

Ogbo-Gebhardt, Ezezi; Ogbo, Oruaro

Conference Paper

Using a Large Language Model-Powered Assistant in Teaching: Stories of Acceptance, Use, and Impact among Ethnic Minority Students

24th Biennial Conference of the International Telecommunications Society (ITS): "New bottles for new wine: digital transformation demands new policies and strategies", Seoul, Korea, 23-26 June, 2024

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Ogbo-Gebhardt, Ezezi; Ogbo, Oruaro (2024) : Using a Large Language Model-Powered Assistant in Teaching: Stories of Acceptance, Use, and Impact among Ethnic Minority Students, 24th Biennial Conference of the International Telecommunications Society (ITS): "New bottles for new wine: digital transformation demands new policies and strategies", Seoul, Korea, 23-26 June, 2024, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/302517>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Using a Large Language Model-Powered Assistant in Teaching: Stories of Acceptance, Use, and Impact among Ethnic Minority Students

Erezi R. Ogbo-Gebhard^a and Oruaro Ogbo^{b,c}

^aNorth Carolina Central University

^bMeta Platforms, Inc.

^cBelieve Better Ltd.

Keywords: Generative Artificial Intelligence, Large Language Models, UTAUT, Digital Inclusion, HBCU, Pedagogy Technologies

Abstract: *Generative Artificial Intelligence (Gen AI), such as pre-trained Large Language Models (LLMs), is being introduced into society and is expected to transform various fields of human endeavor. Although researchers have explored how LLM-powered tools can assist in student learning, limited studies have been conducted on the adoption, acceptance, and use by ethnic minority students, a group often underserved and underrepresented in developing and deploying cutting-edge technologies. This study analyzes qualitative data from graduate students and recent alumni of a US Historically Black College and University. We implement an LLM-powered tool that leverages Retrieval Augmented Generation (RAG) grounded in their teaching material, then ask participants to complete tasks based on the material, with and without the presence of the tool. We conducted semi-structured interviews before and after the tasks to obtain first-hand feedback on using Gen AI tools in a practical educational scenario. This study sheds light on some of the expectations, use, and barriers to acceptance from ethnic minority groups and can inform future designs and implementations of Generative AI tools in pedagogical environments.*

1. Introduction

Since the release of ChatGPT in November

2022, there has been a growing body of literature examining the use, impact, and ethical considerations of tools powered by large language models (LLMs). LLMs can comprehend and generate high-quality human-like text in near real-time and are expected to impact all sectors, with some industries, such as high tech, healthcare, and education more likely to see significant transformative impact (Chui et al., 2023). Within education, LLMs are expected to transform the way students learn by “serving as virtual tutors, offering instant feedback, explanations, and assistance to students” (Jermakowicz, 2023).

Several studies have assessed the use and impact of LLM-powered tools in an authentic classroom setting and found that they promote students’ motivation to learn (Hmoud et al., 2024) and improve learning outcomes such as completion speeds, student satisfaction (Choi et al., 2023), self-efficacy in writing tasks (Pellas, 2023), and higher grades (Sheese et al., 2024). So far, limited research has investigated the acceptance, use, and impact among ethnic minority students. With advancements in LLM’s capabilities and adoption, it is critical to understand the role it could play in digital inequity. Specifically, there’s a need for LLM research that prioritizes groups that are often digitally unserved and underrepresented in technology research.

This pilot exploratory study is conducted with graduate students and recent alumni (n=8) of the School of Library and Information Sciences (SLIS) at North Carolina Central University (NCCU). NCCU is a small, comprehensive, historically black college

¹Erezi Ogbo-Gebhardt
eogbo@nccu.edu

Oruaro Ogbo
oru@meta.com

and university (HBCU) in the United States. The majority of the institution's student body is categorized as members of the "covered population" defined in the Digital Equity Act Sec. 60302(8) as those who are more likely to be disadvantaged in the digital divide, including those that are low-income, from racial and ethnic minority groups, veterans, among others.² This study seeks to investigate the following research questions:

1. What factors influence the acceptance of an LLM assistant for learning among ethnic minority students?
2. What are their lived experiences when using an LLM assistant for learning?
3. What factors influence their use patterns of the LLM assistant for learning?
4. What is the impact of an LLM assistant on their educational outcomes?

