cards alongside their project blueprints, allowing the ethical evaluation to inform the 'Launchpad' project's development in real-time.

Lesson Introduction: Setting the context by connecting AI ethics to students' lived experiences and future aspirations. The curriculum explicitly states that "the technology we create and use today affects real people's lives, from job decisions to healthcare," immediately establishing that ethical reasoning has real-world consequences.

Activity #1: Model Cards Foundation (5 minutes): Students examine Model Cards from Google, including Gemma, MediaPipe FaceMesh, and Object Detection models, to understand how professional developers document ethical considerations.

Activity #2: Project-Specific Ethics Documentation (25 minutes): Students create their own ethical documentation. Students select between two pathways based on their launchpad project characteristics. Responsible Use Cards for creative projects that don't rely on datasets, or traditional Model Cards for projects involving data training. This differentiation ensures that all students can meaningfully engage with ethical evaluation regardless of their project's technical complexity.

The Responsible Use Card Pathway guides students through IBM's Five Pillars framework using a provided template and example (see figure 6.). Students systematically evaluate their AI tool implementations across the 5 pillars – Privacy, Explainability, Robustness, Fairness, and Transparency dimensions, with structured prompts guiding their reflection. The collaborative process requires teams to examine ethical tradeoffs through guided questions.This discussion component ensures that students grapple with real-world constraints rather than viewing ethical implementation as straightforward.

Pedagogical Rationale

Research consistently demonstrates that 'constructivist' pedagogies prove most effective for AI ethics education (Taylor & Francis Online). A systematic review found that "by some considerable margin (10 of 15 studies), the adoption of constructivist pedagogies was identified as the most popular and effective strategy for helping students understand ethical challenges posed by AI"(Taylor & Francis Online). Placing students in the role of AI designers and developers "correlated with an increased understanding of the ethical challenges associated with these technologies and how best to approach them" (Taylor & Francis Online).

The theoretical foundation for Responsible Use Cards also draws from the 'constructionist' learning principles. Seymour Papert's constructionism posits that learning happens most effectively when learners actively construct meaningful artifacts that make abstract concepts tangible and shareable (Papert, 1980). This approach is particularly relevant for ethics education, where abstract principles often remain disconnected from practical contexts.

Additionally, Mitchell Resnick's Four P's of Creative Learning (Projects, Passion, Peers, and Play), provide additional pedagogical justification for the Responsible Use Cards framework (Resnick, 2017). Students engage in meaningful **Projects** by creating personalized cards that document their own AI implementations. **Passion** emerges from the personal relevance of evaluating their own creations rather than abstract case studies. **Peers** contribute through collaborative evaluation and discussion of different ethical perspectives. **Play** occurs through the iterative process of reflection and revision as students discover ethical considerations they hadn't initially recognized.

Lastly, the MIT Media Lab's AI + Ethics curriculum identified three essential design principles for effective resources – (1) **active learning** (where students engage in activities and process information through reflection), (2) **embedded ethics** (teaching technical and ethical concepts in tandem), and (3) **low barriers to access** (using unplugged activities and web-based tools that don't require expensive hardware) (MIT). By adhering to these essential design principles, Responsible Use Cards prove to be an innovative yet simple pedagogical approach for teaching responsible AI principles to students.

Initial Results

Participants were asked, "If you created a Responsible Use Card for your Launchpad Project, did it help you consider the ethical impacts of your project?" with the following response options: "I did not make a Responsible Use Card" or a 5-point Likert-type scale ranging from 1 ("Not at all helpful") to 5 ("Extremely helpful").

Initial survey data from Mark Cuban Foundation AI Bootcamp participants (n=577) reveals a positive reception of the Responsible Use Cards framework and activity. Among students who completed Responsible Use Cards, **74.3% reported finding them helpful to extremely helpful (ratings 3-5)**, with **48.7% specifically rating them as very helpful (31.0%) or extremely helpful (17.7%)**. Only 8.7% of participating students found the cards minimally helpful, with merely 1.2% reporting them as not helpful at all. Notably, 17.0% of survey respondents indicated they did not create a Responsible Use Card, suggesting the student chose the traditional Model Cards pathway for their data-intensive projects or did not complete the activity. The distribution of responses indicates that students found concrete value in the structured ethical evaluation process, with nearly half of participants rating the cards as highly beneficial for their project development. These preliminary results suggest that the framework successfully engages students in meaningful ethical reasoning while providing practical utility for their AI project development, though the self-reported nature of the data and lack of control group comparison limits definitive conclusions about learning outcomes.

