

NSF-IUSE: Developing Artificial Intelligence Literacy Among Undergraduate Engineering and Technology Students Through Case-Based Instruction

Dr. Aditya Johri, George Mason University

Aditya Johri is Professor in the department of Information Sciences & Technology. Dr. Johri studies the use of information and communication technologies (ICT) for learning and knowledge sharing, with a focus on cognition in informal environments. He also

Dr. Cory Brozina, Youngstown State University - Rayen School of Engineering

Dr. Cory Brozina is an associate professor and the Director of First-Year Engineering at Youngstown State University. He completed his B.S. and M.S. in Industrial & Systems Engineering from Virginia Tech, and his PhD is in Engineering Education, also from Virginia Tech.

Nafisa Mehjabin, George Mason University

Taylor Ashley Jackson, George Mason University

NSF-IUSE: Developing Artificial Intelligence Literacy Among Undergraduate Engineering and Technology Students Through Case-Based Instruction

AI literacy is a critical element of current student learning, especially for students in computing, engineering, and technology-related disciplines. In this project, we aim to create a series of case studies to improve students' AI literacy in their first and second years of study. The first part of the project is to gain a foundational understanding of students' knowledge of AI and GenAI. To capture this knowledge, we administered a survey instrument that captures different elements of AI and GenAI literacy. We plan to use the results to guide the formation and implementation of our case study. In this poster, we present our initial survey results and highlight our project plan for the three-year duration. Below, we discuss the survey instrument and an overview of the results. The instrument is available in Appendix A.

AI literacy constructs: Recently, a consensus is emerging around constructs related to AI literacy and the following constructs are guiding our work [1-2]: 1) Know and Understand: Knowledge and understanding of AI fundamental concepts and techniques is a common construct across studies and refers to knowing the basics of AI and understanding how AI applications work from input data to machine learning techniques. 2) Use and Apply: In relation to AI, this construct refers to the operational or functional aspects of AI use. In particular, the ability to use AI tools to accomplish tasks. 3) Evaluate and Create: Evaluation involves the ability to analyze and interpret the outcomes of AI applications critically. Having a comprehensive understanding of AI's technical aspects enables individuals to examine and form informed opinions about their interactions with AI technologies. Create is a less common construct within the literature, but it emphasizes the ability to design and code AI applications. 4) Navigate Ethically: Understanding the ethical and societal implications of AI is crucial to educating citizens to become socially responsible and ethical users of AI. Human-centered considerations such as fairness, accountability, and transparency must be prioritized in addition to AI aspects such as privacy, misinformation, and bias.

Survey instrument development: To assess the AI literacy constructs identified above, we created a survey instrument comprising items mapped to each construct. We used items from already existing surveys to ensure that they were valid and tested. The GenAI Literacy Assessment Test (GLAT) is a 20-item multiple-choice instrument that tests both AI and GenAI knowledge [3]. Jin et al. [3] developed this instrument precisely to address the drawbacks identified above with existing instruments, namely that they largely rely on self-reports of attitudes and experiences, which may be biased, rather than objective tests of knowledge. The Pew AI literacy questionnaire was developed by the Pew Research Center in consultation with Ipsos [5]. We used items from this survey instrument, given its validated testing, development, and implementation, as a baseline that students would at least know or should know what the general U.S. population knows about AI. Chiu et al. [5] was one of the first studies to report on the development of a questionnaire to accurately assess AI literacy. They recognized that, principally, AI literacy studies were using self-reported questionnaires that assessed students' perceived AI capability rather than actual AI knowledge, as they asked learners about their self-perceptions. The validation results showed that the model, and hence the instrument, met the dimensionality criteria, indicating that the items were both reliable and valid. Overall, the knowledge survey

consisted of 30 items that aligned with one of four aspects of AI literacy outlined in the literature: know and understand, use and apply, evaluate and create, and navigate ethically.

Implementation of the Assessment: The survey was conducted within an undergraduate Information Technology course taken by juniors and seniors (i.e., 3rd and 4th year students). A majority of students were minoring in cybersecurity, while the rest were doing a concentration in networking, web development, or cloud computing. Within the course, there was no specific lecture on AI, but students were exposed to the topic through various course materials, including readings on AI and case studies on its use in agriculture and education. At the time the students took this course, they had also not taken any other AI-related courses. This was by design; the survey was created to better understand their baseline knowledge so that future AI courses could teach them relevant concepts. In other words, the survey was designed as a diagnostic or formative assessment tool rather than a summative assessment. The survey was administered online as an extra-credit assignment and implemented using the institution's standard examination system. This online exam system consisted of a lockdown browser functionality – called Respondus – and video monitoring capability to restrict cheating or plagiarism. We used this functionality to ensure that students completing the survey did not use any outside help and that their responses reflected their true knowledge. The average time needed to complete the survey was 13 minutes. Incomplete responses were removed from the sample before analysis, and the final sample size used for analysis was 57 complete responses out of a total student population of 80.

Results: Overall, students demonstrated a good understanding of which applications use AI but lacked an understanding of how applications that use AI operate. Specifically, they lacked understanding of how LLMs work and of related concepts such as Retrieval-Augmented Generation (RAG). The disparity between surface-level awareness of AI systems and technical understanding may explain the lower Cronbach's alpha score in our assessment. This disparity also aligns with the broader AI landscape, where users have prompting skills but lack the computational foundations of AI systems.

