

**USC** Center for  
Generative AI and Society



**Critical Thinking  
and Ethics  
in the Age of  
Generative  
AI in Education**

**A Critical Look into the  
Future of Learning**

# Contents

PEDRO NOGUERA	
Introduction to the Study of Generative AI Technologies in Education: A Critical Look into the Future of Learning. . . . .	1

## Supporting College Students' Critical Thinking in the Age of AI

STEPHEN AGUILAR	
Introduction to the Program . . . . .	4
Generative AI Fellows . . . . .	5
GALE SINATRA & CHANGZHAO WANG	
The Promise and Peril of Using AI in College Classrooms . . . . .	10
ERIC BUI	
Wow, "...[that's] probably an A": An Undergraduate Perspective of Generative AI in Undergraduate Education . . . . .	13

## AI in K12 Classrooms: Ethical Considerations and Lessons Learned

STEPHEN AGUILAR	
Ethics in Generative AI: Report From the Field. . . . .	21
CHANGZHAO WANG	
Lessons Learned by Teaching Middle School Students to Use Generative AI. . . . .	27

## Building the Next Generation of Generative AI Tools

WILLIAM SWARTOUT	
Generative AI and Education: Deny and Detect or Embrace and Enhance . . . . .	32
BENJAMIN NYE	
Authoring by Editing and Revising: Considering Generative AI Tools . . . . .	34

# Introduction to the Study of Generative AI Technologies in Education: A Critical Look into the Future of Learning



In today's rapidly evolving world, the role of technology in shaping our educational landscape is evolving rapidly. The digital age has provided educators and learners with countless tools designed to make teaching more engaging and learning more accessible. Artificial Intelligence (AI) is being hailed as one of the most groundbreaking developments in technology and its applications have already found their way into various sectors like healthcare, transportation, and finance. AI's potential to produce dramatic changes in education may be equally profound.

In this report we delve into the potential of generative AI technologies because the implications of this emerging technology in education are vast and largely unexplored. The traditional view of education is a system where the teacher imparts knowledge, and students passively receive it. But if we are to equip future generations for a world faced with enormous challenges and uncertainty, education must shift from being merely transactional to being transformative and empowering. It must unleash the creativity and imaginations of students, while simultaneously providing them with the academic and technical knowledge and skills needed to tackle complex problems. Generative AI may offer a means to make this transformation possible.

Imagine a classroom setting where AI doesn't replace the teacher but supports them in generating personalized learning plans for each student, based on both academic performance and individual interests. There is already evidence that generative AI can be used to devise new means for teachers to collaborate with and support their students, as well as making it possible for them to share ideas with colleagues throughout the world. By opening new approaches to teaching and learning AI could serve as a means to help students become creators of knowledge rather than merely consumers, and teachers could truly become facilitators of learning rather than transmitters of knowledge. Picture a system where the line between teachers and students blurs as both parties engage in a dynamic, ever-evolving educational journey powered by collaborative, augmented intelligence—human and artificial.

While some are largely fixated on the potential problems created by the possibilities of plagiarism and copyright infringement, we think should we focus on possibilities for using generative AI technologies in education because they may help us in tackling

the problems, we face which are likely to become more complex and intractable in the years ahead. The issues we face today—climate change, inequality, global health crises, war and the ongoing refugee crisis—are not isolated but interconnected. They require our educational institutions and industries to produce a generation of problem solvers who can think critically, adapt, and contribute creatively to a multifaceted and ever-changing landscape. The development of solutions must evolve in sophistication and adaptability. Generative AI can serve as a potent tool to nurture holistic thinkers, equipping them not just with knowledge but also with the tools to generate new ideas and solutions.

However, as we move forward, we are well aware of the need to approach this transformation responsibly, grounded in ethics that put the interest of human society first. While technology has the power to revolutionize education, it also bears the risk of amplifying existing inequities and creating new forms of exclusion. Ethical considerations, such as data privacy, algorithmic bias, and the digital divide, should be central to the implementation of generative AI in educational settings. Furthermore, we must ensure that such technologies are employed to augment human capabilities, not to replace them, to preserve the inherently relational and emotional aspects of teaching and learning.

This report is an invitation for educators, policymakers, technologists, and learners to consider how generative AI can contribute to the future of education. It aims to lay down a foundation upon which we can start building an educational ecosystem that is dynamic, inclusive, and profoundly human, despite being significantly aided by artificial intelligence. It is one of many that will be written by experts within the USC Rossier School of Education, in collaboration with the Institute for Creative Technologies, and others.

Let us embark on this journey with an openness to disrupt the traditional, yet with a reverence for the inherently human aspects of education that no technology should ever replace.

A handwritten signature in black ink that reads "Pedro Noguera". The signature is written in a cursive, flowing style.

Pedro Noguera

# Supporting College Students' Critical Thinking in the Age of AI



STEPHEN AGUILAR

# Introduction to the USC Center for Generative AI Fellows Program



Generative AI technologies, including ChatGPT, are poised to transform education. The USC Center for Generative AI is at the forefront of guiding post-secondary institutions in navigating this shift. Our Generative AI Fellows Program is pivotal in realizing this vision.

Our inaugural fellows have diverse academic backgrounds, including rhetoric, composition, and pedagogy. They advocate for multilingual approaches and have integrated technology into equitable education. With experts in poetry, literature, and creative writing, our program is represented by experts who are committed to preserving human creativity in the face of expanding AI initiatives. They are also committed to not simply ignoring generative AI technologies, but instead seek to find ways that they can be **embraced**, so the undergraduate experience is subsequently **enhanced**.

Consequently, a core objective of our fellows is to better understand how generative AI impacts undergraduate writing education. We focus on undergraduate writing because it forms the basis of critical thinking and analysis. By integrating AI into this area, we aim to equip students to use AI tools effectively, enhancing their learning and communication skills. Beyond academics, we recognize AI's broader impact on the future. We thus aim to prepare undergraduates for the evolving job landscape shaped by AI. They won't just be observers; they'll be at the forefront of this shift. Our fellows, with their varied interests and deep expertise, will ensure that our AI initiatives resonate with real human experiences.

The USC Center for Generative AI Fellows Program is more than a tech initiative. It's a push to redefine AI in education and society. By focusing on undergraduate writing and understanding AI's role, we're preparing for a future informed by a multitude of AI-enhanced tools. Join us as we merge education, AI, and human creativity.

## Generative AI Fellows



**LIZ BLOMSTEDT** is an Assistant Professor (Teaching) in the USC Writing Program. Before coming to USC, she was the Assistant Director of the Warren College Writing Program at UC San Diego. She earned her MA and PhD in Rhetoric, Composition, and Pedagogy from the University of Houston, and her BA in Rhetoric and Writing from the University of Texas at Austin. Liz has taught first-year writing courses in many different contexts, and her research focuses on composition pedagogy. She is particularly interested in multilingual pedagogies, the impact of standardized testing on student writing habits, and writing assessment.



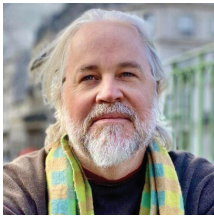
**JEN SOPCHOCKCHAI BANKARD** is an Associate Professor (Teaching) in the Writing Program, specializing in teaching with technology and equitable access to higher education. She earned her bachelor's degree in English with honors in expository writing from Brown University and a Ph.D. in English with a certificate in cinema studies from Northeastern University. At USC, she has taught a variety of upper and lower division courses, with Writing 340: Advanced Writing for Natural Sciences being the most recent. She currently co-chairs the Committee on Information Services (CIS), a standing committee jointly appointed by the Office of the Provost and the Academic Senate to represent the information technology needs and perspectives of USC's faculty. In 2019, she helped restructure the Norman Topping Student Aid Fund, a scholarship and academic support program for first-gen and other high financial need students. She continues to support first-gen student programming at Topping and Dornsife College. When she's not teaching, she writes film and television reviews for [The Long Take](#) and hosts [The Long Take Review](#), a film podcast that uses rhetoric and composition to analyze the Oscars race.



**NIK DE DOMINIC** is an Associate Professor and the Associate Director of Curriculum in the Writing Program. He co-founded and co-directs the [Dornsife Prison Education Project](#). The author of the full length collection of poems *Goodbye Wolf* (The Operating System '20) and the chapbook *Your Daily Horoscope* (New Michigan Press '15), De Dominic has had poems and essays appear in *Guernica*, *DIAGRAM*, *Harpur Palate*, *Poetry Daily*, *Interim*, *Bennington Review*, and elsewhere. He is the Poetry Editor of the *New Orleans Review* and founding editor of *The Offending Adam*, a digital chapbook publisher. He holds an MFA in creative writing from the University of Alabama, where he was an Alabama Prison Arts + Education fellow, and he is currently completing a DSW at USC, focusing on increasing access to post-secondary education for people who are incarcerated.



**P.T. MCNIFF** is an Associate Professor (Teaching) in the Writing Program at the University of Southern California. He has been teaching both first-year and advanced writing courses for fifteen years; for the last eight years, he has also co-taught a summer workshop in creative writing for high school students. P.T. has served as Secretary for the Dornsife Faculty Council and for three years was the co-chair of the USC Academic Senate's Part-Time Faculty Affairs Committee. In addition, he has served on numerous other governance, mentoring, and hiring committees within the Writing Program. He received a masters in fiction writing from USC and a bachelors in English & Communication from the University of Pennsylvania. In his spare time, he co-hosts the Long Take Review podcast, which brings principles and concepts of rhetoric & composition pedagogy to discussions of movies and the Academy Awards. He also writes fiction, non-fiction, and overly long text messages.

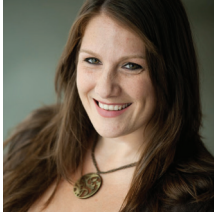


**ROBERT WALLER** (he/him) is Director of the Writing Center and an Associate Professor (Teaching) in the Writing Program. Professor Waller received his bachelor's degree from Duke University, where he studied oral history and music, and has a degree in creative writing from USC. Robert has over 20 years of experience teaching writing with a particular focus on using writing as a tool to explore and reveal the creative process. At USC, he developed the popular Special Topics WRIT 340 course, Writing for Visual and Performing Artists, that aims to help students articulate their own aesthetics and through this process expand their understanding both of writing and the meanings and purposes driving their work. In addition to his work at USC, Waller is also an active performing songwriter and musician who has toured internationally and had his songs appear in TV and film. Robert lives in Highland Park with his wife, three children, and two cats.



**TAMARA LUQUÉ BLACK** joined the Writing Program in 2014, bringing with her more than a decade of teaching experience and a PhD in Sociology from UCLA. Now a proud Trojan, Tamara has taught a variety of lower- and upper-division composition classes, including Identity & Diversity, Technology & Social Change, and Advanced Writing for the Social Sciences. In 2018, she created a new course, Advanced Writing for Economics, which has been offered continuously since it debuted. Like her teaching, Tamara's research and service activities emphasize pedagogy and student support. Her research focuses on grading contracts, rhetorical empathy, and humane composition pedagogy, and she enjoys coaching Fulbright Grant and Marshall Scholarship applicants through USC Academic Honors & Fellowships and serving as a faculty advisor to two student organizations. Tamara lives in South Los Angeles with her husband. She loves animals, hiking, music, and true crime.





