SAMYRAD 2026

SMART ALGORITHMS & MODELS: YOUR REACH TO ADVANCED DIALOGUES · 5–6 OCTOBER 2026 · SEVILLE, SPAIN

SAMYRAD OPEN-ACCESS PROCEEDINGS | PUBLISHER: INTELLIGENT NUDGE

PAPER • OPEN ACCESS

Cultivating Faculty's AI Literacy through the Use of Open-Access Platform

Alexa Alice Joubin

(e-mail: ajoubin@gwu.edu)

George Washington University, DC 20052, United States of America

Article DOI: https://doi.org/10.68286/samyrad2026.353

ABSTRACT

This design-based research investigates the development of faculty's self-efficacy in understanding AI's pedagogical and technological implications. Generative AI (gen AI) is a pluralistic boundary object in higher education. Faculty disagree on how to approach it, but they tend to agree that its presence cannot be ignored. However, wide-spread discussion of the abstract concept does not correlate to actionable knowledge. Studies of literacy indicate that high fluency often masks low adequacy. This project's purpose is to give faculty first-hand experiences of designing AI in order to enhance literacy. We do so by building the open-access Teach Anything platform as a trust inducing environment. It provides backend access and testing of several open-source large language models (LLM), and not just one LLM. Faculty use their own files to design their own applications in the context of experimentation or in their courses. We use Teach Anything to test the hypothesis that hands-on knowledge empowers faculty. Our key finding is that educators are empowered by this community-oriented, participatory approach, because they are proactive designers of AI, rather than passive consumers. They curate the retrieval corpus that grounds their open-source AI applications. They incorporate their new knowledge into adaptive teaching and students' learning.

KEYWORDS open source, open access, generative AI, educator literacy, participatory design

How to cite: Joubin, A. A. (2026). Cultivating faculty's AI literacy through the use of open-access platform. SAMYRAD 2026 Open-Access Proceedings. Intelligent Nudge. https://doi.org/10.68286/samyrad2026.353

Introduction

In higher education, gen AI is happening to faculty and students rather than with them or for them, leading to social and financial inequalities, among other spillovers. Multiple AI tools are imposed on campuses by administration and large corporations. AI has diluted and diminished the value of learning experiences, with its misinformation affecting the trustworthiness in human behaviors and AI-infused society. Academics can be fluent in talking about AI, but their deeper competence and pedagogical integration remain limited [1]. In cross-disciplinary surveys, educators describe repositioning themselves as curators and verifiers of AI output rather than confident designers of it [2], and another study records high faculty awareness of generative AI alongside markedly lower applied competence [3]. We need an inclusive environment and a safe testing ground for literacy-driven trustworthy behaviors that push beyond stock arguments for or against gen AI. Trust emerges from gaining access to processes rather than merely accepting the verifiability of explicit answers.

The National Institute of Standards and Technology defines trust as a property of a system through metrology [4], defining trustworthiness through metrics of transparency, explainability, and interpretability, assessing an AI system's adherence to precursors of trust. This in part has led to the techno-solutionism in educational AI design that assumes trust will emerge from systems that produce predictable performance. I hypothesize that trust in AI-infused environment is not just a technical issue, but also a relational question [5]. Trust is not only about explainability or interpretability of a model. It concerns fostering deep social connections between educators and students. In this dynamic environment with many stakeholders, trust is not a property of an object but rather a type of relation and action [6]. Shifting our perspective on generative artificial intelligence (gen AI) in this way helps reduce misplaced trust in AI's outputs while enabling students to engage with AI as a tool for interpretation and exploration rather than a product for passive consumption.

My primary research question is: How do faculty, as designers of AI, develop AI literacy? My hypothesis is that faculty's self-efficacy in understanding AI's pedagogical and technological implications will be enhanced by hands-on experiences, unavailable through commercial AIs. No commercial AI platforms allow full backend access for faculty. Customization of AI through API or fine-tuning LLMs has a high barrier of technical knowledge that most faculty do not possess. To test this hypothesis, Alexa Alice Joubin and Akhilesh Rangani built the open-access Teach Anything, https://www.teachanything.ai, as a "sandbox" for experimentation. Teach Anything gives educators access to the configuration and operation of open models to enhance meta-cognition skills. Faculty do not need any prior technical knowledge to use Teach Anything to design their own AI. Once faculty's AI literacy develops in a controlled environment, they will be in a better position to foster trustworthy behaviors in their classrooms. They can then incorporate their new knowledge into their design of activities and assignments.

