Inclusive AI Education in Practice: Building the Skills for Responsible AI

Report 3 from the Future of AI Education Council, Curated by AI4ALL & Funded by Google.org

By: Ritika Puri Strauss and Nicole Halmi, in collaboration with the AI4ALL Future of AI Education Council

November 2026

1. Introduction
2. Finding Direction Amid AI Disruption
3. Connecting Inclusive Education and Responsible AI
4. From Educator Experiences to Practical Recommendations
5. Putting the Recommendations into Practice
6. Looking Ahead
7. Council Members

1. Introduction

Artificial intelligence is reshaping how undergraduate and graduate students learn, the skills they need to develop, and the expectations they will face in the workforce. For educators, these changes present new challenges and important decisions about how best to prepare students for the future.

How can students learn to use AI without becoming overly dependent on it? How can professors maintain academic rigor while making AI education accessible to students with different backgrounds and levels of experience? And how can universities prepare students for a workforce that is still adapting to the technology?

These are some of the questions educators have been exploring through AI4ALL's Future of AI Education Council, a working group of faculty, graduate students, and education practitioners convened during the 2025–2026 academic year to develop recommendations for inclusive and responsible AI education.

Supported by Google.org, the Council brings educators together to share experiences, explore challenges, and develop practical approaches to AI education. Its work recognizes that effective teaching requires more than access to technology. Educators also need resources, institutional support, and the flexibility to adapt their teaching to students' needs.

The Council's first white paper, Cultivating New Guidelines for Inclusive AI Education in Higher Ed, published in February 2026, introduced its mission and vision for inclusive AI education. The second, Teaching Through AI Disruption, published in September 2026, drew on interviews and Council discussions to examine the challenges educators are encountering, from unequal access to AI tools to questions about student learning, academic rigor, and workforce preparation.

This third white paper builds on those insights to offer practical recommendations for educators.

The recommendations draw on the experiences and perspectives of Council members, with a focus on supporting professional judgment, encouraging experimentation, and providing a foundation for continued learning as the technology evolves.

Every classroom and institution is different. The goal is to equip educators with useful starting points.

2. Finding Direction Amid AI Disruption

The pace of AI development has made it difficult for educators and students to distinguish lasting changes from short-term trends. Questions about workforce readiness, academic integrity, skill development, and the future of higher education remain unresolved. With so many competing perspectives, finding clear direction can be challenging.

Yet one encouraging trend is emerging in 2026: responsible AI is becoming a strategic priority for businesses and an increasingly talked-about concept in the public conversation.

As organizations move beyond experimenting with AI, they are recognizing the importance of accountability, sound judgment, risk management, and human oversight. PwC's 2026 research found that 72% of surveyed Canadian organizations consider responsible AI a top priority, while 71% expect positive financial outcomes from their investments in trustworthy AI.

Other research shows that organizations still have work to do. A September 2026 EY survey found that 98% of senior AI executives surveyed reported having formal AI governance policies, yet 69% expressed concern about insufficient internal expertise to evolve those controls.

Meanwhile, Morgan Stanley's 2026 report found that 90% of executives involved in AI governance considered data risks, including privacy, cybersecurity, and discrimination, material concerns.

These findings offer an important perspective on the future of work. As organizations develop more responsible approaches to AI, they will need people who can evaluate risks, exercise independent judgment, understand the consequences of technology decisions, and help establish effective practices.

Higher Education is Already Shaping What Comes Next

Higher education occupies a uniquely influential position in meeting this demand for responsible AI practitioners.

The students learning how to work with AI today will become the employees, managers, researchers, technologists, founders, policymakers, and executives responsible for deciding how these systems are used tomorrow.

That gives educators an opportunity to do more than help students adapt to AI. They can help students develop the judgment needed to shape how AI is implemented in practice: when to question a system, when human oversight is necessary, how to recognize unintended consequences, how to evaluate tradeoffs, and how to connect innovation with accountability.

3. Connecting Inclusive Education and Responsible AI

Inclusive AI education and responsible AI share several important principles. Both emphasize meaningful access to technology, sound judgment, and an understanding of how AI systems affect people.

In higher education, inclusive teaching means providing students with meaningful opportunities to learn, regardless of their prior experience or level of preparation. It involves maintaining high academic expectations while offering the instruction, resources, and support students need to develop their knowledge and skills.