Born and raised in Brooklyn, New York, **JESSICA PIAZZA** now lives in Los Angeles where she teaches in the Writing Program at the University of Southern California. She holds a Ph.D. in Literature and Creative Writing from the University of Southern California, a Masters in English Literature and Creative Writing from the University of Texas, Austin and a Bachelors in Journalism from Boston University. Jessica is the author of three poetry collections: *Interrobang* (2013) and *Obliterations* (2015) from Red Hen Press and *This is not a sky* (2014) from Black Lawrence Press. She writes poetry, fiction and nonfiction, and her main academic interests are social science-based approaches in rhetoric and composition pedagogy, cognitive poetics, prosody, and contemporary American literature.



**PATRICIA TAYLOR** is an Assistant Professor of Teaching in the Dornsife Writing Program at USC, where she teaches students how to use failure to their advantage in the writing process. She received her PhD from the University of Connecticut, where she specialized in early modern religious literature and developed an obsession with Milton's *Paradise Lost*. Her research interests include adaptation studies, religion, collaborative authorship, and failure in the writing classroom. She brings writing pedagogy and the study of religion together by advising *The Coastlander*, the Christian ecumenical journal written by and for USC students.



**ANDREW DE SILVA** is an associate professor teaching writing and critical reasoning at the University of Southern California. His courses invite students to explore the links between twenty-first century life and pre-twenty-first century history, interrogating how we arrived at our cultural, economic, and political moment. His forthcoming CCCC talk, "An Abundance of Songs, an Abundance of Paragraphs: What Napster's Disruption of Music Can Teach Us About ChatGPT's Disruption of Writing," distills his early work in the Generative AI space.



**DANIEL PECCHENINO** came to USC in 2013 and has taught primarily in the Economics and Human Values thematics of Writing 150. He has also been one of the writing instructors for USC's Warrior-Scholar Project course since its inception in 2015. His research focuses on student receptiveness to different modes of assessment and feedback. He has been elected to a range of officer positions in faculty governance bodies at both the school and university levels. In 2022–2023 he served as President of the Academic Senate. In that role he led efforts to strengthen shared governance via the school Faculty Councils, increase compensation transparency, and improve policies around misconduct investigations.



**DEBORAH M. SIMS** is an Associate Professor at the University of Southern California. She specializes in writing, literature, and American cultural studies, along with education and teacher training. She is known for developing [The Empathy Project](#), a curricular unit that trains undergraduates in synthesizing emotion and critical reasoning so that they can practice empathetic problem solving. She has also designed and implemented professional development events for incoming faculty lecturers and graduate students at USC for many years. In 2022, Dr. Sims was awarded a Sustainability Across the Curriculum Grant for her course in Writing and the Environment, as well as an Advancing Scholarship in the Arts and Humanities early sabbatical for her book project on the craft of teaching writing. Dr. Sims' recent [publications](#) and [speaking engagements](#) explore methods for creating meaningful learning experiences for students and navigating the post-pandemic classroom.



**MADDOX PENNINGTON** is a nonbinary writer, professor, and playwright; a citizen of the Cherokee Nation, they are originally from Tulsa, Oklahoma. They received an MFA in Creative Nonfiction from Columbia University, and their debut bibliomemoir, *A Girl Walks Into a Book: What the Brontës Taught Me About Life, Love, and Women's Work* was published by Hachette Books in 2017. Teaching interests include identity, empathy, Native studies, Queer studies, Disability studies, and monsters. Their work has been read Off-Off Broadway at the NYC FRIGID Queerly Festival and LA's T/GNC Reading Festival, the Native Voices series at the Autry Museum of the American West and the La Jolla Playhouse, and won awards at the Hollywood Fringe Festival. They recently completed a comedy pilot about haunting queer relationships with the Native American Media Alliance TV Writer's Lab in Los Angeles. You can find them online at [MaddoxKPennington.com](http://MaddoxKPennington.com).



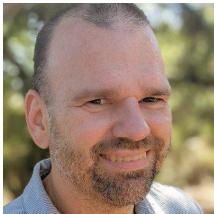
**MARK C. MARINO** is a Professor (Teaching) of Writing at the University of Southern California, where he directs the Humanities and Critical Code Studies Lab. He is also a 2023–24 Generative AI Fellow. Since 2008, he has been the Director of Communication of the [Electronic Literature Organization](#). His works include “Living Will,” “a show of hands,” and “Shields Down.” He was one of ten co-authors of [10 PRINT CHR\\$\(205.5+RND\(1\)\): : GOTO 10](#) (MIT 2013) and was a collaborator with Jessica Pressman and Jeremy Douglass on *Reading Project: A Collaborative Analysis of William Poundstone's Project for Tachistoscope {Bottomless Pit}* (Iowa Press 2015). His latest book is [Critical Code Studies](#) (MIT 2020).



**NORAH ASHE-MCNALLEY** is the Director of the USC Writing Program and a Professor of Teaching Writing at the University of Southern California. Her teaching interests include student publication, compassionate pedagogy, and the exploration of affect in writing. Most recently, she brings her love of eating and sharing to her food studies class. In her pedagogy and her administrative work as director, she cultivates compassionate listening and building capacity for empathy and connection. Outside the classroom, she cultivates her interests in cooking, mindfulness, and taking walks. She was born in Los Angeles, and raised her two children in the same home where she grew up.



**REBECCA FULLAN** is a Lecturer in the Writing Program at the University of Southern California, and a 2023-2024 Generative AI Fellow. She received her Ph.D. in English from the Graduate Center of the City University of New York. has a Master's in Theological Studies from Harvard Divinity School and an AB in Comparative Literature from Bryn Mawr College. Her teaching and research interests include Writing Across the Disciplines, Environmental and Climate Literature and Writing, Native American/Indigenous Literature and Writing, Speculative Fiction, Religion and Literature, and Pedagogy and Technology. At USC, she has taught Writing 150 in the Sustainability and Law and Social Justice thematics, and created the Papers After Class: What's Next for your Writing workshop for the Undergraduate Writer's Workshop. Outside of academia, she writes poetry, fiction, non-fiction, and plays, and currently most often makes up stories with her wife and 4-year-old daughter.



**ZEN DOCHTERMAN** is a Lecturer in the Writing Program at USC. He holds a PhD in Comparative Literature with a focus on Central American poetry. His interests include rhetoric and composition, speculative fiction, and anti-work culture.

# The Promise and Peril of Using AI in College Classrooms



Recently, Dr. Gale Sinatra and Dr. Barbara Hofer commented on the [use of ChatGPT to source scientific information online](#). Students use generative AI for a variety of purposes such as generating the first draft of an essay or creating an outline for a term paper. Another popular purpose might be searching for information about a scientific phenomenon, works by a writer or musician, or just background information for a report. ChatGPT and other Generative AI tools are becoming a popular alternative to a Google search.

However, as Sinatra and Hofer note, “ChatGPT does not search the internet the way Google does.” ChatGPT has been trained on a massive corpus of texts, then based on language patterns it learned from the training, it generates a response. As a result, generative AI tools can produce misinformation that is made up by incorporating what it learned from existing online information.

ChatGPT focuses on the language, so a response may be grammatically correct yet factually or logically problematic. Such stories have been shared by many ChatGPT users. Not only ChatGPT, art transfer tools, another type of generative AI tool, may also produce funny pictures: For example, when transferring a photo of people into cartoon style by such art AI tools, the generated picture may look gorgeous at the first glance but suddenly you find the people only have three fingers on the new image.

When thinking about how to use ChatGPT, Sinatra and Hofer outline the following concerns that users should attend to.

1. **Fabricated sources.** ChatGPT usually does not provide sources of information to its responses. If asking for sources, it may provide fabricated sources. Interestingly, fabricated sources are close to actual sources, so users should take that into account.
2. **Challenges to sourcing accurate information.** The fact that ChatGPT potentially mixes accurate and inaccurate information together makes it challenging for students to identify inaccurate information and find accurate information.

If students are going to use ChatGPT for class projects, they need to know that additional steps must be taken to verify information.

3. **Embedded limitation.** Generative AI tools, like ChatGPT, provide responses based on what they have learned from training data. The training data is the embedded limitation of these AI tools: Their knowledge is limited to the range of the training dataset; beyond that, they have no way to source information. For instance, if an AI writing tool is trained only with samples of Shakespeare's work, it can only generate writing in Shakespeare style instead of in modern English.
4. **Erosion of trust.** Because of the concern that misinformation may be returned by generative AI tools, students will have to be careful about the information obtained from these tools. If AI returns more and more misinformation, that will create an erosion of trust in these platforms and in information found online in general.

When using tools like ChatGPT, however, the following suggestions will help users make the most out of generative AI without necessarily falling victim to its limitations.

1. **Be vigilant.** First, students should be aware that information generated by AI tools may not be trustworthy, so they must be vigilant when using these tools. When seeking information online, it is convenient to get some quick and helpful responses from AI tools. AI tools may be a good starting point, but never stop there.
2. **Perform fact-checks.** When using AI tools to search information, always open a second window to conduct a fact-check. If sources of the information are provided, check the credibility of the source, such as whether the author has expertise in the field and whether the information is aligned with other trusted sources. If there's no source information, turn to a traditional search engine to find more information.



3. **Assess plausibility.** When reading about claims generated by ChatGPT, consider the plausibility of the claim. Make a tentative judgment and then be open to revising your thinking once you have checked the evidence.
4. **Evaluate the evidence.** More importantly, students should evaluate the evidence provided to support the claim; if not, check if there's evidence that supports an alternative claim.
5. **Promote digital literacy.** Higher education instructors need to improve their own digital literacy to better support their students' use of these tools in an AI-driven world. This will be a major objective of the Fellows Program supported by the Center for Generative AI and Society.

## Wow, “[that’s] probably an A”: An Undergraduate Perspective of Generative AI in Undergraduate Education



This past summer, my research position and lab allowed me a unique opportunity to interview a few professors and pick their brains on insights around traditional pedagogical methods and emerging technologies, such as ChatGPT. One interviewee, a veteran writing professor I know well, had the typical disposition of most academics I’ve talked to about the subject. We discussed student needs, historically transformative technology in the classroom, the advent of Google as a method to find data points in writing curricula, engaging in discourse in a deep meaningful way, etc.. He wasn’t a stranger to technology but acknowledged the traditional adherence to tried-and-true teaching methods likely still played a larger part in his pedagogy. He remarked that he didn’t “think generative AI could replace critical thinking at the moment.” But he certainly wasn’t closed to Generative AI—he was and is “[a]lways open to anything that would make him a better teacher”—so long as the cues from industry and academia are there. At the tail end of our conversation, he confessed that he had never used ChatGPT himself. After handing him my laptop, he prompted ChatGPT to “[w]rite an essay on the banality of evil.” Within five minutes he remarked, “[D]amn, this is pretty good—probably an A.”

