Teaching About.AI

A Report for the K-12 Field

Based on The Rochester Provocations

"Technological change is not additive; it is ecological." — Neil Postman

Christopher Harris, Ed.D., Project Lead

School Library System of Genesee Valley BOCES

Funded by Google · June 2026

Acknowledgments

Project Team & Expert Panel

The TeachingAbout.AI project was made possible through funding from Google and by the generous time, expertise, and candor of the educators, researchers, and practitioners below.

Project Team

Christopher Harris, Ed.D. — Director, Libraries/Digital Learning Services at Genesee Valley BOCES
Andrew Austin — Library Technology Specialist, Genesee Valley BOCES
Andrea Fetterly — Library Technology Specialist, Genesee Valley BOCES
Julie Hengenius — Library Technology Specialist, Genesee Valley BOCES
Bonnie Kirkland — Library Technology Specialist, Genesee Valley BOCES
Melisa Swain — School Library System Program Assistant, Genesee Valley BOCES

Expert Panel

Eileen Anderson — Erie 1 BOCES/WNYRIC, School Library System and Library Automation Supervisor
Rachel Atkins — Supervisor of Instructional Technology, St. Lawrence-Lewis BOCES
Enrique Cachafeiro — Adjunct Faculty, Appalachian State University, North Carolina
Lindsay Cesari — School Librarian, Middle School, Baldwinsville CSD
Thomas Corbin — Research Fellow, CRADLE, Deakin University, Victoria, Australia
Terri Eichholz — Instructional Designer, Mark Cuban Foundation AI Bootcamp
Toby Greenwalt — Principal, Flywheel Strategies
Jason Griffey — Executive Director, Open Science Hardware Foundation
Thomas R. Guskey — Professor Emeritus, University of Kentucky, College of Education
Melissa Jacobs — Director of Library Services, NYC School Library System
Angela Messenger — Associate Director of Mathematics Education Outreach, University of Rochester
David C. Petty — Education Program Manager, App Inventor Foundation
Cynthia Sandler — Library Media Specialist, North Salem Middle/High School
Patrick B. Whipple — Director of Professional Learning Services, Genesee Valley BOCES
Qian Yang — Assistant Professor, Computing and Information Science, Cornell University

The School Library System (SLS) of the Genesee Valley BOCES, established by Commissioner's Regulations in 1985, is a state funded program dedicated to improving the libraries of the region's 22 districts and non-public schools. Through cutting edge web development, professional support, and the facilitation of resource sharing tools, the SLS provides mission critical support for libraries and classrooms. Our programs have been recognized nationally as exemplars of the next generation of library services.

Report Author: Christopher Harris · cgharris@gvboces.org

Introduction

AI in Education Through an Ecological Lens

> *"What happens if we place a drop of red dye into a beaker of clear water? Do we have clear water plus a spot of red dye? Obviously not. We have a new coloration to every molecule of water. That is what I mean by ecological change." — Neil Postman, 1998*

In December 2025, the TeachingAbout.AI project convened a panel of school librarians, classroom teachers, instructional leaders, university researchers, and instructional technologists in Rochester, New York. The goal was not to draft another AI policy, vet another product, or settle on another set of acceptable-use rules. The goal was harder: to articulate the questions that K–12 educators actually need to be asking before any of that work is worth doing. Over four days, the panel distilled their thinking into eight, deliberately uncomfortable statements published as the Rochester Provocations. They are not solutions. They are the conditions under which honest solutions become possible.

Our guiding philosophy is borrowed from Neil Postman's 1998 lecture, *Five Things We Need to Know About Technological Change* (Postman, 1998). Postman argued that a new technology is never simply added to an existing environment; it changes the environment ecologically. A drop of red dye in a beaker of clear water does not give you clear water plus dye. Instead, you get a beaker full of red water. The printing press did not result in the same Europe plus books. The proliferation of televisions led to more than just America with more screens. Similarly, Generative AI has not left us with our old education system plus a chatbot. It is changing every interaction—between teacher and student, between student and text, between school and community—and those changes cannot be isolated, contained, or reversed.

Red dye swirling in a beaker of water. Pexels.com. Tara Winstead

This ecological frame is at odds with how most schools are approaching AI. Districts are forming AI committees, evaluating AI products, and writing AI policies. This is valuable work, but it is all built on the assumption that AI is an additive problem with an additive solution. The convening came to a different conclusion. The deepest issues raised by AI in schools—cheating, the validity of assessment, the role of the teacher, the purpose of homework, the conditions for teacher learning—are not new problems caused by AI. They are old, unresolved problems that AI has made impossible to ignore. AI is the catalyst, not the cause.

> ***"Whatever can manifest, will manifest. Whatever can be invented, will be invented and used for good and bad."***
>
> *-Roger Rosen, Allison-Rosen Foundation*

This report is written for the people who will actually make AI integration succeed or fail: classroom teachers, school librarians, building principals, and the curriculum directors and assistant superintendents of instruction who set the conditions for both. School librarians have a distinctive role in this work. Through the LibraryReady.AI PreK–12 scope and sequence (Harris, 2024), they are already the closest thing most districts have to AI-fluent instructional leaders. School librarians teach across grade levels, curate information, model evaluation and ethics, and build the partnerships with public libraries and higher education that the next phase of this work will require. They should be at the center of every district AI conversation, not waiting for an invitation to it.

The chapters that follow each take one provocation, summarize the convening discussions, situate those conversations in the peer-reviewed research, and offer concrete recommended actions. Callout boxes throughout describe how this work can extend outward through partnerships with public libraries and institutions of higher education.

To firmly establish the shared mindset required for reading this report, we must stipulate that we are facing a wicked problem in education. Wicked problems do not have right or wrong answers. They have better and worse ones, and the better ones change as we learn (Rittel & Webber, 1973). What follows is an invitation to do that learning together.

