Agency for Whom and To What Ends: A Plan for Investigating Impacts of Agentic AI in Computing Education

Janice Mak*
Arizona State University
Tempe, Arizona, USA
Janice.Mak@asu.edu

Tony Clear*
Auckland University of Technology
Auckland, New Zealand
tony.clear@aut.ac.nz

Tingting Zhu*
University of Toronto
Mississauga, Canada
tingting.zhu@utoronto.ca

Alison Clear*
School of Business, Eastern Institute of Technology
Auckland, New Zealand
AClear@eit.ac.nz

Oana Andrei
University of Glasgow
Glasgow, UK
oana.andrei@glasgow.ac.uk

Martin Goodfellow
University of Strathclyde
Glasgow, UK
martin.h.goodfellow@strath.ac.uk

Asanthika Imbulpitiya
Otago Polytechnic Auckland International Campus
Auckland, New Zealand
asanthika.imbulpitiya@op.ac.nz

Elizabeth Oladapo
Manchester Metropolitan University
Manchester, UK
e.oladapo@mmu.ac.uk

Aadarsh Padiyath
University of Michigan
Ann Arbor, Michigan, USA
aadarsh@umich.edu

Jose Antonio Pow-Sang
Universidad Católica del Perú
Lima, Peru
japowsang@pucp.edu.pe

Rebecca Williams
University of Maryland Baltimore County
Baltimore, Maryland, USA
rmwillia@umbc.edu

Abstract

Agentic AI, where AI systems have their own forms of agency, will present some of the most critical challenges that computing education will face, including regulatory gaps and amplified cascading social and ethical impacts. This working group proposes a landscape analysis of the ethical and societal implications of using Agentic AI in higher computing education. By exploring emerging literature and use cases of Agentic AI, we aim to contribute a timely landscape study exploring 1) the emerging challenges and opportunities associated with Agentic AI, 2) how higher education is beginning to adopt and use Agentic AI and the resulting ethical and societal impacts, and 3) the implications (e.g. challenges, opportunities, limitations) of integrating Agentic AI in computing education. The expected outputs are: 1) a protocol and literature scoping review of Agentic AI in education, and 2) an analysis of current practices and use cases in computing education. Together, these outputs aim to identify emerging patterns, use cases, key challenges, and principles to support the ethical and responsible integration of Agentic AI in education.

CCS Concepts

• **Social and professional topics → Computer science education**.

Keywords

AI, Agentic AI, Computing Education, Ethical and Societal Impact

ACM Reference Format:
Janice Mak, Tony Clear, Tingting Zhu, Alison Clear, Oana Andrei, Martin Goodfellow, Asanthika Imbulpitiya, Elizabeth Oladapo, Aadarsh Padiyath, Jose Antonio Pow-Sang, and Rebecca Williams. 2026. Agency for Whom and To What Ends: A Plan for Investigating Impacts of Agentic AI in Computing Education. In *Proceedings of the 31st ACM Conference on Innovation and Technology in Computer Science Education V. 2 (ITiCSE 2026), July 10–15, 2026, Madrid, Spain.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3803401.3812047

1 Background and Related Work

The role of Generative AI (GenAI) to date has mostly been conceived as communicative chatbots (e.g., ChatGPT, Gemini, Copilot). The limitations of these narrower scope AI spurred interest in a paradigm shift [7] to Agentic AI, which marks a new era of AI systems that are capable of planning, reasoning, adapting, and interacting with the environment without direct human intervention [1, 9]. This working group will conduct an exploratory analysis to investigate the emergent use and impact of Agentic AI. These rapidly evolving directions in technology, including trends where AI systems have their own forms of agency, will present some of

*Working Group Co-leader. All authors contributed to this research.

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

ITiCSE 2026, Madrid, Spain

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

ACM ISBN 979-8-4007-2633-0/2026/07
https://doi.org/10.1145/3803401.3812047

the most critical challenges that computing education will face in the immediate future.

These challenges include and are not limited to regulatory gaps and amplified cascading effects [1, 4] of Agentic AI in health care, education, manufacturing sectors, legal practices, creative industries and customer-facing services. There are predictions that one-third of industry will use Agentic AI in their workflows by 2028, up from one percent in 2024 [4]. GenAI impacts how we access and use information, particularly in educational settings. These impacts extend to society at large and include impacts on intellectual and creative works. The context for the working group is framed through the *ACM Taskforce on The Ethics and Societal Impacts of GenAI in Higher Education*[1] which investigated the ethical and societal impacts of GenAI within the higher education landscape.

The results of this working group will provide insights into the potential contributions and risks of Agentic AI, building on prior ACM sponsored working group reports [10, 12]. The working group will conduct a landscape analysis on ethical questions arising from Agentic AI tool use in higher education contexts, identifying promising principles, use cases, challenges, and principled ways to navigate the implementation of Agentic AI in ethical and responsible ways.

2 Goals of the Working Group

By exploring the emerging literature and use cases of Agentic AI in higher education, this report will contribute a timely exploratory landscape study to guide higher computing education and the research community. Our research questions (RQ) are:

(1) What emerging challenges and opportunities are associated with the use of Agentic AI in computing education?
(2) How are HE institutions using Agentic AI in computing education?
(3) What ethical and societal implications arise from these emerging practices?