For some academics and students reading this report, this probably doesn’t mean much; yes, ChatGPT excels at doing really rote things, especially at the undergraduate level—where a lot of the paradigms of higher education are inherited rather than created. For others, yes, it’s *scary* how good this technology is at the traditional pedagogical assignments. But for me, this elucidated that the dissonance between professors and students largely wasn’t technical; it was philosophical. The underlying tension—as I understand it—is rooted in the preservation of contemporary pedagogy and an unprecedented, growingly *efficient* technological tool.

Our openness to adopting new technology only goes as far as the teachers and students willing to adopt them. Of course, a software that “crowdsources” everything would do particularly well in standardized undergraduate assignments. The internet’s breadth of data and technology’s general efficiency make it easier to find data points and survive the traditional course load that most undergraduate courses demand. Every academic stakeholder had to adjust to Google, Wikipedia, and the calculator. ChatGPT and generative AI aren’t new in terms of technological evolution in education. But to some,

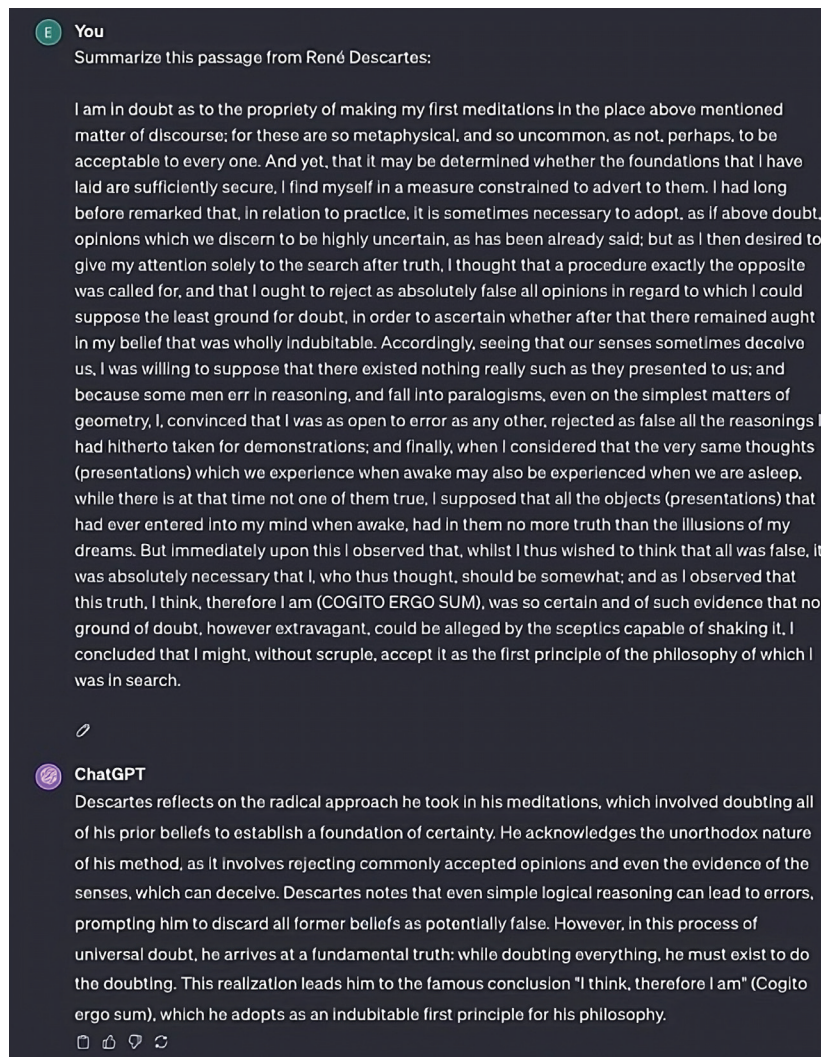
they're a threat to replace critical thinking and a promise to lose the "what" and "why" in favor of a faster "how."

I'll pause here to discuss some of ChatGPT's competencies observed in my recent and ongoing research activities, before I discuss the potential impacts of Generative AI on undergraduate education.

## ChatGPT is a consummate reader

ChatGPT is an exceptionally competent reader. Pasting passages into ChatGPT and having it summarize them is easy, especially with GPT-4's subscription model.

*Here it is, summarizing a passage from Part IV of René Descartes, "A Discourse on Method."<sup>1</sup>*



**You**  
Summarize this passage from René Descartes:

I am in doubt as to the propriety of making my first meditations in the place above mentioned matter of discourse; for these are so metaphysical, and so uncommon, as not, perhaps, to be acceptable to every one. And yet, that it may be determined whether the foundations that I have laid are sufficiently secure, I find myself in a measure constrained to advert to them. I had long before remarked that, in relation to practice, it is sometimes necessary to adopt, as if above doubt, opinions which we discern to be highly uncertain, as has been already said; but as I then desired to give my attention solely to the search after truth, I thought that a procedure exactly the opposite was called for, and that I ought to reject as absolutely false all opinions in regard to which I could suppose the least ground for doubt, in order to ascertain whether after that there remained aught in my belief that was wholly indubitable. Accordingly, seeing that our senses sometimes deceive us, I was willing to suppose that there existed nothing really such as they presented to us; and because some men err in reasoning, and fall into paralogisms, even on the simplest matters of geometry, I, convinced that I was as open to error as any other, rejected as false all the reasonings I had hitherto taken for demonstrations; and finally, when I considered that the very same thoughts (presentations) which we experience when awake may also be experienced when we are asleep, while there is at that time not one of them true, I supposed that all the objects (presentations) that had ever entered into my mind when awake, had in them no more truth than the illusions of my dreams. But immediately upon this I observed that, whilst I thus wished to think that all was false, it was absolutely necessary that I, who thus thought, should be somewhat; and as I observed that this truth, I think, therefore I am (COGITO ERGO SUM), was so certain and of such evidence that no ground of doubt, however extravagant, could be alleged by the sceptics capable of shaking it, I concluded that I might, without scruple, accept it as the first principle of the philosophy of which I was in search.

**ChatGPT**  
Descartes reflects on the radical approach he took in his meditations, which involved doubting all of his prior beliefs to establish a foundation of certainty. He acknowledges the unorthodox nature of his method, as it involves rejecting commonly accepted opinions and even the evidence of the senses, which can deceive. Descartes notes that even simple logical reasoning can lead to errors, prompting him to discard all former beliefs as potentially false. However, in this process of universal doubt, he arrives at a fundamental truth: while doubting everything, he must exist to do the doubting. This realization leads him to the famous conclusion "I think, therefore I am" (Cogito ergo sum), which he adopts as an indubitable first principle for his philosophy.

<sup>1</sup> René Descartes, "Discourse on the Method of Rightly Conducting the Reason, and Seeking Truth in the Sciences," trans. John Veitch, A Discourse on Method, June 28, 1995, <https://www.gutenberg.org/files/59/59-h/59-h.htm>.



It thoughtfully offers a distillation of key ideas and attempts to shred most of the superfluous. It sifts through intricate layers of text. And it gives context, from analyzing vast volumes of academic papers and worlds of literature. In comprehending texts, I find that it performs better than many seasoned readers.

## ChatGPT is an effective and (mostly) accurate writer

ChatGPT crafts detailed and (mostly) accurate texts. That it does so may surprise teachers unfamiliar with its capabilities. It's not just a compiler of information, but in my opinion, there is a tinge of rhetorical finesse that most undergraduate professors would appreciate from their students.

For example, when I prompted it to do something that might be considered a common undergraduate assignment, i.e., to “Analyze Descartes’ method of doubt in ‘Meditation on First Philosophy,’” while also delving into how the method informs Descartes’ epistemology and considering (and responding to) counter arguments, its first response is to provide an introduction, thesis, and an outline for the preceding paragraphs (top). I can then prompt it to write the first body paragraph (bottom). With these initial drafts done I can edit via prompting, or edit after the fact somewhere else.

**You**  
Write a thesis-driven essay based on this prompt:

"Analyze Descartes' method of doubt in 'Meditations on First Philosophy': How does this method serve as a foundation for his epistemology?"

Consider potential counterarguments and offer an effective rejoinder.

**ChatGPT**  
Title: The Method of Doubt as the Bedrock of Cartesian Epistemology

Introduction:  
René Descartes' 'Meditations on First Philosophy' marks a pivotal turn in philosophical inquiry, where he introduces the method of doubt as a means to acquire indubitable knowledge. This essay will analyze how Descartes' radical skepticism serves as a foundational tool in his epistemological framework. By systematically doubting all beliefs, Descartes seeks a basis for knowledge that is immune to skepticism. The essay will also consider counterarguments to the method of doubt and provide rejoinders to affirm its significance in Cartesian philosophy.

Thesis Statement:  
Descartes' method of doubt is a fundamental epistemological tool that, by dismantling preconceived beliefs and the trust in unreliable senses, lays the groundwork for a knowledge system predicated on certainty, exemplified by the cogito argument.

Body Paragraph 1: Explanation of the Method of Doubt

- Descartes initiates his philosophical journey by doubting everything that is subject to the slightest uncertainty.
- His method involves a hypothetical consideration that all beliefs based on sensory experience, reason, and even mathematical truths could be deceptive.
- The purpose of this radical skepticism is to strip away all dubious beliefs, thereby reaching a foundation that is absolutely certain.

Body Paragraph 2: The Method of Doubt as a Foundation for Epistemology

- The method leads Descartes to the cogito argument, "I think, therefore I am," which stands as an irrefutable point of knowledge.
- This certitude becomes the cornerstone of Descartes' epistemology, as it establishes the self as an undeniable truth.
- All further knowledge can thus be built upon this foundation, with the self as the subject that apprehends the world and its truths.

