Investigating User Attitudes Towards and Benefits from Integrating AI Assistants into Research Computing Support

Injila Rasul
University of Massachusetts Amherst
irasul@umass.edu

Georgia Stuart
University of Massachusetts Amherst
gstuart@umass.edu

ABSTRACT

High-Performance computing clusters used for Research Computing, hosted by universities, are an essential part of the ongoing teaching, learning, and research at these institutions. Users must understand myriad scientific, mathematical, and computing concepts. They have a range of experience and comfort with these platforms, requiring regular support as they engage with it for their research. To assist users on the Unity Research Computing Platform, the support team provides the Facilitation Slack channel to get help, find relevant documentation, learn new information, and troubleshoot, requiring significant staff time and funding. This study explores the design and implementation of an AI assistant chatbot augmenting existing support with HPC Facilitator oversight. We investigate the efficacy of AI assistants in extending the productivity and impact of research computing personnel while maintaining a high degree of direct contact with users. We discuss the Human-Centered AI Design and testing process and its significance for large-scale interventions.

KEYWORDS

Human-Centered AI, High Performance Computing Facilitation, Research Computing Support

1 INTRODUCTION

The use of AI-chatbots as assistants for troubleshooting is documented across a range of platforms [4]. These AI chatbots are forms of software that simulate human-like conversations and interactions, understand user prompts, process input, and provide personalized responses [6]. AI chatbots function as digital assistants for learning and teaching, help expand user skills and abilities, and support higher-order executive skills [6].

Research computing is a field with the potential to be significantly improved by AI-assisted support. When interacting with a research computing platform, such as a High-Performance Computing (HPC) cluster, users must understand a wealth of scientific, mathematical, and computing concepts. To assist users, HPC centers employ facilitators or other support staff. However, funding enough user support staff to meet the needs of a typical HPC center is often untenable. In this study, we address initial feasibility questions around using AI assistants in research computing support, specifically in the context of the Unity Research Computing Platform (Unity) at the University of Massachusetts Amherst.

Unity is a collaborative and high-performance research computing platform that supports collaboration across multiple universities and colleges in the northeastern US. Platform users need assistance with issues ranging from simple troubleshooting, such as signing in, managing files, installing software, or accessing datasets, to complex code and workflow optimization questions. The Unity facilitation team offers help through a variety of modes, notably a Slack group. However, more basic queries may be addressed by an AI-enhanced search function, which primarily directs users to relevant areas of the documentation that can help resolve the issues that they encounter. It is imperative that the design and roll-out of an AI assistant within this system is done while being mindful of concerns such as incorrect understanding by the chatbot, user frustration, and lack of trust in AI responses [1].

The design and testing process of this AI-enhanced search/chatbot is the goal of this feasibility study. The following questions guide our work:

(1) In what ways does an AI-enhanced search for Unity users effectively facilitate basic questions that arise in research computing?
(2) What is the role that the AI chatbot can play in assisting the facilitation team in clarifying user questions and filtering out common misunderstandings?

2 RELATED WORKS

We discuss research on user involvement in design, a human-centered framework for studying AI systems, and human-AI interaction breakdowns. We describe Human-centered design in AI and discuss why it is essential to iteratively create, test, and refine the designs of such systems to minimize negative user experiences. Finally, we summarize the relevant literature on AI-assisted user experiences in education, with emphasis on the need for continued human oversight.

2.1 User-Involvement in AI Design

Prior work on integrating Artificial Intelligence into Education explores its potential impact due to its scalability. However, concerns about privacy and agency are highlighted, as end-users are often excluded from the design process [1]. Actively involving the target audience in the design of AI systems fosters trustworthiness, reliability, and a balance between human control/oversight and AI automation [2]. Incorporating end-user insights from the start of the design process helps address the identified need for greater consideration of ethical, methodical, and contextual factors [2].

https://doi.org/10.22369/issn.2153-4136/17/1/12

The viability and feasibility of AI chatbots have been investigated in studies examining dimensions of interaction in comparison with human-human interaction, yielding framework elements such as (i) words per message, (ii) words per conversation, (iii) messages per conversation, (iv) word uniqueness, and (v) use of profanity, shorthand, and emoticons [3]. User studies have also examined the framing of AI responses and their connection to aspects such as frustration, usefulness, appropriateness, and relevance in response to AI chatbot interactions [9]. Key areas of concern are breakdowns in user-LLM interactions that lead to poor user experiences, which users categorize as "irrelevant and useless output," "inability to answer," and "problematic communication style" [7].

