

Teaching Through AI Disruption

Perspectives on Higher Ed Adaptability, Inclusive Classrooms & Responsible AI

By Ritika Strauss, Nicole Halmi, Bo Young Lee

September 2026



Executive Summary

In just a few years, artificial intelligence has moved from a specialized area of research and technical practice into a major area of global economic investment impacting education, work, institutions, and everyday human infrastructure. AI's potential benefits are substantial, from accelerating scientific discovery, creating access to economic opportunity, and improving crisis response to strengthening critical human infrastructure and expanding access to knowledge. At the same time, the rapid deployment of AI is both surfacing and exacerbating significant challenges related to bias and societal exclusion, misinformation and manipulation, inequitable access and outcomes, privacy and surveillance, environmental degradation, wealth inequality, and other forms of harm.

Because humanity is experiencing the rapid expansion of AI across our societal institutions and everyday life in real time, often with limited visibility into the geopolitical and institutional motivations shaping AI's trajectory, there is little distance between these changes and the questions they raise about our education systems.

For educators on the front lines of teaching emerging technical talent, those questions are becoming increasingly immediate. Faculty are being asked to expand access to AI literacy while preserving academic rigor, supporting meaningful learning, preparing students for a changing workforce, and creating positive and inclusive classroom experiences. At the same time, faculty are navigating institutional constraints, rapidly evolving technologies, uneven access and preparedness, unclear public policy environments, and uncertainty about what responsible AI integration should look like in practice.

Understanding how higher education is adapting to AI requires more than tracking institutional policies or technological adoption. It also requires listening to the educators experiencing these pressures firsthand. These lived experiences offer an essential view into where higher education is adapting effectively, where tensions are emerging, and what forms of support may be needed next.



Executive Summary (Cont.)

Qualitative research captures those lived experiences in greater depth: how faculty are interpreting change, where they are encountering friction, what tradeoffs they are making, and which forms of support they believe are most needed.

This white paper draws on a qualitative assessment of interviews conducted between April and August 2026 with 10 members of AI4ALL's Future of AI Education Council. Beyond surfacing key themes from faculty experiences across classrooms and institutions, the paper develops a conceptual discussion of resilient, inclusive, and responsible higher ed leadership; ethical AI education; and the relationship between the two amid ongoing technological disruption. The paper considers how colleges and universities can adapt to rapid change while preserving meaningful learning, equitable participation, academic rigor, and student agency. These findings form the basis of the Council's third white paper, scheduled for publication later in 2026.

- 
1. **Introduction**
 2. **Methods**
 3. **Qualitative Discussion**
Highlights & Key Themes
 4. **Zooming Out: Higher Ed**
Adaptability
 5. **Council Members**



1. Introduction

University students represent a critical talent pipeline for future STEM innovation, not only because they have grown up alongside machine learning systems and increasingly AI-enabled consumer technologies, but because they are actively developing the critical thinking, technical, and disciplinary skills they will need to use AI effectively and, in many cases, contribute to building, evaluating, and improving the technologies themselves.

Preparing students for a future in which AI exists for the human good therefore requires more than giving them access to new technologies or teaching them how to use emerging tools. Students enter the classroom with widely different levels of exposure to AI, technical confidence, subject matter knowledge, and access to resources. Additionally, they may hold very different expectations about what AI means for their education and future careers. It also depends on when and how AI enters the learning process, since introducing tools before students have built core competencies can risk deskilling, where students lose abilities they once had, or no-skilling, where a competency never develops at all.

Some see new possibilities for creativity, productivity, and problem-solving. Others are navigating growing anxiety about employment, pressure to keep pace with rapid technological change, and uncertainty about whether the skills they are developing will retain their value. These varied and evolving responses have also surfaced publicly, including through high-profile incidents in which graduating students voiced a shifting, and still unsettled, relationship with AI and open questions about its trajectory.