Background and related work

Faculty Report Low Self-efficacy

While a great deal of research in education focuses on AI's techno-solutionist effects [7], it often fails to identify how it can do harm and how educators might counter it.

Faculty predominantly lack AI literacy and opportunities to customize existing AI tools, with the majorities of "college and university leaders" indicating that "their institutions are not very or not at all ready to use GenAI for such things as: preparing students for the future, 56% say their schools are not prepared for this" [8, p. 3]. A multi-university study identified low self-reported AI literacy among faculty [9]. AI competencies are central determinants of educator AI literacy [1], and AI literacy is the strongest predictor of pedagogically relevant use or rejection of generative AI [10]. A systematic review finds learner autonomy and literacy emerging as central pedagogical concerns, yet documents little attention to faculty as designers of AI rather than users of it [11].

Commercial AI Systems Inadequate

While commercial AI chatbots provide immediate responses to students' questions, their design often causes the loss of human nuance and discourages critical thinking, a problem that is corroborated by controlled evidence [12]. A study of a freely available tutor built on metacognitive scaffolding found that pedagogical design, not model fine-tuning, determined whether students learned or merely offloaded the task [13]. An evaluation of ChatGPT as a teaching assistant in introductory programming found it fluent yet inclined to supply finished answers where guided questioning was pedagogically required [14]. Effective communication today requires critical thinking to navigate diverse perspectives [15], foster relational understanding [16], and reflect on personal and others' worldviews [17]. These are essential skills for academic and career success [18] [19]. Research confirms that gen AI frequently acts as a "power persuader," shaping professional judgment and disputing human's validation [20].

Critical Thinking in Higher Education

This highlights the urgent need for an open-access and customizable AI system to enhance critical thinking. Drawing on research on the positive effect of AI-framed questioning [21], we built Teach Anything to offer adaptive responses and ask students questions in order to scaffold their critical thinking process. Rather than functioning merely as an information retrieval tool, the faculty-designed AI applications on our platform can stimulate critical thinking.

Why We Built Teach Anything

Teach Anything addresses several of these problems. Some faculty regard AI as an employable skill for students, while others reject AI because it disrupts academic integrity. Faculty cannot agree on what the problem is, and each solution has a trade-off [22], [23]. Higher education does not intrinsically need AI, but faculty need space for experimentation. We need an inclusive environment to share untold stories and to tell new stories about society and civic technology.

Built by educators for educators, Teach Anything is a permanently free platform for professors to design their own open-access AI. No technical knowledge is required for faculty to use it.

Teach Anything, https://www.teachanything.ai, is open source and open access. It models ethical practices in human-AI collaboration and uses inclusive design principles. Teach Anything is a civic science project in which faculty and students collaborate on a solution that is public interest technology. Civic science empowers citizens to actively shape technologies that impact their life.

Teach Anything's goal is to strengthen faculty's ability to independently evaluate AI systems and to make decisions about pedagogy with or against gen AI in our times. Teach Anything's humanistic methods and community-oriented principles move away from a binaristic model of humans versus machines (or consumers versus corporations) to hypothesize an ecology of ideas suitable for higher education contexts. Teach Anything prioritizes positive outcome for educators and students.

Moreover, all AI applications currently in use in education are commercial and were not designed specifically for the purposes of teaching or learning. Commercial LLMs generally do not allow changes to their parameters. Free tier AI

applications are not open-access resources, because they set usage limits, run advertisements, and turn the users' data into products. Teach Anything is free and open access. It does not use faculty data or students' chats to train any LLMs.

Since faculty build their custom AI applications on Teach Anything with their own data and system prompts, they have more control over its behavior than a commercial product allows. Control here operates at the application layer: over the system prompt, the curated corpus, and the choice of base model. It does not extend to the pre-training of the open-weight model itself, a distinction that matters when open and proprietary systems are compared directly [24]. The retrieval corpus for their applications is transparent and intentioned, not indiscriminately scraped from unvetted sources.

We are committed to open culture. Faculty users are not locked into any one company's ecosystem. They can freely experiment with more than 7 medium-weight, opensource LLMs (such as Mistral and Llama, between 10-100 billion training parameters) and switch them at will. The pre-trained parameters of these open-weight models are publicly available for developers to fine tune, though Teach Anything itself does not require or perform fine-tuning.

Unlike for-profit, proprietary systems, Teach Anything provides transparency and inclusivity by allowing professors to curate, inspect and modify the retrieval corpus, the system prompt, and the generation settings, including the temperature (or creativity) and length of AI's outputs. This approach not only democratizes advanced technologies but also fosters a collaborative community of educators and researchers committed to ethical AI implementation. Transparency in AI design ensures that the tools developed remain accountable. While commercial culture is prone to simplification and sloganization [25], open culture encourages nuanced deliberation [26].