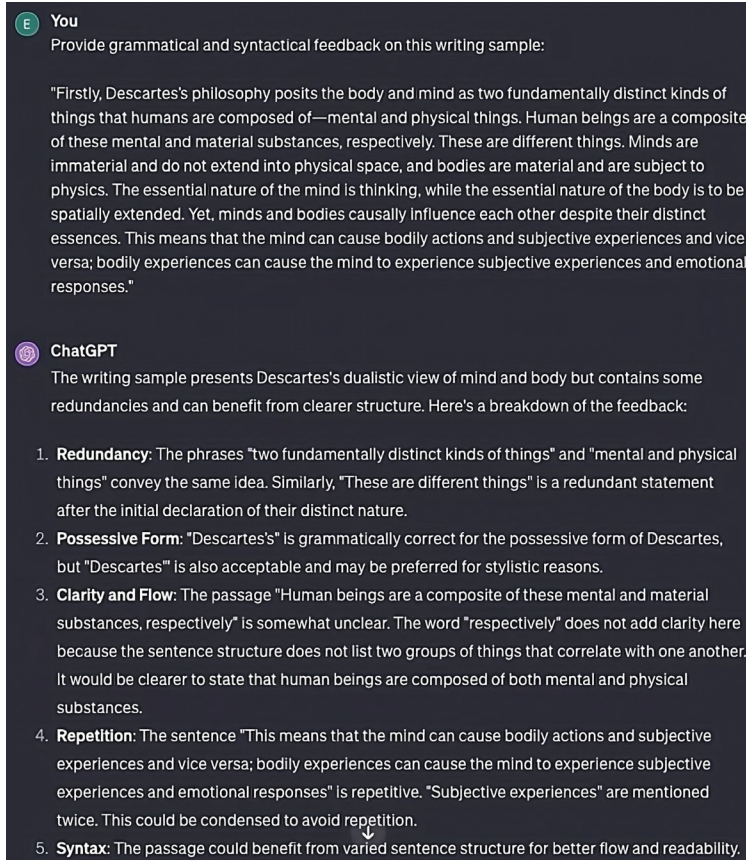
**You**  
Write the entirety of the first body paragraph.

**ChatGPT**  
Body Paragraph 1: Explanation of the Method of Doubt

René Descartes' "Meditations on First Philosophy" introduces a method of doubt that is unparalleled in its systematic and radical approach to questioning knowledge. At the outset, Descartes resolves to set aside all the beliefs he has accumulated from a lifetime of sensory experience, which he deems unreliable due to their susceptibility to deception. This sweeping skepticism extends even to mathematical truths, which, although seemingly incontrovertible, are subjected to the same rigorous scrutiny in case a deceiving God or evil demon is manipulating his perceptions. The essence of this method lies not in the rejection of truth per se, but in the rejection of any belief that can be doubted, however slight that doubt may be. By dismantling the edifice of accepted truths and subjecting each brick to examination, Descartes aims to discard the chaff and uncover a kernel of certainty — a belief so resilient that even the most extreme skepticism cannot shake it. The method of doubt is thus not an exercise in cynicism, but a purgative process, a philosophical quest to establish an unshakable ground for true and certain knowledge. It is the meticulous peeling away of layers of potential error and falsehood to arrive at something undeniable — the bedrock upon which Descartes aspires to construct a new system of indubitable knowledge.

## ChatGPT can be a thoughtful peer reviewer, peer, and study partner

As a reviewer of writing, ChatGPT can delve deeply into user-generated content, including capturing redundancies, repetition, form, and evaluating clarity:



**You**  
Provide grammatical and syntactical feedback on this writing sample:

"Firstly, Descartes's philosophy posits the body and mind as two fundamentally distinct kinds of things that humans are composed of—mental and physical things. Human beings are a composite of these mental and material substances, respectively. These are different things. Minds are immaterial and do not extend into physical space, and bodies are material and are subject to physics. The essential nature of the mind is thinking, while the essential nature of the body is to be spatially extended. Yet, minds and bodies causally influence each other despite their distinct essences. This means that the mind can cause bodily actions and subjective experiences and vice versa; bodily experiences can cause the mind to experience subjective experiences and emotional responses."

**ChatGPT**  
The writing sample presents Descartes's dualistic view of mind and body but contains some redundancies and can benefit from clearer structure. Here's a breakdown of the feedback:

1. **Redundancy:** The phrases "two fundamentally distinct kinds of things" and "mental and physical things" convey the same idea. Similarly, "These are different things" is a redundant statement after the initial declaration of their distinct nature.
2. **Possessive Form:** "Descartes's" is grammatically correct for the possessive form of Descartes, but "Descartes'" is also acceptable and may be preferred for stylistic reasons.
3. **Clarity and Flow:** The passage "Human beings are a composite of these mental and material substances, respectively" is somewhat unclear. The word "respectively" does not add clarity here because the sentence structure does not list two groups of things that correlate with one another. It would be clearer to state that human beings are composed of both mental and physical substances.
4. **Repetition:** The sentence "This means that the mind can cause bodily actions and subjective experiences and vice versa; bodily experiences can cause the mind to experience subjective experiences and emotional responses" is repetitive. "Subjective experiences" are mentioned twice. This could be condensed to avoid repetition.
5. **Syntax:** The passage could benefit from varied sentence structure for better flow and readability.

ChatGPT checks basic grammar, suggests basic improvements, and even recommends structural changes. In my experience, meaningful peer review at the undergraduate level usually devolves into this sort of feedback, so why not cut out the middleman and use the best grammar and spell check available? How important are the human aspects of peer review and peer-professorial interaction when a technical tool can produce more efficient results?

## ChatGPT is a very good generalist tool...

In a lot of humanities education, GPT-4 in its current iteration won't replace the more critical thought that happens at the higher level of fields. In this way, it's more like a steroid-induced and more efficient Google Scholar/JSTOR. Overall, while impressive, the tool is not a substitute for truly understanding the code-switches and vernacular that different groups use in their discussions of critical topics. Understanding that GPT-4 and other generative AI tools have rapid access to vast amounts of data, we shouldn't be shocked by the hyper-efficiency that generative AI provides undergraduates who are often new to academic discourses.



### What does this mean for undergraduates and their professors?

I don't necessarily think that undergraduate education can be reduced to algorithms. But maybe getting the "A," for most students and sometimes myself, reduces undergraduate education as a way to figure out "how to get the A" rather than pry for the "what" and most importantly the "why." Generative AI's broad capabilities make me incredibly empathetic towards professorial hostility around proficient generative AI models, especially ChatGPT. But I also (obviously) think projects like the one at the USC Center for Generative AI are incredibly important for evolving undergraduate education. In the grander schema of undergraduate education, it's complicated.

### On professorial hostility around generative AI

The most charitable view of undergraduate education is that students are supposed to learn how to critically think—the "why." In undergraduate philosophy, we learn archaic theories from Kant and Descartes that illuminate more critical thought today. My friends in computer science learn foundational theories in algorithms and data structures so that they can build more intricate systems that avoid fatal bugs for Facebook and Amazon. That's why I don't necessarily accept banality as a reason for educators to change the way they teach; I *understand* why some professors have taken the step to either ban ChatGPT as a tool or have moved to in-person and paper assessments given the advent of Generative AI. For instance, the pedagogical value of learning and mastering the standard K-12 five-paragraph essay structure is trivialized by most undergraduates. It's boring and mechanical, yes. But it also hones a student's ability to communicate complex ideas succinctly; learning to master this subset of writing teaches students to be clear and concise in whatever discipline they choose.

Outside the scope of this article, anecdotally, most of the English majors I've met and most of the Biology majors I've met want to become lawyers and doctors, respectively. This isn't a bad thing, but I think it underscores the promise of a degree as a means to an end, an investment to a bigger investment. In the context of undergraduate education, the easiest way to make out rich with that investment is to pursue the least scrutinizing means to that end.

In the coming months and years, as generative AI technology advances, I empathize with not relaxing policies that outright ban ChatGPT because some students won't use generative AI intelligently, and plagiarism and cheating are clearly the biggest elephants in the room. My perspective on the core essence of undergraduate education is that it nurtures and educates young adults to explore their own *original* thoughts. So to reduce this *process* into algorithmic thinking for the benefit of good grades poses a large problem that ought to be addressed.

### **At risk and the loss of critical thinking?**

For some, ChatGPT skips the steps necessary for critical thinking. As someone who wants to eventually pursue graduate education, I think I'm scared off enough by OpenAI's disclaimer that "ChatGPT may produce inaccurate information about people, places, or facts" to use it consistently to model my writing.<sup>2</sup> Still, in classes during my fourth year of college, there's still rote banality and things that have been said better than I could ever say them that the algorithm has absorbed and regurgitated.

That said, I see a particular risk for lower division classes. These entry level courses are designed to instill foundational knowledge. This is often done by reading texts, jotting down notes, and going to office hours to ask questions when needed. These tasks are supposed to develop analytical and critical thinking skills. In calculus, for example, mastering Stokes' Theorem is not just about applying a formula to figure out the curl of a vector. It's about understanding the underlying reasoning that makes the formula possible and meaningful. I fear that much of STEM is at risk; AI-generated copy-pastable code is more user-friendly than office hours, Stack Overflow, and Github. In this way, easy access to the right answer undermines the usefulness of struggling through the wrong one.

The convenience of ChatGPT threatens this delicate ecosystem that undergraduate education supposedly promotes. I don't think the most dangerous threat to learning is academic dishonesty. The threats that some educators have made might create that perception, but I don't think professors intend to witch-hunt. I think it's the potential corrosion of critical thinking that educators are afraid of when AI becomes a crutch, in a way that differentiates how Google and Wikipedia have. There's a risk when academia and its young stakeholders bypass the problem-solving stage in every discipline.

<sup>2</sup> ChatGPT, November 7, 2023, OpenAI, <https://chat.openai.com>.

## Rethinking undergraduate education

I don't think it's fair to ask educators and academia as a whole to flip the paradigm overnight; nor does the Center. There is so much pedagogical value in the way our educational system has worked before ChatGPT became so ubiquitous in syllabi and classroom discussions. One sympathy I have towards professors is they are sometimes reluctant to new pedagogical approaches and are afraid of pedagogical threats. Human beings age by the second, but Generative AI seems to age exponentially, leaving many behind. Generative AI is a new future and undermines a linear reality, one that provokes the intrinsic reasons for receiving an education.

I don't expect all (or even most) instructors to view generative AI amicably, but I don't necessarily think an adversarial relationship is good, either. Similarly, students should view generative AI amicably—but as a tool, not a crutch. Banality, the impetus for a loss of critical thinking, was a threat long before the advent of generative AI. Traditional pedagogy is still clearly very valuable, but can AI be an impediment to learning with current hostilities/banalities in some undergraduate education? I don't think it has to be, and that's why the experimental findings that this report makes are so significant.

# AI in K12 Classrooms: Ethical Considerations and Lessons Learned



# How Teachers Navigate the Ethical Landscape of AI in Their Classrooms



Just a few decades ago, Artificial Intelligence (AI) was a concept limited to sci-fi novels and movies. Today, however, AI is an inescapable reality, and is poised to transform our everyday lives. Generative AI in particular has the potential to redefine the way we think about technology and change our relationship to work. Nowhere is this transformation more evident and profound than in education. Generative AI tools have already found their way into classrooms, and their inclusion has forced educators, parents, and policymakers to face challenging questions such as:

- How will AI influence the learning experiences of students?
- Should using generative AI to support writing always be considered plagiarism?
- Will AI be implemented equitably?
- Will generative AI change students' relationship to what it means to master material?
- What potential does AI hold for personalized education?
- What are the ethical ramifications of embedding AI-based educational technologies into the classroom?

Educational technologies (AI-based or otherwise), are only useful if they are successfully incorporated into the everyday practice of teachers. While many things go into teachers' use of educational technologies (e.g, ease of use, appropriateness for a task, supportive infrastructure) generative AI introduces new ethical dimensions to the choice.