To answer these questions, the participants are asked to complete two technical Database administration and analysis tasks and are given access to an LLM assistant for the second task. We conducted two semi-structured interviews with the participants, before and after completing the tasks, to understand their perceptions towards using LLM tools for learning and their lived experiences using the tool.

Through a systematic review of the qualitative data, this paper identifies emergent themes, provides comparisons, and presents insights on effective strategies for using LLM-powered tools to improve the learning experience and outcomes of minority students. This work contributes to the literature on digital equity and inclusion, and emerging pedagogical practices in higher education.

This paper begins by discussing the theoretical background, focusing on the digital divide, models of acceptance and use, and the related literature on the use of LLMs for learning. It will then go on to present the methodology of this study, followed by a discussion of the findings. The paper concludes with a discussion of the main contributions of the paper, the study limitations, and ideas for

future work.

2. Theoretical Background

This section begins with a brief discussion of studies centered on digital inequality and the digital divide. Next, in 2.2, we present an overview of the UTAUT model. We then discuss some of the studies on LLMs when used in pedagogical settings (2.3).

2.1 The Digital Divide

Research on the digital divide has a long history. Defined in the mid-1990s as the differences "between people who have access and use of digital media and those who do not" the digital divide has evolved in scope (Dijk, 2020). Today, a more widely accepted definition is "the gap between those who have affordable access, skills, and support to effectively engage online and those who do not" (NDIA, n.d.).

Evidence from the literature shows that the digital divide varies across demographic and socioeconomic groups, with historically marginalized groups more likely to be disadvantaged. A comprehensive list of digitally vulnerable groups was defined in the Digital Equity Act Sec. 60302(8) and includes low-income families, people with low literacy, racial and ethnic minorities, veterans, and English learners, among others. It is well established from a variety of studies that beyond lack of broadband access, the main barriers to the adoption and use of broadband and related digital technologies are lack of affordability - which includes of services and devices, limited digital literacy and technical skills, and the lack of perceived relevance of connectivity (Horrigan, 2010; Horrigan and Duggan, 2015).

Significant effort has been made to improve digital equity by connecting the unconnected, providing digital literacy training, and offering low-cost devices. However, with the expected ubiquitous impact of Generative AI (Gen AI), there is a concern that it could create a new digital divide frontier.

2.2 The UTAUT Model

² [Digital Equity Act of 2021](#)

Several theoretical models have been developed and applied in studies of diffusion of innovation to understand the factors that lead to the adoption and use of new technologies. Some early theories in this field include the technology acceptance model, the social cognitive theory, and the theory of reasoned action. These theories vary in the extent to which they are able to explain the variance in individual intentions to use technology.

The UTAUT, developed in the early 2000s and based on eight previously established theoretical models, proposes four determinants of technology acceptance and use: performance expectancy, effort expectancy, social influence, and facilitating conditions. Performance expectancy is a user's expectation that the technology improves their performance in their task. Effort expectancy refers to the extent to which people find it easy to use the technology. Social influence is the degree to which an individual believes that people who are important support their use of the new system. Social influence includes the impact of coworkers, peers, friends, and families, as well as institutional policies. Facilitating conditions refers to the availability of resources – such as financial resources, know-how, and technical support – to support the use of the technology.

Models based on the UTAUT explain up to 70% of the variance in user behavior, compared to between 17% and 53% found in models based on any of the individual eight theories of technology acceptance and use. The UTAUT has been widely applied in different disciplines and contexts. For example, it has been used to examine the adoption and use of general-purpose technologies such as broadband (Ogbo et al., 2021); old technologies in a new context such as online learning platforms among hospitality students (Sabri et al., 2023) and health information technologies and eHealth literacy (Magsamen-Conrad et al., 2020); as well as emerging technologies – such as cryptocurrency (Ebizie et al., 2022; Ishak et al., 2022).

2.3 Related work: Use of LLMs in teaching

The adoption of generative AI across many sectors, including higher education, has been rapid,

with limited time to establish policies to guide the effective use of the tool (Crompton and Burke, 2023). Several studies have examined the possible opportunities and challenges in education presented by generative AI tools, with many highlighting the need for a balanced policy that supports innovation yet upholds academic integrity (Alberth, 2023; Xiao et al., 2023). (Xiao et al., 2023) examined the approaches implemented by leading universities around the globe to guide the use of Gen AI tools for teaching and learning. Unsurprisingly, they found significant variation in the policies established. Less than a third of universities in the study had a policy, out of which 67% embraced Gen AI. Among the universities that ban the use of Gen AI, concern about violating academic integrity was a major factor. As universities move to create or update their policies on Gen AI use, according to (Chiu, 2024; Johnston et al., 2024), it is important to consider students' views when developing policies on Gen AI use in education.