2.2 Human-Centered AI Design Process and Principles

Human-Centered AI (HCAI) emerges when Human-Centered Design (HCD) approaches (such as centering human need, values, and perspectives) are applied to the design and development of learner analytics or AI in Education systems [5]. A key tenet is considering AI as a system that serves human interests rather than pursuing a technical goal or aiming to replace humans altogether [8].

Design Phases of the HCD process are defined as follows: Phase 1: Planning, Scoping, and Definition; Phase 2: Exploration, Synthesis, and Design Implications; Phase 3: Concept Generation and Early Prototype Iteration; Phase 4: Evaluation, Refinement, and Production; Phase 5: Launch and Monitor. This study reports results from a system in Phase 4.

Alfredo et al. outline HCAI Principles of Safety, Reliability, and Trustfulness for study design that emerged from their thematic analysis of AI system design research [1]. These principles are interconnected and do not exist in isolation from one another; rather they feed into the design process throughout and work together to foster confidence in the system and therefore its successful adaptation and use.

2.3 Safety

Data Privacy, Data Sharing, Data Collection, and Monitoring and Surveillance emerge as significant themes under the umbrella term of Safety. These elements are essential to consider when designing AI systems so that they are integrated into the DNA of the system and not transplanted in at a later stage. Data related to personal identification, academic performance, or learning progress should be considered sensitive, with the decisions about privacy influencing stakeholder willingness towards adopting the AI system. Data Sharing concerns are usually addressed by anonymizing participant identities, which allows task performance without peer judgement. For Data Collection transparency, studies needed clearer communication between researchers and participants around informed consent. For Monitoring and Surveillance, concerns around student awareness of monitoring and resulting changes in behavior are an important consideration.

2.4 Reliability

The Reliability principle is framed around system accuracy, bias, and strategies to ensure data integrity. System reliability is the accuracy and perception of the machine algorithm and human interpretation of the system. For bias, when researchers/designers are working in isolation, it increases risk of human interpretation bias to be introduced into the system and decisions around the design of the system. Therefore, including diverse perspectives, ideally from including end-users, actively in the design process increases reliability of the system.

2.5 Trustfulness

Trustfulness in the system can refer to stakeholder trust, user perception, and accountability. Trust in human-centered AIED systems is improved by providing an explanation when users request feedback and accuracy of the information it provides. Additionally, it connects to transparency around the way the system gathers and manages information and communicates or documents research.

2.6 Human-in-the-Loop by Design in Self-Regulated Learning

Research around AI in Self-Regulated Learning (SRL) is in its infancy [4]. SRL is very nuanced, with conceptual models and frameworks drawing from different aspects of behavioral, cognitive, and affective theories of learning. Notably, Zimmerman's (2002) SRL model has three phases: Forethought (i.e., goal setting, strategic planning), Performance (i.e., help-seeking, learning monitoring, and instructional strategies), and Self-Evaluation (i.e., self-reflection, feedback) [10]. Therefore, AI applied in this domain will also need a broad and deep capacity for augmenting and supporting SRL.

The distinction between human-centered self-regulation with AI as facilitator (a tool to monitor, control, and support learning) and AI-centered self-regulation with AI making data-based decisions is an important one to make cite4. Most AI applications engage target-specific dimensions of the SRL cycle rather than the whole process holistically and thus support some aspects more than others.

2.7 Breakdowns in User-AI Interaction

Errors and warnings arise in user-AI interactions when user requests cannot be fulfilled due to technical limitations or policies restricting specific responses [9]. When these breakdowns occur in conversations, users might repeat or clarify, emulating a human-human interaction, but for a chatbot interaction, that might not lead to a clarification but rather frustration. Handling user frustration in such instances is crucial for managing these interactions effectively. Wester et al. [9] tested methods to manage the HCI in the case of request denials and found that avoiding baseline denials that do not adequately inform the user as to why their request is not answered. Instead, diverting denial styles were more favorable, which navigates the denial by suggesting alternatives that the AI can assist with [9].

3 METHODOLOGY

For the context of this research, "effective facilitation" is operationally defined as users perceiving themselves as more capable of solving problems with the tool in use than without it. The bespoke instruments were piloted and refined using feedback from the Unity facilitation team, which consists of experts in optimizing High Performance Computing workflows and designing learning technology.