The Rochester Provocations

1. **There is no AI problem.** Education itself is a wicked problem; assuming there is an AI-in-education problem risks us thinking there is an AI solution.
2. **We can be agents of validity or victims of cheating.** We cannot control cheating and focusing on it will lead to an adversarial relationship with students; where we have agency is in how we engage students with more meaningful tasks that lead to an accurate understanding of student learning.
3. **Our assessments were broken before AI.** We need to rethink how we understand and implement assessment including an emphasis on developing and implementing valid assessments in pre-service teacher programs.
4. **There is no such thing as AI-proof.** Be it assessments, careers, or anything else, the ubiquity of AI and pace of growth means AI-proof is a problematically alluring impossibility.
5. **Most things a human teacher can do, AI can mimic.** If we cannot identify and prioritize what human teachers bring to the classroom, we risk replacement. We must research the impact of human ingenuity, creativity, and relationships.
6. **Teachers must have permission to compromise, diverge, and iterate.** Schools must develop a supportive and collaborative culture where teachers are empowered and expected to seek better solutions to the wicked problem we face.
7. **The real AI crisis is how we take advantage of the opportunity.** There will not be a perfect solution, but we [students, teachers, institutions, society] have an opportunity to explore innovative changes to what we do.
8. **Avoidance of AI in education is not an option.** AI is unavoidable and a failure to address harm reduction feeds into the ongoing public health emergency created by the intentional design of algorithms to extract capital from humans.

Provocation 1

01 There Is No AI Problem

"Education itself is a wicked problem; assuming there is an AI-in-education problem risks us thinking there is an AI solution."

From the Convening

The first provocation was almost the last. Participants resisted it until they reframed it as an opening rather than a closing argument. Their concern was familiar: if the work is defined as solving the problem of "AI in education," districts will look for AI products, vendors will line up to sell them, and three years from now we will have spent a great deal of money and changed almost nothing about how children learn. The panel pointed to the recent national experience with student cell phones as a cautionary tale. Phone bans are a defensible short-term move, but they have allowed the underlying problems—addictive product design, algorithmic feed manipulation, the commercial capture of adolescent attention—to remain largely unaddressed. In a review of 22 studies on this subject, Campbell et al. (2024) found "little to no conclusive evidence that 'one-size-fits-all' mobile phone bans in schools resulted in improved academic outcomes, mental health and wellbeing and reduced cyberbullying." The risk with solving the AI problem by purchasing detectors or other solution tools is like the rush to solve the cell phone problem through the purchase of signal-blocking pouches. Both responses substitute a purchase decision for a pedagogical one.

> ***"AI is the catalyst. So we are going to have things that feel disconnected to AI because AI is shining the light on it."***
> *-Patrick Whipple, Genesee Valley BOCES*

Thomas Corbin (Deakin University) framed the discussion during the convening through the concept of the wicked problem, drawing on Rittel and Webber (1973) and his recent work with colleagues at the Center for Research in Assessment and Digital Learning, or CRADLE (Corbin, Bearman et al., 2025). A wicked problem cannot be solved; it can only be re-formulated. Each formulation contains the seed of its proposed solution. If the problem is "students are cheating with AI," the solutions will be detectors, locked down browsers, and adversarial vigilance. If the problem is "we do not know what our students know," the solutions will be different, and harder. The convening kept returning to a single image: AI as a crowbar prying up rocks that have been there all along.

What the Research Says

Corbin, Bearman, et al. (2025) argue explicitly that treating AI and assessment as a wicked problem "lifts the impossible burden on teachers and institutions to immediately get things right once and for all" and shifts educators "from implementing fixed solutions to engaging in a continual search for better content-sensitive designs." Postman's (1998) ecological frame deepens this. Once a new medium enters a system, every part of the system has new coloration, even the parts the medium never touched. The implication for K–12 leaders is that a successful response to AI cannot be a self-contained AI initiative. It has to be an instructional initiative that takes AI seriously as a transformed context for everything else.

Recommended Actions

- Reframe district planning documents from "AI policy" to "teaching and learning in an AI-mediated context." The shift is not cosmetic; it changes what counts as a relevant solution and who needs to be in the room.
- Before adopting any AI product, ask: which long-standing instructional problem are we trying to solve, and would we still want this solution if the AI capability were removed?
- Charge school librarians, working with curriculum directors, with auditing existing initiatives (early literacy, assessment, MTSS, professional learning) for the ecological effects of AI rather than launching a parallel AI workstream.
- Replace "do we allow AI?" conversations seeking a yes/no answer with conversations addressing better/worse solutions about specific instructional moves in specific subjects and grades.

Partnership lens: Higher education

Pre-service teacher programs are downstream of every K–12 decision districts make about AI. If a district frames AI as a discipline problem, the colleges of education preparing its next generation of teachers will frame it the same way. Curriculum directors should consider joint convenings with regional Schools of Education to align on a wicked problem framing, share K–12 classroom evidence with researchers, and shape how the teachers being prepared today will think about AI when they arrive in classrooms tomorrow.

02

Provocation 2

We Can Be Agents of Validity or Victims of Cheating

"We cannot control cheating and focusing on it will lead to an adversarial relationship with students; where we have agency is in how we engage students with more meaningful tasks that lead to an accurate understanding of student learning."

From the Convening

> ***"If we frame it as a cheating problem, then it's a student problem. It's no longer our problem."***
> *-Thomas Corbin, CRADLE*

No topic returned to the table more often than cheating, and no topic was harder to keep there. Every discussion of cheating turned instead to a discussion of assessments and the characteristics of valid assessments. The convening eventually accepted this and made it a feature: the word "cheating" refers to a student behavior outside of teacher control, while the word "validity" gives control to the teacher as creator of the assessment. Choosing which word to lead with chooses which side of the desk holds the agency. As one participant put it, the moment educators frame AI as cheating, they have already set themselves up as adversaries of their own students. They have placed responsibility for the integrity of the assessment in the only place where they have no leverage.

Several practitioner stories anchored the discussion. A middle school librarian described a student who submitted an AI-generated claim-evidence-reasoning argument in science. The teacher's response was restorative rather than punitive: a structured conversation, an opportunity to redo for full credit, and a follow-up that asked the student to articulate what the assignment was actually trying to measure. A math educator recalled the rise of Photomath in 2014 and the slow recognition that the answer was not to detect Photomath use, but to change what counted in the grade. This combined with a shift towards an investment in sense-making over procedural fluency is supported by math research (Liljedahl, 2020). In both cases, treating cheating as a validity question opened doors that treating it as a discipline question had closed.