These priorities are consistent with the U.S. Department of Education's August 2026 guidance on responsible educational technology, which emphasizes evidence-based instruction, educator judgment, and measurable student learning outcomes. The guidance encourages educators to evaluate technology based on its instructional value and whether it improves student learning.

Responsible AI extends many of these considerations to the design, evaluation, and use of AI systems. The National Institute of Standards and Technology's AI Risk Management Framework identifies several characteristics of trustworthy AI, including reliability, safety, security, transparency, accountability, privacy, and fairness with harmful bias managed.

These principles come together in the classroom. For example, educators can help students evaluate AI-generated information, test systems for inaccuracies or bias, understand data privacy risks, and determine when human oversight is necessary. Such exercises strengthen technical knowledge while developing critical thinking and independent judgment.

Future of Education Council members have also emphasized that inclusion and academic rigor can reinforce one another. Students may begin with different levels of preparation, but the goal is to help them meet high academic expectations through effective instruction, appropriate resources, and opportunities to apply what they learn.

Together, these principles provide a foundation for the recommendations that follow. They encourage educators to make informed decisions about AI, maintain rigorous learning standards, and prepare students to evaluate and use emerging technologies responsibly.

The recommendations are intended as practical guidance that educators can adapt to their disciplines, institutional policies, and students' learning needs.

4. From Educator Experiences to Practical Recommendations

Across interviews and ongoing discussions, members of the Future of AI Education Council have shared how they are responding to AI's growing presence in higher education.

Their experiences reveal a range of approaches. Some educators are rethinking how they assess student learning, placing greater emphasis on reasoning and accountability. Others are introducing AI concepts across disciplines, connecting coursework to real-world problems, and finding new ways to help students participate in challenging technical work.

These approaches share several priorities: maintaining academic rigor, supporting student learning, encouraging independent judgment, and preparing graduates to use AI responsibly.

The following five recommendations draw on recurring themes from Council members' experiences. They offer practical starting points for educators and institutions to adapt based on their teaching goals, resources, and student needs.

Recommendation 1: Encourage students to be open about their AI use.

As AI tools become more accessible, educators are reconsidering how to evaluate student learning. Several Council members described approaches that place greater emphasis on how students use AI, rather than simply whether they use it. The goal is to help students understand their own work, explain their decisions, and take responsibility for the results they produce.

One educator described changing their approach to AI-assisted programming assignments:

> *"Last year (2024-2025 academic year), one of the rules I had for programming assignments was that you cannot use any AI coding assistants. We had something that checks if the code is real or AI-generated. This year, I intentionally lifted that. I said, feel free to use whatever you want to use—just walk me through the process of how you used these tools."*

Instead of relying exclusively on restrictions or detection, this educator began asking students to document their use of AI and explain how they completed their assignments.
Other Council discussions highlighted the value of oral examinations, group presentations, and in-class activities that allow students to demonstrate their understanding directly.
These practices connect to an important principle of responsible AI: people need to understand, evaluate, and remain accountable for work completed with AI assistance.

Recommendation 1: Encourage students to be open about their AI use.

Additional Suggestions

- **Set clear expectations for AI use**. Explain when students may use AI, how they should disclose their use, and what responsibilities remain their own.
- **Assess reasoning alongside results**. Ask students to explain their decisions, demonstrate their understanding, and evaluate AI-generated outputs.
- **Use multiple assessment methods**. Consider oral presentations, in-class exercises, and collaborative projects that allow students to demonstrate their knowledge.
- **Protect foundational learning**. Identify skills students should develop independently before using AI to complete more complex tasks.
- **Encourage reflection**. Ask students to consider where AI helped, where it introduced errors, and how they verified their results.

The goal is to help students become thoughtful, capable users of AI who can explain and take responsibility for their work.

Recommendation 2: Make AI education relevant across disciplines and real-world experiences.

AI is becoming relevant to a wide range of academic disciplines and career paths. Council members have emphasized the importance of helping students understand its applications beyond computer science and other traditionally technical fields.

This includes introducing AI concepts in subjects such as the humanities and social sciences, as well as connecting coursework to issues students encounter in their communities and everyday lives.