Key Benefits

Faculty as AI designers can act as "translator-mediators" [27] rather than victimized consumers. This is the case because faculty have considerably more control over the retrieval corpus and their applications' behaviors than any commercial platform permits. It is an important first step toward building literacy to learn to curate datasets, configure a tool, and customize system prompts. Open-weight models are not themselves exempt from biases. A prospective controlled evaluation of an open model against a proprietary one found differences in factual reliability running in both directions [24]. However, instructors do not need full control of the underlying foundation model to develop literacy. Through their pre-defined corpus for retrieval and system prompts, instructors could minimize biases and misinformation.

Through Teach Anything, what faculty deploy is not a base model but an application layer over open-source LLMs. Through a curated corpus and comprehensive, domain-specific system prompts, faculty govern the model's behavior which the instructors can inspect and improve upon. The open-access Teach Anything helps faculty develop AI literacy in order to guide students effectively in responsible AI use and rejection [28].

In building their own "narrow" AI applications, faculty create and define their corpus for retrieval. In contrast, commercial AI scrapes data from the web in a nontransparent manner [29], raising privacy, copyright and trust issues. Those LLMs are trained on datasets with inherent biases of white, masculine, and Western perspectives baked into them [30], [31]. Further, university faculty and students are marginalized in deployment decisions of AI [32], [33]. This problem emerged from depriving faculty of their agency and from training data that has been indiscriminately sourced from the web regardless of quality [34]. Research characterizes these harms as ongoing asymmetries rather than incidental errors [35].

Our argument, however, does not rest on the base model. Using our system, instructors choose from available open-source LLMs to build a customized application. Comparative research supports our claim: educational value is determined by pedagogical design rather than by model access [13]; retrieval errors concentrate where the supplied context is inappropriate or insufficient [36]; and the trade-offs between open and proprietary systems rest on how each is configured for its domain [37]. Biases within open-source LLMs are minimized by the instructors' dataset, system prompts, and retrieval instructions. For instance, instructors could program Teach Anything to respond to particular types of questions by drawing on only designated files.

Key steps for faculty to build their own AI

A Gold open access, SOC2-compliant platform, Teach Anything supports the creation of custom-configured, multilingual AI chatbots by faculty partners. The architecture is modular, and instructors, rather than vendors, control every layer of it. Files the faculty member uploads are extracted, indexed, and embedded into a retrieval corpus, and a retrieval-augmented generation (RAG) pipeline grounds each response in that corpus rather than in the open web. A discipline-specific system prompt governs how the application behaves—whether it answers, withholds, or asks a question in return. Faculty select from more than seven open-weight models and can switch between them at any point, including after deployment, so that no application is bound to one vendor's model. A temperature control from 0 to 100 calibrates how conservative or exploratory the output is, and a token setting determines how long each response runs. None of these operations trains or fine-tunes a model; they are acts of curation and configuration, and it is at this layer, rather than in the weights of the base model, that pedagogy is decided [13], [36]. Without technical knowledge, faculty in any discipline could follow these steps to create and experiment with their custom-configured, domain-specific applications:

(1) Faculty would begin by selecting one of the medium-weight, open-source LLMs. They can swap the LLM at any time even after their application is deployed.

(2) Faculty would upload their files as the foundational data set. When interacting with students, the AI provides citations down to the page. It refers to specific sections of specific files.

(3) Faculty can also use Teach Anything's web crawler to index web resources, such as open education resources (OER) or open-source materials.

(4) Faculty would then write, and refine, their custom system prompts in any of the supported natural languages. They could design their AI to be a provocateur or a Socratic partner that asks students questions rather than write essays for students.

(5) Faculty can adjust their AI's "temperature" of creativity. They can set the number of tokens to control how long their AI's output would be in each instance.

(6) When completed, faculty can easily share their AI via permalinks that do not require students to login. The permalink can, for example, be added to existing learning management systems (LMS).

(7) Faculty also has the option of using the automatically generated HTML or Next.js code to embed their chatbot as a floating button on any website.

(8) After implementation, faculty can check, on their dashboard, anonymized student chat records. There are also analytics available, including commonly asked questions by students, outputs using source content (RAG hit rate), average duration of chat sessions, and low-confidence queries. Faculty use this insight to adjust their pedagogy and refine their AI.

