

# Who Benefits and Who is Excluded? Transformative Learning, Equity, and Generative Artificial Intelligence

TRACIE ADDY Lafayette College

TINGTING KANG Lafayette College

TIM LAQUINTANO Lafayette College

VIVIENNE DIETRICH Lafayette College

#### Abstract

In our essay, we discuss equity implications surrounding the usage of generative artificial intelligence (AI) in higher education. Specifically, we explore how the use of such technologies by students in higher education such as, but not limited to, multi-language learners, students from marginalized linguistic communities, students with disabilities, and low-income students has the potential to facilitate transformative learning. We describe how such tools, when accessible to learners, can help address barriers that prevent students from fully engaging in their learning. Additionally, we explain how the usage of generative AI has the potential to alter the lens through which students view their learning, countering assumptions and broadening what can be considered an "appropriate" use of assistive technologies to support learning for diverse students. We also address various limitations of generative AI with regards to equity such as the requirement to pay to access some of the applications, as well as linguistic and other biases within the outputs produced, reflective of the data used to train the tools. Throughout this piece, we share insights from a study of undergraduate students' perspectives and usage of generative AI and potential future directions for the technologies. This essay aims to increase awareness of the opportunities and challenges around who benefits and who is excluded when generative AI is used within colleges and universities.

#### Introduction

The mission of many colleges and universities is to educate a diverse student body. Essential to accomplishing these goals is creating transformative learning experiences that are both equitable for and inclusive of all students. The increased availability of generative artificial intelligence (AI) directly implicates such equity work. The emergence of more sophisticated tools has the potential to create greater disparities between those who can access them and those who cannot, as well as reinforce societal biases that exist which can lead to marginalized groups being increasingly targeted. Such barriers can impede progress toward a more just and equitable academy.

The widespread integration of generative AI within higher education settings is a paradigm shift that continues to lead to transformative learning, especially at the institutional, instructor, and student

levels (Mezirow, 1994). Institutions are forced to reflect on their values and existing practices to determine how to create significant, equitable learning experiences for their students, while simultaneously defending the value of higher education amidst uncertainty and change. Instructors are rethinking how to best teach students in ways that achieve desired learning outcomes. Educators must now decide how to create equitable generative AI course policies and how to design assessments to support learning. Their assumptions about products produced by a learner (e.g., an essay, etc.) as the sole output that demonstrates learning and skill development are challenged when students can essentially use large language models (LLMs) to create entire assignments. The necessity to apply new criteria for what constitutes academic dishonesty is at the forefront.

Students are challenged by knowing when and how to use these technologies. They also face issues of access and how to navigate usage of tools and their own learning and skill development. Students from diverse backgrounds such as multi-language learners and students with disabilities may encounter experiences interacting with LLMs that impact them in ways that are inequitable.

Generally, generative AI has the potential to transform learning, described by Mezirow (2008) as

The process by which we transform problematic frames of reference (mindsets, habits of mind, meaning perspectives) – sets of assumption and expectation – to make them more inclusive, discriminating, open, reflective and emotionally able to change. Such frames are better because they are more likely to generate beliefs and opinions that will prove more true or justified to guide action (p. 92).

In this essay, we argue that the influence of generative AI in higher education necessitates frame of reference shifts on part of both instructors and students with regard to equitable learning. This aligns with the learning paradigm where "We now see that our mission is not instruction but rather that of producing learning with every student by whatever means work best" (Barr & Tagg, 1995, p. 13). We posit that such means can involve generative AI to enhance equity, but that there are cautions. Subsequent sections explore the influences of generative AI on language learning, minoritized discourses, neurodivergence and disability, in addition to challenges such as bias and access. Throughout this essay we explore who benefits and who is excluded through generative AI.

#### **Students' Perspectives Study**

In this piece we interweave the findings of an undergraduate student survey study conducted with the research goals of better understanding if and how learners utilize generative AI technologies and their viewpoints on the advantages and disadvantages. This investigation, approved by the Institutional Review Board, involved surveying a sample of 1,000 students from a private liberal arts college, with 250 students randomly selected from each of four class years. Participants were emailed via an anonymous survey using Qualtrics software and sent two reminder messages if they had not yet completed the study. Students were given the option of entering a raffle for prizes if they completed the survey. Based on the estimate in Qualtrics software, the total time for completion of the survey was approximately 7 minutes.