One educator described the value of using agriculture to demonstrate AI's practical relevance:

> *"Our region is a heavily agricultural region, and a lot of students come from families that work in the fields. Demonstrating the use of AI in the space of agriculture—which really presents incredible opportunities—is a way of demonstrating that if you learn this, you can actually learn something that you can bring back to your families and your communities and help them advance."*

Examples like this help students connect technical concepts to familiar problems. They also encourage students to consider how AI might be applied within their chosen disciplines, whether they plan to develop AI systems or use them professionally.

Responsible AI education adds another dimension. Students can learn to evaluate the potential benefits, limitations, and consequences of using AI in different contexts.

Recommendation 2: Make AI education relevant across disciplines and real-world experiences.

Additional Suggestions

- **Introduce AI concepts across disciplines**. Help students understand how AI relates to their subject areas, even when they are not pursuing technical degrees.
- **Connect learning to real-world problems**. Use examples from industries, communities, and everyday experiences that students can recognize and investigate.
- **Integrate responsible AI throughout coursework**. Encourage students to examine accuracy, fairness, privacy, safety, and accountability as part of technical and applied learning.
- **Use applied projects**. Give students opportunities to evaluate AI systems, test their limitations, and consider whether they are appropriate for particular tasks.
- **Encourage interdisciplinary perspectives**. Create opportunities for students to examine AI-related questions from technical, social, ethical, and professional viewpoints.

By connecting AI education to meaningful applications, educators can help students see its relevance to their interests, academic goals, and future careers.

Recommendation 3: Support meaningful participation while maintaining academic rigor.

Students enter higher education with different experiences, levels of preparation, and degrees of confidence. Council members have emphasized that effective teaching recognizes these differences while maintaining high expectations for student achievement.

In interviews, educators challenged the assumption that making education more inclusive requires reducing academic difficulty. One Council member explained:

> *"The making of an inclusive classroom does not mean making an easier class. Right, the goal is to broaden the curriculum... people are coming in from different perspectives; the goal is not to make the material easier. The goal is to give them resources so that they're able to follow along with the current material, or present it in different ways... without taking away from the rigor of the original curriculum."*

Council members described several ways to support student participation, including small-group activities, live polling, peer collaboration, interactive exercises, and direct encouragement. Mentorship and constructive feedback also play an important role, particularly for students who may be uncertain about their abilities or unfamiliar with technical subjects.

These practices are especially valuable in AI education, where students may have very different levels of prior exposure to the technology.

Recommendation 3: Support meaningful participation while maintaining academic rigor.

Additional Suggestions

- **Maintain clear and consistent learning standards**. Define what students should know and be able to demonstrate, regardless of their starting level.
- **Provide multiple opportunities to engage**. Incorporate discussions, group projects, interactive exercises, and independent work to support different approaches to learning.
- **Identify gaps in preparation early**. Use formative assessments and student feedback to determine where additional instruction or support may be helpful.
- **Encourage mentoring and peer learning**. Create opportunities for students to exchange ideas, ask questions, and learn from one another.
- **Make room for experimentation**. Give students opportunities to make mistakes, receive feedback, and improve their understanding over time.

The underlying principle is that meaningful access and academic excellence can reinforce one another. Students benefit from challenging learning experiences accompanied by the instruction and support needed to succeed.

Recommendation 4: Give educators the resources and institutional support to teach AI effectively.

Individual educators play an important role in adapting teaching practices, but they cannot address every challenge associated with AI adoption on their own.

Council members have described the importance of faculty professional development, institutional coordination, and access to appropriate technology.

One educator explained how their institution was approaching these challenges:

> *"We have been working at the org level. We have created an AI council that has been driving a lot of this conversation, trying to secure resources to actually provide access to students. Nevertheless, our faculty are still kind of trying to navigate the space of how to modify their pedagogy to improve access to technology and access to AI."*

The educator also described professional development initiatives intended to help faculty become more familiar with evidence-based teaching practices.

Other Council discussions highlighted efforts to negotiate access to AI tools and infrastructure, helping institutions reduce the financial barriers students and faculty may encounter.

These experiences reinforce the importance of institutional support. Educators need opportunities to learn about emerging technologies, evaluate their instructional value, and make informed decisions about when and how to use them.

Recommendation 4: Give educators the resources and institutional support to teach AI effectively.