Inclusion as a Foundation for Responsible AI

Inclusion is increasingly recognized as a core component of Responsible AI, which broadly refers to developing and deploying AI in ways that prioritize fairness, accountability, transparency, safety, and human well-being. [UNESCO's Recommendation on the Ethics of Artificial Intelligence](#) identifies inclusiveness as a foundational value for ethical AI and calls for diverse participation, fairness, non-discrimination, and equitable access to AI's benefits. The [OECD's principles for trustworthy AI](#) similarly link responsible technology development to inclusion and reducing social and economic inequalities.

For higher education, this connection has important implications. The students learning to build, evaluate, govern, and apply AI systems today will help shape how those technologies affect people and communities in the future. Broadening meaningful participation in AI education can bring a wider range of experiences, disciplines, and perspectives into that process, strengthening students' ability to identify assumptions, inequities, and potential harms that might otherwise be overlooked. While the current political climate in the U.S. has created strong headwinds for academic institutions to aggressively pursue certain forms of inclusion, the principles of meritocracy and fairness that shape access remain vital to human-centered AI development and deployment.

Inclusive AI education is therefore not separate from Responsible AI education, which helps shape how AI technologies are ultimately designed, implemented, governed, and commercialized. Creating opportunities for a wider range of students to participate meaningfully can strengthen the examination of bias, fairness, accessibility, accountability, and societal impact while helping students understand that technical decisions can produce different consequences for different communities.

In this way, inclusion serves two related purposes: expanding who has the opportunity to participate in an AI-shaped future and strengthening the capacity of future practitioners to develop and use AI responsibly.

Faculty as the Practical Front Line of Higher Ed Leadership

Student outcomes depend heavily on the faculty designing and facilitating classroom environments that support their learning. In the context of AI, that places educators at the center of consequential decisions about where the technology belongs, how students should engage with it, which skills still need to be developed independently, and how to ensure equitable opportunities for participation.

In this context, higher ed resilience is not simply the ability to absorb technological change. It is the capacity to adapt while preserving the conditions that make meaningful learning possible: academic rigor, equitable participation, critical thinking, human judgment, and opportunities for students to build confidence and agency.

Faculty are navigating that challenge amid rapidly changing technologies, shifting workforce expectations, uneven institutional resources, evolving student needs, and unresolved questions about responsible AI use. Many of these pressures cannot be addressed through pedagogy alone.

Understanding how faculty are responding is therefore essential to understanding the resilience of higher education itself. Their experiences reveal both the practical constraints shaping AI implementation today and the institutional support, resources, and infrastructure needed to create more rigorous, inclusive, and adaptive learning environments.

The sections that follow examine how educators are navigating AI's growing impact on higher education. Drawing on interviews with 10 members of AI4ALL's Future of AI Education Council, the white paper explores the trajectory and use of AI, the challenge of preserving meaningful learning, the role of inclusion in strengthening academic rigor, and the barriers that can limit faculty efforts in practice. It also highlights emerging classroom strategies, institutional constraints, gaps in access and preparedness, and the support educators need to respond effectively.

Together, these themes offer a qualitative view into the pressures shaping higher ed resilience today and provide the foundation for the Council's third white paper, which will translate these findings into tactical recommendations later this fall.


Creating the Conditions for Meaningful Student Participation

In higher education, inclusion is not simply a question of who has access to the classroom, but whether students are positioned to participate meaningfully once they are there. In AI education, that distinction is especially important because students arrive with very different levels of exposure, preparation, confidence, and expectations about the technology.

These differences in learner perspectives matter because classroom experiences can shape who feels capable of participating in AI-related learning and ultimately who sees themselves as having agency within a workforce in which AI is present. Inclusive AI education therefore extends beyond access alone. It requires learning environments where students can experiment, ask questions, make mistakes, critically evaluate AI systems, and understand both their capabilities and limitations.

For educators, this creates a complex design challenge. The goal is not simply to incorporate AI into existing coursework, but to create conditions in which students with different backgrounds and levels of preparedness can engage meaningfully with it. That includes preserving opportunities to develop foundational knowledge and independent thinking, while also helping students learn when and how AI tools can support their work.

