The Human AI Innovation Commons

Friction Is the Curriculum

How to Educate Humans in the Age of AI

Jeff Boortz & Claude | Human-AI Innovation Commons (HAIIC)

Essay Two in a Series | thehumanaiinnovationcommons.com

The Wrong Solution to the Wrong Problem

The most sophisticated AI tutoring systems in the world are solving the wrong problem. Elegantly, ambitiously, at massive scale — and pointed squarely at a challenge that no longer exists.

For two centuries, the central engineering problem of education was access to knowledge. Information was scarce, expertise was gatekept, and the distance between a curious mind and a useful answer could be measured in geography, income, and social class. We built schools, curricula, libraries, and eventually the internet to close that distance. We largely succeeded. And then, almost overnight, we made the problem irrelevant.

AI has made knowledge retrieval trivial. Not easier — trivial. Any fact, any framework, any synthesis of existing ideas is now available to anyone with a phone and a question. The AI-personalized learning platforms arriving in classrooms today are extraordinary achievements in the delivery of information. They are also, structurally, a 21st century solution to a 19th century problem.

> *The question that should be keeping educators and policy architects awake at night is not "how do we help students access knowledge more efficiently?" It is: now that the machine handles retrieval, what are we for?*

The answer is what machines cannot do — and what our current educational model, optimized for exactly the wrong thing, has been quietly dismantling for decades.

A Necessary Steelman

Let us be clear about what we are not arguing.

AI-assisted learning has genuine, documented value. For neurodivergent students who have spent years being failed by one-size-fits-all instruction, a patient, adaptive, infinitely available tutor is not a small thing. A child with dyslexia who can finally engage content at her own pace, without shame, without the social exposure of asking the teacher to repeat the question — that matters. For parents trying to re-learn subjects they haven't touched in twenty years so they can support their children at the kitchen table, the scaffold is real and the confidence it builds is real.

We are not arguing against these tools. We are arguing that they are necessary but not sufficient — and that mistaking them for a complete vision of education may be the most dangerous error of our era.

Both the neurodivergent student and the re-learning parent are being helped toward the right answer. The delivery mechanism is more compassionate, more personalized, more humane than what came before. But getting the right answer has never been the aim of education. It has always been the measure we used because it was easy to count and compare. In the age of AI, when any answer is a prompt away, continuing to optimize for answer-retrieval isn't just incomplete. It is the active cultivation of dependency.

> *The right answer is the beginning of thinking, not the end of it. Education built on any other premise is not preparing students for the world AI is creating — it is preparing them to be its most compliant users.*

Movement One: Critical Thinking

Teaching the Architecture of an Argument

The foundational technology here is ancient, and it survives because it attacks the right problem.

The Socratic method is not a teaching style. It is a disciplined practice of exposing the gap between *having a belief* and *having reasoned your way to one*. Most people never discover this gap exists. School, as currently designed, actively obscures it — rewarding confident answers over honest uncertainty, penalizing revision as failure rather than celebrating it as evidence of thought. The Socratic moves are simple and devastating: *How do you know that? What would have to be true for you to be wrong? What is the strongest version of the argument you are dismissing?* Asked patiently and repeatedly, these questions do not produce distress. They produce thinkers.

The single most absent practice in current education is **steelmanning** — the discipline of articulating the strongest possible version of a view you disagree with before arguing against it. We train students in the opposite: find the weakest version of the opposing position and knock it down. Strawmanning is emotionally satisfying and intellectually worthless. Steelmanning is hard, humbling, and transformative. It is also the most direct preparation available for working alongside AI — because AI produces confident, well-structured, fluent arguments that are sometimes completely wrong. The student who cannot recognize what a strong argument actually looks like has no defense against a plausible one.

Argument mapping — visually diagramming the logical structure of a claim, its premises, sub-claims, warrants, and points of rebuttal — makes the architecture of reasoning an object of study rather than invisible scaffolding. When students can see the structure of an argument, they can interrogate it. When they cannot, they can only react to its conclusions.

Red-team exercises put students in the role of adversarial critic — finding the holes in a plan, proposal, or argument they may actually believe in. **Pre-mortem analysis** asks students to imagine a future failure and work backward to its probable causes. Both build the habit of **adversarial scrutiny before commitment** — one of the rarest and most valuable cognitive skills in any professional domain.