One hundred fifty-four students completed the survey in its entirety. Including those who either fully or partially completed the survey, the response rate was 18%. Of respondents who completed demographic questions, 93% indicated that English was their primary language, 22% described themselves as neurodivergent or a student with a specific learning disability, 19% indicated that they were the first in their family to attend college in the United States, 16% were Pell-eligible, and 53% were receiving aid in other financial amounts (n = 152 responses). We include quotes from study respondents in this essay to support our claims and to provide contextual information regarding how generative AI can lead to transformative learning.

## How Generative AI Can Transform Language Learning Transformative Learning for Multi-Language Learners

In the field of second language learning, educators have been seeking innovative ways to enhance students' learning experiences. By combining cutting-edge technologies with existing pedagogical expertise and educational theory, generative AI tools have the potential to offer unprecedented opportunities for second language learning. However, various equity issues may arise due to language learners' unique language backgrounds, such as language bias and constrained linguistic expressions. In this section, we explore how generative AI tools can transform or hinder the learning of multi-language students and provide recommendations to instructors by applying Mezirow's (1994) Theory of Transformative Learning.

The advent of digital tools and resources has revolutionized language education, providing educators with innovative methods to enhance the learning experiences of multi-language learners. For example, computer-assisted language learning (CALL), mobile-assisted language learning (MALL), and the integration of corpora (a collection of naturally occurring languages) using computers have emerged as some of the prominent areas of study in the field of language education. These approaches offer learners access to authentic language resources and interactive activities, ultimately promoting autonomous learning and individualized progress over time (e.g., Deignan & Potter, 2004; Liu & Yu, 2022; Roussel, 2011).

Recent advancements in generative AI tools have further amplified these benefits owing to their exceptional natural language processing capabilities and vast knowledge base. Many language educators started to implement AI tools in their teaching (e.g., *Making English Fun*). Table 1 summarizes how generative AI tools like ChatGPT can be applied to almost all aspects of English language learning and teaching including grammar, vocabulary, writing, reading, pronunciation, conversation, culture, and translation.

Areas	Learning Activities
Grammar	Correct English sentences or provide explanations on English grammar rules.
Vocabulary	Define or provide examples of English words and phrases; generate sentences or paragraphs that include specific vocabulary words; distinguish academic and non- academic words.
Writing	Provide writing prompts; give feedback on essays; provide suggestions or corrections.
Reading	Generate English text on various topics, such as news articles, stories, or essays; create comprehension questions
Pronunciation	Pronounce words or sentences; provide feedback on pronunciation; dictate speech
Conversation	Real-life interactions; initiate conversations with ChatGPT by asking questions, discussing topics, or engaging in role-play exercises

Table 1: Sample 1	Language	Learning Acti	vities Using ChatGP	Γ
т	<b>T</b> •	A 10 010		

Language Areas	Learning Activities
Culture	Ask about English-speaking countries' culture, customs, and traditions; learn about English idioms, expressions, and social norms to gain a deeper understanding of English language and culture
Translation	Translate words, sentences, or paragraphs into different languages.

Students in our study also seemingly saw the advantages of these tools in supporting their learning. Nine out of ten students who reported that their primary language was not English indicated that they used generative AI tools. A few respondents noted the utility of these technologies in supporting language learning as indicated below:

It [generative AI] can also be very helpful for students who haven't had as many English classes or [for] whom English isn't their first language.

AI can definitely be used as an asset, but it should not be submitted for anything word for word; rather, it's better for getting ideas. I also used it to help me understand Spanish word choice since it's often hard for the average speaker to explain why they use one word/phrasing as opposed to another to convey specific points.

The comments of these respondents highlight a broadened understanding of how languages can be learned by leveraging generative AI.

However, the use of generative AI tools by multi-language learners also raises several equity concerns. Many of these tools have been developed primarily by English-speaking countries, resulting in potential accuracy issues when processing text in languages with limited training data. Secondly, students who are not proficient in the language used by generative AI tools may encounter difficulties in comprehending and accurately evaluating the generated text. Lastly, multi-language learners might be targeted more for academic dishonesty. A study highlighted the presence of bias in AI detectors against multi-language learners with limited linguistic expressions (Liang et al., 2023). In this case, the AI detectors were more prone to categorizing writings from multi-language learners as AI-generated, while accurately identifying native writing samples. Such are barriers to equitable, transformative learning.