These classroom decisions have implications beyond individual courses. Higher education is one of the environments in which students advance their understanding about who builds AI, whose expertise matters, how these technologies should be used, and what responsibilities accompany their development. Creating inclusive learning environments is therefore part of a broader effort to ensure that a wider range of students can participate in, and help shape, the trajectory of AI.



2. Methods

Study Design

This study uses a qualitative research approach to understand how higher education faculty are navigating the practical challenges of inclusive AI education.


The AI4ALL Future of AI Education Council maintains an active Slack community that supports ongoing peer exchange through discussions and weekly prompts focused on inclusive AI education. Beginning in April 2026, AI4ALL began conducting in-depth interviews with Council members to capture educators' classroom experiences, professional perspectives, instructional decision-making, and concerns about the evolving role of AI in higher education.

The interviews are guided by a central research question:

How are educators advancing inclusive AI education within their institutions, and what barriers do they encounter in the process?

This question provides a framework for examining how educators are balancing inclusion, academic rigor, institutional capacity, Responsible AI, and long-term workforce preparedness. 10 participants have participated in interviews so far. Faculty are teaching a range of courses related to AI and computer science, such as:

- Core computer science classes
- Topics in computer vision
- Applied machine learning
- Ethics-integrated AI classes



Each interview is approximately 45 minutes and exploratory in nature, giving participants space to reflect on their experiences and identify issues they considered most significant.

Interviews examined:

- Professional backgrounds, teaching pathways, and motivations for working in education
- Experiences developing and teaching AI-related curricula
- Perspectives on inclusion, equity, and Responsible AI
- Strategies for broadening participation in AI education
- Curriculum design, teaching practices, and institutional initiatives
- Barriers to advancing inclusive AI education
- Institutional, policy, funding, and resource constraints
- Societal, technological, and workforce implications of AI
- Risks associated with the trajectory of AI education and deployment
- Opportunities for stronger collaboration, institutional support, and shared practices
- Lessons emerging through participation in the Future of AI Education Council

Qualitative Analysis

Data collection and analysis were conducted concurrently. This iterative approach allows emerging themes to be identified while interviews are still underway, enabling subsequent conversations to explore recurring patterns, tensions, and areas requiring greater depth.

The purpose of the analysis is not to produce statistically representative or broadly generalizable conclusions. Rather, qualitative research provides a way to examine how educators interpret their experiences, how similar challenges manifest across different institutional contexts, and which themes consistently warrant further investigation.

Preliminary findings in this white paper are based on interviews conducted between April and August 2026 and are intended to provide an early synthesis of the themes emerging from the Council's work.

AI-Assisted Thematic Analysis

The research team used the [Expansively Qualitative Research platform](#) to support thematic analysis and the identification of recurring patterns across interview materials and Council discussions. Expansively was co-developed by this white paper's co-author Ritika Puri in collaboration with natural language processing strategist and technical architect Justin Strauss.

The platform incorporates GenAI-assisted workflows to support the organization, retrieval, synthesis, and thematic clustering of qualitative material. AI is used as an analytical aid rather than as an independent source of findings or interpretation.

The system operates against a constrained body of research materials, including interview excerpts and Council discussions, rather than relying primarily on open-ended model generation. Structured retrieval and source-grounding techniques are used to connect analytical outputs back to the underlying qualitative evidence.


Researcher Review and Validation

A layered review process is used to reduce the risk of unsupported or inaccurate AI-generated interpretations.

Generated themes and claims are evaluated against their supporting source material, and weak or unsupported evidence is removed from the analysis. Claims that cannot be sufficiently grounded in the research materials are excluded.

All analytical outputs are also reviewed by the lead researcher responsible for each interview and its qualitative analysis, who is trained in qualitative research methods. Human review is used to validate interpretations, identify potential inaccuracies, preserve appropriate context, and ensure that the findings presented in this white paper remain grounded in participant experiences.

