

AI and the Next Layer of Human Work

A Position Paper for Students, Faculty, and Administrators

James M. Hyman

Department of Mathematics, Tulane University

New Orleans, LA 70118, USA

`mhyman@tulane.edu`

May 25, 2026

Abstract

Concern that artificial intelligence will displace human work and weaken human capacity is widespread across higher education. This position paper argues that AI fits a long historical pattern: every powerful tool removes some forms of labor while shifting the value of judgment, taste, design, and responsibility. Photography did not end painting; calculators did not end mathematical thinking; programming libraries did not end software development; spellcheck did not produce good writers. In each case, skill did not disappear, it relocated. Drawing on six historical analogies and contemporary evidence on generative AI's effects on productivity, learning, and labor markets (Noy and Zhang, 2023; Brynjolfsson et al., 2025; Dell'Acqua et al., 2023), the paper develops a single thesis: as AI becomes more capable, the human responsibility to direct it, evaluate it, and accept consequences for it grows rather than shrinks. The paper closes with concrete guidance for three audiences. Students learn five practices that preserve learning while using AI well. Faculty receive three principles for assessment redesign and a framework for deciding what AI changes about a discipline. Administrators receive a four-part framework covering policy, infrastructure, faculty development, and academic integrity. Across all three audiences, the paper warns against two equal and opposite mistakes: prohibition that protects the old form of the work without asking what the work is for, and permissiveness that treats a powerful tool as a feature to deploy rather than a responsibility to govern. The danger ahead is not only that AI will replace people. The deeper danger is that people will use AI passively and let their own capacity to think weaken.

Opening

Many people are afraid that artificial intelligence will take over their work, their creativity, and eventually large parts of their lives. Artists worry that AI-generated images will make human art unnecessary. Programmers worry that AI-generated code will make human developers obsolete. Teachers worry that students will stop learning. Writers worry that language itself will become cheap.

These fears are not foolish. AI is genuinely disruptive, and it will change how many people work (Frey and Osborne, 2017; Acemoglu and Restrepo, 2019).

But AI is not the first technology to frighten people by redefining skill. Again and again, powerful tools have removed some forms of labor while making judgment, taste, design, and responsibility more important, not less. The real question is not whether AI will change the work. It will. The better question is what kind of human skill will matter more once it does.

This paper develops that question through six historical analogies and then turns to what the answer means in practice for students learning with AI, faculty teaching with AI, and administrators setting policy for AI in higher education.

1 Photography Did Not End Art

When photography became widely available in the nineteenth century, painters had reason to worry. If a camera could capture a likeness quickly and cheaply, why would anyone pay for a painted portrait ([Newhall, 1982](#))?

That concern had some validity. Photography did reduce demand for certain kinds of realistic visual copying, and portrait painters whose livelihoods depended on commissioned likenesses faced genuine economic hardship. The transition was not painless, especially for those caught in it. But painting did not disappear. Instead, it transformed. Painters moved more deeply into color, abstraction, emotion, and personal vision, into territory a camera could not follow. Photography also became an art form of its own, and today many painters use photographs as part of their process, studying light, testing composition, and using the camera as a tool for seeing.

The camera did not destroy the artist. It changed what the artist was needed for.

AI may do something similar. It will likely reduce the market for quick, routine commercial images and flood the world with cheap visual material, which is a real and legitimate concern. But it cannot replace what art actually requires: attention, memory, pain, humor, style, taste, and the human need to make meaning. An artist's value has never been simply the ability to produce an image. It is the ability to see.

2 Calculators and the Question of Understanding

The calculator offers a lesson that educators know well ([Cuban, 1986](#)).

When calculators became common, people worried that students would stop learning arithmetic. That concern was not wrong. A student who reaches for a calculator before developing number sense may never fully acquire it ([Sweller, 1988](#)). But calculators did not end mathematics. They shifted the level at which mathematical work takes place.

We do not value mathematicians because they can multiply long numbers by hand. We value them because they can reason, model, prove, and interpret, and because they know when an answer makes sense even before they verify it. A calculator can give an answer without understanding the question. That gap is exactly where human judgment lives.

AI creates the same tension. Used carelessly, it allows people to skip the very learning that would let them evaluate what the tool produces ([Sparrow et al., 2011](#)). Used well, it frees them from routine work and helps them think at a higher level. The goal is not to use AI instead of understanding. The goal is to use AI only after building enough understanding to judge what it gives back.