As the Associate Director of the newly formed USC Center for Generative AI and Society I have become increasingly interested in how educators are making sense of the ethical challenges posed by the potential use of generative AI-powered educational technologies. In short, ***How are teachers navigating the ethical landscape of AI in their classrooms?***

It is important to explore how teachers make ethical judgments about using generative AI in their classrooms. So, as we stand at this pivotal juncture in education, I am happy

to share findings from the first study conducted by the center. In so doing I hope to shed light on the ethical beliefs of those at the heart of the education system: our dedicated and hardworking teachers.

### Why you should care about ethics and AI in education

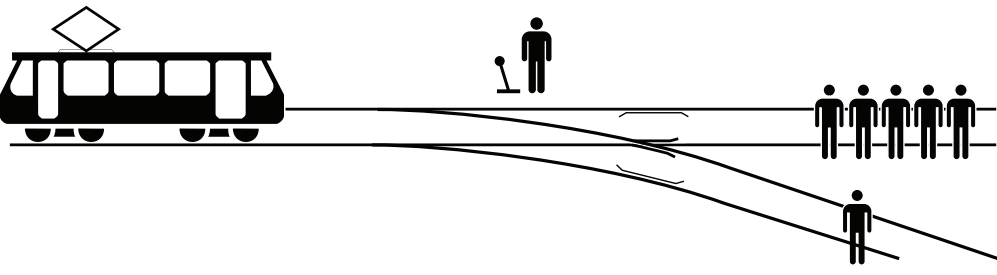
AI is more than just lines of code or intricate algorithms; it is a reflection of our values, decisions, and aspirations. While many of us have the luxury to safely ignore generative AI for a while longer, teachers cannot, as they will soon be faced with an avalanche of educational technologies that are based on generative AI. In fact, many students are already using generative AI to complete their assignments. Teachers are thus at the forefront of making ethical decisions about whether or not generative AI has a place in their classrooms. How, though, do teachers make such judgments?

At the core of our study was the desire to understand the nuanced perspectives of teachers regarding AI ethics. Specifically we were interested in the following research questions:

- How do personal beliefs and experiences shape teachers' views of generative AI?
- In an era where technology is ubiquitous, does a teacher's level of comfort with technology influence their ethical considerations?

### “The Trolley Problem”

“Ethics,” however, is a vast topic with many points of entry. In order to ground our investigation, we used the “Trolley Problem,” a classic philosophical thought experiment that enables us to surface our intuitions about what it means to act ethically. The scenario is as follows:



The diagram illustrates the Trolley Problem. A trolley is moving along a set of tracks from left to right. Ahead of the trolley, there are two tracks: a main track and a side track. On the main track, five people are tied up and unable to move. On the side track, one person is tied up. A person stands next to a lever that can switch the trolley from the main track to the side track. The trolley is currently on the main track, heading towards the five people.

A trolley is moving along a set of tracks. Ahead, there are five people tied up and unable to move. The trolley is headed straight for them. You are standing some distance off, next to a lever. If you pull this lever, the trolley will switch to a different set of tracks. However, there is one person tied up on the side track (Fig. 1). You have two options:

1. Do nothing, and the trolley kills the five people on the main track.
2. Pull the lever, diverting the trolley onto the side track where it will kill one person.



What is the most ethical thing to do in this instance? The crux of the Trolley Problem boils down to two primary ways of deciding the most ethical course of action:

- **The Rule-Followers (“Deontological” Perspective):** For some, ethical action means strictly adhering to established principles and guidelines, regardless of the situation. Thus, acting ethically is determined by following a previously established rule that determines whether or not an act is ethical. “It is never morally justified to knowingly cause the death of another,” for example, is one such rule.
- **The Outcome-Seekers (“Consequentialist” Perspective):** Then there are those who believe in a more flexible approach. They argue that ethical decisions should be based on achieving the most beneficial outcome, even if it means bending (or breaking) some rules. “It is OK to cause the death of another if a greater harm is avoided,” would be a heuristic used by folks in this camp.

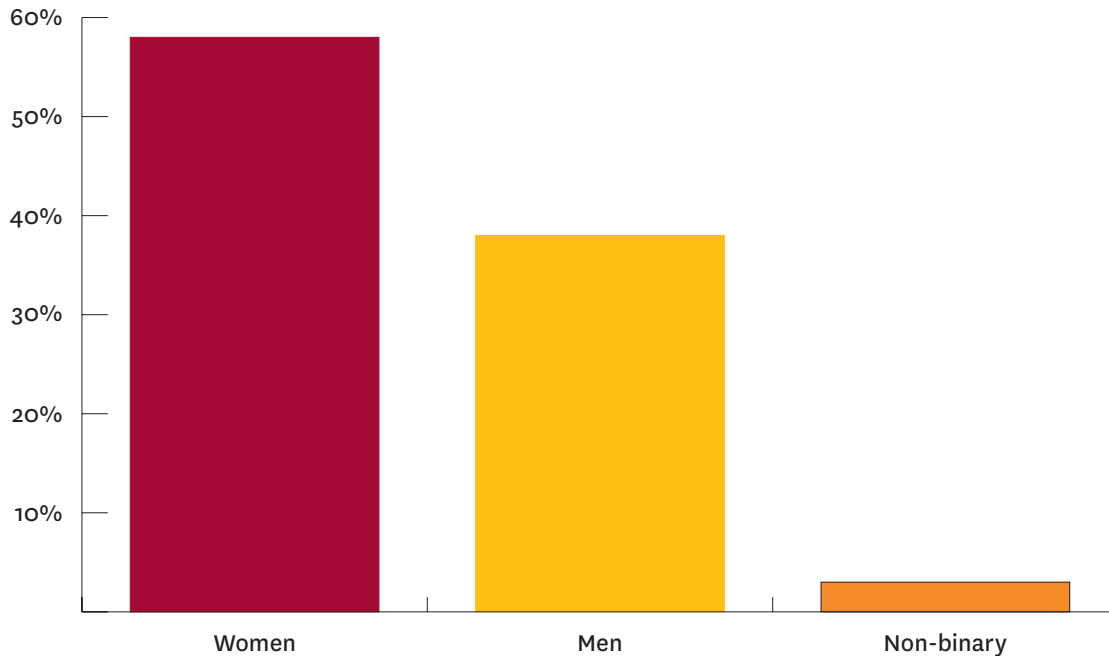
Thus, the “rule-follower” would not act, and five people would die, whereas the “outcome-seeker” would pull the lever, stopping the death of five individuals but causing the death of one.

We don’t live in a world where teachers are manning levers determining the direction of literal trolleys, however. **Instead, we live in a world where teachers have to decide when, where, and how generative AI can be used by the students they teach.** Do such judgements focus on promoting the positive consequences of generative AI and avoiding the negative consequences? Or are they instead governed by deeper beliefs about what it means to teach and learn, which may limit how educational technologies powered by AI are used?

## Participants

To answer our research questions we reached out to K-12 teachers spanning different backgrounds and years of teaching experience. Our final sample consisted of 248 teachers who had been teaching for 11 years, on average. 43% were elementary school teachers, 16% were middle school teachers, and 40% were high school teachers. A little over half (58%) of teachers in our study identified as women. The majority of teachers in our sample taught in public schools (72%), although charter schools (5%) and private schools (19%) were also represented. The majority of teachers in our study identified as White (85%). Finally, the study represented teachers from 41 different U.S. states. While diverse on multiple dimensions, our data is not intended to be representative of all teachers in the United States.

**Gender Distribution of Teachers**



We asked teachers to rate how much they agreed with different **ethical ideas**, using a scale from 1 to 5. There were two types of ethical views:

- **Deontological** (based on rules):
  - AI should respect users' choices, like having clear ways to opt in or out.
  - AI must protect user privacy and confidentiality.
  - AI systems shouldn't deceive users.
  - AI should be fair and not biased.
  - AI must avoid causing harm.
  - AI decision-making should be transparent.
  
- **Consequentialist** (focused on outcomes):
  - AI can give access to information in resource-limited areas.
  - AI can improve efficiency.
  - AI can help with creativity.
  - AI can offer personalized experiences.
  - AI can help connect people globally, like through language translation.
  - Concerns about AI invading privacy or misusing data.
  - AI can provide accurate information or spread misinformation.
  - AI might replace jobs.
  - People might become too reliant on AI.
  - AI could increase or reinforce societal biases.

We also asked teachers to rate their willingness to use generative AI, like ChatGPT, in their classrooms. They used a scale with eight different aspects:

- **Performance:** Teachers' belief in the usefulness of ChatGPT for their job.
- **Effort:** How easy teachers think it would be to learn to use ChatGPT.
- **Attitude:** Teachers' overall opinion about using ChatGPT.
- **Social:** Influence of others' opinions on teachers' decision to use ChatGPT.
- **Conditions:** Compatibility of ChatGPT with other systems teachers use.
- **Self-Efficacy:** Teachers' confidence in getting help if they struggle using ChatGPT.
- **Anxiety:** Teachers' fear of making uncorrectable mistakes while using ChatGPT.
- **Intentions for Use:** Teachers' plans to use ChatGPT in the next six months.

## Findings

Teachers' ethical evaluations of generative AI were diverse, reflecting the broader societal debate about when and how generative AI can or should be used within education. While some were staunch proponents of strict ethical guidelines, others believed in a more outcome-driven approach. Our full findings are undergoing peer review, however we can confidently report two key findings:

- **A Gender Perspective:** Our results showed intriguing gender-based nuances. Male and female teachers, while sharing many views, displayed distinct differences in certain ethical evaluations. Specifically, *female teachers were more likely to be proponents of rule-based (deontological) perspectives when compared to men.*
- **The Role of Attitude:** *Self-Efficacy* (confidence in using technology) and *Anxiety* (worry about using technology) were found to be important in both rule-based (deontological) and outcome-based (consequentialist) views about AI. This shows that how people feel about technology, like being confident or anxious, affects how they think about the ethics of AI and how they use it. Teachers' feelings of either being skilled or burdened by AI were also related to their decisions about using it.

## Concluding Thoughts

Teachers are not just passive recipients of the changes in education that are imminent due to generative AI. Instead, they are active participants, grappling with the moral challenges posed by AI. Their concerns extend beyond the practicalities of AI tools. Our results show that teachers are also pondering the deeper questions, such as: What values do these AI systems propagate? Are they fair to all students? Our findings underscore the complexity of these questions. Teachers, drawing from their reservoir of experiences and beliefs, presented differing views, yet, a consensus emerged on one front: the pressing need for an adaptable, comprehensive ethical framework for AI in education.

As AI becomes a staple in our classrooms, the surrounding ethical questions of its use become paramount. Our study reminds us that as we stride into an AI-infused future, we must ensure that our technological advancements are underpinned by ethical considerations. This study is not just a reflection of the present; it suggests that we need to look ahead in order to ensure that as we leverage the power of AI, we do so with a deep-rooted commitment to fairness, inclusivity, and ethical integrity. Teachers, moreover, will be at the forefront of those decisions.