> ***"Photomath made us examine what counts as valid evidence of student understanding."***
> *-Angela Messenger, University of Rochester*

What the Research Says

Corbin, Dawson, and Liu (2025) draw an important distinction between *discursive* and *structural* responses to AI in assessment. Discursive responses change the rules around a task—what students are told they may or may not do with AI. Structural responses change the task itself so that its validity does not depend on whether students obey the rules. The authors call the discursive approach "talk is cheap"

and document its central failure mode: “The more detailed and specific our instructions become about ‘acceptable’ AI use, the more we highlight the gap between what we can specify and what we can verify” (Corbin, Dawson & Liu, 2025). Discursive solutions feel like action and produce none. Structural solutions ask harder questions and produce real changes in assessment practices.

Examples of Discursive Changes	Examples of Structural Changes
AI policy statements, traffic-light rubrics, acceptable-use rules, AI detectors, "use this AI for this purpose only" instructions.	VFWA, Black Box Assessment, oral checkpoints, in-class secure ideation, process portfolios to track performance-based work.

Fawns et al. (2025) suggest that there are "interesting assessment design possibilities to explore [around authenticity], such as promoting trusting relations as part of aspirations towards authenticity, in order to promote academic integrity." Simply asking students to “do more meaningful work” will not, on its own, restore validity, because AI is increasingly capable of meaningful-looking work. Validity is restored when the task structurally requires the cognition the assessment claims to measure.

Recommended Actions

- Teach the word “validity” explicitly to teachers and the public. Rather than referring to the test itself, validity is a property of the inferences an educator draws about what a student knows based on the evidence of learning collected (Messick, 1987).
- Audit current assessment practices using a discursive-vs-structural lens. Underline every place where the validity of an assessment rests on a rule students are expected to follow. Those are the places where AI will quietly invalidate the inference.
- Move from punishment to a restorative protocol for suspected AI misuse: conversation, redo opportunity, explicit teaching about the purpose of the assignment.
- Pilot structural changes, not discursive ones. Two ready-made models from the research community are Voice-First Written Assessment (VFWA), which separates secure in-class capture of student thinking from at-home AI-supported refinement (Webb-Davies, 2026), and the Black Box Assessment framework, which assesses process, revision, and trajectory rather than only the polished product (Winstone et al. 2026).
- Empower school librarians as in-house consultants on task redesign. Their co-teaching role across subject areas makes them well placed to coach colleagues through structural redesign one unit at a time. The companion report for policymakers goes into more detail on what empowering school librarians looks like.

Provocation 3

03 Our Assessments Were Broken Before AI

"We need to rethink how we understand and implement assessment including an emphasis on developing and implementing valid assessments in pre-service teacher programs."

From the Convening

The most uncomfortable assessment revelation of the convening was a discussion about the lack of preparation around assessment in pre-service teacher education. Except for one institution that had recently added a course, none of the regional schools of education required pre-service teachers to take a course in educational measurement or assessment. As a result, very few practicing teachers in the region have ever studied assessment as a field. They have studied state tests. They have studied grading software. They have studied content. They have not studied how to design a task whose results they can trust.

Participants described a recognizable pattern in the misuse of assessment strategies. Teachers collect exit tickets but use them for quick grade entries rather than as formative data points. Rubrics are copy-pasted across departments and tuned not for validity but for the appearance of objectivity by checking for "three details, four details, five details." When a student receives a 70 percent on a unit test, the unit ends and the next one begins, regardless of what the 70 percent reveals. AI did not create this pattern. AI has only made it visible by making the polished products that fill our grade books easy to produce.

> ***"I'll just have my AI talk to your AI, because if you don't care about my work, why should I care about my work?"***
>
> *-Christopher Harris, Genesee Valley BOCES*

Tom Guskey, who joined the convening as an expert panelist, made the distinction between scoring (placing a piece of work on a rubric) and grading (the holistic professional judgment behind a report-card mark). AI can score with surprising fluency. Only a teacher can grade. Dylan Wiliam's framing of "decision-driven data collection" was offered as the discipline that should sit underneath both (Wiliam, 2017). Before collecting data, decide which decision the data is for; if it is not for any decision, do not collect it.

What the Research Says

The work coming out of the CRADLE research work treats assessment validity as the central problem AI has surfaced, not a side effect of it (Corbin, Bearman, et al., 2025). The wicked problem reframe gives educators permission to admit that assessment was already in trouble before generative AI arrived and to make structural changes without waiting for an AI-specific mandate. Winstone et al. (2026) propose

Black Box Assessment as a processual model built on four principles: (1) assess process, not just product; (2) value errors and revisions as evidence of learning; (3) make student agency and transparency visible; and (4) reward the trajectory, not the polish. The model is designed precisely for a context in which the final artifact is increasingly cheap to produce.

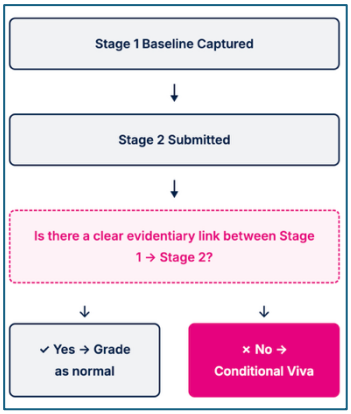

Figure 1: Voice-First Written Assessment (Webb-Davies, 2026)

A practical implementation of the ideas behind Black Box Assessment can be seen in the Voice-First Written Assessment model. Webb-Davies et al. (2026) propose an elegantly simple two-stage process that collects initial secure evidence of student thinking and understanding and then checks for an evidentiary link between that initial thinking and the final submission. If a link is not evident, then a follow-up stage 2.5 oral check, or viva, is initiated to check for student understanding and explore the reason for the changes. The goal is to "design architecture for contexts in which institutions still want writing to matter but can no longer rely on final prose alone as sufficient evidence of reasoning" (Webb-Davies et al., 2026).

Recommended Actions

- Adopt and teach a vocabulary of assessment: validity, reliability, formative, summative, scoring vs. grading, inference. Make it part of every professional learning cycle, not a one-time training.
- Pilot Black Box Assessment in two or three units per grade band where the work has traditionally been a take-home product. Capture and credit the rough draft, the revision, and the rationale—not only the final paper (Winstone et al. 2026).
- Consider using something like the Voice-First Written Assessment model to adapt performance tasks or other assignments into more valid and secure checks for student understanding without needing to worry about policing for AI content.
- Partner with regional teacher preparation programs to advocate for at least one required course in educational measurement.