As Egan (1999) wrote, "The best of technology does not by itself create a productive learning environment" (p. 281). Mezirow's (1994) transformative learning theory is defined as "an orientation which holds that the way learners interpret and reinterpret their sense experience is central to making meaning and hence learning" (p. 222). This learning theory emphasizes that through self-awareness, self-directed learning, and critical theory, students can grasp the underlying meaning structures and engage in a process of reevaluating their assumptions. At the same time, self-awareness strategy and self-directed learning have also proven to be effective in second language learning (e.g., Vohs & Baumeister, 2004; Roussel, 2011). Therefore, when designing language learning activities, it is crucial to invite students to critically reflect on the content, process, and premise of their interaction with AI tools as recommended by Mezirow's transformative learning theory. The following are some questions that instructors can use to ask students to reflect on their language learning experiences using AI tools:

• **Content:** What do you like or dislike about the output produced by this AI tool? What specific language skills or knowledge are you gaining from this AI tool? Can you identify any cultural biases present in the AI-generated content? Can you guess how the AI output content was generated?

- **Process:** What strategies have you used to get the expected output from the AI tool? Are there any difficulties or challenges you have encountered in using the AI tool, and how have you addressed them?
- **Premise:** What assumptions or expectations did you have about using AI tools for language learning, and have they been met or challenged? How do you think future advancements in AI technology might change the way we integrate AI into language learning?

## Transformative Learning for Speakers from Minoritized Discourses

In addition to multi-language learners, the usage of generative AI can impact speakers from minoritized discourses. For decades, those interested in the politics of language in higher education have been grappling with a core problem: Standard Edited English (SEE) is ostensibly the currency of the writing economy of American higher education and the professional world. SEE enables groups to communicate across space and time in a way that facilitates economic exchange, legal systems, and public safety. It has also facilitated global empires, and it has been deployed in the service of linguistic imperialism, linguistic discrimination, and linguistic racism (Baker-Bell, 2020; Canagarajah, 2006; Smitherman, 1977). It is what James Paul Gee (1989) would call a "dominant discourse," and its American iteration has largely emerged from the language practices of the upper-middle and wealthy classes. As Gee argues, the extent to which one's home discourse aligns with the dominant discourse of higher education is the extent to which a student has linguistic privilege. This situation has created a conundrum for students whose linguistic practices and home discourses do not align well with the academy (and a conundrum for the instructors who teach them). They can learn (and sometimes struggle to learn) the standard conventions of written English, a discourse of power, prestige, and cultural capital, or they can work to code mesh and integrate their home discourse(s) into academic debate, which allows them to draw on a new reservoir of meaning-making practices that have largely been excluded from SEE (for an example of this see Young, 2010).

With their ability to transform even the most garbled of sentences into SEE, language models offer a solution to this issue, but that solution is unsatisfactory. LLMs have numerous faults as text generators (early models create somewhat generic prose and the models operate with no ground truth), but they excel at re-writing and copyediting prose. Those concerned with the politics of language in higher education now face a core question similar to that facing multilingual writers: Do students who come to higher education with fluency in minoritized discourses deploy language models to help them write in SEE and avoid the pernicious effects of linguistic racism? Or do we collectively continue to make the case that SEE is only one kind of English and a kind that tends to limit the full written expression of many kinds of people, which might allow for the adoption of a fuller spectrum of language practices in the academy? And, if we prohibit students from using LLMs for linguistics tasks such as copyediting, to what extent are we foreclosing on the opportunity for groups of students to avoid the penalties of linguistic racism that have long been a part of many assessment systems?

The benefit to LLMs is that we now have what promises to be a decent technological solution for writers who struggle to shape their writing into the kind of SEE that is acceptable in the academy. Linguistic discrimination is deeply enmeshed in systems of white supremacy and our systems of writing assessment, for that matter (Inoue, 2015). From a pragmatic perspective, writers who would have otherwise faced judgment for writing with non-standard language practices have access to a tool that will provide high-quality and effective copy editing that will potentially enable them to circumvent (at least partially) problematic assessment systems.

The benefit, however, is also the challenge. When we talk about having students rewrite their sentences and paragraphs using a language model, we're effectively talking about a new form of linguistic colonization. Bender et al. (2012) have argued in a now-famous paper that the worldviews of many LLMs reflect the voices most strongly represented in data on which the language model was trained; or, as they put it: "the voices of people most likely to hew to a hegemonic viewpoint are also more likely to be retained" (613). We can add to this that the linguistic dialects and linguistic constructions of those most well-represented in the data are more likely to be retained and output by the LLM. Asking LLMs to

copyedit a student text equates to taking voice from a marginalized discourse, which should be represented in academic discourse, and homogenizing it into a dominant discourse via automation. In other words, clicking a button to use a machine to change one's writing to SEE could be seen as capitulating to a racist linguistic system and perpetuating a relatively limited linguistic domain.

