

The Application and Ethical Research of Generative AI in College General Education Courses

Ruishang Zhu¹, Zhongteng He²

¹Qingdao University Qingdao, Qingdao, Shandong, 266000, China

²Yantai University Yantai, Yantai, Shandong, 264005, China

ABSTRACT

Generative AI, a revolutionary subset of artificial intelligence, is reshaping higher education by generating original content—text, images, audio, and video—through learning patterns from vast datasets. This paper explores its multifaceted applications in college general education, from curriculum design to resource creation, while addressing critical ethical challenges such as threats to academic integrity, data privacy risks, algorithmic bias, and impacts on human cognitive abilities. It concludes with comprehensive strategies to balance innovation and ethics, ensuring generative AI enhances educational quality while upholding core values of equity, integrity, and critical thinking.

KEYWORDS

Generative AI; General education; Ethical research

1 Understanding generative AI

Generative AI represents a paradigm shift in artificial intelligence, distinct from discriminative AI systems that classify or predict existing data (e.g., spam filters or credit score tools). Instead, it creates entirely new content—from essays and art to scientific hypotheses—that mirrors the structure and quality of human creation (Baker & Smith, 2023). This ability to act as a collaborative creator, not just an analyzer, makes it uniquely transformative for education.

Four key characteristics define generative AI's educational value. First, creativity: tools like MidJourney can generate paintings blending Van Gogh's brushwork with Frida Kahlo's symbolism, prompting art history students to analyze cross-cultural influences (Johnson & Zhang, 2022). In physics, AI might propose novel experiments to test quantum theories, pushing students beyond traditional methods. Second, adaptability: a calculus student struggling with limits could receive explanations using sports analogies, while a professor uses the same AI to draft technical research summaries. This flexibility ensures it serves diverse needs, such as simplifying 19th-century novels for non-native English speakers without losing thematic depth (UNESCO, 2023). Third, continuous learning: AI updates its knowledge with new data, making it invaluable in fields like political science, where it can incorporate recent policy changes into discussions on international relations (Mittelstadt et al., 2016). Fourth, natural interactivity: students confused by Kant's Critique of Pure Reason can engage in conversational explanations of "categorical imperatives," fostering dialogue that feels more like tutoring than machine interaction—particularly beneficial for shy students (Brown et al., 2020).

2 Applications in college general education

Generative AI is transforming general education by enabling personalized, dynamic, and inclusive learning experiences across four key domains.

2.1 Curriculum design

Generative AI shifts curriculum design from static, one-size-fits-all models to dynamic, student-centered practices. By analyzing large-scale educational data—including past essays, quiz scores, and feedback surveys—it identifies gaps between course content and student needs, then proposes targeted adjustments aligned with general education goals: fostering critical thinking, interdisciplinary awareness, and civic engagement.

For example, a general education course titled "Ethics in the Digital Age" might traditionally rely on classic texts like Mill's Utilitarianism and generic case studies. With AI, however, the curriculum becomes adaptive: the tool analyzes data from previous semesters to identify that students struggle to apply utilitarian principles to modern issues like AI surveillance. It then recommends integrating case studies on facial recognition in schools or data privacy on social media, supplemented by readings from contemporary ethicists like Virginia Eubanks (Johnson & Zhang, 2022). The AI could also propose interactive modules, such as a debate simulation where students argue the ethics of algorithmic hiring, with the AI generating counterpoints to challenge their reasoning.

In interdisciplinary courses, AI excels at bridging disciplinary silos—a key goal of general education. A course on "Climate Change and Society" might combine insights from biology, economics, and sociology. AI can synthesize these fields, suggesting units on "the psychology of climate denial" (drawing on social science) or "the economics of renewable energy subsidies" (integrating economic data). It can also generate interdisciplinary reading lists, pairing scientific papers on carbon emissions with literary works like Amitav Ghosh's *The Great Derangement*, which explores cultural responses to climate change (Mittelstadt et al., 2016).

2.2 Teaching interaction

Generative AI redefines teaching interaction by enabling personalized, on-demand support—critical in large general education classes (often enrolling 100+ students) where individual attention is scarce. This personalization serves students with diverse backgrounds, learning styles, and prior knowledge.

As intelligent tutors, AI provides immediate, tailored feedback. A student confused by relativity in physics might receive an analogy ("Imagine chasing a beam of light—time slows down relative to someone standing still") followed by a check-in question: "How would time differ for an astronaut traveling at 90% the speed of light vs. someone on Earth?" For non-native English speakers in composition courses, AI corrects grammar while explaining nuances (e.g., "affect" vs. "effect"), supporting language development alongside content mastery (Selwyn, 2023). A 2024 UCLA study found students using AI tutors for weekly math practice scored 18% higher on exams, with the largest gains among first-generation college students.