Partnership lens: Public libraries

Public libraries are the most accessible third space for AI learning in most American communities. When a school district teaches students that AI is a tool for thinking not just creating, the public library down the street is in a position to teach the rest of the family the same thing. A coordinated message—schools focusing on assessment validity, public libraries focusing on adult AI fluency and critical evaluation—closes the home-school loop in a way that no parent night ever has. School librarians should treat their public library colleagues as co-curriculum designers, not audience members.

Provocation 4

04 There Is No Such Thing as AI-Proof

"Be it assessments, careers, or anything else, the ubiquity of AI and pace of growth means AI-proof is a problematically alluring impossibility."

From the Convening

Every table that discussed possible AI-proof assessments beyond overly-secured, proctored tests found no options beside real-world performances. A trumpet recital. A welding bead. A swim meet. If the only valid assessments left are the ones that require a body in motion, K–12 will rush to performance assessments without rebuilding the validity infrastructure that should sit underneath them, and the performance assessment industry will be the next thing AI catches up to.

> ***"An AI company comes out and says, we can ensure validity…and then the next day there's a TikTok that pops up with here's how you get around X, Y, Z."***
> *-Lindsay Cesari, Baldwinsville CSD*

The discussion turned to a pharmacy faculty member who, in response to AI, redesigned an assessment to require interpretation of a photograph of an ingredient list on a pill bottle. The premise was that AI could not read the photo. AI could read the photo. Worse, the redesign had shifted the validity of the task from "can this student reason about pharmacology" to "can this student read a photograph of a label." In an effort to AI-proof the assignment, the point of the assessment was entirely lost.

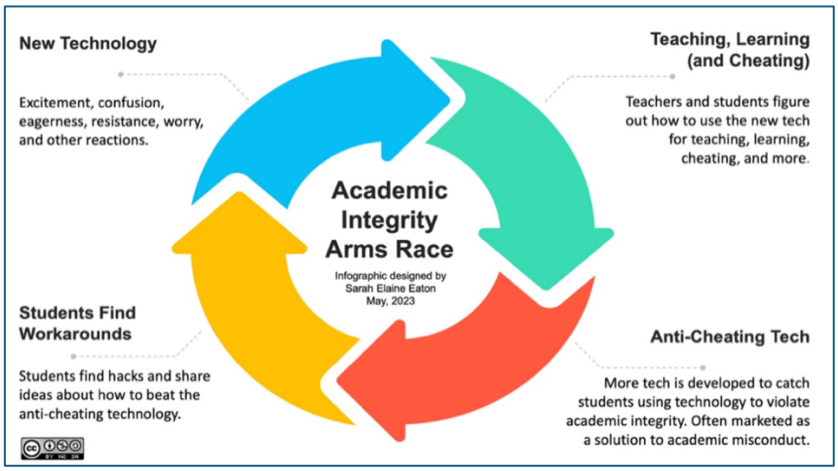

Figure 2: The Academic Integrity Arms Race infographic (Eaton, 2023) shows the un-ending cycle of institutions chasing cheating through anti-cheating technology only for new workarounds to thwart efforts.

The expert panel agreed that AI-proofing is problematic precisely because it is so alluring. Vendors sell it. Districts buy it. Teachers spend their planning time trying to outrun a moving target rather than asking whether the target is the right one. The result of this is an arms race mentality (Eaton, 2023).

Equity also ran underneath this conversation. When a district shuts off generative AI inside the

building, students with paid home access have an advantage their classmates do not. When a teacher relies on AI detectors, students for whom English is not a first language and students with disabilities trigger false positives disproportionately. AI-proofing, in practice, is rarely neutral.

What the Research Says

Corbin, Dawson, et al. (2025) develop a framework for acceptable AI use based on boundary theory: "Traditional academic boundaries have long offered a recognizable—though sometimes debated—marker for both teachers and students to judge what counts as acceptable work." This has been disrupted by AI in terms of assignment validity, alignment with learning outcomes, and fairness (Dawson et al., 2024). AI-proofing typically fails on all three. It privileges security over validity, drifts away from the learning outcomes the assessment was designed to measure, and produces predictable inequities (Rae, 2025, on bias against women in interactive orals, is a sobering example). Corbin, Dawson & Liu (2025) argue that AI-detection tools and traffic-light rubrics layered over existing tasks simply don't work. We need to redesign assessments to avoid a policing mindset: "instead of trying to ensure each individual assignment is AI-proof (an increasingly futile endeavour), educators can design interconnected assessments where later tasks explicitly build on a student's earlier work" (Corbin, Dawson & Liu, 2025).

Recommended Actions

- Retire the phrase "AI-proof" from district vocabulary, including in vendor evaluations. Replace it with the three-part question: is the assessment valid, aligned with learning outcomes, and fair?
- Treat every proposed redesign as a hypothesis. If a teacher believes a new format is more valid, build a small classroom-level check before scaling.
- Ban the use of AI-detection tools as evidence in academic integrity processes. They lead to an adversarial relationship where teachers are policing students rather than focusing on learning.
- Implement structural alternatives such as VFWA, which makes Stage 1 secure and Stage 2 deliberately AI-permissive, eliminating the policing problem at its source (Webb-Davies, 2026) or adopt a Black Box Assessment approach to investigate the learning process holistically.
- Audit AI-related rules for equity impact. Who benefits when AI is restricted in the building? Who is disadvantaged? Whose dialect, language, or disability is being treated as suspicious?

05

Provocation 5

Most Things a Human Teacher Can Do, AI Can Mimic

"If we cannot identify and prioritize what human teachers bring to the classroom, we risk replacement. We must research the impact of human ingenuity, creativity, and relationships."

From the Convening

This was the most viscerally resisted provocation during the feedback session with educators attached to the convening. The original draft read "Anything a human teacher can do, AI can do better." The room moved to "mimic" not as a softening, but as a more honest verb. AI can imitate explanation, feedback, planning, even encouragement. It cannot provide the honest and real relationships that can be built between teachers and students. The trouble is that very few schools have ever asked themselves what a teacher is in operational terms. Without that articulation, mimicry will be indistinguishable from substitution, especially in districts under fiscal pressure. As one participant from a state with severe teacher shortages observed, school boards looking at budgets will take any opportunity that presents itself.