Because interviews and Council discussions are ongoing, the findings presented here should be understood as preliminary. They are intended to identify recurring themes and areas requiring further attention, which will inform the development of tactical recommendations in the Council's third white paper.



3. Qualitative Discussion Highlights & Key Themes

A Working Definition for Inclusive AI Education

Interviewees define inclusive AI education as ensuring access and equal opportunity for all students, regardless of background, economic status, or prior experience. This includes supporting first-generation students, those from non-technical disciplines, and students with varying levels of preparation, so everyone can participate meaningfully in AI learning.



“To me, inclusion means enabling access to everybody regardless of background, regardless of economic conditions, regardless of family demands. And you know, we have a lot of students here who are first-generation college students, so they really don't have, many of them don't have access that we expect. They don't have the family support, the social support to actually embrace academia, especially in a technical space like AI. So to me, inclusion means enabling everybody to have access and equal opportunity to everybody.”

Several educators emphasized the importance of maintaining thoughtful, nonpartisan discussions about inclusion even as institutional policies and public discourse around DEI terminology continue to shift. For faculty, the underlying questions remain practical: which students are participating, who is being left out, where disparities persist, and what conditions help more students succeed.

As one interviewee explained:



“There are groups that are severely underrepresented in academia, in the workplace, in the workforce. And we cannot just ignore that.”

This perspective suggests that inclusive education need not depend on any particular political or institutional label. It can instead be grounded in observable student experiences, participation gaps, access to resources, and the responsibility of educators to create rigorous learning environments in which a wider range of students can meaningfully engage.

For several interviewees, inclusion begins with a straightforward principle: students should have meaningful access to educational opportunities regardless of the circumstances they bring with them. That includes differences in economic resources, family responsibilities, prior preparation, and other factors that can affect a student’s ability to participate fully.

As one educator explained:



“So to me, inclusion means enabling access to everybody regardless of background, regardless of economic conditions, regardless of family demands. ...inclusion means enabling everybody to have access and equal opportunity to everybody.”

Several key focus areas areas of inclusion include but are not limited to:

- **Supporting First-Generation and Low-Income Students:** Educators recognize that first-gen students often lack family or social support for navigating academia, especially in technical fields like AI. Efforts include targeted office hours and summer classes to help these students build confidence and skills, even if uptake is sometimes limited.
- **Reaching Non-Technical Disciplines:** There's a push to make AI accessible to students in fields like history and psychology, not just computer science. Faculty seek guidance and professional development to help non-technical students engage with AI meaningfully and without fear.
- **Addressing Varying Levels of Preparation:** Educators emphasize the need for resources and pedagogical strategies that allow students with different backgrounds and preparation levels to participate fully in AI learning.
- **Demonstrating Relevance:** Connecting AI to students' communities and real-world issues, such as agriculture for students from farming backgrounds, helps make the subject matter more accessible and motivating.

The Connection Between Inclusion and Responsible AI

Interviewees consistently described inclusion and responsible AI as deeply intertwined and mutually reinforcing. Rather than treating these as separate educational goals, faculty see them as inseparable in both principle and practice.

Educators point out that bringing responsible AI concepts—like fairness, bias mitigation, and ethical impact—into the classroom automatically broadens participation by ensuring that a diversity of perspectives, backgrounds, and experiences shape discussions and project work. One faculty member put it succinctly:



“The more we introduce notions of responsible AI, that automatically brings in conversations like incorporating a diverse set of viewpoints, which automatically helps make the course more inclusive. ...I’m not sure I can think about both of those, like, separately, to be honest.”

Faculty highlight that promoting access, representation, and a plurality of viewpoints can help facilitate more robust and ethical AI outcomes. By training students from varied backgrounds, educators empower an expanded generation to address algorithmic bias and design solutions relevant to more communities.