Recent evidence has accumulated to show that Gen AI improves learning outcomes. Drawing on qualitative data from 15 graduate students with at least one semester's experience using Gen AI, (Hmoud et al., 2024) argue that generative AI could promote students' motivation to learn and thus improve their learning outcomes. Similarly, (Pellas, 2023) found that Gen AI improved students' writing efficacy. However, they note no significant impact on students' perception of creative identity in their writing. A randomized control trial conducted with 59 law school students at the University of Minnesota Law School, (Choi et al., 2023) found that students saw improved completion speeds and were more satisfied with using GPT-4 to complete legal tasks. Interestingly, they also reported that the lowest skilled participants saw the largest improvements in the quality of their output with AI assistance. Relatedly, an analysis of 52 students' use of an LLM-powered tool in a computer programming class showed that increased frequency of use is associated with higher grades (Sheese et al., 2024).

Many recent studies have investigated how students perceive Gen AI tools, with several

reporting mixed emotions regarding the benefits for learning versus concern about academic dishonesty. Although students report that Gen AI improves personalized learning, supports research, and provides writing assistance, they are worried about accuracy, privacy, and plagiarism (Chan and Hu, 2023; Walczak and Cellary, 2023). Students are also concerned about being prepared for a changing work environment that relies more on Gen AI (Chan and Hu, 2023; Chiu, 2024). A handful of studies examine the acceptance, use, and impact of Gen AI among minority students. For example, a study of ChatGPT use among 4 low-income Chinese language learners in high school suggests that its use improves student writing and empowerment (Li et al., 2023). In another study that examined how Gen AI could improve educational outcomes of international students in OECD countries, the authors note that: “it is vital to recognize that AI is not a one-size-fits-all solution and should be implemented in ways that align with international students’ unique needs and cultural differences” (Wang et al., 2023).

3. Methodology

This study utilizes a mixed methods approach. However, only the results from the qualitative study are presented in this paper. Participant recruitment and data collection took place over four weeks in April and May 2024. Participants were recruited from current students and recent alumni of the School of Library and Information Sciences (SLIS) and North Carolina Central University (NCCU).

To be eligible to participate in the study, subjects needed to be 18 years or older and enrolled in or previously taken the Database Systems class offered in SLIS. To recruit participants, we emailed the entire student body of NCCU’s SLIS, allowing participants in the target population to opt in. Eight graduate students were recruited into this study—*names anonymized for confidentiality*. Participants' summary characteristics and prior experience with generative AI tools are presented in Table 1.

Table 1: Characteristics of interview participants and their prior experience and reported trust with generative AI tools

Participant	Age	Gender	Race	Trust in Gen AI*	Frequency and Type of Gen AI Use over the Past Three Months
Participant 1	21 - 30	Male	Black	5	Once a week. Searching for additional information during informal discussions.
Participant 2	31 - 40	Female	Black	4	Two to three times a week. Understanding new concepts in easy-to-understand terms.
Participant 3	50+	Female	Black	3**	Once a week. Diverse occupational use cases based on need e.g., editing or brainstorming.
Participant 4	31 - 40	Female	Asian	3.5	Once a week. Diverse use cases from academic research to brainstorming in personal contexts.
Participant 5	21 - 30	Male	Asian	2	Once a month. Brainstorming in personal and informal contexts.
Participant 6	21 - 30	Male	Black	5	Daily. Used at work and in personal contexts.
Participant 7	41 - 50	Female	Black	4.5	Twice a month. Testing out the tools to understand its capabilities.
Participant 8	50+	Female	Black	4	Twice a month. Assistant in automating processes and retrieving information for informal discussions.

* Responses were based on a scale of 1 (least) to 5 (most)

** Neutral response without committing to a figure.

Participants were asked to spend about an hour completing virtual research tasks. On obtaining informed consent from the participants, they were asked to complete a brief questionnaire to collect basic demographic information. Next, we conducted a 15-minute semi-structured interview to collect information about their default approach to searching for information, prior use of generative AI tools, and factors that impact their prior use, such as their perspectives on the impact of generative AI tools for learning, required skills to use these tools, and concerns with the use of these tools for learning. Sample questions included “What do you think would be the impact of the LLM-powered assistant on your learning?” and “What are some concerns that you have with using the LLM-powered assistant for learning, if any?”