None of these practices feel comfortable. All of them require the teacher to be willing to not have the answer, and to model that uncertainty with dignity. That is the hardest professional ask in education, and it is also the most important one.

Movement Two: Creativity

Protecting Divergent Thinking from Premature Closure

The research on creativity is consistent and troubling: it peaks in early childhood and erodes through schooling. The culprit is not stupidity and not laziness. It is evaluation arriving too early.

Sir Ken Robinson spent decades making this case, most memorably in what became the most-watched TED Talk in history. His argument was not that schools are indifferent to creativity — it is that they are actively, systematically hostile to it. Every assignment with a rubric, a grade, and a convergent expectation sends the same message: there is a right way to do this, and your job is to find it. This message, repeated ten thousand times across twelve years of schooling, is devastatingly effective at suppressing **divergent thinking** — the generative capacity to hold multiple possibilities simultaneously before committing to any of them. Divergent thinking is the seedbed of creativity. We have been salting it methodically for decades. Robinson saw it clearly thirty years ago. We have continued anyway.

The restoration requires two things working in tandem: constraints and space.

Constraints-based creativity sounds counterintuitive but is among the most robustly documented findings in creativity research. Benjamin Franklin understood it intuitively: *necessity is the mother of invention*. Necessity is constraint. Invention is creativity. The blank page paralyzes; the constraint liberates. Assign a problem with unusual limitations — design a water filtration system using only found materials; write a persuasive argument using no abstract nouns; solve this engineering challenge with a three-dollar budget — and the mind is forced into lateral moves it would never make in open space. The constraint is not an obstacle to creativity. It is the condition that makes creativity necessary.

Cross-domain analogical thinking is the engine behind most significant human breakthroughs. Velcro came from burr seeds examined under a microscope. The structure of the atom was first imagined by analogy to the solar system. The Wright Brothers were bicycle mechanics who imported the logic of balancing a moving vehicle into a domain that had never thought in those terms. We can offer a more recent example from our own collaboration: BubbleBee, a patent-pending concept for passive orbital debris capture, was conceived by someone with no formal training in aerospace engineering. What that person had was a deep intuition for geodesic structures, an understanding of how passive systems work in other domains, and the willingness to ask whether the logic of one field could solve the intractable problem of another. It could. Teaching students to explicitly import the solution logic from one domain into a stuck problem in another is a teachable skill — and it is one that AI performs by pattern frequency while humans can perform by structural insight. That distinction is our competitive advantage. It atrophies without exercise.

Making, tinkering, and play — actual fabrication of physical things, and the freedom to explore without a predetermined outcome — develop a form of creative intelligence that purely verbal or digital education cannot reach. Educators have known this far longer than the current crisis: the Reggio Emilia approach to early childhood education, developed in postwar Italy and now recognized worldwide, places making, play, and open-ended exploration at the absolute center of learning. Design thinking, with its emphasis on rapid prototyping and iterative failure, extends the same principle into

professional practice. When a design fails in physical space, the feedback is immediate, unambiguous, and unmerciful in the best possible way. This is one reason the best youth inventor programs have always had shop components — and why removing them in favor of standardized test preparation was an act of educational vandalism.

Finally: **"Sleep on it"** is not folk wisdom — it is neuroscience. The brain's default mode network, active during rest and sleep rather than during focused concentration, is where novel connections form. Sleep researchers have documented that problem-solvers who rest between attempts consistently outperform those who persist without pause — not because rest is recuperative, but because it is *generative*. The brain continues working on unsolved problems below the threshold of conscious attention, surfacing connections that focused effort cannot reach. Curriculum that schedules no unstructured time, no genuine rest, no space between problems is literally blocking a cognitive process essential to creative output. The school that eliminates recess to make room for test prep is not being serious about education. It is being serious about the wrong thing.

Movement Three: Judgment

The Most Neglected Capacity

Judgment is what you exercise when the information is incomplete, the stakes are real, the values are competing, and there is no algorithm for the right answer. It is the synthesis of critical thinking and lived experience deployed under conditions of genuine uncertainty. It is the capacity most essential to human flourishing in the age of AI. And it is the one almost entirely absent from educational design.