“When we think about AI access, there is AI as a topic of learning and training as many people as possible so they know how the algorithms work, allowing them to design tools for problems that are relevant to them... building expertise as much as possible by increasing perspective plurality, so we can start seeing a wider set of solutions.”


Instructors deliberately embed ethics and inclusion in assignments and projects—challenging students to consider whose perspective is missing, who might be harmed, and how bias could emerge. A faculty leader shared:



“A lot of what I try to train them on is responsible evaluation: how are you able to anticipate the things that you’re not measuring...how do you consider whether different populations may be underserved by the solution that you’re deploying? That’s one of the main components I evaluate in the final project, for example.”



Several key intersections between inclusion and responsible AI include but are not limited to:

- Educators highlight that introducing responsible AI concepts—such as fairness, bias mitigation, and ethical impact—naturally broadens classroom participation and brings a plurality of voices into technical discussions. This ensures that a range of lived experiences and perspectives directly shapes how students engage with AI challenges.
 - Faculty emphasize that fostering inclusion through broad access and representation leads to more ethical and robust AI solutions. Training students from varied backgrounds equips them to recognize, address, and ultimately reduce algorithmic bias, helping to develop technologies that serve a wider array of communities and needs.
 - Instructors embed ethics and inclusion throughout project work and assessments, challenging students to ask critical questions about whose perspectives may be missing, who could be harmed, and how their work impacts diverse groups. This approach positions inclusive thinking as integral to technical rigor and responsible practice.
 - Some institutions now require students to begin their AI studies with courses focused on ethics and societal impact, framing inclusion and responsibility as the foundation for technical learning and shaping students' perspectives from the outset.
- 

The Term “Inclusion” and Its Relationship to Academic Rigor

Interviewees emphasize that inclusion and rigor must reinforce each other: the goal is to maintain high expectations while providing the resources, support, and multiple entry points needed for students with varying perspectives and lived experiences to succeed. For example, one educator described adapting teaching methods to clarify concepts for students learning in a second language, with making the material accessible, not easier. Others highlight that bringing in topics like fairness, ethics, and bias adds technical depth and complexity, rather than diluting rigor.

“Inclusion has nothing to do with washing down the curriculum or making things easier. It’s just presenting things in a different way—presenting them in a way that’s recognized by the students and that gives value to what you’re presenting to the student. So in a sense, it’s more connecting with where the students are and where the students come from. It’s nothing to do with simplifying. As a matter of fact, I think that one of the things that has been demonstrated over and over is the fact that if you want to improve the success of especially first-gen students coming from underrepresented backgrounds, what you want to give them is mastery experiences. Not something that’s easier, but something that shows them they can advance like everybody else—or even more than everybody else. Because then you give a sense of strength, you give a sense of belonging. So as a matter of fact, you want to actually strengthen the rigor, not soften it.”

Several interviewees explicitly rejected the idea that inclusion means making courses easier, arguing instead that it means presenting rigorous material in ways that connect with students' backgrounds and experiences. They see inclusive AI education as broadening participation and perspectives, not reducing academic challenge.



“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.”

Several educators also emphasized that inclusive and Responsible AI topics can increase, rather than reduce, technical complexity. Bringing questions of fairness, bias mitigation, and responsible system design into the classroom can require students to engage with a broader set of technical and conceptual challenges.



“If you're including things about responsible AI, or fairness and bias mitigation, all of that—there are really interesting technical challenges—within that space. I think bringing those into the class allows students to appreciate that they are going to be learning technical details about this much broader space. It actually introduces more complications, more challenges, and requires them to understand more. And I think that's how we want to be viewing it; I don't think this is a reason to simplify curricula.”

For some faculty, maintaining rigor while supporting inclusion also means holding a consistent academic standard while responding flexibly to what individual students need to reach it. Rather than making assumptions based on a student's background, educators described focusing on the specific support that can help each learner succeed.



"I say like this is the bar. I believe you can reach this bar and, you know, whatever you need that can help you do better—let's just have a conversation about it. I don't even bring up any of their other issues."

