Proceedings Paper:
Pereira Cipriano, B., Petrovska, O., Pombo, N. et al. (2026) Towards improving CS students' generative AI literacy. In: Proceedings of the 31st ACM Conference on Innovation and Technology in Computer Science Education V. 2. ITiCSE 2026: ACM Conference on Innovation and Technology in Computer Science Education, 10-15 Jul 2026, Madrid, Spain. ACM, pp. 789-790. ISBN: 9798400726330.

https://doi.org/10.1145/3803401.3812055

Towards Improving CS Students' Generative AI Literacy

Bruno Pereira Cipriano*
Universidade Lusófona, INESC INOV Lab
Lisbon, Portugal
bcipriano@ulusofona.pt

Olga Petrovska*
Swansea University, IoC in Wales
Swansea, Wales, United Kingdom
olga.petrovska@swansea.ac.uk

Nuno Pombo*
Universidade da Beira Interior
Covilhã, Portugal
ngpombo@ubi.pt

Lina Battestilli
North Carolina State University
Raleigh, North Carolina, United States
lbattestilli@ncsu.edu

Laura Farinetti
Politecnico di Torino
Turin, Italy
laura.farinetti@polito.it

Richard Glassey
KTH Royal Institute of Technology
Stockholm, Sweden
glassey@kth.se

Maria Kasinidou
Open University of Cyprus
Nicosia, Cyprus
maria.kasinidou@ouc.ac.cy

Olakunle Olayinka
University of Sheffield
Sheffield, England, United Kingdom
o.olayinka@sheffield.ac.uk

Anshul Shah
University of California, San Diego
La Jolla, California, United States
ayshah@ucsd.edu

Alexander Steinmaurer
Interdisciplinary Transformation
University (IT:U)
Linz, Austria
alexander.steinmaurer@it-u.at

Ramalakshmi Vaidhiyanathan
Cardiff University
Cardiff, Wales, United Kingdom
vaidhiyanathanr@cardiff.ac.uk

James Weichert
University of Washington
Seattle, Washington, United States
jpw@cs.washington.edu

Abstract

The widespread adoption of Generative AI (GenAI) tools by students across different educational levels highlights the need for them to develop robust GenAI literacy, including a working understanding of these systems' fundamental concepts, their limitations, and implications for responsible use. However, misconceptions about GenAI, such as perceiving these systems as mere search engines or database lookup systems, are commonly observed among students, while the availability of teaching resources remains fragmented, and learning objectives lack alignment. This Working Group aims to design pedagogical resources for computing science instructors, enabling them to develop students' GenAI literacy. To achieve this, the Working Group will first identify a concise set of GenAI literacy learning objectives informed by instructor experience, research literature, and community input, and subsequently design pedagogical resources aligned with these objectives.

CCS Concepts

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

Keywords

GenAI, Large Language Models, computing education, instructional design

*Co-leader

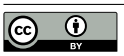

ITiCSE 2026, Madrid, Spain

ACM ISBN 979-8-4007-2633-0/2026/07
https://doi.org/10.1145/3803401.3812055

ACM Reference Format:
Bruno Pereira Cipriano, Olga Petrovska, Nuno Pombo, Lina Battestilli, Laura Farinetti, Richard Glassey, Maria Kasinidou, Olakunle Olayinka, Anshul Shah, Alexander Steinmaurer, Ramalakshmi Vaidhiyanathan, and James Weichert. 2026. Towards Improving CS Students' Generative AI Literacy. 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/10.1145/3803401.3812055

1 Motivation and related work

Generative AI (GenAI) systems, such as ChatGPT or GitHub Copilot, are changing professional workflows. Their ability to generate a wide range of outputs, from detailed textual explanations to functional code, offers significant potential for productivity gains. For instance, Google's research has shed some light on these tools' capacity to reduce the required time for code migrations [11]. Furthermore, they estimated a 21% increase in development speed attributable to GenAI [10].

The 2025 DORA report [4] documents this evolution in the professional software development practices, highlighting positive gains in both productivity and code quality. Their global survey has shown that 90% of professional developers reported using AI at work, while 80% perceived that AI has increased their productivity, and 65% reported relying on AI at least "a moderate amount".

Simultaneously, misconceptions about how GenAI systems operate appear to be common. These include perceiving GenAI systems as search tools [2] or attributing them actual intelligence [6]. Imprecise perceptions and inaccurate mental models can significantly influence how an individual approaches GenAI. On the other hand, lack of awareness or insufficient practice with GenAI can also be problematic for students' future careers.

The education sector, including the computer science education (CSE) community, is struggling to keep up with the rapid evolution of these technologies. The integration of GenAI across computing courses remains fragmented; for example, non-determinism is covered in [12], while [3] covers probabilistic behavior, sensitivity to training data and biases. In addition, reliable educational resources on GenAI integration are difficult to find [1] and lack of knowledge about GenAI among educators has been reported as a barrier to adoption [1, 5]. This is likely problematic, since GenAI tools have certain characteristics, such as their probabilistic nature, their dependence on training data which can lead to biases of technical [8] and social [7] nature, limited explainability, and potential of generating incorrect information.

Collectively, these observations highlight the need for a more systematic and shared definition of what students should understand about GenAI, and a shift from broad AI literacy frameworks to more fine-grained GenAI literacy [9] that is adapted for computing science students and relevant for their future careers. We see this as a key step towards the development of mental models for working with generative systems: recognising models as stochastic generators of plausible continuations rather than retrievers of facts, treating prompt sequences as a form of programming refined iteratively over time, and adopting a verification mindset in which outputs are hypotheses requiring judgment and validation. Students must also understand how system performance is shaped by their training data and where this leads to failures, manage context and interaction history, and view LLMs as components within broader complex workflows.