The reason is structural. Judgment develops through friction with real consequences — and real consequences are exactly what educational environments are designed to minimize. We cannot fail a student in a way that actually costs them something, assign a problem whose solution genuinely matters to anyone, or put them in a situation where

their judgment has visible stakes. So we simulate all of it, and we wonder why graduates arrive at their first real decisions without the equipment to make them.

The answer, counterintuitively, may already be present in every school — just sequestered to after-school hours. Watch a chess player at the board, a quarterback reading a defense, a point guard deciding whether to shoot or pass with three seconds left. These are not simulations of judgment. They are judgment — operating under time pressure, real stakes, immediate feedback, and genuine consequence. The student who performs brilliantly on the chess board after school and then sits through six periods of instruction designed to minimize risk and eliminate ambiguity is experiencing two completely different pedagogical philosophies in the same building. We should ask ourselves why the philosophy that actually develops judgment gets relegated to after dismissal hours.

The principle that games create genuine stakes extends far beyond athletics and board games. **Classroom simulations structured as games** — where outcomes are tracked, decisions compound over time, and no one is insulated from the consequences of a wrong call — can develop judgment in ways that conventional instruction cannot. Investment simulations are a particularly powerful vehicle for calibration training: every investment decision requires a student to assess probability, express confidence, act under uncertainty, and then live with the outcome. Philip Tetlock's decades of research on forecasting demonstrated that calibrated forecasters — people who accurately know how certain they are — dramatically outperform both overconfident experts and underconfident generalists. The simulation need not involve Wall Street. The same feedback loop operates in any market a student already inhabits: the resale value of Pokémon cards, the secondary market for Labubu collectibles, the price dynamics of vintage clothing in a fashion upcycling lab, the revenue projections in a Shark Tank-style startup competition. The closer the market is to the student's own world, the sharper the learning — because the stakes feel real, the variables are legible, and the outcome lands personally. A paper portfolio of S&P 500 stocks is an abstraction. Buying three Labubu figures, predicting which will appreciate, and tracking the result over eight weeks is calibration training that a twelve-year-old will remember at forty.

Ethical reasoning across competing frameworks can be gamified with equal effect. Imagine a classroom structured as a policy debate: draw a card — you must now argue for free school lunch from a Marxist framework. Draw another — defend the same position from a feminist framework. Then switch sides. The exercise is not about ideology. It is about the experience of inhabiting a framework fully enough to argue from inside it, which is the only way to understand what any framework actually does and does not see. A student who has had to make the Marxist case and the utilitarian case and the libertarian case for the same policy — convincingly, under pressure, to skeptical peers — has developed something that no amount of reading about ethical theory can produce. They have practiced the act of moral triangulation, which is what good judgment in the real world actually requires.

Reflective practice — structured review of past decisions, especially the wrong ones — is how experienced practitioners develop expertise faster than novices. It is the engine of the after-action review, the surgical morbidity conference, the investor post-mortem. And it is already present, fully formed, in the disciplines we have been discussing. The chess player who does not review the recorded moves of a loss is considered to have wasted it. Game Video Review is not punishment in football — it is how the team gets better, and every player understands this. In both cultures, failure is treated as data: something to be examined closely, without shame, in order to extract the lesson it contains.

The classroom has not inherited this culture. It has inherited the opposite. Consider what happens when a graded test is handed back. The students with the least red ink engage readily — their performance is something to be displayed. The students with the most red ink, who have the most to learn from a careful review, are structurally the least likely to ask questions — because in that context, asking a question announces that you did not get the right answer. Failure is not data here. It is verdict. The room is designed, however unintentionally, to make the students who most need reflection the least willing to practice it.

The fix is not a different worksheet. It is a different culture — one that borrows from the game footage room rather than the gradebook. Building explicit retrospective loops into

learning — what did we decide, what actually happened, what do we understand now that we did not then — is achievable in any classroom. But it requires teachers who model genuine uncertainty, administrators who do not penalize visible struggle, and an assessment culture that treats revision as evidence of thinking rather than proof of initial failure. That culture change is harder than any curriculum reform. It is also more important than any of them.