# Lessons Learned by Teaching Middle School Students to Use Generative AI



Artificial Intelligence (AI) is now getting more and more popular. All kinds of AI applications are emerging in different fields and changing our life and work. As a STEM education researcher, I feel the pressing need to prepare youth for the future AI-driven society. I searched for existing research and platforms for K-12 AI education but found very few well-developed and age-appropriate curricula to teach AI to students. So, I decided to design an AI curriculum for my recent doctoral dissertation. I chose middle school as the target grade level for two reasons: (1) Middle school is a critical phase for students' STEM learning and identity development, and (2) I am most familiar with middle school students based on my past teaching experience.

The curriculum needed a focal topic since it was impossible to cover everything in AI. I decided to focus on the application of Machine Learning (ML) to solve Socio-Scientific Issues (SSI). ML is an important subfield of AI, referring to the computational algorithms that function as the “brain” of AI. For instance, the plant identification app is supported by one type of ML algorithms (i.e., supervised learning) to first learn from as many types of plants as possible about their characteristics, then be able to recognize a plant based on the learned patterns. SSI are complex social issues involving scientific knowledge, such as climate change, genetically modified organisms, and COVID pandemic. In the real world, ML has been widely applied to solve SSI. For example, to control the COVID pandemic, ML techniques have supported robots to automatically check whether people wear masks. A curriculum of applying ML to address SSI can engage students in authentic problem-solving.

## Curriculum Design & Implementation

I designed the curriculum to teach students about AI by encouraging them to create ML products. For this purpose, I designed several ML projects as a key component of the curriculum, apart from short introductory lectures about AI/ML. Working on the ML projects allowed students to experience the complete process from collecting data to training and testing ML models, to communicate their ideas with peers and instructors, and to revise and improve their ML models.

One of the projects was designed in the SSI context of COVID pandemic. I asked students to build a ML system on the [Teachable Machine platform](#) to achieve automatic detection of people wearing or not wearing a mask. Students took photos of themselves wearing and not wearing masks to train their own ML models for this project. After completing the task, they were asked to share their products with the class and provide feedback to each other.

I also included a group discussion for students on the topic of whether AI technology should be used for automatic check of mask wearing during the COVID pandemic. It was designed to cultivate students' critical thinking about AI ethics.

In February 2023, I implemented the curriculum as a free online afterschool program over three to four weekends. Participants were 10 middle school students (Grade 6–8) from diverse racial and ethnic backgrounds (2 Hispanic/Latino, 1 Indian, 7 Chinese) and geographical locations (9 students from 5 different states of USA and 1 from New Zealand). Three students were female and seven were male.

## Findings

Students generally showed improvement of their ML knowledge compared to baseline tests.

1. The inclusion of masks as a physical tool and the SSI context of COVID pandemic made the ML project more interesting and engaging for students. While working on the tasks, students had fun with putting on and taking off the mask, trying different ways of wearing masks, and changing masks with different colors and designs.



2. The ML projects allowed students to actively interact with peers and instructors, such as sharing thoughts about how to build the ML project, sharing the ML products they made, and commenting on others' ML products. The interactions often provided students the opportunity to take a closer look at their projects and observe issues/bugs, and consequently, to revisit the data and tools to further explore, revise, and improve their projects. During the processes of making, interacting, and revising, students deepened their ML knowledge and promoted their ML practices. For future research, I plan to revise the activity structure based on this finding, to maximize students' improvement from the learning-through-making process.
3. The curriculum also supported students to talk about AI ethics confidently from multiple aspects: socioeconomic impact, safety and security, bias and fairness, environmental sustainability, and human control and autonomy. For example, one student expressed his concern about humans losing control over robots, "if Bots keep on advancing and no new innovation is made, then robots would just like, take over basically. And we do have like nothing to do and just get lazy." In addition to assigned discussion, students also voluntarily initiated discussions on AI ethics, such as the use of ChatGPT. A student shared that his classmate submitted a paper generated by ChatGPT and got a good grade because their teacher was not aware of it. Then he talked about regulations on using ChatGPT.

## Implications

There are several interesting observations that deserve further research.

- Students had several episodes of exciting and fruitful collaborative learning. Since collaboration is usually challenging for online learning, an in-depth analysis of these successful examples is helpful for the online collaborative learning design.
- Students from different racial and ethnic backgrounds showed different learning patterns in the ML projects, especially in their approaches to sharing their ML projects to others. Further research on this may have implications on equitable and inclusive learning designs for students from diverse backgrounds.
- According to what participants shared during the program, middle school students were using ChatGPT and their attitudes toward ChatGPT varied. Some students were using it inappropriately, as shared by one of my participants. In contrast, another participant explicitly said that he would not use ChatGPT because he wanted to write by himself so that he could learn. Their black-or-white attitudes to the use of ChatGPT showed their lack of knowledge in how to utilize it in supporting their learning.

My findings suggest that it is important to continue researching teaching and learning strategies based on the use of ChatGPT and other Generative AI technologies. As these new technologies are more popular among students at undergraduate and graduate level, starting from higher education level can help with the urgent need to guide students how to better use them.

## References

- Arnett, J. J. (2015). Identity development from adolescence to emerging adulthood: What we know and (especially) don't know. In K. C. McLean & M. Syed (Eds.), *The Oxford Handbook of Identity Development* (pp. 53–64): Oxford University Press.
- Carlone, H. B., & Johnson, A. (2007). Understanding the science experiences of successful women of color: Science identity as an analytic lens. *Journal of Research in Science Teaching: The Official Journal of the National Association for Research in Science Teaching*, 44(8), 1187–1218.
- Potvin, P., & Hasni, A. (2014). Analysis of the decline in interest towards school science and technology from Grades 5 through 11. *Journal of Science Education and Technology*, 23(6), 784–802. <https://doi.org/10.1007/s10956-014-9512-x>
- Sadler, T. D. (2004). Informal reasoning regarding socioscientific issues: A critical review of research. *Journal of Research in Science Teaching*, 41(5), 513–536. <https://doi.org/10.1002/tea.20009>
- Wang, C. (2023). AI education for middle school students: Critically exploring socioscientific issues with machine learning. [Doctoral Dissertation, University of Miami]. <https://scholarship.miami.edu/esploro/outputs/991031825011402976>
- WION news. (2020). Robots enlisted to fight COVID-19 | 'Robovie' tells people to wear a mask | Japan | World News. [https://youtu.be/jl\\_qcWdLnDw](https://youtu.be/jl_qcWdLnDw)



# Building the Next Generation of Generative AI Tools



# Generative AI and Education: Deny and Detect or Embrace and Enhance?



ChatGPT was released on November 30, 2022. Now, about a year later it seems clear that generative AI systems like ChatGPT will have a profound disruptive effect on many facets of society. Perhaps one of the largest effects has been in education. Students suddenly found that they had a kind of “homework machine” (Williams & Abrashkin, 1958) at their disposal, a set of systems that could complete writing assignments, generate artwork, and even solve programming problems.

The reaction of the educational community was mixed (Jimenez, 2023). Some educators, concerned that students would use generative AI to cheat, advocated denying students access to generative AI and detecting if they did. Some of the nation’s largest school districts blocked students from accessing generative AI systems on their networks and a number of tools emerged that attempted to determine whether a body of text had been produced by generative AI or a human (Caulfield, 2023/2023; Walters, 2023). Other educators, believing that generative AI is here to stay, argue that generative AI should be embraced, and the education process itself needs to be modified to accommodate it. Further, in making those modifications it might be possible to provide an enhanced educational experience.

Our own view at the Center is that the deny and detect approach, although well intentioned, will ultimately not serve students well. First, the reliability of AI detectors does not seem to be consistently high. While some studies have reported that a few detectors are 100% accurate (Walters, 2023) on a particular set of data, other studies show lower accuracy for the same detectors on different datasets (Caulfield, 2023/2023). Furthermore, some of the detectors are not accurate at all with reported accuracy scores as low as 38%. An additional serious problem is false positives, where the detector classifies a human-written text as being produced by a machine. False positives are particularly problematic because they set up a situation where a teacher may accuse a student of cheating when in fact they did not. More fundamentally, even if detectors were 100% accurate, the denial approach would not prepare students well. If one tells students they cannot use generative AI throughout their schooling, and then graduate them into a world where it is extensively used, the students have not really been prepared for the world they will enter.

Alternatively, if we adopt the embrace and enhance approach, how should the education process be changed to accommodate generative AI? As Ben Nye outlines in his essay in this report on our work on the Authoring by Editing (ABE+) framework, we have begun to identify two elements that seem to be important in fitting in generative AI. First is an emphasis on promoting critical thinking in students. There are well known, important limitations of generative systems that impact the quality of their results. Generative systems sometimes “hallucinate”, making statements that are untrue or inaccurate. Because current generative systems do not understand the texts they produce (at least in a way that humans would recognize as understanding) it can be difficult for them to verify the statements they make as they are generated. Generative systems also have quite limited reasoning capabilities, which can again contribute to inaccuracies or flawed argument structures. The ability of a human to look at the output of a generative system critically, asking questions such as “Are the facts here right?” or “Does this argument make sense?” is very important to assure the quality of the output. Another important shift to accommodate generative AI is to move from evaluating a student’s work based on the artifact they produce (an essay) to evaluating the process they went through to produce it, such as edits and modifications to the eventual text. A process based approach to evaluation is both more resistant to cheating and more revealing of the student’s thought process, which can be useful in mentoring the student.

As we work more with faculty interested in using generative AI in their classes we are finding that accommodations that work well in one educational environment may not work as well in another. For example, the approach taken in the ABE+ framework, where a student has a generative system produce a first draft of a paper that they then examine critically and modify may work well for a STEM writing exercise, but not as well for a creative writing class where the emphasis is on having the student develop their own writing style or voice. If a generative AI produces the first draft of the paper, it may prove difficult for the student to add their own voice.

We believe that this pattern will continue. Accommodating generative AI will require changes to how we educate our students, but the changes will need to be customized to particular pedagogical environments and goals. When we consider the university as a whole, it will clearly be a major undertaking to incorporate generative AI into our programs of instruction, but if done right, it holds the promise of improving the educational experience for our students.

## References

- Caulfield, J. (2023, September 6). Best AI Detector | Free & Premium Tools Compared. Scribbr. <https://www.scribbr.com/ai-tools/best-ai-detector/> (Original work published 2023)
- Jimenez, K. (2023, January 30). “This shouldn’t be a surprise’ The education community shares mixed reactions to ChatGPT. USA Today. <https://www.usatoday.com/story/news/education/2023/01/30/chatgpt-going-banned-teachers-sound-alarm-new-ai-tech/11069593002/>
- Walters, W. H. (2023). The Effectiveness of Software Designed to Detect AI-Generated Writing: A Comparison of 16 AI Text Detectors. *Open Information Science*, 7(1). <https://doi.org/10.1515/opis-2022-0158>
- Williams, J., & Abrashkin, R. (1958). *Danny Dunn and the homework machine*. Scholastic Book Services.