Then, participants were asked to complete two database administration and analysis tasks. To reduce bias, each task was a multiple-choice test of 4 questions each, delivered using Microsoft Office Forms, timed, and auto-graded. Students were asked questions that required knowledge of database concepts and writing queries and given up to 10 minutes to complete each task. During the second task, participants were allowed to use the LLM-powered assistant (accessible through WhatsApp, a secure and free instant messaging app).

The final stage of the study consisted of a follow-up 15-minute semi-structured interview. During the follow-up interview, we asked questions about their experience using the tool, the impact of the tool in assisting to complete the task, and their confidence in the support provided by the tool. Sample questions include “What are some examples of ways that [the LLM-powered assistant] was most useful?” and “On a scale of 1 to 5, where a higher number indicates more trust, to what extent do you trust the responses given by the tool?” Study participants received a \$10 Amazon gift card to compensate for their time.

All interviews were conducted over Zoom, transcribed using Otter.ai—with a quality manual check by the research team—and entered into Atlas.ti for coding and content analysis. One researcher used

Atlas.ti to mark sections of the transcript that related to the core research questions in this paper. The full research team conducted the final review of the coded segments.

Ethical approval was obtained from the North Carolina Central University Institutional Review Board, approval number 1201780.

4. Designing AI Assistants for Minority Students

We developed the AI assistant for this study using Anthropic PBC's Claude 3 Sonnet language and vision model with Retrieval Augmented Generation (RAG), a technique introduced by (Lewis et al., 2021) to improve NLP generative tasks. We limited the use of the model to language tasks, using it like a traditional LLM.

To generate responses, the LLM queried and retrieved embeddings of relevant course material, including reading material, lecture notes, and recordings of the synchronous sessions. We used Amazon's Titan Embeddings G1 - Text v1.2 model to generate 1536-dimensional vector space embeddings of the material and stored them in an Amazon OpenSearch Serverless vector database for the LLM's use. We hosted the LLM and vector database on Amazon Bedrock, ensuring efficient retrieval and scalability.

Participants engaged with the LLM assistant through a familiar and accessible interface that requires minimal digital skills: WhatsApp chat. This allowed for real-time educational support and interaction.

4.1 Pedagogical quality control (data sourcing)

We designed a RAG-based LLM-assistant system that retrieves and uses course-related material because:

1. By limiting responses to course material, we sought to maintain quality control over the responses most likely returned to study participants. While LLMs are known to sometimes ignore newly retrieved information that differs from their priors (Wu et al., 2024), by constraining the study to a well-established

Computer Science field of study, we aimed to limit the effect of this tension.

2. LLMs are known to hallucinate when a model generates responses that are random and not grounded in facts or relevant to the input. RAG has been demonstrated as an effective method to mitigate hallucinations (Kang et al., 2024; Shuster et al., 2021). RAG systems have also been demonstrated to outperform fine-tuning LLMs with small datasets (Li et al., 2024). Given the limited size of teaching material, RAG was thus a more suitable solution for this study.

4.2 Technological considerations

We made some design decisions because the study was conducted on a “covered population”.

- *Accuracy*: To improve the accuracy of generated responses, we used a simple RAG architecture to augment the capabilities of the chosen LLM.
- *Accessibility*: To reduce the barrier to adoption, we chose to use WhatsApp, which is free and has a familiar chat interface.
- *Latency*: We located the LLM inference servers and RAG database in AWS’s North Virginia data centers, less than 300 miles from NCCU, for lower latency responses.
- *Privacy*: End-to-end (e2e) encryption of messages and pseudonymization of personally identifiable information was a key consideration in the study design and informed technology choices, e.g., using WhatsApp (which provides e2e message encryption) and only storing information necessary to complete the study.
- *Ethics, anthropomorphism, and creativity*: It has been argued that some anthropomorphic generation from Advanced AI Assistants, such as the LLM-powered tool, can foster trust and lead to increased adoption (Gabriel et al., 2024). However, this could also lead to overreliance on the responses generated by the assistant. Given the educational context of the study, we limited anthropomorphism by

deciding not to give the assistant a name and limiting responses to the specific learning content.