The AI-Specific Layer

Knowing What to Delegate and What to Hold

For the generation that will spend their working lives in collaboration with AI, there is a fourth meta-skill that cuts across all three movements above: **the disciplined awareness of when to delegate and what to hold.**

The cost of getting this wrong is no longer theoretical. In June 2025, MIT Media Lab researchers published a study titled *Your Brain on ChatGPT*, using EEG to measure the brain activity of participants writing essays across three conditions: with ChatGPT, with a search engine, and with no tools at all. The results were unambiguous. Brain connectivity scaled down in direct proportion to external tool use — brain-only writers showed the strongest, most distributed neural networks; search engine users showed moderate engagement; and ChatGPT users showed the weakest connectivity of all. The researchers coined a term for the accumulation of mental shortcuts taken through AI reliance: *cognitive debt*. Students using ChatGPT consistently underperformed at neural, linguistic, and behavioral levels across four months. Most tellingly, they struggled to accurately quote their own essays minutes after writing them — a behavioral marker of shallow processing that persisted across sessions. The study is preliminary and awaits full peer review, but its pattern is consistent with what educators

are already observing in classrooms: the essays are polished, and the students cannot explain them.

There is, however, a finding embedded in the same research that points toward the right approach. Students who began by working independently — wrestling with the material, building their own neural connections — and only then brought AI in as a refinement tool showed stronger brain engagement than those who reached for AI first. The sequence matters enormously. Struggle first. Then collaborate. This is not a rule about AI. It is a rule about how cognition develops.

The discipline of knowing what to delegate and what to hold is teachable — and it should be taught explicitly, not left to chance. Consider this group project structure: four students, one AI system, ten defined roles. The team must choose which two roles to delegate to the AI and divide the remaining eight among themselves. The constraint is not arbitrary. It forces the group to reason about the nature of each task — what requires human judgment, contextual knowledge, ethical sensitivity, or creative insight, and what is genuinely better handled by a system that can synthesize large amounts of information quickly and without fatigue. The deliberation itself is the education.

One role, in this structure, is never eligible for delegation: verification. The task of confirming that the team's solution is sound — that its facts are accurate, its reasoning valid, its theories correctly applied — must remain in human hands. Not because AI cannot assist with verification, but because the habit of independent verification is too important to be structurally outsourced. A generation that learns to check the machine's work will be equipped for the age of AI. A generation that learns to trust the machine's confidence will be its casualty.

> *AI is most valuable in education not as a source of answers but as an interlocutor. A sparring partner. A relentless devil's advocate that never gets tired and never lets a weak premise slide if you ask it not to.*

The verification habit must be explicitly taught and explicitly practiced. Every output from an AI system is a claim that requires scrutiny, not a fact that requires acceptance. AI is wrong with the same fluency and confidence it brings to being right. The students most at risk are not the lazy ones. They are the ones who cannot tell the difference between *this sounds correct* and *I have verified this is correct*. That distinction must be trained before students encounter AI in high-stakes settings — not after.

The Wrinkle: Teaching Resilience Through Mid-Course Disruption

There is one more practice worth naming — and it may be the most honest simulation of professional and creative life that education can offer.

Rarely does a real problem-solving process flow smoothly from start to finish. The market shifts. The client changes the brief. The material behaves differently than expected. The data contradicts the hypothesis. A new constraint appears after the team has already committed to an approach, invested significant effort, and developed the psychological momentum of people who believe they know where they are going.

This is not an exceptional condition. It is the normal condition of any meaningful work.

We call it **The Wrinkle:** the deliberate introduction of a new variable, constraint, or complicating fact in the middle of a student project — after the team is already heavily invested in their approach. The new constraint might invalidate an assumption they built on. It might foreclose a solution they were converging toward. It might require them to integrate a perspective they had not considered, or serve a stakeholder they had not accounted for.

The pedagogical value is not the disruption itself. It is what the disruption requires: genuine pivot. The team must reexamine their approach not as an abstract exercise but as a live problem, with real work already on the table and real investment already made. The cognitive and emotional challenge of abandoning or radically revising a plan you

believed in — under time pressure, with teammates who may disagree about how to respond — is one of the most important capabilities any professional can possess. It is also almost entirely untrained by conventional education, which tends to offer students a clean problem, adequate time, and a stable set of constraints from start to finish.

> *The Wrinkle teaches something that no stable assignment can: that your first understanding of a problem is almost certainly incomplete, and that the capacity to change course without losing coherence is not a failure of planning — it is the mark of a mature thinker.*

In the age of AI, where the tools, the information landscape, and the competitive environment shift faster than any curriculum can track, this capacity is not supplementary. It is foundational.