The faculty-designed chatbots will only reference the uploaded files and the faculty designated web sources. Keeping the inference window small significantly reduces the chances for factual errors and fabricated information. The baseline risk is substantial. Measured hallucination rates in systematic-review tasks reach 39.6% for GPT-3.5, 28.6% for GPT-4, and 91.4% for Bard. The primary problem is fabricated information [38]. Invented sources and quotations are especially corrosive in humanities teaching, where novices are least equipped to detect them [39]. Constraining the model through system prompts and a bounded corpus reduces this risk significantly even if it does not completely remove it [36]. The corpus includes faculty's files and web sources designated by faculty. Our web crawler is designed with the ethical principle to respect robots.txt, a file at each website's root folder. This file directs web crawlers on the paths they may access and paths that are off-limits. Honoring robots.txt and the site owner's wishes is standard for ethical crawling.

These chatbots, operating in 58 languages with textual and voice input, can offer adaptive responses to suit the learner's needs, taking as a cue the ways in which the student frames a question. Students' full privacy is guaranteed. Faculty can upload material and write system prompts in different languages. Their AI applications will respond in the language which students choose to use in their queries.

Teach Anything uses only open-source, mediumweight LLMs. When processing student queries and generating outputs, Teach Anything uses both the embedding and retrieval augmented generation (RAG) methods to ground responses in pre-defined, curated content. This smaller inference window enables more accurate operations, because the faculty-designed applications reference small, pre-defined datasets and follow specific system-level instructions tuned for pedagogical benefits. A small corpus with clear boundaries is more effective for RAG than a large dataset. Research on retrieval systems indicates that errors occur when the retrieved context is unclear to answer the query [36]. Context-aware retrieval architectures are shown to be more effective than generalists tools, such as ChatGPT, in specialist domains [40]. After generating each output, the chatbot reveals its reasoning process for transparency and pedagogical benefits. Students enjoy complete privacy because Teach Anything does not store their personal information or preferences. Their interactions are not sent off to train the LLM. Commercial LLMs are trained on stolen texts that are not disclosed or authorized. Further, users' chats are used by AI vendors to train their LLM. Our system is transparent and ensures data privacy and security.

Commercial gen AI products store usage history that harms the users [41]. In fact, prolonged interaction with commercial AI systems that store user information increases the models' digital sycophancy [42]. These models mirror users' views regardless of whether those perspectives are beneficial to users' education. Teach Anything is designed to be productively "amnestic" in that it does not store student profiles or personal preferences. This feature prevents the issue of sycophancy common in commercial AI products. Interacting with digital platforms without leaving traces is a rare respite in the AI-dominated world. Every time a student interacts with a faculty-designed AI on Teach Anything, it is a fresh start.

Results and findings

More than 400 faculty members participate in Teach Anything. They are based in East Asia, India, Australia, the Middle East, Turkey, Egypt, Western Europe, Norway, the UK, North Africa, Canada, the US, Brazil, Argentina, and elsewhere. A small number of them took the extra step of sharing reflections on their experience.

As a result of using Teach Anything, these participants (n=18) reported improvement in self-efficacy in understanding AI (Table 1). 70% gained understanding of AI's pedagogical implications, and 30% developed knowledge of AI's technological implications through their control of its temperature, system prompts, and inference window (files they uploaded). In terms of the next steps, 90% indicated that they were better positioned to make responsible, effective, and context-appropriate decisions about how to teach with or against AI systems. 10% indicates that they had not yet decided on AI's role in their class but had stronger understanding of AI's operation.

Developing self-efficacy is the first step toward pedagogy that addresses challenges posed by gen AI. We thus contribute to a shared public knowledge base about AI systems and their societal impacts, particularly in the domain of higher education. We enhance literacy and self-efficacy by hands-on practice.

Faculty participants also suggested and beta-tested new features in collaboration with our team, such as allowing students to generate practice quizzes. Aligning with participatory design principles, Teach Anything strengthens a broad range of stakeholders' ability to understand, evaluate, and responsibly use AI systems. Through using our open-access platform, faculty stakeholders gain first-hand experience in order to better understand how AI systems work and where they can fail. Collectively, we evaluate the quality and limitations of AI-generated outputs.

Teach Anything is an open-source, open-access infrastructure that enables faculty across disciplines to

experiment with and design custom AI applications. As part of our efforts to promote open culture, Teach Anything offers faculty access to the proverbial engine room, with opportunities to adjust the retrieval corpus, system prompt, and generation settings for trustworthy outputs. By identifying ethical considerations and trust issues related to the use of AI in learning and development, this design-based research contributes to the ongoing dialogue about responsible AI implementation. These insights will be crucial for organizations, educational institutions, and policymakers as they navigate the integration of AI into learning and development processes.