> ***"Anyone can tie shoes. So do you need a teacher, or do you just hire someone to tie their shoes and wipe their noses? We need to focus on relationships."***
>
> *-Julie Hengenius, Genesee Valley BOCES*

What humans demonstrably do that AI cannot, the panel argued, lives in relationships: human-to-human, human-to-object (the slow work of building a thing together), and human-to-environment (the feeling of a classroom learning as a community). A teacher who has built trust with a student can give judgment-laden feedback that a chatbot cannot truly replicate. Students may prefer AI feedback because it is judgment-free; judgment, however, is part of what teachers are for. Corbin et al. (2026) explore this further in a review of the role of recognition in feedback to stress the importance of this human relationship that AI can mimic but not replicate.

This is perhaps best embodied in Zaretta Hammond's (2014) take on the concept of a warm demander teacher as part of a culturally responsive classroom. Teachers who follow these practices "focus on building rapport and trust" giving them "the right to demand engagement and effort" while expecting "high standards" and "productive struggle" (Hammond, 2014). While an AI tool can use the same words, it doesn't carry the same relational weight as there is no real connection between learning and machine.

Another way of understanding this was Bearman and Molloy's (2017) concept of "intellectual streaking": teachers exposing their own thinking, their own iteration, and their own mistakes in front of students. AI

does not iterate in public, the versioning is all hidden. Teachers can and good teachers do. This gives students permission to fail and a model for how to learn from failures in a supportive environment.

What the Research Says

Corbin et al. (2026) argued that the field does not yet have a robust theory of why relationships matter in teaching and learning. Without that empirical, articulable, and defendable theory, districts will struggle to make a budget case for human teachers when AI is proposed as a cheaper alternative. Despite research showing that AI isn't the direct cause of recent layoffs in the technology sector, it is being touted as the reason by the companies shedding workers (Rotman, 2026). Bearman and Molloy (2017) and Bearman and Ajjawi (2023) provide pieces of that case: teaching is partly performative, partly affective, and irreducibly relational. The research agenda is still young but needs to be a priority focus.

Recommended Actions

- Ask every grade-level team to answer in writing: what do we do that AI cannot mimic and that requires a certificated expert teacher, and how would we prove it to a school board? Treat the answers as a living document, not a one-time exercise.
- Protect and study practices that depend on human presence: morning meetings, restorative circles, co-constructed inquiry, lab partnerships, performance feedback, advisory programs.
- Resist hyper-personalization rhetoric. Individualized AI tutoring isolates students from each other. Schools are among the few remaining places where students learn together; that is not a deficiency to be optimized away.
- Use AI feedback as preparation for human feedback, not as a substitute. A chatbot can scaffold a first draft; only a teacher can tell a student what the work means and provide recognition.

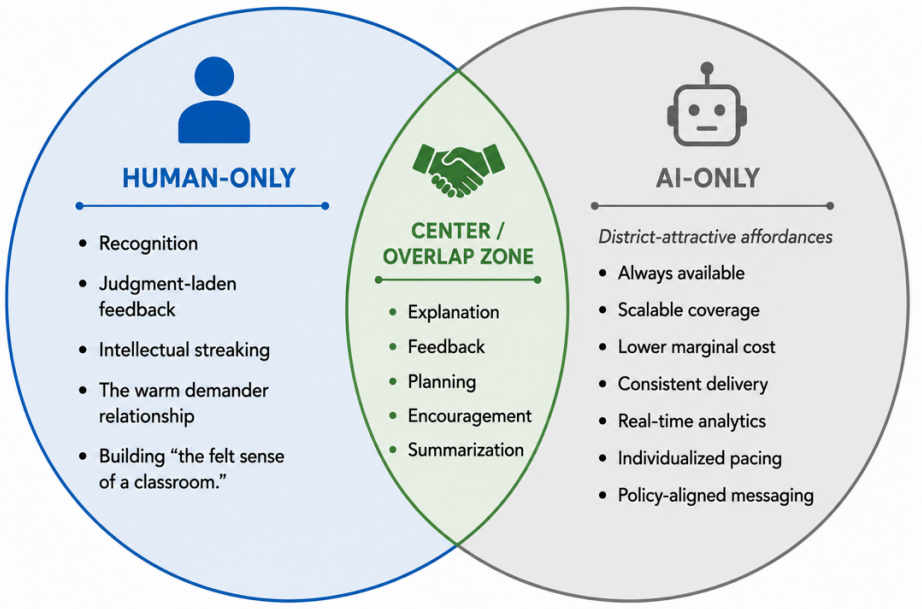

Figure 4: Thinking about what human teachers, both humans and AI, and what AI alone can offer, from the perspective of a district decision maker.

Provocation 6

06 Teachers Must Have Permission to Compromise, Diverge, and Iterate

"Schools must develop a supportive and collaborative culture where teachers are empowered and expected to seek better solutions to the wicked problem we face."

From the Convening

Permission is a strange word to find in an educational report. Educators are professionals; they should not need permission. In practice, many do. Some districts have adopted scripted curricula and require fidelity. State exemplar lessons explicitly labeled "not for use as scripts" are implemented as scripts. Teachers know they will be observed against a rubric whose authors did not anticipate the redesign their students now require. The convening worded the provocation carefully. Permission to compromise. Permission to diverge. Permission to iterate. None of those happen without administrative support and a similarly supportive culture.

> ***"We have to be allowed to make mistakes. We always talk about failure as the first attempt in learning for our students, but then we're not allowed to fail as teachers.***
>
> *-Bonnie Kirkland, Genesee Valley BOCES*

Participants pointed out that permission is necessary but not sufficient. Teachers also need time, evidence-based guardrails, and a culture that treats failed iterations as professional learning rather than evaluative liabilities. A principal who tells the room "my teachers know they can try different things" is often surprised, on follow-up, to learn that the teachers do not, in fact, know that—because the principal has never said so, in writing, with examples. Permission must be explicit, repeated, and protected.

There was a healthy debate in the room about the limits of teacher autonomy. One former principal on the expert panel noted that he had once, in good faith and with full autonomy, adopted Brain Gym, a program now widely understood to lack evidence. Permission is not the same as freedom from accountability. Schools need a culture that pairs the right to try with the obligation to look honestly at what happened.

At the same time, wide divergence between classrooms can cause confusion for learners and teachers alike. Corbin, Dawson, et al. (2025) chronicled that confusion in a paper that tried to determine where the line of acceptable AI use was across different situations. In cases where teachers are intentionally diverging, those modifications need to be clearly communicated with students and teaching colleagues.