3 Proposed Methodology

The overarching goal of this working group is to produce an exploratory analysis of the emerging use and impacts of Agentic AI in higher education through a literature review and qualitative study. The scope of our literature review will include empirical practices and research involving Agentic AI and their implications for computing curricula, instruction, pedagogy, assessment, and competencies. We acknowledge from the outset the critical role that grey literature (GL) will play in our work. Including GL in a rigorous way will support the identification of emerging practices, challenges, and opportunities, especially given the rapid and emerging role that Agentic AI currently has in industry. Therefore, we will approach the literature review by taking a Multi-Vocal literature review approach (MVLR) [2, 6, 13, 14], which represents multiple voices and perspectives on emerging topics. As such, MVLR embodies diverse information sources beyond academic settings, including professional societies, government, research centers and trade press [3, 5, 8, 11].

[1]https://www.acm.org/binaries/content/assets/about/annual-reports/ed-board-fy25.pdf

For the qualitative research component, we will interview individuals who are using Agentic AI in computing education across a range of institutions (e.g. public, private, urban, rural) representing diverse countries. We will recruit participants through the ACM listserv and working group members. Arizona State University will serve as the lead institution for the IRB process and ensure compliance across institutions. We will use a combination of inductive and deductive coding. Inductive coding will identify emerging practices, challenges, and opportunities, while deductive coding will apply the ESI-Framework [10] to analyze ethical and societal implications.

4 Expected Deliverables

The expected deliverables will include: 1) development of a protocol and a scoping literature review of the emerging challenges and opportunities associated with the use of Agentic AI in computing education, and 2) an analysis of current practices and use cases of Agentic AI in computing education through a socio-ethical lens.

References

[1] Deepak Bhaskar Acharya, Karthigeyan Kuppan, and B Divya. 2025. Agentic AI: Autonomous intelligence for complex goals–a comprehensive survey. *IEEe Access* (2025).

[2] Priyanka Antil. 2020. Requirements engineering in global scaled agile software development environment: a multi-vocal literature review protocol. *arXiv preprint arXiv:2004.12647* (2020).

[3] ACM Europe Technology Policy Committee. 2025. *Systemic Risks Associated with Agentic AI: A Policy Brief.* Report. ACM. https://www.acm.org/binaries/content/assets/public-policy/europe-tpc/systemic_risks_agentic_ai_policy-brief_final.pdf

[4] Yogesh K Dwivedi, Mohamed YI Helal, Ibrahim A Elgendy, Rasha Alahmad, Paul Walton, Ayoung Suh, Vinay Singh, and Il Jeon. 2025. Agentic AI Systems: What It Is and Isn't. *Global Business and Organizational Excellence* (2025).

[5] World Economic Forum. 2025. *AI Agents in Action: Foundations for Evaluation and Governance.* Report. World Economic Forum and CapGemini. https://www.weforum.org/publications/ai-agents-in-action-foundations-for-evaluation-and-governance/

[6] Vahid Garousi, Michael Felderer, and Mika V Mäntylä. 2019. Guidelines for including grey literature and conducting multivocal literature reviews in software engineering. *Information and software technology* 106 (2019), 101–121.

[7] Soodeh Hosseini and Hossein Seilani. 2025. The role of agentic AI in shaping a smart future: A systematic review. *Array* 26 (2025), 100399. doi:10.1016/j.array.2025.100399

[8] A.I News Hub. 2025. AI Inference Costs 2025: Why Google TPUs Beat Nvidia GPUs by 4x. https://www.ainewshub.org/post/ai-inference-costs-tpu-vs-gpu-2025

[9] Georgios Kostopoulos, Vasileios Gkamas, Maria Rigou, and Sotiris Kotsiantis. 2025. Agentic AI in Education: State of the Art and Future Directions. *IEEE Access* 13 (2025), 177467–177491. doi:10.1109/ACCESS.2025.3620473

[10] Janice Mak, Joyce Nakatumba-Nabende, Tony Clear, Alison Clear, Ibrahim Abluwi, Oana Andrei, Lorenzo Angeli, Stephen MacNeil, Solomon Sunday Oyelere, Matthew Hale Rattigan, Judy Sheard, and Tingting Zhu. 2026. Navigating the Ethical and Societal Impacts of Generative AI in Higher Computing Education. 313–359 pages. doi:10.1145/3760545.3783975

[11] S. Mylius, P. Slattery, Y. Zhu, M. Narayanan, A. Thinnyun, A. K. Saeri, J. Graham, M. Noetel, and N. Thompson. 2025. Mapping the AI Governance Landscape: Pilot Test and Update. https://airisk.mit.edu/blog/mapping-the-ai-governance-landscape-pilot-test-and-update

[12] Claudia Szabo, Nickolas J. G. Falkner, Mbodila Munienge, Judithe Sheard, Mabberi Enock, Tony Clear, Delali Kwasi Dake, Omowumi Ogunyemi, Oluwakemi Ola, Tsholofetso Taukobong, and Bimlesh Wadhwa. 2025. Ethical and Societal Impacts of Generative AI in Higher Computing Education: An ACM Task Force Working Group to Develop a Landscape Analysis - Perspectives from the Global Souths and Guidelines for CS1/CS2/CS3. 371–372 pages. doi:10.1145/3736251.3749526

[13] Ángel Jesús Varela-Vaca and Antonia M Reina Quintero. 2021. Smart contract languages: A multivocal mapping study. *ACM Computing Surveys (CSUR)* 54, 1 (2021), 1–38.

[14] Xiaofan Zhao, Tony Clear, and Ramesh Lal. 2024. Identifying the primary dimensions of DevSecOps: A multi-vocal literature review. *Journal of Systems and Software* 214 (2024), 112063.