Table 1: Pre and Post-Test Questionnaire

	QUESTION COMPONENTS
LEARNER PROFILE	1. I am a frequent user of research computing clusters (for example, Unity). 2. I am comfortable using research computing clusters to conduct my research. 3. I consider myself well-versed in solving computing issues. 4. When I face an issue, I can comfortably navigate the internet or documentation to help solve my issue. 5. I am comfortable navigating through software documentation to resolve issues I face. 6. When I feel stuck in a platform I am using, I utilize any available AI-assistance. 7. I prefer AI-assistance to human-assistance for questions around computing/software. 8. If I get stuck on a technical problem during research computing, there is no chance I'll figure it out on my own.
AI-USE	1. I am not comfortable prompting AI chatbots for assistance when I am stuck on a technical problem. 2. If I get stuck on a technical problem during research computing, there is no chance I'll figure it out on my own. 3. There is usually only one correct approach to solving a technical issue while working on research computing. 4. AI chatbots are confusing to use. 5. The instructions that AI-chatbots give are clear or easy to follow. 6. When I use AI chatbots to problem-solve, I feel frustrated if the answer is unclear to me. 7. AI chatbots can give incorrect and misleading answers to my question prompts. 8. AI chatbots are useful for preliminary assistance for simpler problems. 9. I prefer human assistance when the problems are more complex or multi-faceted. 10. I am better able to solve problems when I have a human assisting me. 11. I prefer human assistance when I am facing software/computing issues, compared to AI assistance.
COGNITIVE WALKTHROUGH	The participant will be given a series of tasks, one at a time, to complete on the Unity HPC Cluster. Once a task is completed to the satisfaction of the participant, the researcher observer will give them the next task. Tasks are intentionally open-ended, but can be answered from existing documentation. Tasks are roughly ordered by difficulty. Task list (example): 1. Log into Unity OnDemand and upload a test file to your home directory (test file is an ipynb file for subsequent tasks) in a subdirectory named "Unity ai_study". 2. Start a JupyterLab session with 4 CPU cores and 4 GB of Memory for a 1-hour time limit in the "cpu" partition. 3. Open the notebook (ipynb) file uploaded in Task A with the default Python kernel and execute all cells. 4. Create a non-interactive job to run the notebook (ipynb) uploaded in Task A without the JupyterLab Unity OnDemand app.
INTERVIEW: SEMI-STRUCTURED QUESTIONS	1. Which tasks, if any, were you familiar with in advance? 2. In general, what is your opinion on AI-chatbots being used to assist in searching for answers or clarification? 3. In general, what is it like for you to interact with an AI-chatbot? 4. Can you think of a time when you felt frustrated while using an AI-chatbot? 5. Overall, would you consider AI-chatbots to be useful in your experience, and why? 6. In your experience, have AI-chatbot responses typically been appropriate to the situation and/or your question? 7. In your experience, was the content of the AI-chatbots' answer clear and relevant to your question? 8. Have you asked questions in the user Slack before? 9. If YES, in what ways was your experience different? Was it better or was it worse? 10. If NO, what are the reasons why you have not used the user facilitation slack channel?

3.1 Data Types and Collection

There are four categories of data collected for this project. The collection and analysis of data is currently ongoing. Instrument components are given in Table 1.

(1) Pre- and Post-test (qualitative in nature) taken by users on user experience and self-perception of computing ability, administered through Qualtrics.
(2) One-on-One semi-structured interviews with users on user experience
  (a) Cognitive walkthrough exercises,
  (b) Experience-related questions.
(3) Pre- and Post-test (quantitative in nature) taken by the user on content knowledge before and after the bot facilitation experience, administered through Qualtrics.
(4) Qualitative analysis of the question-and-answer log from the active Unity help-desk channel.

3.2 Participants

Our target audience is the existing Unity users, typically research teams utilizing the cluster for their research projects. This includes PIs, research associates, and student research assistants who are using the system for their active research projects. User consent is obtained at the start of the testing regarding the use of their data for developing and refining the design of the AI chatbot.

3.3 Data Analysis

The collected data is analyzed for themes around user experience, the degree of facilitation the AI-chatbot provides, and the ways in which it can improve. The qualitative coding process involves classifying major and minor themes that emerge from the data, as well as developing a codebook. The initial codebook uses the HCAI framework of themes around Safety, Reliability, and Trustfulness.

4 PRELIMINARY RESULTS

Data collection and analysis are currently ongoing and scheduled to conclude by early Fall of 2025. Data analysis is an iterative process conducted in collaboration with Unity team facilitators and administrators, who participate in work sessions and presentations through which emergent and key findings and patterns are discussed. Experts in high-performance computing and computational science education raised questions, shared insights, and surfaced contextual knowledge that helped shape the ongoing AI assistant development and strategies for data collection and analyses. In line with the centers' goals, decisions about which analyses to conduct and what data would be most useful were made by researchers.