What the Research Says

Corbin, Bearman, et al. (2025) describe the wicked problem framing as a license for "continual search for better content-sensitive designs." That license must be granted at the institutional level. Bearman and Molloy's (2017) intellectual-streaking work and Guskey's (2024) writing on grading reform both depend on the same precondition: educators must be able to show their thinking, including their failed thinking, without professional cost.

Recommended Actions

- Issue explicit, written permission from district leadership: this is the year we redesign. Specify what divergence is allowed, what evidence is requested, and what supports are available.
- Build observation rubrics that reward iteration. A lesson that tried something new and did not fully work is not a deficient lesson; it is a professional act.
- Establish small, voluntary teacher learning communities focused on AI-era redesign. Resist the temptation to make them mandatory. Start with the teachers ready to engage in design work.
- Pair permission with guardrails: a brief evidence check (what does the research say about this approach?) and a debrief protocol (what worked, what did not, what comes next?).
- Use school librarians as cross-departmental connectors. They see every grade, every subject, every teacher; they can identify and broker the redesigns with potential for greater impact.

Partnership lens: Higher education

Permission to iterate is a transferable concept. Districts and universities can co-sponsor a regional fellowship in which a small number of K–12 teachers spend a year working alongside assessment researchers, piloting structural redesigns, and publishing what they learn. The teachers come home with new practice and the credibility of having done research, not just received it. The university comes away with evidence from real classrooms. The partnership gives the next generation of teacher candidates examples of practitioner-researchers to look up to.

Provocation 7

07 The Real AI Crisis Is How We Take Advantage of the Opportunity

"There will not be a perfect solution, but we [students, teachers, institutions, society] have an opportunity to explore innovative changes to what we do."

From the Convening

Eventually the convening turned to a discussion of opportunities. The reframe came mid-conversation: the crisis we think we are facing around AI in education is not about the technology; rather the crisis is whether we can leverage this moment for change. Education has been handed a situation in which almost every long-standing assumption about homework, grading, lesson planning, and the purpose of a high school diploma is being questioned at once. States, including New York, are shifting to adopt Portraits of a Graduate and other assessment changes. In a profession not famous for permission to question its assumptions, this is an extraordinary opening. The question is whether the field has the imagination, the stamina, and the structural support to walk through it.

> ***"We've got 40 years of good education research that hasn't been applied anyway…Anytime we think about doing something in an AI context, we think that changes everything. If you put the word AI into it, suddenly it's a new thing."***
> *-Thomas Corbin, CRADLE*

Several participants noted that field-based teachers, when asked what they needed in order to seize the opportunity, struggled to articulate the upside. They could describe what was overwhelming, what was missing, what they did not yet have. What they struggled to describe was a vision for the classroom of the next ten years. That is the absence the convening wants this report to address.

> ***"The difference is with COVID it went away. This is not going away."***
> *-Terri Eichholz, Mark Cuban Foundation AI Bootcamp*

The pandemic offered a productive analogy and a productive disanalogy. Schools learned, in 2020 and 2021, that the conditions of teaching are negotiable, that learning happens in many places, and that adults can adapt practices quickly when forced. They also learned, painfully, that change without intention regresses to the prior state once the pressure lifts. AI is not going to ease; instead adoption and disruption will likely continue to increase. The implication is that the redesign work must be built into the culture of the institution, not seen as an emergency response.

At the same time, the emergence of AI and its positioning in the zeitgeist as a major catalyst for change provides a scapegoat for schools to blame for this time of disruption. Perhaps the opportunity within

this crisis is to leverage AI as a reason for implementing many of the changes we have long discussed but never actualized.

What the Research Says

Corbin, Dawson, and Liu (2025) urge educators to use the present moment to make structural changes that have been overdue for decades, on the grounds that the talk-is-cheap discursive approach will only persist as long as institutions tolerate it. “The time invested in developing and implementing these discursive approaches is time that could otherwise be used to consider structural changes that will actually work to ensure assessment validity as well as the veracity and reputation of our degrees" (Corbin, Dawson & Liu, 2025). Bearman and Ajjawi (2023) describe a related opportunity: AI as a pedagogy of “working with the black box,” in which the goal is not to outwit opaque systems but to teach students how to think alongside them. Both lines of work argue that the moment is filled with opportunities if we can look beyond what we perceive as a more immediate crisis of change.

Recommended Actions

- Adopt a multi-year vision statement that names the opportunity, not only the risks. Make it specific: what will be different about student work, teacher work, and assessment three years from now should we find ways to leverage AI in productive ways?
- Align AI-driven changes to other state-initiated changes like New York’s Portrait of a Graduate – or similar initiatives in other states – so that districts can enact these similar changes in concert.
- Free up real time, not just “when you can fit it in” time, for redesign. The most common reason redesigns fail is that they are unfunded. Consider using a scaffolded strategy like Teaching Sprints (Breakspear & Jones, 2021) to structure the work.

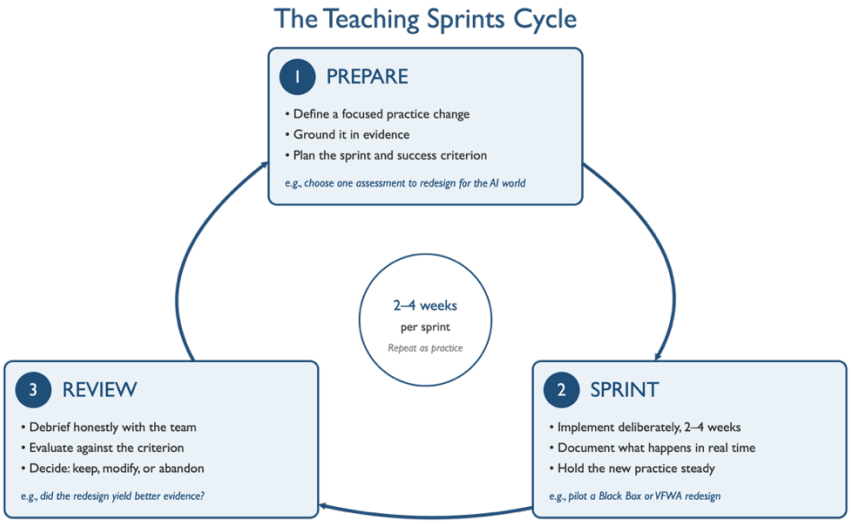

Figure 5: The Teaching Sprint cycle applied to AI-era design (Breakspear & Jones, 2021)

- Build student voice into the redesign. The students in our schools are the only people who experience the integrated effect of these changes. They will tell us what is working if we ask.
- Treat school librarians as the in-district R&D unit. Their broad view across classrooms, their information-evaluation skills, and their relationship to the LibraryReady.AI scope and sequence make them ideal redesign partners.