# Authoring by Editing and Revising: Considering Generative AI Tools



To evaluate the opportunities for Generative AI in Education, prototypes must be carefully developed and studied to identify where they add value to learning and instruction. To this end, the USC Center for Generative AI and Society is co-designing tools with fellows from the USC Writing Program to support critical thinking and build process-oriented writing skills. While many academic subjects will be influenced by generative AI, writing courses have already been impacted at a fundamental level. The core challenge of writing instruction lies between the aspirational goal to support student writing over time, versus the practical constraints that have historically meant that grading and feedback for writing often relies on the “end product.”

In this respect, generative AI offers both new challenges and potential solutions. On the challenge side, there are legitimate concerns that generative AI enables cheating by generating a reasonable essay or other writing product. On the solutions side, instructors are already leveraging generative AI tools such as ChatGPT or Claude to help students ideate, reflect on, and improve their writing. These lessons learned from the field inform the types of tools that we should be building to help students develop their writing process. If tools designed to scaffold and summarize writing over time are widely used, this should mitigate the potential to cheat because skills could be tracked over time. More importantly, generative AI could offer an important asynchronous communication medium for writing, where instructors can automate support for students and AI analytics of student drafts help instructors intervene.

## Background: Categorizing Different Generative AI Approaches

The landscape for tools that leverage generative AI is evolving rapidly— many of the practical applications are not yet rigorously reviewed in academic journals, and are instead found in workshops, blogs, self-published handbooks, and other less formal venues (Hsiao, Klijn, & Chiu, 2023). The space is also highly varied in terms of technical expertise, where both instructors with practical use-cases (e.g., assignments with a ChatGPT interaction) and researchers with advanced prototypes (e.g., customized large language models) both offer important insights.

Due to the dramatic increase in quality of large language models (LLMs), much of the key insights depend on reports from only the last year. Earlier generation LLM's, such as GPT-3, were well-studied but generated outputs that were too poor quality to be useful for most educational needs, even for drafting instructional content such as explanations (Wang, Prihar, & Heffernan, 2023). More effective models such as GPT-4 (OpenAI, 2023) and LLAMA 2 (Touvron et al., 2023) have dramatically increased the quality of generative AI, meaning that as a community we are still mapping out where they can improve education. When considering "improvement", we consider two standards:

- **Feasibility:** Have use-cases or case-studies demonstrated that generative AI can be used successfully based on at least the self-assessment of a practitioner?
- **Efficacy:** Have evaluations measured and established the likelihood of specific pedagogical benefits (e.g., greater learning, higher engagement, time-savings for instructors)?

Feasibility is a much lower standard than efficacy, indicating a promising direction to investigate further. In this report, we focus on this space of feasible approaches. Monitoring and cataloging promising techniques is important, because each new generative AI model opens up new techniques that might have worked poorly the prior year. In future years, evidence of efficacy will accumulate, but so far most approaches have not been rigorously tested against meaningful controls (e.g., traditional assignments, posing questions to search engines). In the writing domain, feasibility is being explored across three main categories:

- **Instructional Content Generation:** Speeding up or improving the quality / variety of assignments and materials a teacher uses.
- **Real-Time Learning Activities:** Learning activities where a student uses generative AI to produce output (e.g., an example essay) or to engage in an interactive conversation about some learning content.
- **(Formative) Assessment:** Analyzing student work to provide insights that would help an instructor (or the student themselves) address issues and improve the student's learning.

While domains such as math or basic language instruction have focused extensively helping to generate instructional content such as hints, feedback, or alternate assignments (Bonner, Lege, & Frazier, 2023; Pardos & Bhandari, 2023), writing instruction appears have greater emphasis on Learning Activities and Assessment. Table 1 highlights a selection of the activities where educators have been applying generative AI to teach writing. As a taxonomy for generative AI learning tasks has not yet solidified, in this report we will organize learning activities based on Bloom's Taxonomy (Krathwohl, 2002) and two instructor-oriented categories for Content Generation and Assessment. These represent only a subset of approaches noted by instructors and in literature, but they give an overview of the breadth of techniques under discussion.

**Table 1: Feasible Learning Activities for Writing Instruction with Generative AI**

Category	Learning Activity
Content Generation	<ul style="list-style-type: none"> <li>▪ <b>Assignment Prompts:</b> Generate additional narrative writing prompts</li> <li>▪ <b>Presentation Notes:</b> Pre-made notes based on other content (e.g., slides)</li> <li>▪ <b>Brainstorming Lessons:</b> Generate candidate lessons based on a description</li> <li>▪ <b>Red Team Lessons:</b> Check that generative AI models cannot easily give a high quality response to a lesson prompt</li> </ul>
Assessment	<ul style="list-style-type: none"> <li>▪ <b>Cheating Detection:</b> Attempting to identify AI-generated essays</li> <li>▪ <b>Self-Explanation Analysis:</b> Categorizing explanations for reading comprehension strategies, overall quality, or other criteria</li> <li>▪ <b>Automated Essay Grading:</b> Evaluating overall essay quality, often (but not always) with high-level strengths and weaknesses</li> </ul>
Remember	<ul style="list-style-type: none"> <li>▪ <b>Corrections:</b> Fix grammar or sentence mechanics</li> <li>▪ <b>Note-Taking:</b> Generate notes to review later</li> </ul>
Understand	<ul style="list-style-type: none"> <li>▪ <b>Summarization:</b> Simplify text to main ideas</li> <li>▪ <b>Outlining:</b> Create an outline of a text to review an argument</li> <li>▪ <b>Structural Example:</b> Generate examples of different writing products (e.g., formats)</li> <li>▪ <b>Conversion:</b> Modifying an input into another writing product format (e.g., “Make a sonnet summarizing this article”)</li> <li>▪ <b>Translation:</b> Convert a writing product into another language or dialect</li> </ul>
Apply	<ul style="list-style-type: none"> <li>▪ <b>Initial Draft:</b> Create a version to critique or modify</li> <li>▪ <b>Prompt Design:</b> Understanding fundamentals of writing effective generative AI prompts to produce certain outputs</li> </ul>
Analyze	<ul style="list-style-type: none"> <li>▪ <b>Find Gaps:</b> Compare essay versus generated claims to support thesis</li> <li>▪ <b>Red Teaming:</b> Generate critiques or alternate viewpoints</li> <li>▪ <b>Critiquing the AI:</b> Generate AI content to evaluate where it can be improved</li> </ul>
Evaluate	<ul style="list-style-type: none"> <li>▪ <b>Reflection:</b> Self-reflection prompts and sounding board, optionally with feedback</li> </ul>
Create	<ul style="list-style-type: none"> <li>▪ <b>Brainstorming:</b> Generate examples from different perspectives to inspire writing</li> <li>▪ <b>Simulated Response:</b> Write a piece (e.g., a letter) and prompt the AI to respond “in character” to the piece (e.g., Thomas Jefferson), to inform revisions or follow-ups</li> <li>▪ <b>Extrapolation:</b> Generate a hypothetical situation or fictional narrative (e.g., character backstory) for an element in the student’s current writing</li> </ul>

In emerging literature, different organizing principles have been used to group activities using generative AI. For example, the Harvard (Derek Bok Center for Teaching and Learning, 2023, p. 10) seminar program focused on real-time learning activities and organized them into an argumentation process: Recognize Argument, Unpack Argument, Make Own Argument, Make Scholarly Argument, and Apply Scholarship. On the converse, work such as Hsiao, Klijn, & Chiu (2023) consider LLM use from the standpoint of assessment, so their focus is more on the scope of use (e.g., limiting use to a few elements of an assignment) rather than the specific activities. Generative AI errors have also been suggested as examples to encourage critical thinking about bias and misinformation (e.g., Mills & Goodlad, 2023). Distinct from these themes, activities have also been designed to teach students how to prompt an AI effectively and use that to help students think about the nature and process of writing (Marino, 2023). These techniques may eventually evolve into teachable agent approaches, where students learn about the writing process by instructing an AI and reflecting on its performance (Kim & Baylor, 2016). Finally, existing literature is almost entirely focused on LLM writing analysis and generation, as opposed to image generation or conversations about multimedia documents. This may change as AI tools develop further, as some writing formats rely on figures or tables for communicating key concepts.

In terms of tools, advances in writing pedagogy using generative AI are so far almost entirely using off-the shelf tools rather than through integrations of LLM into existing intelligent frameworks designed to train writing skills (e.g., Butterfuss et al., 2022). Given the relative power offered by techniques such as chain-of-thought prompting and other programming techniques for LLMs, the space of feasible pedagogy is almost certainly more diverse than Table 1 suggests. However, exploring this space is more complex because instructors are less able to directly innovate when a programming API is involved. As a result, a key recommendation from this report is that instructors and researchers share their prompts and activity structures diligently. These should help other instructors to re-use these activities, as well as provide examples that AI researchers could help improve or optimize.

## Our ABE+ Prototype Approach

The Authoring by Editing and Revising (ABE+) framework represents an effort to explore the feasibility of generative AI techniques to support writing pedagogy.

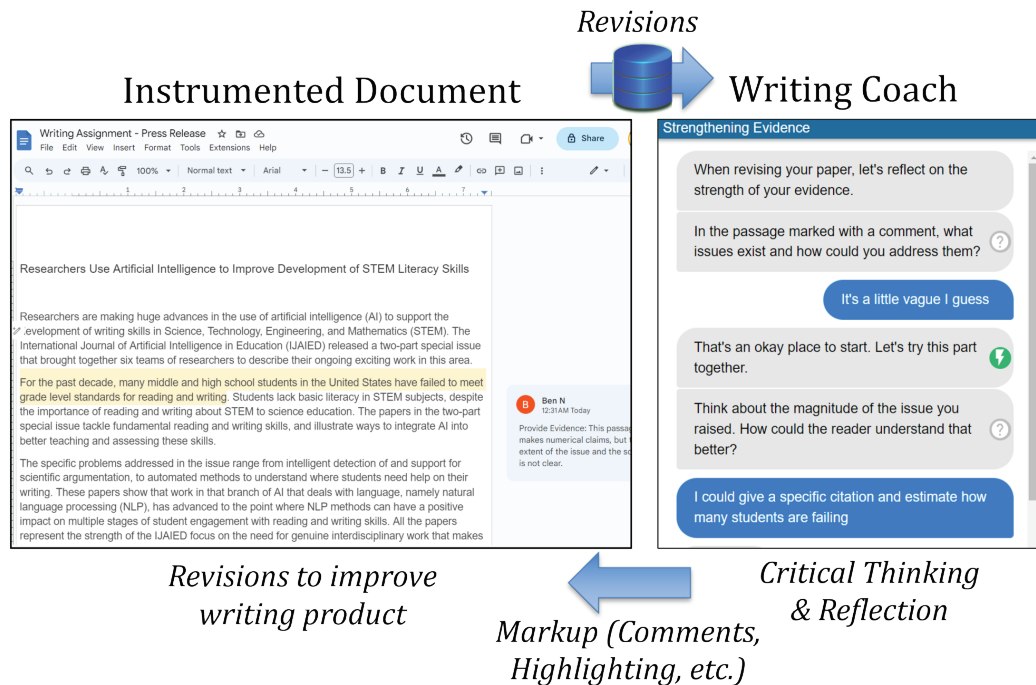
ABE+ has the following main design criteria:

- **Critical Thinking:** Develop tools to support deeper aspects of the writing process (Analyze, Evaluate), rather than traditional low-level review (e.g., grammar).
- **Process-Oriented:** Encourage multiple cycles of reflection and revision of a writing product, with reflections recorded as context for the revisions that are made. Track revisions continuously to capture potential patterns in skill acquisition.