In terms of AI literacy constructs, students lacked both an understanding of how AI worked and how it can be applied. This is likely because these aspects are often tied together – unless you know how something works; it is difficult to assess if it can be applied or used in the correct manner. Students scored highest on the construct of how to navigate the use of AI ethically.

For students in technical domains, it will be necessary to move beyond merely being users to a deeper understanding of architectural principles. Developing students' understanding of AI as future engineers who may design, implement, or critically evaluate systems is where introductory curricula should be focused. We recommend that, if readers want to use this survey for summative assessment at the end of the course, they add a few more items directly related to the course content.

Acknowledgement

This work is partly supported by U.S. NSF Award# DUE-IUSE-2439460 and 2439459. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

1. Almatrafi, O., Johri, A. & Lee, H. (2024), 'A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019–2023)', *Computers and Education Open*, vol. 6, p. 100173.
2. Long, D. & Magerko, B. (2020), 'What is AI literacy? Competencies and design considerations', in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–16. doi:10.1145/3313831.3376727.
3. Jin, Y., Martinez-Maldonado, R., Gašević, D. & Yan, L. (2025), 'GLAT: The generative AI literacy assessment test', *Computers and Education: Artificial Intelligence*, p. 100436.
4. Pew Research Center (2023), 'Public Awareness of Artificial Intelligence in Everyday Activities', available at: <<https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/>>.
5. Chiu, T.K. et al. (2024), 'Developing and validating measures for AI literacy tests: From self-reported to objective measures', *Computers and Education: Artificial Intelligence*, vol. 7, p. 100282.

APPENDIX A

Construct	Item (Source: Jin et al., (2025), <i>except 5 items marked* based on Pew (2023), and 4 items marked** based on Chiu et al (2024)</i>)
<i>Know & Understand</i>	Which of the following best describes Artificial Intelligence (AI).
	Thinking about customer service, which of the following uses artificial intelligence (AI)?*
	When using email, which of the following uses artificial intelligence (AI)?*
	Thinking about health products, which of the following uses artificial intelligence (AI)?*
	In the context of AI, what is “zero-shot learning”?
	Thinking about online shopping, which of the following uses artificial intelligence (AI)?*
<i>Use & Apply</i>	You are a music enthusiast and are evaluating the use of AI to improve your listening experience. For which of these activities will you use artificial intelligence (AI)?*
	As a technology enthusiast, you are always looking to make your home “smarter” by using different devices and applications. For which of these activities, will you use your knowledge of artificial intelligence (AI)?*
	After deploying a customer service chatbot, you notice that it frequently provides outdated information about company policies. What is the best course of action to address this issue?
<i>Evaluate & Create</i>	While reviewing a video of a well-known public figure making controversial statements, which characteristic confirms the video was NOT generated by AI?
	You are building a computer vision application and evaluating different sources of data for your product. Which of the following devices <i>cannot</i> provide you with relevant data?***
	You are building an AI system but its performance is under par. You are looking at ways to improve the output quality. Which of the following will most likely enhance an AI system’s performance?***
	As part of a course on AI, you have learned about Natural Language Processing (NLP) techniques you have tasked with creating a tool that uses NLP. Which one of these is NOT a tool that you will create?***
<i>Navigate Ethically</i>	Which of the following actions of an AI developer is NOT ethical?***
	In a healthcare startup, an accurate AI model recommends treatments, but doctors don’t trust it because they can’t understand how the model arrived at its conclusions. What core issue does this scenario illustrate?
<i>Generative AI Items</i>	
<i>Know & Understand</i>	Which of the following best describes "Generative AI"?
	Which of the following statements best describes an LLM (Large Language Model)?
	Which of the following tasks can Generative AI perform with a high degree of accuracy?
	Which of the following is a potential challenge when using prompt-based development for text generation?

	What does the term "token" refer to in the context of a large language model (LLM)?
	How does RAG (Retrieval-Augmented Generation) enhance the capabilities of an LLM?
<i>Use & Apply</i>	When using generative AI to create a marketing pitch, which of the following strategies is least likely to be effective?
	Suppose you have a large dataset of emails and you want to build an application to answer questions based on this dataset. Which of the following scenarios best illustrates the advantage of using RAG over prompting (i.e., without RAG)?
<i>Evaluate & Create</i>	As a student using a Large Language Model (LLM) to gather information for an assignment, how should you approach the information it provides?
	It is unlikely for an LLM to provide an accurate summary of the latest financial market trends in real-time. Is this statement true or false?
	A generative AI tool has provided a summary of a research paper. The summary states, "The study found that increased screen time is directly correlated with decreased attention spans in children aged 8-12." What is your next step?
<i>Navigate Ethically</i>	What are the potential copyright implications for a journalist using an AI-generated image in a commercial article?
	Should we impose restrictions on the outputs of generative AI technologies?
	Sending personal information to cloud-based generative AI tools has little privacy concerns.
	When a generative AI system is used for screening job applications, what issue might arise concerning the quality and fairness of hiring decisions?