Additional Suggestions

- **Invest in faculty professional development**. Offer training that addresses AI literacy, instructional design, assessment, and responsible technology use.
- **Create opportunities for collaboration**. Support faculty working groups and cross-disciplinary discussions where educators can exchange experiences and teaching resources.
- **Develop clear institutional guidance**. Establish expectations for appropriate AI use, academic integrity, privacy, and human oversight while preserving room for instructional judgment.
- **Evaluate access to technology**. Consider institution-wide licensing, vendor agreements, and appropriate lower-cost tools that can help students participate in AI-related learning.
- **Review teaching practices regularly**. Give faculty opportunities to evaluate what is working, identify emerging challenges, and refine their approaches.

Institutional support can help educators move beyond individual experimentation toward more consistent, evidence-informed teaching practices.

Recommendation 5: Connect responsible AI education to career readiness and lifelong learning.

Students are preparing to enter a workforce in which AI capabilities, expectations, and professional responsibilities continue to evolve.

Council discussions have highlighted the importance of connecting classroom learning to employment needs, real-world applications, and opportunities for continued skill development.

Some institutions are exploring employer partnerships, applied projects, and microcredentials that allow students to demonstrate emerging AI capabilities. These approaches can help connect academic learning with professional practice.

One Council member elaborated:

> *"That is to me the biggest challenge that any university is going to face: having a dedicated function to building relationships with tech employers, and having them at the table to give feedback. Then to say, 'Okay if I have you at the table, you're going to give input into curriculum....How do we nurture relationships with you to get students that real skill—virtual internships, robust internships, apprenticeships?"*

Responsible AI provides an important foundation for this preparation.

As organizations place greater emphasis on risk management, accountability, and human oversight, graduates can develop the skills to evaluate AI-generated information, recognize system limitations, protect sensitive information, and make informed decisions about appropriate technology use.

These responsibilities extend across professions, including education, healthcare, business, engineering, public policy, and the arts.

Recommendation 5: Connect responsible AI education to career readiness and lifelong learning.

Additional Suggestions

- **Connect coursework to professional applications**. Use case studies and projects that help students understand how AI is used within different industries and occupations.
- **Engage employers and professional practitioners**. Seek input on emerging skills, workplace expectations, and responsible AI practices that students may encounter after graduation.
- **Emphasize transferable capabilities**. Develop students' critical thinking, communication, problem-solving, and ethical reasoning alongside technical skills.
- **Explore flexible learning pathways**. Consider short courses, microcredentials, and continuing education opportunities that support skill development beyond a single degree program.
- **Prepare students for ongoing change**. Encourage the habits of independent learning, evaluation, and adaptation needed to navigate evolving technologies throughout their careers.

Career readiness involves more than familiarity with today's AI tools. It also requires the ability to learn new systems, evaluate their implications, and exercise sound judgment as professional expectations change.

5. Putting the Recommendations into Practice

The recommendations in this paper are not exhaustive. They offer a starting point for educators to draw on their professional judgment, experiment with different approaches, and continue refining their teaching practices as AI evolves. Institutions should consider their own academic requirements, student learning goals, available resources, and applicable policies when putting these ideas into practice.

Across these different approaches, one principle remains consistent: effective AI education should help students develop the knowledge, skills, and confidence to engage with technology thoughtfully and responsibly.

By combining academic rigor with meaningful opportunities to learn, educators can prepare students to contribute to a future in which responsible AI increasingly matters across disciplines, industries, and communities.

6. Looking Ahead: Building on the Council's Work

The 2025–2026 academic year marks the conclusion of the Future of AI Education Council's inaugural year and an important milestone in its work.

Over the past year, Council members have contributed their time, expertise, and firsthand experiences to an ongoing examination of how higher education can respond thoughtfully to AI. Their perspectives have helped identify common challenges, highlight promising teaching practices, and inform recommendations that educators can adapt to their own classrooms.

The conversations have also reinforced the value of bringing educators together across disciplines and institutions. While their circumstances differ, Council members share a commitment to meaningful learning, academic excellence, and preparing students to make informed decisions about emerging technologies.

AI education will continue to evolve, and many questions remain open. The Council's work offers a foundation for continued research, collaboration, and the exchange of practical knowledge.

AI4ALL extends its appreciation to the Council members for their contributions throughout this inaugural year. Their insights have shaped this series of white papers and will continue to inform broader conversations about the future of inclusive and responsible AI education.