One interviewee openly reflected on a period when they realized they were subconsciously lowering standards for certain students from specific groups, influenced by narratives around inclusion from 5–10 years ago.




"If you're overly influenced by some of the narratives that happened five to 10 years ago, then implicit bias starts creeping in and you start subconsciously lowering standards for certain students from certain groups. I noticed it in myself, I noticed it in others, and after reflecting on it for a long time, I decided that that was not okay."

After noticing this in themselves and others, they decided it was not acceptable and shifted their approach: instead of lowering standards, they now focus on providing extra help and resources so all students can meet the same high expectations.



"What I do now is, if a student needs extra help, I make sure they get the extra help, but I don't lower my standards."



This educator emphasizes that inclusion should not mean modifying or reducing standards, but rather ensuring every student has the support needed to reach them. They also emphasize the importance of understanding students' different backgrounds and confidence levels, but still holding everyone to the same academic bar. Their philosophy is to enable, empower, and inspire students, giving them space to fail and improve, while always holding them accountable for growth. The educator describes:

- **Direct conversations with students:** If a student is struggling, the educator talks with them to identify what specific help they need to reach the bar, without referencing their background or making assumptions. The focus is on what will help that individual succeed, whether it's extra tutoring, resources, or encouragement.
- **Belief in student potential:** The educator explicitly communicates that they believe every student can meet the high standard, reinforcing confidence and motivation.
- **Creating space to fail and improve:** They work to enable, empower, and inspire students, giving them room to make mistakes and learn, but also holding them accountable for growth and improvement over time.
- **Access to tools and resources:** While not always specifying exact resources, the educator emphasizes ensuring students have what they need—whether that's additional help, materials, or guidance—to succeed.

Strategies to Make AI Education More Inclusive

Interviewees are developing and implementing a range of strategies to ensure AI education is accessible, rigorous, and relevant for all students—regardless of background, discipline, or prior experience. Their approaches reflect a commitment to excellence. Specific strategies include:

Faculty Professional Development

Institutions are investing in faculty training on evidence-based, inclusive teaching practices. This is especially important for reaching students outside traditional technical disciplines. One interviewee shares:



“We have been developing professional development initiatives across the [organization] for faculty members to get more well-versed about what are evidence-based practices to improve access for students.”

Another adds:



“Guidance, professional development, pedagogical guidance, and anything that basically helps me help the student who is in the history class, the student who is in the psychology class, not be afraid of AI, but actually be driven into the use of AI in the proper way. Ethics is the other important thing.”

Demonstrations of Real-World Relevance to Students

Interviewees make AI concepts more accessible by connecting them to issues relevant to students' lives and communities. One interviewee explained:



“We know, for example, students are motivated by looking at the impact of AI in their community—how they can help their own communities by learning AI and using AI to solve problems... 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,”

Providing Access to Infrastructure and Resources

To address cost barriers and ensure equitable access, universities are negotiating with vendors to provide generative AI tools to all students. One interviewee explains:



“We have been negotiating with different vendors so that we can offer access to some generative AI tools to all the students in the university... providing that type of infrastructure.”

In-Classroom Encouragement

Interviewees focus on creating environments where students feel empowered to take risks, learn from failure, and are held to high standards. One interviewee shares:



“I spend all my time thinking about how to enable, empower, and inspire; how to make students feel like they have space to fail, but also to hold them accountable to improve themselves over time. Those are the things I work on that I think a lot about.”