5. Results and Discussion

In this section, we discuss the interview findings, with participants’ perspectives before using the LLM-powered assistant to complete the tasks interspersed with their reflections after using the tool.

5.1 Acceptance of LLM-Powered Assistants

After using the tool to complete the tasks in the study, the general consensus among participants is that they recommend using the LLM-powered assistant for teaching. For a couple of participants, this reflects a shift in their perception before using the tool in the study. While all the participants had prior experience using a generative AI tool, only three reported using it to support their learning during their degree program. We begin our discussion on the findings on the factors that impact the acceptance of LLM-powered assistants with the UTAUT model, and then highlight other factors on acceptance that emerged from the interviews.

5.1.1 UTAUT Acceptance Factors

A common view that emerged is that LLM tools improve the learning process and outcomes. Within the UTAUT model, this falls under the Performance Expectancy construct. Participants noted that LLM tools could complement learning and help parse information, save time, and assist with brainstorming. However, some stated that it is important that students are not too dependent on the tool.

“In some cases, it could help you with getting a paper started, or getting an idea off the ground, not for it to write the entire paper, but to generate some thoughts.” (Participant 8)

The Effort Expectancy construct in the UTAUT model measures the influence of technical skills on the acceptance and use of technology. As noted by Participant 6, “it takes a good understanding

to use generative AI for learning.” While all the participants in the study reported that they have the skills to use an LLM-powered assistant for learning, their perspectives on the most important required skill varied. For example, some participants identified a clear understanding of what they need the tool to do and the ability to construct the prompt properly as the primary required skill. A couple of participants reported that patience is required. The views on a need to understand how the tool works varied; one participant stated that it is important to understand how the tool works, while several argued that it is not.

The Social Influence construct in the UTAUT model captures the impact of family and friends, as well as institutional policies. Half of the participants reported that a member of their family or a friend used generative AI and either introduced them to the tool or assisted them with using it. When asked about their perception of the acceptance of generative AI tools across the institution, the majority of participants argued that there’s an absence of a centralized policy and that acceptance of the tools is determined by the individual faculty member. However, most participants believe that SLIS faculty are generally more accepting of using generative AI for learning compared to the rest of the university, attributing it to SLIS being a technology-focused school.

“I believe SLIS is very accepting of [generative AI]. They're more like, ‘okay, we understand ...the importance of trying to get students to use it, make it like a part of the repertoire, and try to be a better or more productive individual with it.’ I believe SLIS has a different viewpoint than traditional academia on ChatGPT.” (Participant 5)

5.1.2 Other Acceptance Factors

Other factors that impacted acceptance of the tools, beyond UTAUT constructs, emerged from the interviews. For example, out of the eight participants, five reported that curiosity was a major factor. For some, it was a technology they “stumbled onto”

online. For others, it was more nuanced and driven by a need to understand how it works or to introduce this “new technology” into their daily lives. For example, Participant 7 said she used it because it “popped up on a website.” Participant 5 explained that while the media attention on generative AI was one of the initial motivations for using the tool, understanding the downsides and limits of the tool was also a key factor.

“The motivation for using the tool [was] just to see how it worked, and to see kind of to test the accuracy of the results... I like to be on the front end and try to figure things out and see if it can be useful to me down the road.” (Participant 8)

A fear of social isolation was also identified as a factor that could impact acceptance of LLM tools. Participant 7 explained that she felt a bit of anxiety considering how increased adoption of LLM tools could lead to increased social isolation. However, after completing the tasks with the LLM tool designed for this study, she noted that her fears were reduced:

“After using the tool, I can understand that it probably will not replace [social connection] ...it depends on what you may need to speak with the professor about, but for basic things, maybe assignment-based [questions] ...I think it is very useful. And back to the point of social interaction. [The tool] will eliminate some aspects of [social interaction], but I guess we are in a technological era.” (Participant 7)

5.2 Use of LLM-Powered Assistants

5.2.1 Type of Use

Across the study participants, generative AI had previously been used for personal, work, and educational activities. Multiple participants shared that they had used it to retrieve fundamental knowledge on a topic—to understand specific concepts, get a broad overview, or search for

secondary sources or articles on a topic. A couple of participants reported using generative AI for personal or work-related event planning, while some noted that they used it to create templates for a Word document or presentation or to check the language in a work email.