Discussions with the HPC facilitators and review of the existing Slack channel messages indicated a level of redundancy in existing questions, prior to deployment of the AI chatbot. The facilitation team found that the developed AI assistant reasonably answered many of the common questions. However, data collection from participants who are not subject matter experts is forthcoming. The facilitator panel further provided feedback to refine the tasks for the cognitive walkthrough to cover tasks at a range of complexity, shown in Table 1.

5 SIGNIFICANCE

Rolling out the AI-chatbot to the entire audience of researchers at the six different institutes that use this research cluster and eventually integrating it within the User Facilitation Slack is a large-scale change to the system and requires a learner-focused approach to its design and implementation. The study focuses on ensuring AI is integrated intentionally and intelligently, not for the sake of an AI tool but rather with the intention of centering learner needs and utility. A human-centered approach to developing and deploying new AI tools ensures they are tailored to address gaps in human-provided support and support humans who are the lead facilitators.

Retaining the Human-in-the-Loop also ensures that learner frustration is managed, learners are understood and supported in their work, and that custom facilitation and support are readily available. The automation of facilitation in the research computing environment is essential; however, exercising caution around the design, development, and large-scale deployment of such a tool is necessary to be mindful of the limitations of such technology. The bespoke instrument we developed provides fellow researchers with a starting point for assessing learner experience, identifying pain points, and addressing the issues that are identified. Positioning the AI tools in the facilitator role, while incorporating ongoing monitoring and evaluation by the human ensures that the end-user's learning is managed while lowering the teaching burden on the team.

Ultimately, this design approach strengthens the process for scaling high-quality learner facilitation and learning in a more informal SRL context. It reiterates that AI-enhanced learning technology tools should be developed through an iterative process, driven by existing and evolving literature on design principles, and be contextually intelligent by integrating the needs of the end-learners from the start.

ACKNOWLEDGEMENTS

Our thanks to NSF for supporting this research effort [Award Number: 2430001, Project Title: CC* Compute-Region: Harmony for Research Computing: A Green Compute Solution Designed to Enable Emerging Workflows and to Maximize Access for New England Researchers and Educators].

REFERENCES

[1] R. Alfredo, V. Echeverria, Y. Jin, L. Yan, Z. Swiecki, D. Gašević, and R. Martinez-Maldonado. 2024. *Human-centred learning analytics and AI in education: A systematic literature review.* Artificial Intelligence, 6, 100215, Computers and Education. https://doi.org/10.1016/j.caeai.2024.100215

[2] M. Bond, H. Khosravi, M. De Laat, N. Bergdahl, V. Negrea, E. Oxley, P. Pham, S. W. Chong, and G. Siemens. 2024. A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education* 21, 1 (2024), 4. https://doi.org/10.1186/s41239-023-00436-z

[3] J. Hill, W. Randolph Ford, and I. G. Farreras. 2015. Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations. *Computers in Human Behavior* 49 (2015), 245–250. https://doi.org/10.1016/j.chb.2015.02.026

[4] M. Lan and X. Zhou. 2025. A qualitative systematic review on AI empowered self-regulated learning in higher education. *Npj Science of Learning* 10, 1 (2025), 21. https://doi.org/10.1038/s41539-025-00319-0

[5] C. Lang and L. Davis. 2023. Learning Analytics and Stakeholder Inclusion: What do We Mean When We Say "Human-Centered"?. In *LAK23*. 411–417. https://doi.org/10.1145/3576050.3576110

[6] P. Pergantis, V. Bamicha, C. Skianis, and A. Drigas. 2025. AI Chatbots and Cognitive Control: Enhancing Executive Functions Through Chatbot Interactions: A Systematic Review. *Brain Sciences* 15, 1 (2025). https://doi.org/10.3390/brainsci15010047 Article 1.

[7] M. Skjuve, A. Følstad, and P. B. Brandtzaeg. 2023. The User Experience of ChatGPT: Findings from a Questionnaire Study of Early Users. In *Proceedings of the 5th International Conference on Conversational User Interfaces*. 1–10. https://doi.org/10.1145/3571884.3597144

[8] U. A. Usmani, A. Happonen, and J. Watada. 2023. Human-Centered Artificial Intelligence: Designing for User Empowerment and Ethical Considerations. In *5th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*. https://doi.org/10.1109/HORA58378.2023.10156761

[9] J. Wester, T. Schrills, H. Pohl, and N. van Berkel. 2024. As an AI language model, I cannot. In *Investigating LLM Denials of User Requests. Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/10.1145/3613904.3642135

[10] B. J. Zimmerman. 2002. Becoming a self-regulated learner: An overview. *Theory into practice* 41, 2 (2002), 64–70.