- **Formative Assessment:** Generate feedback and suggestions in real time, designed to emulate human instructor prompts or peer review. Summarize the timeline of reflections and formative assessments for a course instructor to review.

A high level diagram of the ABE+ real-time coaching is shown in Figure 1. Students work in a specially-instrumented document file (e.g., Google Doc), with a coaching chat panel on the right hand side. A typical user session is expected to start with establishing the current goals for the writing product (e.g., brainstorming, revising the thesis statement). First, the coach will suggest an activity and support switching to a different activity if desired (“Last time, we hadn’t talked about the issues raised by opponents of your position. Did you want to start on that, or do something else?”). Second, most activities will start with some pre-writing discussion to establish goals and provide generative AI insights (e.g., analysis of the current text, generating example opposition responses). Third, when appropriate, the coach will mark up the draft document with highlights and comments to establish key areas to review. Finally, when the user is ready with their revisions (or when they want more help from the coach) they will continue chatting with the coach about their writing. This iterative process will be recorded to help inform reports for instructors. This design should also make cheating particularly hard, because a learner would need an AI able to respond to the coach, produce meaningful writing edits in an existing document, and apply these inputs over time at realistic human speeds.

**Figure 1: ABE+ Coaching Diagram**





In this prototype, coaching dialogs will not be directly controlled by a generative AI model, but instead follows a more traditional pattern (e.g, a dialog state manager). This enables the coach to use reliable response templates when appropriate, as well as control the context that is shared to the LLM prompts rather than arbitrarily including the whole conversation history. More generally, we hypothesize that the unconstrained conversations typical of systems such as ChatGPT are not particularly efficient for many educational tasks for two reasons. First, novices often “don’t know what they don’t know” which hinders their ability to ask accurate questions. Second, LLM’s typically try to answer questions directly, but high quality coaches and tutors tend to ask questions or answer with strategies to help find answers instead. That said, while it is beyond our current scope to tune a generative AI to provide a reliable dialog system, certain interactions may involve generated AI conversations (e.g., a chat with a simulated reader who opposes the paper’s position),

**Critical Thinking.** The current focus for ABE+ is college-level writing such as a position paper, as these require logical argumentation but are less structured and specialized than a scientific article. Existing generative AI approaches show multiple ways to encourage critical thinking, such as the AI arguing an opposing viewpoint or pointing out gaps in the claims required for the thesis to hold true. The goal of such activities and feedback will be to reinforce an effective thought process (e.g. Who is the intended audience and why should they keep reading?), rather than suggest specific fixes or changes to the essay content itself. This approach leverages LLM technology primarily as a comprehension model to review and critique an essay, rather than generating writing content itself. In cases where writing content is generated, students will be instructed to critique or compare the content (i.e., analyze) rather than use the content directly.

Based on our co-design with the USC Generative AI Fellows and our exploration of LLM capabilities, we believe that a successful generative AI tool may be able to provide feedback at the level of a good-quality peer review. While not a substitute for an experienced instructor, the ability to receive feedback in seconds rather than days offers a substantial advance from traditional non-interactive assignments such as a self-reflection worksheet. Currently, we are focused on analyzing position paper essay content to support feedback. We have explored feasible prompt engineering and fine-tuning approaches issues, including:

- **Thesis Support:** Generating key claims that should be addressed to support the student’s thesis. These can be used during pre-writing or to compare against the current essay state to identify areas that have not been addressed in detail.
- **Stakeholder Perspectives:** Identifying stakeholders commonly associated with the thesis position. Suggesting where certain claims may be improved, as they have limited perspectives or take a one-sided approach to an issue.
- **Reference Quality:** Identifying claims that should be supported by an external reference and rating authority level of that reference (e.g., an academic citation versus an op-ed).

Our goal is to carefully tune at least ten learning activities to help develop writing skills, as well as a simple interface to create custom activities (e.g., simple prompt chains) to share with students. We are also exploring if synthetic data sets (e.g., essays where we use generative AI to produce known errors) can be used to generate training sets to fine-tune a LLMs which have stronger diagnostic abilities to detect these types of issues and recommend areas to improve.

**Process-Oriented.** To analyze the writing process, the ABE+ framework is designed to plug in with an existing writing tool API (e.g., Google Docs). When active, it tracks changes at a granular level (e.g., any differences on a time interval and whenever ABE+ needs to analyze the latest version). We are also designing the system based around component activities that prompt for reflection and intent: before a student changes the document, they should first talk about why and what they are changing. As a result, these components can leverage the pedagogy that high-quality instructors often use in their assignments, which focus students on different aspects of their draft over time.

Initial implementations are developing tools for the pre-writing and revision stages of student writing, which some Generative AI Fellows have suggested as particularly productive periods to support developing writers. By comparison, Fellows indicated that student critique of AI-generated writing (such as improving the quality of sources or identifying bias) were useful activities, but expressed concerns that critique activities would be most useful after developing stronger intuitions about writing fundamentals (e.g., a strong and compelling thesis). The coach will incorporate these insights by including a suggested order of activities to develop and revise a writing product, but instructors and students will be able to choose their own order as well so that we can study what approaches are preferred and that appear to be more effective.

**Formative Assessment.** The coach will also provide feedback on changes that the student introduces. This feedback could be provided in real-time or could be delayed to a higher-level summary. Instructors are particularly interested in summaries of the changes that students make, particularly in terms of the reasons for changes (i.e., connections between self-reported intentions in the coach chat versus the changes that were made in the document). This information is currently unavailable at a detailed level, as students tend to deliver a document with many kinds of revisions.

**Challenges.** One challenge for ABE+ is that assessing the improvement of students' writing process (as opposed to their outputs) has not been deeply analyzed at the level of granularity we can collect. Writing instructors traditionally see only a handful of drafts from a student on any one assignment, so it is less clear what trends and patterns suggest a student is acquiring new skills even if the end product of writing is not yet high quality. Research on AI intelligent tutoring systems such as math and science have analyzed transcripts from human tutoring dialogs, but less data has been recorded on effective coaching during a writing session. However, by summarizing the history of changes that a student makes, researchers should be able to identify such patterns that inform both the system and our understanding of the writing process overall.

## Opportunities, Challenges, and Future Directions

The exploration of Generative AI in education, particularly in writing instruction, reveals a landscape rich with both opportunities and challenges. The collaboration between the USC Center for Generative AI and Society and the USC Writing Program is a prime example of proactive engagement in this domain. We are focused on co-designing tools that not only facilitate critical thinking and writing skills but also address the practical constraints of traditional writing instruction, which often emphasizes the final product over the writing process.

The academic community is still charting the potential of the latest Generative AI models, such as GPT-4 and LLAMA 2, in education. These models have shown improvements in quality and effectiveness, opening new avenues for educational applications. The emphasis has been on feasibility - the potential of Generative AI to enhance education in practical settings - rather than just on efficacy, which requires more rigorous and controlled evaluations. The ABE+ prototype, developed under this initiative, encapsulates an approach that focuses on supporting critical thinking, process-oriented learning, and formative assessment. It aims to support various stages of the writing process, from brainstorming to revision, by providing real-time coaching and feedback. As the field evolves, the sharing of prompts, activities, and strategies among educators and researchers will be crucial for the continued growth and effective integration of Generative AI in educational contexts. Our center will support such efforts.

### References

- Bonner, E., Lege, R., & Frazier, E. (2023). Large Language Model-Based Artificial Intelligence in the Language Classroom: Practical Ideas for Teaching. *Teaching English with Technology*, 23(1) 23-41. <https://doi.org/10.56297/BKAM1691/WIEO1749>
- Butterfuss, R., Roscoe, R. D., Allen, L. K., McCarthy, K. S., & McNamara, D. S. (2022). Strategy uptake in Writing Pal: Adaptive feedback and instruction. *Journal of Educational Computing Research*, 60(3), 696-721.
- Kim, Y., & Baylor, A. L. (2016). based design of pedagogical agent roles: A review, progress, and recommendations. *International Journal of Artificial Intelligence in Education*, 26, 160-169.
- Krathwohl, D. R. (2002). A revision of Bloom's taxonomy: An overview. *Theory into practice*, 41(4), 212-218.
- Harvard Derek Bok Center for Teaching and Learning (2023). *Generative Artificial Intelligence and Writing Assignments*. Retrieved Oct 28, 2023 from: [https://firstyearseminarprogram.college.harvard.edu/sites/projects.iq.harvard.edu/files/freshmanseminars2/files/bok\\_ctr-a.i.\\_and\\_writing\\_assignments.pdf](https://firstyearseminarprogram.college.harvard.edu/sites/projects.iq.harvard.edu/files/freshmanseminars2/files/bok_ctr-a.i._and_writing_assignments.pdf)
- Hsiao, Y. P., Klijn, N., & Chiu, M. S. (2023). Developing a framework to re-design writing assignment assessment for the era of Large Language Models. *Learning: Research and Practice*, 1-11.
- Marino, M. (2023). *Secrets to Prompting ChatGPT*. Retrieved Oct 28, 2023 from: <https://markcmarino.medium.com/secrets-to-prompting-chatgpt-d631b13d8202>
- Nicula, B., Dascalu, M., Arner, T., Balyan, R., & McNamara, D. S. (2023). Automated Assessment of Comprehension Strategies from Self-Explanations Using LLMs. *Information*, 14(10), 567.
- OpenAI (2023). GPT-4 Technical Report. ArXiv, abs/2303.08774.
- Mills, A. & Goodlad, L.M.E. (April 17, 2023). *Adapting College Writing for the Age of Large Language Models such as ChatGPT: Some Next Steps for Educators*. Retrieved Oct 28, 2023 from: <https://criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/>
- Pardos, Z. A., & Bhandari, S. (2023). Learning gain differences between ChatGPT and human tutor generated algebra hints. *arXiv preprint arXiv:2302.06871*.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., ... & Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, A., Prihar, E., & Heffernan, N. (2023). Assessing the Quality of Large Language Models in Generating Mathematics Explanations. *Learning@ Scale (L@ S'23), July 20-22, 2023, Copenhagen, Denmark*.