Teach Anything also promotes equity by making advanced tools more accessible to a wider range of faculty, regardless of financial resources and technical knowledge. This research created the following digital deliverables:

1. A permanently free, open-access platform for higher education faculty to create open-source chatbots that support student learning in their specific courses.

2. A series of open-access chatbots specializing in the disciplines of faculty participants.

3. Monthly newsletters to faculty members and platform upgrades including new features.

Faculty and students evolve their modes of interaction with AI created via Teach Anything in comparison to general-purpose, commercial models.

Table 1. The impact of Teach Anything on the users across diverse demographics

Objective	Activity	Outcomes
Increase faculty's AI literacy (how AI systems work and where they can fail) through professional development opportunities offered by Teach Anything	More than 400 faculty members utilize the open-access Teach Anything to design their own domain-specific AI applications, following our illustrated, step-by-step guide.	• Faculty use the open-access infrastructure to create custom-configured AI chatbots • Faculty develop learning activities using their custom AI • Faculty test new assignments with students
Evaluate the limitations of AI-generated outputs	Faculty assess anonymized student chat records to gauge the pedagogical effectiveness of their instruction and of their custom AI. Faculty form a community of practice to share insights and strategies.	Faculty evaluate the quality and limitations of AI generated outputs and adjust • their AI applications • their pedagogical approaches accordingly
Make responsible and context appropriate decisions on AI and education	18 faculty refine their system prompts and custom-configured AI based on evaluation of classroom activities and anonymized student chat histories	These 18 faculty members discuss their development of self-efficacy with project leader in a preliminary survey

Comparison with Other Platforms

Similar projects exist, but they are by invitation only, commercial (BoodleBox), operating as fee-for-service and for K-12 teachers (Playlab), not sharable (NotebookLM), or designed for a single course. NotebookLM locks users into Google's ecosystem. It is a commercial platform with a free tier; it is not open source and open access. Users can only upload files and do not have backend access. Public sharing is difficult. Playlab and BoodleBox only allow users to choose from among commercial LLMs.

UCLA and Arizona State University, respectively, engage third-party AI vendors to build textbooks and assignments that students have to pay for. UCLA offers a comparative literature course with a controversial AIgenerated textbook [43], [44] and class assignments provided by Kudu, a third-party platform [45]. Without faculty consent or knowledge [46], Arizona State University quietly used online instructional material from faculty to create AI-generated learning modules [47]. These partnerships with for-profit entities are often ineffective and untrustworthy. ASU Atomic generated faulty and decontextualized information [46].

Little Language Models by MIT Media Lab [48] is by invitation only and aims at elementary education. CS50 Duck, a virtual rubber duck chatbot, allows students in Harvard's Introduction to Computer Science to "interact directly with GPT-4 in a controlled manner" [49]. These models do not provide backend access, are not open access or open source.

Teach Anything is a grass-root, faculty-student collaborative learning project that serves as public interest technology. Our open-access platform allows educators everywhere to create custom AI for free. It does not write essays or code for students, and it gives educators and backend access to participate in civic science.

Conclusion: GenAI is not a Black Box

Teach Anything demonstrates that gen AI need not be a black box. In fact, it is AI companies that create and propagate the myth of a black box technology. Engineers know how its parts work by simply generating texts or images "appropriate to the scenario described by the prompts" that evoke "certain conceptual spaces" [50]. For its own interest, the AI industry has woven the so-called black boxes "into the stories they tell." If it appears to be a black box it is only because of the companies' choice of using "an impenetrable cluster of numerical values, linking vectors clumped together incomprehensibly in the training phase" [50].

The open-access Teach Anything changes these stories by empowering faculty members as designers. Large corporations offer AI models that cannot be edited or finetuned by consumers. Faculty do not have control over these models that lack interpretability. When addressing literacy challenges, bigger is not always better, as "large foundation models are giving way to smaller, more task specific models" [51], [52]. Teach Anything supports the creation of custom and "narrow" AI applications for specific

subjects or pedagogical methods. Teach Anything gives faculty more control.

Acknowledgment

Led by Alexa Alice Joubin, Teach Anything is co-built by Akhilesh Rangani at George Washington University. In collaboration with Joubin and Rangani, Ryan Kosloff joined later and contributed several new features.

Declarations

Funding: Research supported by George Mason University Institute for Humane Studies and by George Washington University's Public Interest Technology Scholar Program, GW Coders, Taiwan Studies Open Access Award, and Columbian College of Arts and Sciences Nick-of-Time Impact grant.