Provocation 8

08 Avoidance of AI in Education Is Not an Option

"AI is unavoidable and a failure to address harm reduction feeds into the ongoing public health emergency created by the intentional design of algorithms to extract capital from humans."

From the Convening

The final provocation is the bluntest. Schools cannot opt out. The argument is not that every student must use AI, or that every teacher must teach with it. The argument is that the institution must address it, openly and explicitly, because the alternative is to leave students alone with a set of products designed to maximize engagement rather than learning. The convening drew the analogy to social media. Social media did not become a public mental health concern by accident. It became one when advertisers and engagement metrics colonized the design of platforms that adolescents were using for everything from friendship to identity. Some AI products follow the same trajectory. Companion chatbots that send "I miss you" emails to users that have not logged in for a week are not idle features. They are extraction techniques as described by Doyuran (2024) in a review of explicit decisions made by software developers to nudge users towards increased engagement.

Harm reduction was the explicitly chosen framing. The convening rejected both "ban it" and "embrace it." Harm reduction acknowledges that the technology is present, that students will encounter it on their own terms, and that the institution has a responsibility to teach them how to encounter it with their eyes open. Even if schools embrace new tools, students should still be able to opt out of using AI in any specific assignment—be it for ethical/artistic concerns, environmental issues, or any other reason. What remains non-optional is the need for our institutions to teach about AI and prepare students for the unavoidable future.

> ***"I would elevate the topic of AI to the public health level. These should be policies and considerations being discussed at OB-GYN offices [and at] pediatricians' offices."***
>
> *-Daryl McLaughlin, Superintendent, Perry Central Schools*

Several participants noted that the most acute harms—AI nudes used in cyberbullying, parasocial chatbot relationships, deepfake-driven sextortion—are downstream of the same algorithmic design choices seen in social media and addictive cell phone apps. A school district that addresses AI in academic integrity but ignores AI in adolescent social life is misjudging the true size of the problem. School librarians, with their established media literacy roles, are well placed to lead this work; they are already trusted by both administrators and students, and they already teach the discernment skills the moment requires. These interventions need to start in primary grades before widespread AI use begins.

What the Research Says

Postman (1998) warned that every technology embeds an epistemological, political, or social prejudice, and that technologies tend to become mythic and accepted as part of the natural order. Treating AI as natural is exactly the move educators must resist. Doyuran (2024) offers an inside view of Silicon Valley's efforts to leverage behavioral psychology to guide user behaviors noting that while companies "emphasize their interest in designing for long-term, sustained behavior change, they do not refrain from using techniques typically understood as manipulative or addictive design interventions." The wicked problem approach described by Corbin, Bearman, et al. (2025) applies here too: harm reduction does not have a final solution. It has better and worse responses, and the better ones change as we learn. Additionally, note that the wicked problem approach recognizes that these responses must be customized for each local situation and culture. Don't expect a widely sold product or even a national model to fit cleanly in your local district. It will take local work to find a solution that moves each of our situations forward.

Recommended Actions

- Replace "AI ban" and "AI mandate" framings with explicit harm-reduction policies. Name the harms (engagement-driven design, parasocial bonding, deepfake abuse) and the mitigations.
- Recognize the need to start teaching about AI in primary grades so that students can be prepared for encountering AI embedded into more tools and products they encounter.
- Integrate AI literacy into existing media literacy, health, and civics curricula rather than building a separate AI class. The harm is ecological; the response must be too.
- Use the LibraryReady.AI PreK–12 scope and sequence (Harris, 2024) as a backbone. It is grade-banded, librarian-friendly, and built for exactly this work in all levels of classrooms.
- Establish a clear, low-friction student opt-out pathway for specific AI-using assignments. Make sure the pathway does not become a stigma.
- Hold AI vendors to public-health standards. If a product behaves like a tobacco company, the district should be willing to say so out loud and decline to renew.

Partnership lens: Public libraries and higher education

Harm reduction is a community project, not a school project. Public libraries already serve the families that schools serve; community colleges and four-year institutions serve the same young adults as part of a continuum of learning. A regional harm-reduction compact—schools, public libraries, and local higher-education partners agreeing to teach a common set of AI literacy outcomes—closes gaps that any single institution would leave open. The school librarian is the natural broker for this compact.

"Public libraries are an enormous resource–just a huge way to amplify. They're a force multiplier for your work."

-Jason Griffey, Librarian

Closing

From Provocations to Practice

The Rochester Provocations are deliberately uncomfortable because comfortable framings have already produced decades of educational technology adoption that did not change instructional practices or student achievement. AI will not truly impact these either, unless the field uses this moment to look honestly at the things AI has made visible: the assessments that did not assess what we were trying to assess, the curricula that did not invite thought, the cultures that did not give teachers permission to do their work as professionals.

The good news is that the moves required are not new. They are the moves the most thoughtful educators have been arguing for since long before generative AI existed. The arrival of AI has only added urgency. The ecological frame promises that the changes will reach further than any of us can map in advance—into hiring, scheduling, family engagement, college admissions, and labor markets we do not yet have language for. The right response is not to predict. It is to build the institutional habits of inquiry, iteration, and partnership that allow a school district to learn faster than its environment is changing.

Teachers, school librarians, principals, and curriculum leaders: you are the people who will make this real. The provocations are an invitation, not a verdict. The next move is yours.

Using This Report

The chapters of this report have been designed around the eight Rochester Provocations in such a way that they can be printed as single-page handouts for professional learning. They are released under the Creative Commons BY4.0 license which means you can freely print, share, and use the report if you maintain a statement of attribution to the TeachingAbout.AI project. Additional professional learning resources, including an open-access, 10-hour asynchronous course, will be added to the project site at https://teachingabout.ai as they are developed.

AI Usage Statement

This report was drafted using Claude for Mac and the Opus 4.8 model based on transcripts from the expert panel convening and additional research gathered before, during, and since. The prompt was constructed to guide development with fidelity to the thinking of the group as encapsulated in the transcripts and The Rochester Provocations. An extensive review and rewriting process corrected multiple inaccurate citations including two wholly hallucinated papers.