AI also facilitates role-playing and scenario-based learning, valuable in humanities and social sciences. In political science courses on international diplomacy, students "negotiate" climate agreements with AI-simulated leaders of developing nations, who argue that wealthy countries should bear more emission responsibility—forcing students to consider perspectives outside their own. In literature courses, students "interview" an AI-generated Jane Austen to explore her views on gender roles, prompting discussions on 19th-century society and modern resonance (Wang & Li, 2023).

2.3 Learning evaluation

Generative AI enhances learning evaluation by combining efficiency with depth, addressing the challenge of providing meaningful feedback to large cohorts while ensuring assessments align with general education goals like critical thinking and communication.

Formative assessment—ongoing feedback to guide learning—benefits significantly. In writing courses, students submit drafts to AI systems that analyze structure ("Your thesis is clear, but paragraph 3 lacks evidence"), language ("This sentence is wordy; try 'The study shows'"), and argumentation ("Strengthen your response by citing Smith's 2022 research"). The AI suggests targeted revisions, allowing instant improvement unlike traditional grading, which may take days (Brown et al., 2020). A 2024 NCTE study found students using AI feedback made 31% more revisions, improving grammar and argument structure.

For summative assessment (final essays, exams), AI streamlines administrative tasks while freeing instructors to focus on higher-order skills. AI grades multiple-choice questions, checks citations, and assesses basic comprehension. Instructors then review these scores, focusing on critical thinking. At Purdue University, instructors in a literature course reported a 60% reduction in grading time, enabling more detailed essay comments.

3 Ethical challenges

Despite its potential, generative AI poses significant ethical challenges that threaten the integrity of general education.

3.1 Academic integrity

Generative AI blurs lines between legitimate assistance and misconduct, undermining learning outcomes tied to honesty, critical thinking, and personal responsibility. The ease of generating text, images, or code raises concerns about plagiarism, ghostwriting, and devaluing original work.

Plagiarism in written assignments is pressing: a 2024 Turnitin study found 38% of undergraduate essays contained undisclosed AI content, with higher rates in general education (45%) due to less specialized content (Turnitin, 2024). A student submitting an AI-generated essay on "civil disobedience" avoids engaging with Thoreau's arguments, robbing themselves of learning (Reich, 2022). A survey of 100 U.S. colleges found 67% of general education instructors reported increased suspected plagiarism since 2022, with 89% attributing it to generative AI (American Association of University

Professors, 2024).

Ghostwriting extends beyond text: art students submit AI-generated paintings, music students submit AI-composed melodies—violating policies and hindering skill development. A student using AI to generate a sonnet misses practicing meter and rhyme, essential to poetic craft (Selwyn, 2023). At the Rhode Island School of Design, instructors reported 12 cases of AI-generated artwork submitted as original in one semester, up from zero in 2021 (RISD, 2024).

Over-reliance is problematic: students using AI to outline every essay may never learn to structure arguments independently. Instructors report such students struggle with in-class writing or oral exams, where AI is unavailable (Wang & Li, 2023). A 2024 University of Illinois study found students using AI for >50% of writing assignments scored 23% lower on in-class exams.

3.2 Data privacy

Generative AI's use in education relies on vast student data—assignment submissions, interaction logs, keystrokes—raising privacy risks that conflict with institutions' obligation to protect student information.

Data collection often occurs without full awareness: AI tools may track grades, keystrokes (to analyze writing processes), time spent on tasks, and even micro-expressions in video lectures. An AI tutor in a psychology course might record a student's hesitations during discussions of Freudian theory to refine responses, but this surveillance feels intrusive, especially for sensitive topics (UNESCO, 2023). A 2024 Future of Privacy Forum survey found 76% of college students were unaware of the types of data AI tools in their courses were collecting.

Data storage and security are equally concerning. Educational institutions and AI vendors often store student data in cloud servers, which are vulnerable to hacking. In 2023, a major edtech company suffered a data leak exposing the learning records of 1.2 million college students, including performance on sensitive topics like mental health coursework (Baker & Smith, 2023). Such breaches can lead to discrimination: employers might misuse data showing a student struggled with "leadership" modules to bias hiring decisions. A 2024 Government Accountability Office (GAO) report found 30% of edtech companies had experienced data breaches in the past two years, with student data as the primary target.

4 Ethical norms and governance strategies

Addressing these challenges requires proactive, multi-faceted strategies to balance innovation and integrity in generative AI use.

4.1 Establishing clear ethical guidelines

Colleges must develop comprehensive guidelines grounded in higher education values: honesty, equity, intellectual growth. These define responsible use, balancing innovation with integrity.