During the technical tasks, multiple participants reported using the LLM-powered assistant to clarify technical concepts required to answer the question, noting that using the tool to generate the answers was inappropriate or lazy. For example, Participant 1 shared that rather than asking the tool to generate the SQL query, he asked it “how to use the limits and threshold statements in SQL.” Relatedly, some participants used the tool sparingly - for 1 out of 4 questions. Some participants indicated that they copied and pasted the technical question into the tool to generate queries. One participant commented that this was done as a quality check after writing the query individually, while others used the query generated by the tool as a starting point for comparison with the options given in the task.

When asked if they would change the way they used the tool, given their score on the task, a few participants said they would use it more. One participant said they would use it more for other questions, while another said they would have asked more follow-up questions to clarify concepts.

5.2.2 Factors that Impact Use

Half of the participants reported having access to resources - including online platforms - to support their use of Gen AI tools. This falls under the Facilitating Conditions construct in the UTAUT model. It is worth noting that among participants who indicated that they had used Gen AI as part of a class, they stated that they received no support in using the tool. For example, Participant 1 explained that in one of his classes, while they were told they could use generative AI to ensure their submitted work was original, they were not told how to actually use the tool.

The themes of trust and the quality of responses - both actual and individual perception - recurred throughout the interviewees. After using the

tool to complete the task, most participants indicated a high level of trust in the tool. Some participants attributed their trust to an informal assessment of the accuracy of the response based on their prior knowledge. A few participants commented that they were concerned about the source of the information used to train the model and the risk of algorithmic bias. Other metrics that participants used to evaluate the quality of responses include the formatting and depth of the responses. Interestingly, participants who reported a significant lack of trust in generative AI before completing the technical task with the tool commented that they built trust based on validating some responses generated by the tool with prior knowledge.

“Some of the information I'm saying that is actually given to me is something I can resonate with, and when I can resonate with that automatically, I'm building trust with [the tool]” (Participant 6)

A couple of participants were particularly concerned with the ethical issues with using generative AI for learning. Particularly with regards to generating a complete solution or essay using the tool, privacy concerns, and how the information is curated, including copyright infringement concerns.

There were some comments about the user experience. A number of participants discussed the response speed, the formatting of the text, and the mobile-friendly system. The majority of participants expressed that the easily digestible information generated by the tool was a significant use factor. The overwhelming majority of participants viewed the tool's accessibility as important.

“If this is going to be used for class, I think it's perfect. Because we all have devices, and, thumbing through a book sometimes, or even, on the internet, you're just bombarded with so much information, it's like, 'where do I go?' So I think having something like that in your hands is going to be helpful.” (Participant 3)

5.3 Impact of LLM-Powered Assistants

Insights from the interviews suggest that LLM-powered assistants will change how students learn and are assessed. The use of generative AI in learning has raised questions about ensuring the ethical use of the tool. Particularly how to ensure students are not submitting plagiarized work. Several participants discussed this issue, and their experience could provide insights into how generative AI may result in a new way of learning and assessment. For example, from the study, we found evidence that some students will attempt to copy and paste the entire question into the tool. However, the participants who used the tool in this way reported that the SQL query generated did not match the choices presented in the task. As a result, they needed to compare and contrast the options in the task with the query generated by the tool. As Participant 4 stated:

“An answer is given [to you], it's elaborating and explaining to you ...but you cannot apply [it] exactly 100% to your quiz. ...it's giving you hints and explanations, but you have to learn and apply.” (Participant 4)

Generative AI is expected to impact all sectors, with information-related roles particularly likely to see significant impact. Surprisingly, only two participants acknowledged the expected global shift and increased reliance on these tools, particularly in the workplace. One of these participants noted that generative AI is helpful and would be fully integrated into the workplace within the next 10 years, particularly to improve worker productivity. Another participant echoed the impact of generative AI on productivity and argued that it helps speed up the brainstorming process.

6. Conclusion

In this study, we conducted semi-structured interviews and analyzed responses to identify some of the key factors affecting the use and acceptance of LLM-powered assistants amongst ethnic minority students.

We observed recurring themes that suggest

design, ethical considerations, and technology-aware institutions will be vital in diverse groups adopting LLM-powered tools and assistants. Participants recurrently reported individual curiosity and professional necessity as drivers for increased use of LLM-powered tools for learning. Participants overwhelmingly reported positive experiences using the LLM assistant for learning and educational outcomes but also expressed concerns that the tool could exacerbate social isolation and increased human dependence on technology.