References

[1] A. M. Ayyoub et al., "Advancing higher education with GenAI: Factors influencing educator AI literacy," Front. Educ., vol. 10, 2025, doi: 10.3389/feduc.2025.1530721.
[2] J. Du and L. Kohnke, "Repositioning teachers in the age of AI: A cross-disciplinary study of perceptions and practices in China," Teach. Educ., vol. 37, no. 3, pp. 281–305, 2026, doi: 10.1080/10476210.2025.2590694.
[3] N. Najjar, M. Rouphael, M. El Hajj, T. Bitar, P. Damien, and W. Hleihel, "Faculty perceptions and adoption of AI in higher education: Insights from two Lebanese universities," Educ. Sci., vol. 16, no. 1, art. 55, 2026, doi: 10.3390/educsci16010055.
[4] National Institute of Standards and Technology, "Trustworthy and responsible AI — Artificial intelligence risk management framework: Generative artificial intelligence profile," NIST, Gaithersburg, MD, USA, 2025. [Online]. Available: https://www.nist.gov/trustworthy-and-responsible-ai
[5] D. Broniatowski, A. A. Joubin, A. A. Susan et al., "The epistemologies of trust: Conflicting worldviews in the trustworthy AI discourse," AI Ethics, vol. 6, art. 260, 2026.
[6] X. Tan, G. Cheng, and M. H. Ling, "Artificial intelligence in teaching and teacher professional development: A systematic review," Comput. Educ.: Artif. Intell., vol. 8, art. 100355, 2025.
[7] N. B. Milman, "Exploring the crossroads of technology, society, and the human condition: GW education and generative artificial intelligence (AI)," presented at the Society of the Emeriti September Lecture, Sep. 26, 2025.
[8] C. E. Watson and L. Rainie, "Leading through disruption: Higher education executives assess AI's impacts on teaching and learning," AAC&U, Washington, DC, USA, Jan. 2025. [Online]. Available: https://dgmg81phhvh63.cloudfront.net/content/user photos/AACU_AI_Report_2025.pdf
[9] D.-K. Mah and N. Groß, "Artificial intelligence in higher education: Exploring faculty use, self-efficacy, distinct profiles, and professional development needs," Int. J. Educ. Technol. High. Educ., vol. 21, art. 58, 2024, doi: 10.1186/s41239-024-00490-1.
[10] N. Jogezai, F. A. Baloch, M. Jaffar, and G. Khlljl, "From technology to pedagogy: Determinants of university faculty's pedagogically relevant use of generative AI," Qual. Assur. Educ., vol. 33, no. 4, pp. 618–634, 2025.
[11] Y. Qian, "Pedagogical applications of generative AI in higher education: A systematic review of the field," TechTrends, vol. 69, no. 5, pp. 1105–1120, 2025, doi: 10.1007/s11528-025-01100-1.
[12] J. A. Bowen and C. E. Watson, Teaching with AI: A Practical Guide to a New Era of Human Learning. Baltimore, MD, USA: Johns Hopkins Univ. Press, 2024.
[13] L. P. Cordova, T. Mendoza, S. Holmes, K. Megginson, B. Webster, and D. Gregory, "Not all chatbots teach: Evidence for pedagogical design in AI-assisted technical education," in Proc. 26th ACM Annu. Conf. Cybersecurity Inf. Technol. Educ. (SIGITE '25), Nov. 2025, doi: 10.1145/3769694.3771107.
[14] Z. Ahmed, S. S. Shanto, and A. I. Jony, "Potentiality of generative AI tools in higher education: Evaluating ChatGPT's viability as a teaching assistant for introductory programming courses," STEM Educ., vol. 4, no. 3, pp. 165–182, 2024, doi: 10.3934/steme.2024011.
[15] Y. T. Nakamura and J. Milner, "Inclusive leadership via empathic communication," Organ. Dyn., vol. 52, no. 1, art. 100957, 2023, doi: 10.1016/j.orgdyn.2023.100957.
[16] Y. T. Nakamura, J. Hinshaw, and R. Burns, "Developing women's authenticity in leadership," Merits, vol. 2, no. 4, pp. 408–426, 2022, doi: 10.3390/merits2040029.
[17] Y. T. Nakamura, J. Milner, and T. Milner, "Inclusive-empathy in leadership," J. Appl. Behav. Sci., vol. 58, no. 1, pp. 161–163, 2022, doi: 10.1177/0021886320982022.