3 Programming Libraries Did Not End Programming

The history of software development shows the same pattern.

In the early days of computing, programmers often had to write everything themselves. A fast Fourier transform might require weeks of work: learning the algorithm, writing it by hand, testing, debugging, and carrying it forward on punched cards (Cooley and Tukey, 1965). Today, an FFT is one line of code. Linear solvers, optimization routines, statistical packages, and machine-learning tools are now standard building blocks.

That did not make programmers unnecessary. It changed what good programming means.

Syntax is necessary, but it is not the heart of the work. A good programmer formulates the problem, chooses the right tools, tests the answer, understands failure modes, and builds something that works reliably under real conditions. AI is the next step in this progression. It can draft code, summarize documentation, find bugs, and produce prototypes quickly (Brynjolfsson et al., 2025). But the human still has to decide what problem is being solved, what assumptions are being made, whether the output is trustworthy, and what consequences follow from getting it wrong.

The skill does not disappear. It relocates.

4 Power Tools and the Limits of Force

A power saw does not make someone a carpenter. It lets a skilled carpenter work faster. It also lets an unskilled person make mistakes faster.

The tool has force but no judgment. It does not know where the cut belongs. It does not know how the pieces fit together, or whether the structure will hold. Skill is what fills that gap, and the more powerful the tool, the more consequential the judgment becomes.

This may be the most practical frame for thinking about AI. AI is a power tool for thought. It can move quickly, produce a great deal, and save real time (Noy and Zhang, 2023). But speed is not wisdom. Fluency is not truth. A confident answer is not necessarily a correct one (Bender et al., 2021). Dell’Acqua and colleagues found that knowledge workers using AI inside its capability range produced substantially better work, while workers using AI on tasks just outside that range produced substantially worse work without realizing it (Dell’Acqua et al., 2023). The boundary between the two is rarely marked. As the tool grows more capable, the human responsibility to direct it, evaluate it, and correct it grows with it (Russell, 2019).

5 Operational Skill vs. Mechanical Knowledge

When cars were new, drivers often needed to understand the engine. Mechanical knowledge was part of driving, because if the car broke down, the driver was often the person who had to fix it.

Today, excellent drivers need not know how a transmission works. Driver education does not begin with pistons and fuel injection. It begins with steering, braking, merging, watching the road, and judging distance: the operational competencies that keep people safe. Mechanical knowledge still matters deeply, but it belongs to a different kind of practitioner.

AI is moving many fields in this direction. Not every user needs to understand how a large language model works at the architectural level (Bommasani et al., 2021). But every serious user needs to know how to ask good questions, detect errors, check outputs, protect privacy, and recognize when the stakes are high enough to require human judgment rather than automated response (Weidinger et al., 2021). Knowing how to use the tool safely is becoming as important as knowing how it is built.

6 Spellcheck Did Not Make Everyone a Writer

Writers have already lived through a smaller version of this change (Carr, 2010).

Spellcheck did not make everyone a good writer. Grammar tools did not make everyone clear, honest, or persuasive. They removed surface errors while leaving the real work intact. A sentence can be grammatically correct and still be dull. A paragraph can be polished and still say nothing.

AI writing tools are far more powerful than spellcheck, but the underlying principle holds. They can assist with drafts, structure, and revision. They cannot decide what is true, supply what the writer has actually lived, or provide courage, taste, or moral clarity. A good writer is not someone who knows grammar. A good writer knows how to create pressure, rhythm, image, and emotional force. AI can help manage language. It cannot replace the need to mean something (Marcus and Davis, 2019).

7 What This Means for Institutions

The analogies above describe individual skill. But students, faculty, and administrators all operate within institutions, and institutions face a version of the same challenge at larger scale.

An institution that simply bans AI use is making the same mistake as a school that took away calculators. It protects the old form of the work without asking what the work is actually for. An institution that allows unrestricted AI use without building the surrounding judgment skills is like a driver education program that hands students the keys without teaching them to watch the road.

The more productive question is what institutions must now teach that they did not need to teach before: how to evaluate AI outputs critically, how to detect confident-sounding errors, how to maintain academic integrity not by avoiding the tool but by understanding its limits, and how to use AI to go further than the work would otherwise allow rather than to substitute for it (Miao and Holmes, 2023; Cotton et al., 2024).