2 Goals

There are two primary goals for our working group:

(1) Identify a set of GenAI literacy learning objectives which give students the essential knowledge and understanding required to use these tools responsibly (e.g., fundamental behavior, limitations, ethical implications);
(2) Prepare pedagogical materials that can be shared with the computing education community. This might include a range of media such as documents, slides, videos, and/or pedagogical games (e.g., card games).

3 Methodology

The proposed work focuses on gathering and synthesizing perspectives and insights from the CSE community, and publications on GenAI literacy. The first will take the form of surveys and the latter will be done using a "snowballing"-based literature review. Following this, we will use community feedback to validate, prioritize and improve the identified GenAI literacy learning objectives based on their perceived importance and relevance across course contexts. Finally, we will design pedagogical materials aligned with the learning objectives (e.g., slides, short activities, games). We aim to design these resources such that they remain current with scientific and technical advances as well as their real world impacts and implications, whilst preserving the intent of the core learning objectives established by this working group.

4 Expected Deliverables

The expected deliverables are:

- An overview of GenAI literacy landscape based on the CSE community input and research literature.
- A concise set of community-validated learning objectives capturing the core knowledge students need to responsibly use GenAI tools and systems.
- A collection of open source teaching materials (to maximise community adoption) aligned with the proposed learning objectives, such as slides, short activities, and/or pedagogical games (e.g., card-based formats).

5 Acknowledgments

Partially funded by national funds through FCT — Fundação para a Ciência e a Tecnologia, I.P., under projects UID/6486/2025 (https://doi.org/10.54499/UID/06486/2025), UID/PRR/6486/2025 (https://doi.org/10.54499/UID/PRR/06486/2025) and UID/PRR2/06486/2025 (https://doi.org/10.54499/UID/PRR2/06486/2025), and by the Cyprus Research and Innovation Foundation Programme under agreement No. BRIDGE2HORIZON/0823E/0203 (PINNACLE) and the EU's Erasmus+ Programme under agreement No. 2025-1-CY01-KA220-HED-000364525 (AI4Everyone).

References

[1] 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 *Proc. of the 2025 Working Group Reports on Innov. and Tech. in Comp. Sci. Ed.* (Netherlands) *(ITiCSE-WGR 2025)*. ACM, New York, NY, USA, 61–107. https://doi.org/10.1145/3760545.3783970

[2] Bruno Pereira Cipriano and Pedro Alves. 2025. "ChatGPT Is Here to Help, not to Replace Anybody": An Evaluation of Students' Opinions on Integrating ChatGPT in CS Courses. In *Proc. of the 17th Int. Conference on Computer Supported Education, CSEDU 2025*. Sci. and Tech. Pub., Lda, 236–245.

[3] Bruno Pereira Cipriano and Lúcio Studer Ferreira. 2025. Programmers Aren't Obsolete yet: A Syllabus for Teaching CS Students to Responsibly Use Large Language Models for Code Generation. In *17th Int. Conference on Computer Supported Education, CSEDU 2025*. Sci. and Tech. Pub., Lda, 412–420.

[4] Google Cloud. 2025. DORA State of AI-Assisted Software Development. https://dora.dev/research/2025/. (2025).

[5] Steven Gordon, Paul Denny, Hieke Keuning, Natalie Kiesler, Amruth N Kumar, Viraj Kumar, Juho Leinonen, and James Prather. 2026. ACM Task Force on Generative AI and Programming Assessment. (2026). https://acm-education-genai-task-force.github.io/ACM_Taskforce_GenAI_Report_16Feb26.pdf

[6] Aditya Johri, Ashish Hingle, and Johannes Schleiss. 2024. Misconceptions, Pragmatism, and Value Tensions: Evaluating Students' Understanding and Perception of Generative AI for Education. In *2024 IEEE Frontiers in Education Conf. (FIE)*. 1–9. https://doi.org/10.1109/FIE61694.2024.10893017

[7] Hadas Kotek, Rikker Dockum, and David Sun. 2023. Gender bias and stereotypes in Large Language Models. In *Proc. of The ACM Collective Intelligence Conference* (Delft, Netherlands) *(CI '23)*. ACM, New York, NY, USA, 12–24. https://doi.org/10.1145/3582269.3615599

[8] Rrezarta Krasniqi, Depeng Xu, and Marco Vieira. 2025. SE Perspective on LLMs: Biases in Code Generation, Code Interpretability, and Code Security Risks. *ACM Comp. Surv.* 58, 5, Article 137 (2025), 16 pages. https://doi.org/10.1145/3774324

[9] Davy Tsz Kit Ng, Jac Ka Lok Leung, Xianghan O'Dea, and Samuel Kai Wah Chu. 2025. *From AI Literacy to Generative AI Literacy: Conceptualization and Frameworks*. Springer Nature Singapore, Singapore, 21–38. https://doi.org/10.1007/978-981-95-0950-8{_}3

[10] Google Research. 2024. How much does AI impact development speed? An enterprise-based randomized controlled trial. (2024). https://arxiv.org/abs/2410.12944.

[11] Google Research. 2025. How is Google using AI for internal code migrations? (2025). https://arxiv.org/abs/2501.06972.

[12] Annapurna Vadaparty, Daniel Zingaro, David H. Smith IV, Mounika Padala, Christine Alvarado, Jamie Gorson Benario, and Leo Porter. 2024. CS1-LLM: Integrating LLMs into CS1 Instruction. In *Proc. of the 2024 on Innov. and Tech. in Comp. Sci. Ed. V. 1* (Milan, Italy) *(ITiCSE 2024)*. ACM, New York, NY, USA, 297–303. https://doi.org/10.1145/3649217.3653584