6.1 Study Limitations

Some study limitations include selection and volunteer bias using a self-opt-in option from the graduate and alumni population. This may skew our participants to be overrepresented by individuals with some interest in Generative AI. We also made use of a text-only Generative AI model. Since the study was developed and administered, OpenAI, Inc. and Anthropic PBC have released multi-modal models capable of receiving and generating text and images. Such models may be more capable of responding to the tasks in the study.

6.2 Next Steps

Our next work will focus on analyzing the revealed preferences, use, and impact data from the LLM-powered tool from the scores on the database administration and analysis tasks. Analyzing acceptance factors across other populations will also be a promising area for further work.

Declaration of Conflict

Oruaro Ogbo is the sole director of the private company that created the Generative AI tool used in this study.

Acknowledgments

The authors wish to thank the participants in this study for sharing their experiences and perspectives.

References

- Alberth, 2023. The Use of ChatGPT in Academic Writing: A Blessing or a Curse in Disguise? *TEFLIN J.* 34, 337–352.
<https://doi.org/10.15639/teflinjournal.v34i2/337-352>
- Chan, C.K.Y., Hu, W., 2023. Students' voices on generative AI: perceptions, benefits, and challenges in higher education. *Int. J. Educ. Technol. High. Educ.* 20, 43.
<https://doi.org/10.1186/s41239-023-00411-8>
- Chiu, T.K.F., 2024. Future research recommendations for transforming higher education with generative AI. *Comput. Educ. Artif. Intell.* 6, 100197.
<https://doi.org/10.1016/j.caeai.2023.100197>
- Choi, J.H., Monahan, A., Schwarcz, D.B., 2023. Lawyering in the Age of Artificial Intelligence. *SSRN Electron. J.*
<https://doi.org/10.2139/ssrn.4626276>
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., Zimmel, R., 2023. The Economic Potential of Generative AI: The Next Productivity Frontier.
- Crompton, H., Burke, D., 2023. Artificial intelligence in higher education: the state of the field. *Int. J. Educ. Technol. High. Educ.* 20, 22.
<https://doi.org/10.1186/s41239-023-00392-8>
- Dijk, J. van, 2020. *The Digital Divide*. John Wiley & Sons.
- Ebizie, P., Nkamnebe, A., Christian, O., 2022. Factors Influencing Cryptocurrency Adoption Among Nigerian University Fintech Entrepreneurs: An UTAUT perspective. *Br. J. Mark. Stud.* 10, 25–37.
<https://doi.org/10.37745/bjms.2013/vo10.n3pp2537>
- Gabriel, I., Manzini, A., Keeling, G., Hendricks, L.A., Rieser, V., Iqbal, H., Tomašev, N., Ktena, I., Kenton, Z., Rodriguez, M., El-Sayed, S., Brown, S., Akbulut, C., Trask, A., Hughes, E., Bergman, A.S., Shelby, R., Marchal, N., Griffin, C., Mateos-Garcia, J., Weidinger, L., Street, W., Lange, B., Ingerman, A., Lentz, A., Enger, R., Barakat, A., Krakovna, V., Siy, J.O., Kurth-Nelson, Z., McCroskery, A., Bolina, V., Law, H., Shanahan, M., Alberts, L., Balle, B., de Haas, S., Ibitoye, Y., Dafoe, A., Goldberg, B., Krier, S., Reese, A., Witherspoon, S., Hawkins, W., Rauh, M., Wallace, D., Franklin, M., Goldstein, J.A., Lehman, J., Klenk, M., Vallor, S., Biles, C., Morris, M.R., King, H., Arcas, B.A. y, Isaac, W., Manyika, J., 2024. The Ethics of Advanced AI Assistants.
<https://doi.org/10.48550/arXiv.2404.16244>
- Hmoud, M., Swait, H., Hamad, N., Karram, O., Daher, W., 2024. Higher Education Students' Task Motivation in the Generative Artificial Intelligence Context: The Case of ChatGPT. *Information* 15, 33.
<https://doi.org/10.3390/info15010033>
- Horrigan, J.B., 2010. Broadband adoption and use in America. Federal Communications Commission Washington, DC.
- Horrigan, J.B., Duggan, M., 2015. Home broadband 2015. *Pew Res. Cent.* 21.
- Ishak, M.F., Wan Hussain, W.M.H., Ayob, A.H., Buja, A.G., Ahmad, R., 2022. The Used of UTAUT Model in Initial Coin Offerings: A Conceptual Paper. *Int. J. Acad. Res. Bus. Soc. Sci.* 12, Pages 748-760.
<https://doi.org/10.6007/IJARBS/v12-i7/13200>
- Jermakowicz, E., 2023. The Coming Transformative Impact of Large Language Models and Artificial Intelligence on Global Business and Education. *J. Glob. Aware.* 4, 1–22.
<https://doi.org/10.24073/jga/4/02/03>
- Johnston, H., Wells, R.F., Shanks, E.M., Boey, T., Parsons, B.N., 2024. Student perspectives on the use of generative artificial intelligence technologies in higher education. *Int. J. Educ. Integr.* 20, 1–21.
<https://doi.org/10.1007/s40979-024-00149-4>
- Kang, H., Ni, J., Yao, H., 2024. Ever: Mitigating Hallucination in Large Language Models through Real-Time Verification and Rectification.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., Kiela,