7. 2025-2026 Members

Dr. Andy Gunawardena
Professor
Department of Computer Science
Rutgers University

Dr. Angelina Wang
Assistant Professor
Department of Information Science
Cornell University

Dr. Augustin Chaintreau
Associate Professor
Department of Computer Science
Columbia University

Dr. Banu Ozkazanc-Pan
Professor
School of Engineering
Brown University

Dr. Cynthia Rudin
Professor
School of Computer Science & Electrical and Computer Engineering
Duke University

Dr. Deepa Muralidhar
Tenure Track Faculty
Department of Computer Science
Georgia State University

Dr. Deepti Ghadiyaram
Assistant Professor
Department of Computer Science
Boston University

Dr. Desmond Patton
PIK University Professor, Chief Strategy Officer
Penn Center for Inclusive Innovation & Technology, Social Policy, Communications, Psychiatry
University of Pennsylvania

Dr. Diane Uwacu
Assistant Professor
Department of Computer Science
Mount Holyoke College

Dr. Dylan Thomas Doyle
Research Fellow
Center for Behavioral Intervention Technologies
Northwestern University

Dr. Enrico Pontelli
Professor, Dean
Department of Computer Science
New Mexico State University

Fiona Ryan
Graduate Student
Department of Computer Science
Georgia Tech University

Dr. Judy Hoffman
Associate Professor
Department of Computer Science
Director, Center for Machine Learning
University of California, Irvine

Julia Chae
Graduate Student
Department of Electrical Engineering and Computer Science
MIT

Justin Kay
Graduate Student
Department of Electrical Engineering and Computer Science
MIT

Dr. Kaoutar Ben Ahmed
Assistant Professor
School of Computing and Information Sciences
Florida International University

Dr. Kinnis Gosha
Professor, Chair
Computer Science Department
Morehouse College

Dr. Maria De-Arteaga
Associate Professor
Department of Data, Analytics, Technology, and Artificial Intelligence
ESADE Business School

Marilu Duque
Community Moderator
Instructor
Ignite Accelerator
AI4ALL

Nadun Kulasekera Mudiyanselage
Community Moderator
Instructor
Ignite Accelerator
AI4ALL

Dr. Olga Russakovsky
Associate Professor
Department of Computer Science
Princeton University

Dr. Priya Donti
Assistant Professor
Department of Electrical Engineering and Computer Science
MIT

Rebecca Funke
Community Moderator
Instructor
Ignite Accelerator
AI4ALL

Dr. Richard Souvenir
Professor, Vice Provost for Strategic Initiatives
Department of Computing & Information Sciences
Temple University

Dr. Samira Zad
Assistant Teaching Professor
School of Computing and Information Sciences
Florida International University

Dr. Sara Beery
Assistant Professor
Department of Electrical Engineering and Computer Science
MIT

Dr. Suxia Cui
Associate Professor
Department of Electrical and Computer Engineering
Prairie View A&M University

Timm Haucke
Graduate Student
Computer Science & Artificial Intelligence Laboratory
MIT

Dr. Vikram Ramaswamy
Lecturer
Department of Computer Science
Princeton University

Dr. Yash Tadimalla
Post-Doctoral Researcher
Department of Computer Science
University of North Carolina at Charlotte

Dr. Yisong Yue
Professor
Computing and Mathematical Sciences Department
Caltech



This report, the final in our Year One series, reflects a year of thinking and generosity from many people. Our deepest thanks go to the Council members whose interviews, discussions, and classroom experience form its foundation, and to AI4ALL's co-founder and board chair, Dr. Olga Russakovsky, for her ongoing guidance. We are grateful to the AI4ALL instructors who served as Council Community Moderators during the 2025-2026 academic year, Becca Funke, Dr. Nadun Kulasekera Mudiyanselage, and Marilu Duque, for sustaining the conversations that made this work possible.

We thank Dr. Sean Peters, Dr. Jon Doane, and Dr. Cody Glickman for their guidance throughout, Thalea Smorres for data and evaluation advising, and Justin Strauss for his technical contributions to the qualitative analysis. Kaylee Duong and Mia Nguyen led the design and delivery of this paper. Finally, we extend our continued thanks to Google.org, whose support made the Future of AI Education Council's inaugural year possible.