AI Can Do Your Homework. Now What?

Developing a Taxonomy of Computing Assessments That Accommodate GenAI Use

Geoffrey Challen*
challen@illinois.edu
University of Illinois
Champaign, IL, USA

Kevin Lin*
kevinl@cs.uw.edu
University of Washington
Seattle, WA, USA

Ranysha Ware*
rware@cs.swarthmore.edu
Swarthmore College
Swarthmore, PA, USA

Muhammad Sajjad Akbar
muhammad.akbar@sydney.edu.au
University of Sydney
Sydney, NSW, Australia

Fraida Fund
ffund@nyu.edu
New York University
New York, NY, USA

Casey Hopkins
c.l.hopkins@swansea.ac.uk
Swansea University
Swansea, Wales, UK

Oscar Karnalim
oscar.karnalim@it.maranatha.edu
Maranatha Christian University
Bandung, West Java, Indonesia

James W. McGuffee
jwmcguffee@ccis.edu
Columbia College
Columbia, MO, USA

Shubbhi Taneja
staneja@wpi.edu
Worcester Polytechnic Institute
Worcester, MA, USA

Abstract

For decades, computing instructors have given students a specification, asked them to translate it into code, and evaluated the result against the specification. Coding agents can now complete most of these assessments with only a small number of prompts. Students who use AI this way learn neither what the assessment was meant to teach nor how to collaborate with AI effectively. Meanwhile, the rapid adoption of coding agents by developers is reshaping how technology is created. Working with them multiplies individual productivity, accelerates teams, and expands what computing can build. Agentic development is an important new skill. Students must learn it. Updated assessments are needed to teach it.

This working group will map how computing educators are responding. In July 2026, we organized a public workshop convening computing educators to discuss how they intended to redesign assessments in response to AI and prepared a public report from those discussions. Then, we will survey computing educators to identify assessments that have been redesigned, replaced, or deliberately preserved, along with the rationale behind each choice. Finally, we will assemble what we discover into a taxonomy that can help future computing education researchers and practitioners design assessments that resist AI delegation and assessments that teach students to direct it.

CCS Concepts

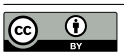
• **Social and professional topics → Computing education**.

Keywords

Agentic Engineering, Assessment Redesign, Curriculum Design, Generative Artificial Intelligence, Pedagogy, Scaffolding

ACM Reference Format:
Geoffrey Challen, Kevin Lin, Ranysha Ware, Muhammad Sajjad Akbar, Fraida Fund, Casey Hopkins, Oscar Karnalim, James W. McGuffee, and Shubbhi Taneja. 2026. AI Can Do Your Homework. Now What?: Developing a Taxonomy of Computing Assessments That Accommodate GenAI Use. In *Proceedings of the 2nd ACM Virtual Global Computing Education Conference V.2 (SIGCSE Virtual 2026), November 12–15, 2026, Virtual Event, USA*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3795860.3831042

*Co-leads contributed equally; remaining authors are listed alphabetically by surname.



This work is licensed under a Creative Commons Attribution 4.0 International License.

SIGCSE Virtual 2026, Virtual Event, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2491-6/2026/11
https://doi.org/10.1145/3795860.3831042

1 Background and Related Work

Monthly generative AI use among computing students rose from 66% in 2023 to 94% in 2024 [7]. Students who delegate traditional assessments to AI learn neither what the assessment was meant to teach nor how to collaborate with AI effectively. At the same time, AI tools risk eroding the social interactions of help-seeking that traditionally support learning [5, 6].

Outside the classroom, the same tools are reshaping professional software development, where working with coding agents changes individual productivity, team velocity, and the scope of what can be built. Effective practice now spans a wide range, from rapid prototyping with AI models to building production systems with AI agents. Across that range, professionals direct AI tools, evaluate their output, and integrate their work into reliable systems. Agentic engineering shifts the bottleneck from writing code to architectural design, triage, and review [8]. In this new paradigm, software engineers work more like team leaders, collaborating with human and AI agents to architect solutions and evaluate risks. This becomes an essential professional skill, and we need new assessments that prepare students to collaborate with AI.

Three recent working groups have mapped this landscape. Prather et al. [12] surveyed students and instructors across twenty countries and reviewed the early LLM literature. Prather et al. [13] systematically reviewed seventy-one papers through May 2024, categorizing

how instructors deploy generative AI by tool type, purpose, and student guidance. Bouvier et al. [2] focused on upper-level computing courses, combining a systematic literature review with a survey of computing instructors to present an overview of current practices. Together these efforts characterize generative AI as an in-context assistant—a tutor, hint generator, or code explainer—integrated into otherwise traditional courses. Agentic coding tools such as Cursor, Claude Code, and OpenAI Codex have moved generative AI from in-context assistant to autonomous worker.