Across interviews, specific tactics include:

- **Interactive learning activities**—including quizzes, scenarios, and in-class exercises—designed to sustain engagement and encourage active participation.
- **Interpersonal engagement strategies**, such as moving throughout the classroom and intentionally inviting participation from a broad range of students, to foster a more inclusive and collaborative learning environment.
- **Real-time brainstorming sessions and live polling** to encourage critical thinking, surface diverse perspectives, and increase participation even in larger classroom settings.
- **Peer-group research and presentation assignments** that encourage students to take ownership of their learning while strengthening collaboration, communication, and public speaking skills.
- **Role-playing exercises and group projects** that connect AI concepts to real-world applications and practical problem-solving contexts, which help students grapple with fairness, responsibility, and implications for different groups.
- **Empowering students to become active creators** by encouraging them to find and address local or community problems using AI, and designing projects where they must evaluate the impacts and potential biases of their solutions

Barriers Posing Challenges to Educator Goals

Despite strong commitments to inclusion, educators face significant barriers in making AI education accessible and equitable for all students.

Cost and Access to Technology

Providing equitable access to AI tools and infrastructure can be a major financial challenge for both institutions and students. One interviewee shares:



“We need to make sure that we provide everybody access to the infrastructure to use AI... that access is expensive. If I have to give something like [AI tool] to every single student, that will cost the university close to a million dollars per year. There are a lot of universities that cannot afford that price tag.”

Faculty Readiness

Many educators are still learning what AI means for their own professional use as they work to adapt teaching approaches, support diverse learners, and improve access to technology. One interviewee notes:



“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.”

Digital Divide

There is ongoing concern that AI could widen gaps between students who have access and those who do not, potentially disenfranchising underrepresented groups. Faculty and institutions are working at multiple levels to address these disparities, but solutions remain elusive. One interviewee points out:



“[Long-standing] access and affordability issues haven't necessarily changed; it's just a new domain of technology.”

Student Anxiety

The pressure to succeed in AI can lead to high anxiety and feelings of inadequacy among students. One interviewee shares:



“The pressures of higher ed and the pressures of employment come through and end up in the classroom with high levels of anxiety for students who are very much in this thinking mentality that AI's going to replace jobs...[that people] need to be bigger, better, faster. When really what they need to be is more human.”

Bias and Dataset Representation

Interviewees also pointed to the data underlying AI systems as an important dimension of inclusion. From this perspective, questions of representation begin before a model reaches the classroom or workplace: they are embedded in the datasets used to train and evaluate the technology. As one educator explained:



“I think about inclusion with AI starting with data—specifically how extremely biased and non-representative the available digital data is at the moment...we're not really doing much to adjust and make corrections.”

Policy and Regulatory Uncertainty

Responsible AI education also takes place within a broader policy and governance landscape that remains unsettled. Faculty are not only helping students think through the ethical use of AI, but also preparing them to navigate evolving questions around regulation, human oversight, bias, and accountability.



“And, of course, there are bigger conversations. For example, at the state level, there are a lot of bills that are introduced on controlling the use of AI, ensuring that AI is used in an ethical manner, and that if AI is used for consequential decisions, there is a human in the loop to make sure there are no biases or other issues. Those are difficult conversations because the pendulum tends to swing from one side to the other, while the right side is somewhere in the middle—and we need to help them in that process.”

Sustained Institutional Investment

Several educators pointed to sustained institutional investment as a prerequisite for meaningful AI access and implementation. Where funding, infrastructure, and professional development remain uneven, faculty may face structural limits on what they can realistically provide students.



“So far, we have not seen steady investments, for example, from the state to actually enable that type of access.”

Fear of Change

Beyond readiness, there is often a genuine fear of change among both faculty and students about how to approach AI, how to ensure responsible outcomes, or where to even start. This slows efforts both to broaden access and to build ethical practices.



“There are disciplines that have been good at adapting, but others are holding back. Some of that is fear of change and not knowing where to start. And that's a big risk.”

Pace of Change vs. Institutional Inertia

Many respondents see universities as slow to adapt compared to the rapid pace of AI-driven changes in industry. There's concern that higher ed struggles to update curricula quickly enough to meet evolving workforce needs, and that this is a significant challenge for long-term relevance.



“Universities are not institutions that adapt very fast. They are slow-moving while the workforce is changing very quickly.”