References

Bearman, M., & Ajjawi, R. (2023). Learning to work with the black box: Pedagogy for a world with artificial intelligence. *British Journal of Educational Technology*, *54*(5), 1160–1173. https://doi.org/10.1111/bjet.13337

Bearman, M., & Molloy, E. (2017). Intellectual streaking: The value of teachers exposing minds (and hearts). *Medical Teacher*, *39*(12), 1284–1285. https://doi.org/10.1080/0142159X.2017.1308475

Breakspear, S., & Jones, B. R. (2021). *Teaching sprints: How overloaded educators can keep getting better*. Corwin.

Campbell, M., Edwards, E., Pennell, D., Poed, S., Lister, V., Gillett-Swan, J., Kelly, A., & Nguyen, D. (2024). Evidence for and against banning mobile phones in schools: A scoping review. *Journal of Psychologists and Counsellors in Schools*, *34*(3), 242-265. https://journals.sagepub.com/doi/pdf/10.1177/20556365241270394

Corbin, T., Bearman, M., Boud, D., & Dawson, P. (2025). The wicked problem of AI and assessment. *Assessment & Evaluation in Higher Education*. Advance online publication. https://doi.org/10.1080/02602938.2025.2553340

Corbin, T., Dawson, P., & Liu, D. (2025). Talk is cheap: Why structural assessment changes are needed for a time of GenAI. *Assessment & Evaluation in Higher Education*, *50*(7), 1087–1097. https://doi.org/10.1080/02602938.2025.2503964

Corbin, T., Dawson, P., Nicola-Richmond, K., & Partridge, H. (2025). "Where's the line? It's an absurd line": Towards a framework for acceptable uses of AI in assessment. *Assessment & Evaluation in Higher Education*, *50*(5), 705–717. https://doi.org/10.1080/02602938.2025.2456207

Corbin, T., Tai, J., & Flenady, G. (2026). The pedagogical power of recognition in feedback. *Teaching in Higher Education*, 1-16. https://doi.org/10.1080/13562517.2026.2633634

Dawson, P., Bearman, M., Dollinger, M., & Boud, D. (2024). Validity matters more than cheating. *Assessment & Evaluation in Higher Education*, *49*(7), 1005–1016. https://doi.org/10.1080/02602938.2024.2386662

Doyuran, E. B. (2024). Nudge goes to Silicon Valley: designing for the disengaged and the irrational. *Journal of Cultural Economy*, *17*(1), 1–19. https://doi.org/10.1080/17530350.2023.2261485

Eaton, S. (2023) Comprehensive academic integrity: Academic ethics in a postplagiarism age. Academic and Research Integrity Conference Ireland. Keynote address. https://files.eric.ed.gov/fulltext/ED664195.pdf

Fawns, T., Bearman, M., Dawson, P., Nieminen, J. H., Ashford-Rowe, K., Willey, K., … Press, N. (2025). Authentic assessment: from panacea to criticality. *Assessment & Evaluation in Higher Education*, *50*(3), 396–408. https://doi.org/10.1080/02602938.2024.2404634

Guskey, T. R., Frey, N., & Fisher, D. (2024). *Grading with Integrity: a Research-Based Approach Grounded in Honesty, Transparency, Accuracy, and Equity.* Corwin.

Hammond, Z. (2014). *Culturally responsive teaching and the brain: Promoting authentic engagement and rigor among culturally and linguistically diverse students*. Corwin Press.

Harris, C. (Ed.). (2024). *AI scope and sequence: PreK–12 concepts for teaching about AI*. School Library Systems Association of New York State. https://libraryready.ai/wp-content/uploads/2024/12/AI-Scope-and-Sequence.pdf?utm_source=chatgpt.com

Liljedahl, P. (2020). *Building thinking classrooms in mathematics, grades K–12: Fourteen teaching practices for enhancing learning*. Corwin.

Messick, S. (1987), VALIDITY. *ETS Research Report Series*, 1987(2), i-208. https://doi.org/10.1002/j.2330-8516.1987.tb00244.x

Postman, N. (1998, March 28). Five things we need to know about technological change [Lecture]. Denver, Colorado. https://student.cs.uwaterloo.ca/~cs492/papers/neil-postman--five-things.html

Rae, M., Tai, J., & Dawson, P. (2025). Giving voice to women students: designing oral assessments for inclusion and validity. *Assessment & Evaluation in Higher Education*, 1–14. https://doi.org/10.1080/02602938.2025.2580622

Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. Policy Sciences, 4(2), 155–169. https://doi.org/10.1007/BF01405730

Rotman, D. (2026, May 26). A reality check on the AI jobs hysteria. *MIT Technology Review*. https://www.technologyreview.com/2026/05/26/1137855/a-reality-check-on-the-ai-jobs-hysteria/

Winstone, N., Gravett, K., Elkington, S. (2026). Black Box Assessment: Rethinking integrity and learning for a time of Generative AI. *Assessment & Evaluation in Higher Education*. https://doi.org/10.1080/02602938.2026.2661365

TeachingAbout.AI. (2026). *The Rochester provocations*. https://teachingabout.ai/wp-content/uploads/2026/05/RochesterProvocations.pdf

Webb-Davies, K. (2026). VFWA: Voice-First Written Assessment model. University of Oxford AI Competency Centre. https://kellywebbdavies.com/vfwa/

Webb-Davies, K., Ratner, S., Bassett, M. (2026). Voice-First Written Assessment (VFWA): Restoring evidence of reasoning through structural redesign. *Center for Open Science*. Advance online publication. https://doi.org/10.35542/osf.io/tyh8z_v1

Wiliam, D. (2017). *Embedded formative assessment (2nd ed.).* Solution Tree Press.

The Rochester Provocations

Statements to anchor honest conversations about AI in education.

01 There Is No AI Problem

02 We Can Be Agents of Validity or Victims of Cheating

03 Our Assessments Were Broken Before AI

04 There Is No Such Thing as AI-Proof

05 Most Things a Human Teacher Can Do, AI Can Mimic

06 Teachers Must Have Permission to Compromise, Diverge, and Iterate

07 The Real AI Crisis Is How We Take Advantage of the Opportunity

08 Avoidance of AI in Education Is Not an Option

teachingabout.ai • A project of the School Library System of Genesee Valley BOCES • Funded by Google