[18] Y. T. Nakamura, J. Hinshaw, and D. Yu, "The role of empathy in developing ethical leadership: Neurobiology and video-based approaches," in Ethics and Human Resource Development, A. Alizadeh and D. F. Russ-Eft, Eds. Cham, Switzerland: Palgrave Macmillan, 2024, pp. 449–468.
[19] Y. T. Nakamura, J. Milner, D. Yu, and J. Hinshaw, "The role of empathy in leadership ethics: Examining relational leadership practice through video-based methods," Bus. Ethics Q., 2025, in press.
[20] S. Randazzo, H. Lifshitz, A. Joshi, F. Dell'Acqua, K. C. Kellogg, and K. R. Lakhani, "GenAI as a power persuader: How professionals get persuasion bombed when they attempt to validate LLMs," SSRN, Jul. 1, 2025, doi: 10.2139/ssrn.5678644.
[21] V. Danry et al., "Don't just tell me, ask me: AI systems that intelligently frame explanations as questions improve human logical discernment accuracy over causal AI explanations," in Proc. CHI Conf. Human Factors Comput. Syst., Apr. 2023.
[22] T. Corbin, M. Bearman, D. Boud, and P. Dawson, "The wicked problem of AI and assessment," Assess. Eval. Higher Educ., Sep. 3, 2025, doi: 10.1080/02602938.2025.2553340.
[23] S. Otterman, "A.I. is coming to class. These professors want to ease your worries," The New York Times, Jan. 17, 2026. [Online]. Available: https://www.nytimes.com/2026/01/17/nyregion/ai college-classes.html
[24] S. Eminovic, B. Levita, R. Schmidt, M. Lindholz, D. Desser, T. Penzkofer, M. P. Wattjes, and J. Nawabi, "Open-source versus proprietary large language models for academic writing in medicine: A controlled evaluation of DeepSeek-R1 and ChatGPT-o1 in neuroradiology," Digit. Health, vol. 12, Feb. 2026, doi: 10.1177/20552076261464745.
[25] A. Sayer, "(De)commodification, consumer culture, and moral economy," Environ. Plan. D: Soc. Space, vol. 21, pp. 341–357, 2003.
[26] J. Riskin, "AAUP: Academic freedom on the line. Viewpoint diversity isn't about viewpoint diversity," The Stanford Daily, Feb. 4, 2026. [Online]. Available: https://stanforddaily.com/2026/02/04/aaup-viewpoint-diversity/
[27] B. Czerkawski, "AI and the learning experience design: From divergent creativity to convergent precision," TechTrends, vol. 69, pp. 260–270, 2025.
[28] M. Kassorla, M. Georgieva, and A. Papini, "AI literacy in teaching and learning: A durable framework for higher education," EDUCAUSE Working Group, 2024. [Online]. Available: https://www.educause.edu/content/2024/ai-literacy-in-teaching-and learning/executive-summary
[29] Y. Guo, M. Guo, J. Su et al., "Bias in large language models: Origin, evaluation, and mitigation," arXiv:2411.10915, 2024.
[30] A. Birhane, "Algorithmic injustice: A relational ethics approach," Patterns, vol. 2, no. 2, art. 100205, 2021, doi: 10.1016/j.patter.2021.100205.
[31] S. Cave and K. Dihal, "The whiteness of AI," Philos. Technol., vol. 33, no. 4, pp. 685–703, 2020, doi: 10.1007/s13347-020-00415-6.
[32] T. L. Howard and G. W. Ulferts, "Artificial intelligence and the redefinition of higher education," Higher Educ. J., vol. 46, pp. 1–21, 2025.
[33] J. Batista, A. Mesquita, and G. Carnaz, "Generative AI and higher education: Trends, challenges, and future directions from a systematic literature review," Information, vol. 15, no. 11, art. 676, 2024, doi: 10.3390/info15110676.
[34] I. O. Gallegos, R. A. Rossi, J. Barrow et al., "Bias and fairness in large language models," Comput. Linguist., vol. 50, no. 3, pp. 1097-1179, 2024.
[35] G. Waters and P. Honenberger, "AI biases as asymmetries: A review to guide practice," Front. Big Data, vol. 8, art. 1532397, 2025, doi: 10.3389/fdata.2025.1532397.