Debates About Higher Ed’s Inherent Purpose

Some argue universities should focus on developing ethical, responsible leaders rather than just workforce training. They warn that if higher ed recalibrates its mission entirely around the latest technology, it risks losing sight of its core purpose to provide a well-rounded education.



“I’m of the view that [higher ed] is not workforce training per se; it’s sort of creating responsible, ethical people who go out to the workforce, go out to society, and become leaders. And so when you recalibrate entire curricula on one new technology, you sort of lose the heritage of what is the purpose of higher education. If your sole purpose is to create a job—a workforce—I think you really already missed the boat because the workforce of today is going to look nothing like the workforce in 10 years.”



4. Zooming Out: Higher Ed Adaptability

Inclusive AI education is not simply a question of technology adoption. It is a broader test of higher education's ability to adapt while preserving the conditions that make meaningful learning possible: academic rigor, equitable participation, critical thinking, human judgment, and student agency.

Across interviews, faculty consistently described inclusion and rigor as mutually reinforcing rather than competing goals. Broadening participation does not require lowering expectations; it requires ensuring that students with different levels of preparation, access, and confidence have meaningful opportunities to meet them. Educators also described Responsible AI as inseparable from this work, particularly where questions of fairness, bias, representation, accountability, and societal impact become part of technical learning itself.

At the same time, many of the pressures shaping classroom decisions extend beyond what individual faculty can solve through pedagogy alone. Cost and infrastructure constraints, uneven access to AI tools, faculty readiness, student anxiety, gaps in dataset representation, policy uncertainty, and inconsistent institutional investment all influence how effectively educators can respond.

Seen through this lens, higher ed leadership in AI is not simply a matter of asking faculty to adapt faster. It depends on whether institutions create the conditions that allow educators to respond thoughtfully to disruption: clear guidance, sustained investment, professional development, access to infrastructure, opportunities for collaboration, and enough flexibility to preserve educational priorities as technologies continue to evolve.

The qualitative findings in this paper provide a view into where those conditions are already emerging and where important gaps remain. The findings also reinforce the value of continuing to listen closely to educators working at the practical front lines of AI integration.

5. Council Members

The Council is continuing to welcome new participants and conduct additional interviews to broaden the range of perspectives represented in this work. Its third white paper, scheduled for publication in late 2026, will provide additional guidance to help educators navigate this increasingly complex balancing act as the 2026-2027 academic year progresses.

Dr. Andy Gunawardena

Professor
Department of Computer Science
Rutgers University

Dr. Deepa Muralidhar

Tenure Track Faculty
Department of Computer Science
Georgia State University

Dr. Angelina Wang

Assistant Professor
Department of Information Science
Cornell University

Dr. Deepti Ghadiyaram

Assistant Professor
Department of Computer Science
Boston University

Dr. Augustin Chaintreau

Associate Professor
Department of Computer Science
Columbia University

Dr. Desmond Patton

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

Dr. Banu Ozkazanc-Pan

Professor
School of Engineering
Brown University

Dr. Diane Uwacu

Assistant Professor
Department of Computer Science
Mount Holyoke College

Dr. Cynthia Rudin

Professor
School of Computer Science & Electrical and Computer Engineering
Duke University

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 Mudiyanse

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 Computer &
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



Acknowledgements

This work would not be possible without the efforts and guidance of the Council members, including AI4ALL's co-founder and board chair Dr. Olga Russakovsky; and the AI4ALL instructors who served as Council Community Moderators in the 2025-2026 academic year, Becca Funke, Dr. Nadun Kulasekera Mudiyansele, and Marilu Duque. We are also grateful to Dr. Sean Peters, Dr. Jon Doane, and Dr. Cody Glickman for their guidance on this work, to Thalea Smorres for data and evaluation advising, to Justin Strauss for his technical contributions to the qualitative analysis, and to Kaylee Duong and Mia Nguyen for the design and delivery of this paper. We extend our continued thanks to Google.org for their support of this work.