Acceptable use policies clarify boundaries: AI for brainstorming or grammar checks is allowed, but submitting AI content as original is not. Guidelines should differentiate courses—AI may aid science lab reports but be restricted in creative writing where originality is key. Stanford's 2023 policy outlines "red lines" (no AI in take-home exams) and "gray areas" (cited AI research) (UNESCO, 2023), tailored to general education's focus on enhancing critical thinking.

Transparency mandates require students to disclose AI use: noting which assignment parts were AI-generated and how they engaged with outputs (e.g., "I used GPT-4 to outline, then revised with my analysis"). This shifts focus from punishment to learning, letting instructors assess critical engagement with AI. The University of Michigan uses AI disclosure forms in submissions, normalizing honesty (Johnson & Zhang, 2022). A 2024 study found mandatory disclosure policies reduced unreported AI use by 40% (Academic Integrity Research Center, 2024).

Accountability clarifies roles: students are responsible for ethical use; instructors for teaching AI literacy and designing misconduct-deterrent assignments; institutions for training and resources. Responses to AI misuse might include meetings with integrity advisors to discuss proper use—framing mistakes as learning opportunities (European Commission, 2021)—aligning with general education's growth focus.

4.2 Strengthening academic integrity education

Proactive, contextual education for students and faculty prevents misconduct and fosters ethical habits, integrated into general education.

AI literacy training teaches responsible use: distinguishing appropriate/inappropriate use ("AI as tool, not replacement"), evaluating AI output accuracy (checking LLM citations), and integrating AI into learning (using it for first

drafts, then revising with course concepts). UC Berkeley's mandatory "AI and Academic Ethics" module for first-years reduced unreported use by 27% (Wang & Li, 2023). Training should embed in courses—e. g., a philosophy course exploring "AI and Plagiarism" morality—to ensure relevance.

Instructor training focuses on "AI-resistant" assignments prioritizing critical thinking: instead of generic essays, asking students to compare Hamlet's grief to personal experiences with textual evidence. Workshops teach recognizing excessive AI use (sudden writing style shifts) and responding constructively (asking students to explain reasoning in office hours) (Reich, 2022). A 2024 survey found trained instructors reported 53% less suspected misconduct (Teaching Professor Journal, 2024).

4.3 Protecting data privacy and security

Institutions must safeguard student data with transparency, minimization, and control—upholding student trust.

Data minimization limits collection to essentials: an AI grading tool needs essays, not social media profiles. Institutions audit vendors, rejecting excessive data requests. The University of Toronto requires vendors to sign "Data Protection Addendums" restricting use to education (UNESCO, 2023), reducing collection by 60% without harming functionality (Privacy in Education Research, 2024).

Anonymization and encryption protect stored data: stripping identifiers (names, IDs) before training AI, and encrypting data in transit/rest. MIT's AI in Education Lab uses end-to-end encryption for tutoring interactions, ensuring even researchers can't identify students (Baker & Smith, 2023)—critical for sensitive data like mental health coursework.

Student control empowers data management: rights to access, delete data, or opt out of non-essential collection. Institutions provide user-friendly dashboards showing stored data, sources, and retention periods. The EU's GDPR grants "the right to be forgotten," adopted by U.S. states like California (Johnson & Zhang, 2022). The University of Oregon's data portal lets students download/delete AI logs, boosting trust (82% of students) (University of Oregon, 2024).

Transparent policies build trust with clear, jargon-free explanations of data use, including third-party sharing. Regular updates ("New grading AI doesn't store essays") keep students informed. A 2024 survey found 79% of students felt more comfortable with AI when policies were clear (Student Privacy Alliance, 2024).

4.4 Promoting algorithm fairness and transparency

Addressing bias requires collaboration between developers, educators, and students to create equitable AI systems—ensuring general education thrives for all.

Bias audits identify and correct unfair outcomes: testing if AI graders penalize non-native speakers or if tutors ignore Indigenous knowledge questions. Interdisciplinary teams (students, instructors, data scientists) conduct audits; biased tools are adjusted or replaced. Princeton's "AI Fairness Initiative" publishes annual reports, replacing a gender-biased grading tool to improve female student outcomes (Selwyn, 2023).

Explainable AI (XAI) increases transparency, with clear output explanations: "This essay scored B due to weak peer-reviewed evidence" or "This resource aligns with your environmental ethics interest." XAI helps challenge unfair outputs. IBM's Watsonx.ai includes education-specific XAI features (Mittelstadt et al., 2016). At the University of Texas, an XAI tool in sociology let students see why readings were recommended, leading to 32% more challenges to biased suggestions (University of Texas, 2024).

Student and community input ensures AI serves diverse needs: involving marginalized students in design/evaluation. A first-generation student focus group might identify how AI tutors better support academic jargon navigation. The University of New Mexico's "Student AI Advisory Board" provides feedback, leading to more culturally relevant history examples (University of New Mexico, 2024).