Faculty designing courses and administrators setting policy face the same underlying problem: how do we preserve the learning that makes powerful tools useful rather than dangerous? That is not a technology question. It is an educational question, and it is one that human judgment must still answer.

The three sections that follow translate that general claim into concrete practice for the three audiences this paper addresses.

8 Implications for Students

Students are the audience with the most immediate stake in the change, and also the audience that bears the greatest risk if the tool is used without thought. A confident-sounding draft, a finished problem set, a polished email arrive in seconds. The hidden cost is that the learning that should accompany the work does not happen, and the student is left holding a credential without the underlying capacity it was meant to certify.

The argument here is direct: AI is a legitimate part of learning, but only after the foundational work is in place. The earlier in the curriculum a topic sits, the more important it is to build the relevant skill before reaching for the tool. The arithmetic, the proof technique, the simple program, the first essay, these are not artifacts to be produced as cheaply as possible. They are scaffolding for everything that follows.

Five concrete practices follow from this.

- **Use AI after, not instead of, the first attempt.** Write the draft, work the problem, struggle with the proof, then bring AI into the revision. The first attempt is where the learning lives. The AI revision is where the work improves.
- **Treat AI output as a hypothesis, not an answer.** Generative models produce confident-sounding text whether they are correct or not. A student who accepts an answer because it sounds right has stopped being a student.
- **Cross-check with sources.** If AI gives a citation, find the actual paper. If it gives a fact, find a corroborating source. The tool is fluent, not accurate.
- **Ask the tool to explain its reasoning, then evaluate the explanation.** Where the explanation is weak, the answer is suspect. Where the explanation is clear, the student has gained something the answer alone would not have given.
- **Be honest about use.** The credibility of the work depends on the credibility of the worker. Hidden AI use is a small lie that compounds into a larger problem about what the credential represents.

The student who learns to use AI well will leave school with two capacities employers and graduate programs actually value: the ability to do the underlying work without the tool, and the ability to use the tool to go further than that work alone would allow. The student who uses AI to avoid learning will arrive at the next stage of life with the credential and without the capacity, which is a position that becomes harder to recover from with each passing year.

9 Implications for Faculty

Faculty face a different problem. They must decide what AI changes about their courses, what it does not change, and how to redesign assessment so that the learning they care about can still be observed.

The first decision is what AI changes about the discipline itself. In a programming course, the answer is significant. Routine code is now cheap, so the value of the course shifts toward problem formulation, debugging strategy, system design, testing, and judgment about when to trust the tool.

In a writing course, the answer is more subtle. AI produces fluent prose, but it does not produce a perspective, a question, or the specific judgments of a writer who has actually thought about the topic. The course can lean into exactly those capacities. In a mathematics course, the answer is again different. AI can produce derivations of variable quality and frequently introduces errors of reasoning. The course remains responsible for the conceptual understanding that lets a student detect those errors.

The second decision is assessment. Three principles help.

- **Assess process, not only product.** A polished final document tells you little about who wrote it. Drafts, in-class work, oral defense, and incremental submissions tell you more. Where high-stakes evaluation is needed, in-person assessment carries weight that take-home work no longer does.
- **Build assignments that require integration with experience.** Ask students to apply concepts to their own data, their own observations, their own community, their own context. AI is good at general patterns and weak at specific situated knowledge.
- **Make AI use explicit and reflective.** Rather than ban the tool or ignore it, ask students to document how they used it, what they accepted, what they rejected, and what they learned. The reflection is itself an assessable learning outcome and a form of digital literacy that will matter in their professional lives (Mollick, 2024).

The third decision is what to teach explicitly that previously could be left implicit. Critical evaluation of confident-sounding text, the difference between fluency and truth, the discipline of source verification, and the ethics of attribution are now first-class learning outcomes in nearly every field. They were always part of an educated person's toolkit. They are now part of the syllabus.

Faculty who treat AI as a discipline-specific question, rather than a generic technology problem, will find that the redesign is less disruptive than it first appears. The deep learning goals do not change. The pathways to them do.

10 Implications for Administrators

Administrators set the conditions inside which students and faculty operate. Their decisions about policy, infrastructure, faculty development, and academic integrity shape what happens in every classroom on campus, often more than any individual syllabus.

Three categories of decision deserve particular attention.