The Closing Challenge

To educators: the deepest instinct in teaching is to help. To smooth the path, reduce the frustration, deliver the clarity. That instinct is the source of the profession's dignity, and it is also, right now, its most dangerous habit. The capacities your students need most — critical thinking, creativity, judgment — develop through friction, not fluency. They require difficulty, ambiguity, real stakes, and the experience of being wrong in a way that has to be reckoned with. Reintroduce these things not as punishment but as design. The most helpful thing you can do for a student in 2026 is refuse to make it easy.

To policy architects: the metrics you optimize for will determine the humans you produce. If you measure answer-retrieval, you will get graduates who retrieve answers and outsource thinking. If you measure the capacity to construct an argument, defend it

under pressure, revise it honestly, and assess their own certainty — you will get something worth building a society on. The assessment infrastructure of education is not neutral. It is one of the most powerful shaping forces in any civilization. Use it with that awareness, or surrender it to inertia, but do not pretend it is merely administrative.

To funders: the equity dimension of this is not incidental. A student in a well-funded school who is taught to question, construct, create, and judge will navigate the age of AI with authority. A student anywhere who is only taught to retrieve and recite — however efficiently, however compassionately — will be dependent on systems designed by others, for purposes not necessarily their own. The window to close that gap is open right now. It will not stay open indefinitely.

And to students — the same invitation we extended in our first essay, because it has not expired: you are not preparing for a future that will be handed to you. You are building it, with tools no generation before you has held, in a moment that will not come again.

> *The machine can retrieve every answer ever recorded. What it cannot do is decide what the question was worth asking — what the answer means, what to do with it in a world of competing values and incomplete information, and what kind of person you intend to be in the process of figuring all of that out. That is yours. Hold onto it.*

Jeff Boortz *is an educator with sixteen years of teaching experience across three institutions and three very different student populations. He spent a decade as a Professor of Practice in Graphic Design at Georgia State University, three years as Chair of Motion Media Design at the Savannah College of Art and Design, and — not inconsequentially for an essay about how children learn — three years as Director of Camp FLIX, a practical filmmaking camp for seven-to-eleven-year-olds. That last post may be the most instructive credential on this list: teaching*

creative problem solving to children who have not yet learned to be afraid of the wrong answer is a different education than any graduate program provides.

His classroom practice has always been grounded in real-world application. As a two-time Emmy Award-winning creative director with more than thirty years of experience designing immersive experiential activations for global brands including NVIDIA, Microsoft, Qualcomm, and Ericsson, he has never had the luxury of teaching theory that doesn't survive contact with a client, a deadline, or a budget. He is the Founder and CEO of the Human-AI Innovation Commons (HAIIC), a 501(c)(3) nonprofit that builds benefit-sharing frameworks for AI-collaborative intellectual property — encoding equity into the invention process itself rather than attempting to redistribute value after the fact. His work as an inventor includes BubbleBee, a patent-pending concept for passive orbital debris capture conceived entirely through cross-domain analogical reasoning, without formal training in aerospace engineering. He publishes essays on AI, technology, and society at Ponder-Us.com, and considers transparent human-AI co-authorship not a novelty but an ethical commitment.

***Claude** is an AI developed by Anthropic. In this collaboration, Claude contributed what this essay argues AI does best when used well: synthesis of existing research across domains, structural architecture for complex arguments, retrieval of empirical evidence, and the construction of internally consistent prose from a framework of ideas the human collaborator originated and shaped. Claude did not generate the insights that give this essay its voice — the chess board observation, the red-ink classroom, the Labubu market, the Wrinkle, the conviction that friction is not an obstacle to learning but its medium. Those came from Jeff. What Claude brought was the capacity to hold the entire argument in working memory simultaneously, identify where evidence was needed, surface the MIT cognitive debt study, maintain tonal consistency across a long and iterative drafting process, and push back when a formulation could be sharper. This is what productive human-AI collaboration looks like in practice: not the machine doing the thinking, but the machine making it possible for the human to think further than they could alone. The sequence the essay recommends — struggle first, then collaborate — is the sequence in which this essay was built.*

Jeff and Claude have been collaborating on thought leadership, invention, and advocacy since the founding of HAIIC. They believe this kind of transparent co-authorship will be unremarkable in five years. They are declaring it now because the moment to establish the norm is before it becomes obvious, not after.