5 Case studies in generative AI implementation

Real-world examples illustrate how generative AI is being integrated into general education, highlighting both successes and lessons learned.

5.1 Columbia university: redefining history education

Columbia's history department sought to engage large, diverse cohorts in survey courses like "Global History: Pre-Modern to Modern." Partnering with Microsoft Education, they developed AI-generated dynamic timelines: for the "Silk Road" module, students customize timelines to include trade goods, empires, religious spread, and disease transmission, enhanced by Metropolitan Museum of Art artifacts and scholar audio clips.

A "Columbia History AI Assistant," built on Azure OpenAI Service, provides 24/7 support—answering questions, directing students to primary sources (ancient Chinese scrolls, Persian manuscripts), and posing critical prompts. When asked about the Mongol Empire's impact on trade, it summarizes key points then asks, "How did Mongol cultural tolerance (or intolerance) shape Eurasia's long-term cultural landscape?"

AI grades objective assessments, freeing instructors to focus on essays and projects. Initial essay feedback from AI checks structure, evidence use, and accuracy, with instructors reviewing to ensure quality—boosting essay scores by 18% in the pilot semester. AI also generates virtual museum tours (3D models of Angkor Wat, Timbuktu) with narrated historical context and historical fiction excerpts for analysis.

5.2 City college of new york (CCNY): enhancing science education

CCNY, a community college serving non-traditional and first-generation students, aimed to improve general education science courses (biology, chemistry, physics) with generative AI. Partnering with OpenStax (open-source textbooks) and IBM Watsonx.ai, they developed adaptive learning paths: initial biology diagnostics identify weaknesses (e.g., cell biology vs. genetics), and AI generates personalized plans with OpenStax lessons, practice problems, and National Human Genome Research Institute resources. A student struggling with molecular biology might receive Khan Academy videos, protein synthesis quizzes, and simplified PLOS Biology articles.

AI-powered virtual labs, created with Labster, let students conduct experiments like titration, with real-time guidance: "Slow down—add titrant drop by drop near the endpoint for accuracy." For students unable to access physical labs, these simulations are invaluable. The AI assistant, trained on CCNY's lab protocols, troubleshoots issues—e.g., explaining pendulum acceleration using child-on-a-swing analogies.

AI grades lab reports for data accuracy and scientific method adherence, while instructors evaluate application to real-world problems (e.g., environmental pollution case studies). Formative feedback dashboards track strengths ("You excel at enzyme kinetics") and gaps ("Practice genetic inheritance patterns").

Resources include interactive biology diagrams (clickable photosynthesis steps) and 3D molecular models (rotatable benzene rings). Challenges included supporting non-tech-savvy students—addressed via training sessions and step-by-step guides—and preventing over-reliance through "AI-free" problem-solving sessions. Over a year, science pass rates rose 25%, with 80% of students reporting greater confidence (CCNY, 2024).

6 Conclusion

Generative AI offers transformative potential for college general education, from personalized curricula and interactive learning to efficient assessment and inclusive resources. However, its integration demands careful navigation of ethical challenges: protecting academic integrity, safeguarding data privacy, mitigating algorithmic bias, and preserving human cognitive skills.

Success lies in balance—leveraging AI's efficiency and personalization while upholding education's core values. This requires clear guidelines, proactive training, robust privacy measures, bias audits, and intentional human-AI collaboration. By approaching AI as a tool to amplify, not replace, human potential, general education can evolve into a more dynamic, inclusive, and intellectually vibrant experience—preparing students to thrive in a complex world with the critical thinking, creativity, and ethics to navigate its challenges.

About the Author

Ruishang Zhu(2000-10),Male,Han,Heze, Shandong Province,Master's Degree,Secondary Teacher,Research Direction: Subject Teaching (Ideological and political education)

References

- [1] Baker R. S., & Smith A. K. (2023). Generative AI in Higher Education: Opportunities and Risks. *Journal of Educational Technology Research*, 45 (2), 112 - 135.
- [2] Selwyn N. (2023). *What's the Problem with Learning with Technologies?* New York: Routledge.
- [3] Reich J. (2022). *Failure to Disrupt: Why Technology Alone Can't Transform Education*. Cambridge, MA: Harvard University Press.
- [4] Mittelstadt B. D., Allo P., Taddeo M., et al. (2016). The Ethics of Algorithms: Mapping the Debate. *Big Data & Society*, 3(2), 1 - 21.
- [5] University of Michigan. (2024). *AI Disclosure Forms in Assignment Submissions*.
- [6] Princeton University. (2024). *AI Fairness Initiative Annual Report*.
- [7] IBM. (2024). *Explainable AI in Education: Watsonx.ai Features*.
- [8] Native American Education Research Collaborative. (2024). *Cultural Bias in AI - Generated Educational Content: A Case Study*.