Policy. A blanket ban on AI is brittle, hard to enforce, and increasingly out of step with the working world students are preparing to enter. A blanket permission abdicates the institution's responsibility to define what a degree means. The productive middle ground is a framework that distinguishes contexts: where AI is appropriate and how its use should be acknowledged, where AI is not appropriate and why, and what happens when use crosses those lines. The framework must be discipline-aware. A computer science department, a writing program, and a clinical training program reasonably differ on what constitutes appropriate use, and the institutional policy should accommodate that variation rather than flatten it.

Infrastructure. Equity is a real concern. Paid AI tools outperform free versions, and students with resources can access capabilities that students without resources cannot. Institutions that provide centralized access to capable AI tools, with appropriate privacy and data protections, help equalize what is rapidly becoming a baseline expectation rather than an enrichment. The cost is real, but the cost of doing nothing is a quiet stratification of student outcomes by family means.

Faculty development. The single most consequential investment an administration can make is in helping faculty understand what AI does in their specific discipline, what assessment redesign is needed, and what conversations to have with their students. This is not a one-time training event. It is an ongoing conversation that needs sustained institutional support, including release time, peer-learning communities, and recognition that course redesign for AI is a legitimate and substantial form of academic work.

A fourth concern cuts across all three categories. Academic integrity processes built for pre-AI plagiarism do not work well for AI use. Detection tools are unreliable. Hyatt et al. (2025) reported false-positive rates in the low single digits for the detectors they tested, but rates from other studies have run substantially higher, particularly for short work and for writers whose prose falls outside the detector’s training distribution. False positives create injustice, and the underlying question, what counts as the student’s own work, is genuinely harder than it used to be. Institutions need to revisit their integrity frameworks with the same seriousness they applied to internet-era plagiarism a generation ago. The goal is not perfect detection. It is a fair and educational process that treats students as people learning to navigate a new technology rather than as suspects.

The administrator’s hardest job is to resist two equal and opposite mistakes: prohibition that protects the old form of the work without asking what the work is for, and permissiveness that treats a powerful tool as a feature to deploy rather than a responsibility to govern. The middle path is harder than either extreme, but it is the only path that takes seriously what universities are actually for.

11 The Real Risk

None of this means AI is harmless.

Some jobs will be disrupted, and routine tasks will lose economic value. People whose livelihoods depend on that work will be genuinely displaced. Artists may lose commissions; programmers may find that entry-level work has changed (Dell’Acqua et al., 2023). Some students will use AI to avoid the very learning that would let them use it well, and some institutions will let them. The point is not that everything always works out for everyone.

The displacement is not symmetric. When painting moved from realistic representation to abstraction, a painter who had built one set of skills could plausibly build another. When mathematical work moved from manual computation toward modeling and proof, mathematicians who had developed number sense were positioned for the higher-level work that followed. The historical pattern is encouraging at the aggregate level, but it does not guarantee that the same individuals displaced from the old work can access the new work. The new skills AI rewards, such as problem formulation, evaluation, and judgment under uncertainty, often require more education and more time to acquire than the routine tasks AI replaces. The aggregate pattern can be true and the individual outcome still grim, and a society that takes the historical pattern as reassurance without investing in the

transition will see the inequality the pattern conceals.

The lesson from earlier tools is not that everything always works out. It is that the people and institutions that adapt thoughtfully often move to a higher level of practice, while those that refuse to engage may fall behind. But there is a risk on both sides. Refusing to adapt is one failure. Adopting tools so uncritically that judgment atrophies is another, and in some ways the harder one to see.

Early empirical evidence already points in that direction. [Gerlich \(2025\)](#) surveyed 666 participants and found that frequent AI use correlates with reduced critical-thinking performance, mediated by cognitive offloading. Younger users were most affected, and higher educational attainment served as a protective buffer. A preprint EEG study from MIT Media Lab compared essay-writers using ChatGPT, a search engine, or no tools; the LLM group showed the weakest brain connectivity ([Kosmyna et al., 2025](#)). The reduction in cognitive engagement persisted in a follow-up session even after the AI was removed.

The danger is not only that AI will replace people. The danger is that people will use AI passively and let their own capacity to think weaken.

12 What Matters More Now

As AI improves, some skills will matter less. Routine first drafts, boilerplate code, mechanical summaries, formulaic writing: all of these become cheaper. That is a real change, and not a trivial one for people whose work currently consists of producing them.