Implications & Limitations

Rationale for Approach

The Responsible Use Cards framework addresses gaps in K-12 AI ethics education by providing a concrete tool that makes abstract principles actionable for student developers. Unlike traditional ethics education approaches that rely on theoretical discussions or external case studies, this framework engages students in evaluating their own creations, increasing personal relevance and motivation for deep ethical reasoning. The combination of individual reflection and collaborative discussion ensures that students encounter diverse ethical perspectives while developing personal responsibility for their technology design decisions.

By grounding the activity in IBM's Five Pillars framework and Model Cards documentation practices, students develop familiarity with professional ethical evaluation processes that will serve them in future academic and career contexts. This professional relevance distinguishes the framework from purely academic ethics exercises, providing students with transferable skills for responsible technology development.

Scalability Potential for Broader Audiences

Initial implementation suggests significant scalability potential across different educational contexts including higher education contexts. The template-based structure reduces implementation barriers for educators while maintaining flexibility for different project types and student populations. The framework requires minimal specialized training for facilitators while providing sufficient structure to support meaningful ethical reasoning among students with varying levels of technical expertise. Educators can incorporate Responsible Use Cards into existing AI projects, computer science courses, or maker education programs with minimal additional resources.

Potential applications extend beyond formal AI education to include widespread industry adoption. Tech developers who have already integrated Model Cards into their deployment workflows could adopt Responsible Use Cards as a complementary documentation practice, particularly for AI applications with significant user interaction or creative outputs where traditional model documentation might be insufficient.

Current Limitations and Next Steps for Research

Several limitations characterize this preliminary implementation and point toward necessary research directions. The intensive bootcamp format may not be representative of a typical classroom environment, potentially overestimating student engagement and learning outcomes in more constrained contexts. The self-selected nature of bootcamp participants (students who voluntarily engage in intensive AI education) may not reflect the broader student population's readiness for ethical reasoning activities.

The current evaluation relies primarily on data from one question in the bootcamp post survey. Future research requires development of validated instruments for measuring ethical reasoning development, technical understanding, and long-term retention of responsible AI principles. Pre- and post-assessments should examine both knowledge acquisition and behavioral changes in students' approach to technology design and evaluation.

Implementation scalability requires investigation across diverse educational contexts, including traditional classrooms, different age groups, and varying levels of technical infrastructure. Research should examine how the framework performs with different facilitation models, from expert instructors to peer-led implementation, and identify necessary adaptations for successful deployment in resource-constrained environments.

To move beyond this initial implementation, future research may employ performance-based assessments. This includes a study of students who completed projects to track 'ethical transfer'—evaluating if students who used the cards in the bootcamp apply similar ethical constraints to independent projects one year later. Additionally, we plan to create validated ethical reasoning rubrics to analyze the depth of student responses on the cards, measuring shifts from compliance to a deeper understanding of limitations and trade-offs.

Conclusion

The Responsible Use Cards framework demonstrates promise as a practical approach to bridging the documented gap between AI ethics principles and K-12 educational practice. By combining the documentation structure of Model Cards with the comprehensive ethical coverage of IBM's Five Pillars framework, this pedagogical tool makes abstract ethical principles tangible and actionable for students. Grounded in constructionist learning theory and implemented through collaborative project-based activities, the approach engages students in meaningful ethical reasoning about their own AI implementations.

While preliminary observations suggest positive outcomes in terms of student engagement and ethical reasoning development, significant research remains necessary to validate the approach's effectiveness and scalability. As K-12 students increasingly interact with AI technologies and

enter careers where ethical technology design matters, practical tools like Responsible Use Cards could prove to be essential for developing the ethical reasoning capabilities that responsible AI development requires.