- D., 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.
- Li, J., Yuan, Y., Zhang, Z., 2024. Enhancing LLM Factual Accuracy with RAG to Counter Hallucinations: A Case Study on Domain-Specific Queries in Private Knowledge-Bases.
- Li, X., Li, B., Cho, S.-J., 2023. Empowering Chinese Language Learners from Low-Income Families to Improve Their Chinese Writing with ChatGPT's Assistance Afterschool. *Languages* 8, 238.
<https://doi.org/10.3390/languages8040238>
- Magsamen-Conrad, K., Wang, F., Tetteh, D., Lee, Y.-I., 2020. Using Technology Adoption Theory and a Lifespan Approach to Develop a Theoretical Framework for eHealth Literacy: Extending UTAUT. *Health Commun.* 35, 1435–1446.
<https://doi.org/10.1080/10410236.2019.1641395>
- NDIA, n.d. Definitions [WWW Document]. Natl. Digit. Incl. Alliance. URL <https://www.digitalinclusion.org/definitions/> (accessed 5.4.24).
- Ogbo, E., Brown, T., Gant, J., Sicker, D., 2021. When Being Connected is not Enough: An Analysis of the Second and Third Levels of the Digital Divide in a Developing Country. *J. Inf. Policy* 11, 104–146.
<https://doi.org/10.5325/jinfopoli.11.2021.0104>
- Pellas, N., 2023. The Effects of Generative AI Platforms on Undergraduates' Narrative Intelligence and Writing Self-Efficacy. *Educ. Sci.* 13, 1155.
<https://doi.org/10.3390/educsci13111155>
- Sabri, M., Baba, N., Sulaiman, W., 2023. Investigating Hospitality Student's Acceptance in Online Learning Platform: Utilising UTAUT Model. *Int. J. Acad. Res. Bus. Soc. Sci.* 13.
<https://doi.org/10.6007/IJARBSS/v13-i5/17047>
- Sheese, B., Liffiton, M., Savelka, J., Denny, P., 2024. Patterns of Student Help-Seeking When Using a Large Language Model-Powered Programming Assistant, in: *Proceedings of the 26th Australasian Computing Education Conference*. pp. 49–57.
<https://doi.org/10.1145/3636243.3636249>
- Shuster, K., Poff, S., Chen, M., Kiela, D., Weston, J., 2021. Retrieval Augmentation Reduces Hallucination in Conversation.
- Walczak, K., Cellary, W., 2023. Challenges for higher education in the era of widespread access to Generative AI. *Economics and Business Review* 9.
- Wang, T., Lund, B.D., Marengo, A., Pagano, A., Mannuru, N.R., Teel, Z.A., Pange, J., 2023. Exploring the Potential Impact of Artificial Intelligence (AI) on International Students in Higher Education: Generative AI, Chatbots, Analytics, and International Student Success. *Appl. Sci.* 13, 6716.
<https://doi.org/10.3390/app13116716>
- Xiao, P., Chen, Y., Bao, W., 2023. Waiting, Banning, and Embracing: An Empirical Analysis of Adapting Policies for Generative AI in Higher Education. *SSRN Electron. J.*
<https://doi.org/10.2139/ssrn.4458269>