New educational infrastructure has begun to address these shifts, including textbooks on AI-assisted Python programming [11] and courses on agentic engineering.[1] But work beyond programming-centered courses is limited, and the adaptations instructors are making across the curriculum have not been systematically organized. As Lau et al. [10] describe, computing instructors are currently engaged in "emergency pedagogical design," grappling with an "assessment misfit" between what traditional assignments measure and how students use generative AI. Instructors lack the visibility, scalable resources, and shared community practices needed to effectively translate high-level AI policies into actionable, forward-looking assessment designs.

This working group takes the assessment itself as the unit of analysis, aiming to collect, organize, and share assessments (e.g., assignments, quizzes, exams, and the rubrics or syllabi that frame them) that instructors are actively using or planning to use. The scope spans the full computing curriculum, from CS1 specify-and-implement assignments and algorithms problem sets to systems labs, HCI design projects, and ethics seminars. This includes emerging forms of authentic assessment, such as "prompt problems" that evaluate students' ability to direct LLMs [4] and even assessment designs that do not directly engage AI, such as oral assessments that probe student reasoning in real time [3, 9]. Central to these efforts is the goal of fostering learner agency: ensuring that students engage in mastery-oriented help-seeking rather than merely outsourcing their learning to AI [14].

2 Goals of the Proposal

This working group aims to systematically map and categorize the diverse ways computing educators are adapting—or intentionally not adapting—their assessments in the era of generative AI. Because the technology and instructional responses are moving faster than the publication cycle, we focus on what educators are doing now and planning for the year ahead rather than on what has already settled into the literature. Our goal is to surface the range of pedagogical approaches currently deployed across different contexts in the hope of building awareness of what is possible in supporting our students' learning. We will document the concrete changes instructors are making, including assessment redesigns, new interaction patterns, and AI-centric learning objectives. Specifically, we will investigate how these changes support the shift toward agentic engineering, the implementation of more authentic forms of assessment, and the preservation of learner agency. We will identify which aspects of traditional assessment are being held steady, and evaluate how those decisions hold up against advancing AI capabilities. We will surface the questions that remain, so that subsequent work can build on a shared map of what is still open.

3 Proposed Methodology

Our methodology draws on three activities: a public workshop, a community survey study, and a systematic literature review. We organize what we find using a taxonomy that we develop through the public workshop.

An open, one-day synchronous online workshop brings together the working group and the broader computing educator community to share assessment designs with each other. Held on July 28, 2026, this workshop drew 118 registrants and 73 attendees: our nine authors and 64 participants. The workshop's public report [1] records the approaches participants have tried, the results they reported, and the questions each room left open, reflecting the community's thinking at this moment in time.

A survey study, conducted with computing educators recruited through SIGCSE channels, catalogs the assessment modifications, preservation strategies, and instructor motivations educators are working with right now.

A systematic literature review sets a baseline. We scan the published computing education literature for assessments designed for the AI era, recognizing that anything indexed today reflects practice from twelve to eighteen months ago.

4 Expected Deliverables

This working group will produce two deliverables. The workshop is itself a deliverable: a synchronous community event during which contributing instructors present their assessment redesigns to one another and to the working group, creating a shared forum at a moment when most adaptation work is happening locally and informally. A comprehensive final report, published in the working group proceedings, will document the taxonomy of assessment adaptations—drawing on the workshop's public summary report, the survey study, and the literature review—along with the pedagogical rationales behind them and the questions that remain open.

5 Human Subjects Research

The survey study collects data from computing educators about their teaching practices and is the only activity in this protocol involving human subjects. The U.S.-based work will undergo Institutional Review Board review, where we expect it to qualify as exempt; the Australia-based work will undergo Human Research Ethics Committee review. Both processes are currently underway.

The workshop is a public community event, not a data-collection activity conducted under this protocol: we do not recruit or collect data from attendees as human subjects. Its public summary report is a publicly available document that we analyze post hoc, alongside the literature review, as secondary source material rather than as human subjects research.