But other capacities become more valuable. The ability to ask a question precisely enough that a powerful tool can help answer it. The ability to evaluate an answer skeptically rather than accepting it because it sounds right. The ability to revise with purpose, knowing not just that something is wrong but why, and what would make it better. The ability to notice what is missing, to bring genuine experience to a problem, and to accept responsibility for the result.

AI can produce. Human beings must still choose.

13 Conclusion

Each technology in this history carried a version of the same fear, and each time the fear was partly right and partly wrong. Photography changed what painters needed to do. Calculators changed the level at which mathematical thinking begins. Programming libraries moved developers from writing every routine to building larger systems. Power tools made skilled judgment more consequential, not less. Spellcheck removed surface errors and left the real difficulties of writing exactly where they had always been.

AI belongs in this history. It is not simply a threat and not simply a miracle. It is a powerful tool that removes some kinds of work and shifts the value of others. Like every powerful tool before it, it rewards the people and institutions that learn to use it while keeping their judgment, their standards, and their responsibility intact.

The future will not belong to those who pretend AI does not exist. It will not belong to those who let AI think for them. It will belong to those who understand what powerful tools can do, know

what they cannot do, and remain clear about which decisions must stay in human hands.

Acknowledgments

The author thanks colleagues at Tulane University and the broader applied mathematics community for conversations that shaped this position paper. Earlier drafts were refined through the Manuscript Review Enhancement Protocol cycle, including SPRE prose refinement, narrative-thread audit, and multi-panel harsh review.

References

- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019. doi: 10.1257/jep.33.2.3.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’21)*, pages 610–623, 2021. doi: 10.1145/3442188.3445922.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, et al. On the opportunities and risks of foundation models. arXiv preprint, 2021. arXiv:2108.07258.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025. doi: 10.1093/qje/qjae044. Earlier version: NBER Working Paper No. 31161 (2023), arXiv:2304.11771.
- Nicholas Carr. *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton & Company, New York, 2010. ISBN 978-0-393-07222-8.
- James W. Cooley and John W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–301, 1965. doi: 10.1090/S0025-5718-1965-0178586-1.
- Debby R. E. Cotton, Peter A. Cotton, and J. Reuben Shipway. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2):228–239, 2024. doi: 10.1080/14703297.2023.2190148.
- Larry Cuban. *Teachers and Machines: The Classroom Use of Technology Since 1920*. Teachers College Press, New York, 1986. ISBN 978-0-8077-2792-3.
- Fabrizio Dell’Acqua, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraye, François Cadelon, and Karim R. Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Working Paper 24-013, Harvard Business School, 2023.
- Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, 2017. doi: 10.1016/j.techfore.2016.08.019.

- Michael Gerlich. AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1):6, 2025. doi: 10.3390/soc15010006. Correction at 10.3390/soc15090252 (Sep 2025); conclusions unaffected.
- Jon-Philippe K. Hyatt, Elisa Jayne Bienenstock, Carla M. Firetto, Elizabeth R. Woods, and Robert C. Comus. Using aggregated AI detector outcomes to eliminate false positives in STEM-student writing. *Advances in Physiology Education*, 49(2):486–495, 2025. doi: 10.1152/advan.00235.2024.
- Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint, 2025. arXiv:2506.08872; MIT Media Lab; preprint, not peer-reviewed.
- Gary Marcus and Ernest Davis. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Pantheon Books, New York, 2019. ISBN 978-1-5247-4825-8.
- Fengchun Miao and Wayne Holmes. Guidance for generative AI in education and research. Technical report, UNESCO, Paris, 2023. URL <https://unesdoc.unesco.org/ark:/48223/pf0000386693>.
- Ethan Mollick. *Co-Intelligence: Living and Working with AI*. Portfolio / Penguin, New York, 2024. ISBN 978-0-593-71671-7.
- Beaumont Newhall. *The History of Photography: From 1839 to the Present*. The Museum of Modern Art, New York, 5th edition, 1982. ISBN 978-0-87070-381-2.
- Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023. doi: 10.1126/science.adh2586.
- Stuart Russell. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, New York, 2019. ISBN 978-0-525-55861-3.
- Betsy Sparrow, Jenny Liu, and Daniel M. Wegner. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043):776–778, 2011. doi: 10.1126/science.1207745.
- John Sweller. Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2): 257–285, 1988. doi: 10.1207/s15516709cog1202_4.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, et al. Ethical and social risks of harm from language models. arXiv preprint, 2021. arXiv:2112.04359.