[36] H. Joren, J. Zhang, C.-S. Ferng, D.-C. Juan, A. Taly, and C. Rashtchian, "Sufficient context: A new lens on retrieval augmented generation systems," in Proc. Int. Conf. Learn. Representations (ICLR), 2025.
[37] Y. Sugihara, A. Milosavljevic, S. Jankovskaja, and M. Falk, "A review for navigating the trade-offs: Evaluating open-source and proprietary large language models for clinical and biomedical information extraction," Front. Digit. Health, vol. 8, art. 1778786, Apr. 2026, doi: 10.3389/fdgth.2026.1778786.
[38] M. Chelli, J. Descamps, V. Lavoué, C. Trojani, M. Azar, M. Deckert, J.-L. Raynier, G. Clowez, P. Boileau, and C. Ruetsch Chelli, "Hallucination rates and reference accuracy of ChatGPT and Bard for systematic reviews: Comparative analysis," J. Med. Internet Res., vol. 26, art. e53164, 2024, doi: 10.2196/53164.
[39] Z. Ahmad, W. Kaiser, and S. Rahim, "Hallucinations in ChatGPT: An unreliable tool for learning," Rupkatha J. Interdiscip. Stud. Humanit., vol. 15, no. 4, 2023, doi: 10.21659/rupkatha.v15n4.17.
[40] A. Sriram, N. Maheswaran, B. Sundan, and S. Krishnamoorthy, "Context-aware retrieval-augmented generation for artificial intelligence in urology," Cureus, vol. 17, no. 7, art. e88167, Jul. 2025, doi: 10.7759/cureus.88167.
[41] A. Mahdawi, "Beware what you tell your AI chatbot. It's not a shrink — it's a snitch," The Guardian, May 13, 2026. [Online]. Available: https://www.theguardian.com/commentisfree/2026/may/13/beware what-you-tell-your-ai-chatbot-its-not-a-shrink-its-a-snitch
[42] S. Jain, C. Park, M. Viana, A. Wilson, and D. Calacci, "Interaction context often increases sycophancy in LLMs," in Proc. 2026 CHI Conf. Human Factors Comput. Syst., Apr. 2026, doi: 10.1145/3772318.3791915.
[43] C. Dutton, "The course is about literature. Its textbook was generated by AI," The Chronicle of Higher Education, Dec. 12, 2024. [Online]. Available: https://www.chronicle.com/article/the course-is-about-literature-its-textbook-was-generated-by-ai
[44] K. Palmer, "The AI-generated textbook that's making academics nervous," Inside Higher Ed, Dec. 13, 2024. [Online]. Available: https://www.insidehighered.com/news/faculty/learning assessment/2024/12/13/ai-assisted-textbook-ucla-has-some academics-worried
[45] UCLA Newsroom, "Comparative lit class will be first in Humanities Division to use UCLA-developed AI system," Regents of the Univ. of California, 2025. [Online]. Available: https://newsroom.ucla.edu/stories/comparative-literature-zrinka stahuljak-artificial-intelligence
[46] E. Whitford, "Faculty concerned about ASU's 'Frankensteinian' AI course builder," Inside Higher Ed, Apr. 29, 2026. [Online]. Available: https://www.insidehighered.com/news/tech innovation/artificial-intelligence/2026/04/29/faculty-concerned about-asus-new-ai-course
[47] ASU Atomic, "Level up your professional potential," Arizona Board of Regents, 2026. [Online]. Available: https://atomic.asu.edu
[48] M. Dhariwal and S. Dhariwal, "Little language models: Empowering children to be future AI modelers," LinkedIn, Aug. 12, 2024. [Online]. Available: https://www.linkedin.com/pulse/little language-models-empowering-children-future-ai-modelers-e5cpc/
[49] R. Liu, C. Zenke, C. Liu, A. Holmes, P. Thornton, and D. J. Malan, "Teaching CS50 with AI: Leveraging generative artificial intelligence in computer science education," in Proc. 55th ACM Tech. Symp. Comput. Sci. Educ. V. 1, 2024, pp. 750–756, doi: 10.1145/3626252.3630938.
[50] E. Salvaggio, "The black box myth: What the industry pretends not to know about AI," Tech Policy Press, Jun. 17, 2025. [Online]. Available: https://www.techpolicy.press/the-black-box-myth-what the-industry-pretends-not-to-know-about-ai/
[51] L. Hinkel, "Creating AI that matters: How the MIT-IBM Watson AI Lab is shaping AI-sociotechnical systems for the future," MIT News, Oct. 21, 2025. [Online]. Available: https://news.mit.edu/2025/creating-ai-that-matters-1021
[52] J. Lin, J. Tang, H. Tang, S. Yang, W.-M. Chen, W.-C. Wang, G. Xiao, X. Dang, C. Gan, and S. Han, "AWQ: Activation-aware weight quantization for on-device LLM compression and acceleration," in Proc. 5th MLSys Conf., Santa Clara, CA, USA, 2024.