[1] https://www.teachcswithai.org/news-and-events/genai-cs-education-workshop-practice-and-research-aicsepar

References

[1] Muhammad Sajjad Akbar, Geoffrey Challen, Fraida Fund, Casey Hopkins, Oscar Karnalim, Kevin Lin, James McGuffee, Shubbhi Taneja, and Ranysha Ware. 2026. *AI Can Do Your Homework. Now What? Report from an Online Workshop on Computing Assessment in the Age of Generative AI.* Technical Report. SIGCSE Virtual 2026 Working Group. https://go.illinois.edu/ai-can-do-your-homework

[2] Dennis J. Bouvier, Bruno Pereira Cipriano, Richard Glassey, Olga Petrovska, Emma Anderson, Anastasiia Birillo, Ryan Dougherty, Raymond Pettit, Nuno Pombo, Ebrahim Rahimi, Charanya Ramakrishnan, Alexander Steinmaurer, Shubbhi Taneja, Muhammad Usman, and Annapurna Vadaparty. 2026. The Rest of the Robots: Generative AI in Post-introductory Computing Education. In *Proceedings of the 2025 Working Group Reports on Innovation and Technology in Computer Science Education* (Netherlands) *(ITiCSE-WGR 2025)*. Association for Computing Machinery, New York, NY, USA, 61–107. https://doi.org/10.1145/3760545.3783970

[3] Kelly Cochran and Chris Piech. 2026. Personalized Exam Prep (PEP): Scaling No-Stakes, No-LLM Dialogue-Based Assessments in a Large CS Course. In *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2026)*. Association for Computing Machinery, New York, NY, USA, 225–231. https://doi.org/10.1145/3770762.3772663

[4] Paul Denny, Juho Leinonen, James Prather, Andrew Luxton-Reilly, Thezyrie Amarouche, Brett A. Becker, and Brent N. Reeves. 2024. Prompt Problems: A New Programming Exercise for the Generative AI Era. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3626252.3630909

[5] William G. Griswold. 2024. Experience Report: Meet the Professor—A Large-Course Intervention for Increasing Rapport. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3626252.3630844

[6] Irene Hou, Sophia Mettille, Zhuo Li, Owen Man, Cynthia Zastudil, and Stephen MacNeil. 2024. The Effects of Generative AI on Computing Students' Help-Seeking Preferences. In *Proceedings of the 26th Australasian Computing Education Conference (ACE '24)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3636243.3636248

[7] Irene Hou, Hannah Vy Nguyen, Owen Man, and Stephen MacNeil. 2025. The Evolving Usage of GenAI by Computing Students. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 2 (SIGCSE TS 2025)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3641555.3705266

[8] Matthew Kam et al. 2025. What Skills and Knowledge do Software Developers Need to Succeed in an Age of Artificial Intelligence?. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (FSE 2025)*. Association for Computing Machinery, New York, NY, USA.

[9] Suhas Kannam, Yuri Yang, Aarya Dharm, and Kevin Lin. 2025. Code Interviews: Design and Evaluation of a More Authentic Assessment for Introductory Programming Assignments. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2025)*. Association for Computing Machinery, New York, NY, USA.

[10] Sam Lau, Kianoosh Boroojeni, Harry Keeling, and Jenn Marroquin. 2026. Barriers that Programming Instructors Face While Performing Emergency Pedagogical Design to Shape Student-AI Interactions with Generative AI Tools. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3772318.3790682

[11] Leo Porter and Daniel Zingaro. 2023. *Learn AI-Assisted Python Programming: With GitHub Copilot and ChatGPT*. Manning Publications, Shelter Island, NY. https://www.manning.com/books/learn-ai-assisted-python-programming

[12] James Prather, Paul Denny, Juho Leinonen, Brett A. Becker, Ibrahim Albluwi, Michelle Craig, Hieke Keuning, Natalie Kiesler, Tobias Kohn, Andrew Luxton-Reilly, Stephen MacNeil, Andrew Petersen, Raymond Pettit, Brent N. Reeves, and Jaromir Savelka. 2023. The Robots Are Here: Navigating the Generative AI Revolution in Computing Education. In *Proceedings of the 2023 Working Group Reports on Innovation and Technology in Computer Science Education (ITiCSE-WGR '23)*. Association for Computing Machinery, New York, NY, USA, 108–159. https://doi.org/10.1145/3623762.3633499

[13] James Prather, Juho Leinonen, Natalie Kiesler, Jamie Gorson Benario, Sam Lau, Stephen MacNeil, Narges Norouzi, Simone Opel, Vee Pettit, Leo Porter, Brent N. Reeves, Jaromir Savelka, David H. Smith IV, Sven Strickroth, and Daniel Zingaro. 2024. Beyond the Hype: A Comprehensive Review of Current Trends in Generative AI Research, Teaching Practices, and Tools. In *Proceedings of the 2024 Working Group Reports on Innovation and Technology in Computer Science Education (ITiCSE-WGR '24)*. Association for Computing Machinery, New York, NY, USA, 300–338. https://doi.org/10.1145/3689187.3709614

[14] Stephanie Yang, Melissa Chen, and Bertrand Schneider. 2025. Navigating Appropriate Help-seeking with LLMs: Associations with Motivation and Beliefs in a CS course. In *Proceedings of the 24th Koli Calling International Conference on Computing Education Research (Koli Calling 2025)*. Association for Computing Machinery, New York, NY, USA.