In teaching AI literacy for the purposes of transformative learning, it is important to be keenly aware of the language practices that LLMs facilitate. Linguistic imperialism is not limited to the output of LLMs. It also includes the input, or the data on which a language model has been trained. We are entering an era of "model proliferation," where we will see most major technology companies release their own language models along with a variety of open-source versions. All these models will be trained on different datasets. The data used for training has become a closely guarded trade secret, and for many language models, academics and users do not have access to training data, which renders it impossible to assess its bias and the extent to which certain kinds of voices dominate the training data. Consequently, although many models can produce SEE, the models struggle when asked to produce convincing versions of other kinds of dialects and world Englishes (see Owusu-Ansah, 2023). Antonio Byrd (2023) has argued that the refusal of companies to share their training data means that such LLMs may not be ethical tools to work with as writers. He suggests that instructors adopt open-source tools that allow users to inspect the data used to train the models.

When it comes to the output of LLMs, which tends to happen in SEE unless the machines are prompted otherwise, students need to acquire a critical awareness of how SEE has been developed and its history of being deployed via linguistic violence against marginalized populations. There is a voluminous literature on critical language awareness, a linguistic movement that has advocated the position that instructors need to help students learn about dominant and marginalized discourses, although April Baker-Bell's (2021) research makes clear that many students who speak marginalized discourses are already well aware of the status of their discourse. Instructors who work with LLMs and AI literacy need to make clear to students that its outputs do not represent some sort of neutral dialect but rather the dialects of those who have had the opportunity to contribute the most data to the training set. Therefore, instructors can facilitate transformative learning experiences for students by having them critically examine the outputs produced by generative AI when marginalized discourses are used and have them explore and consider issues surrounding training the tools with SEE.

#### **Opportunities for Neurodivergent Students & Students with Disabilities**

Generative AI tools have the potential to support the learning of neurodivergent students and students with disabilities, but also have their limitations (McMurtrie, 2023). Neurodiversity is "a biological truism that refers to the limitless variability of human nervous systems on the planet, in which no two can ever be exactly alike due to the influence of environmental factors" (Singer, n.d.). The term describes the many ways that individuals encode and process information. Neurodiversity brings to the forefront how people with conditions such as ADHD, dyslexia, autism spectrum disorder, and bipolar disorder, are not different but rather exist within the continuum of all humans that vary in how we process information. Such challenges the concept of normality. Individuals with disabilities, whether they be mental, physical, sensory, or learning, exist within the diversity of people in our world, and within college and university courses. When considering how to create more equitable, transformative learning environments for such learners, generative AI models present distinct possibilities.

### Transformative Learning for Neurodivergent Students and Learners with Disabilities

Neurodivergent students and learners with disabilities might already use assistive technologies that reduce barriers that they confront to their learning. These might involve obtaining or creating recordings of class sessions and converting text-to-speech, using a screen reader, obtaining notes, or having a notetaker. LLMs can expand students' approaches for learning in a variety of ways such as through organizing course notes in ways that they may never have been able to do previously to allow them to better grasp the material, querying their notes to increase their understanding, and generating initial ideas and topics when they need added support to guide them in the next steps of their work. Such

can lead to transformative learning for these students and their instructors as they work through the disorienting dilemma of how to use such tools to support equitable learning (Mezirow, 1994).

Issues might arise if instructors prohibit their students from using LLMs in their courses. Such restrictions may pose a barrier for neurodivergent students and students with disabilities who might benefit from the tools to support their learning. Additionally, restricted access due to the necessity to pay for various AI tools can also limit their availability to such learners; a tool might exist that can serve the needs of students but be cost prohibitive. Further, students might be stigmatized for their usage of these tools or increasingly targeted for academic dishonesty if they use them regularly.

Instructors can take a variety of steps to support neurodivergent students and learners with disabilities in their usage of generative AI models for learning. A preliminary step is to learn which assistive technologies students use in their courses and how they support learning, to gain a general awareness. Additional individual conversations with students who have accommodations can focus on how and whether generative artificial intelligence tools can support learning. Many colleges and universities also have accessibility offices. Their staff can be valuable resources for instructors deliberating upon how to incorporate generative AI for students with accommodations. Lastly, when instructors adopt policies that are equity-minded around generative AI, they can lower barriers to support diverse students in achieving learning outcomes. While much of this essay focuses on the transformative learning of students, considering generative AI an assistive technology for neurodivergent learners presents a shifted viewpoint by instructors.

## **Cautions of Biased Output**

LLMs, despite their promise, have additional limitations that can hinder transformative learning. The outputs of LLMs reflect the datasets from which they draw information. Therefore, they are subject to reproducing and reinforcing bias. Biased output can be psychologically damaging for members of marginalized groups who are more likely to experience threats to their identities in their everyday experiences. Such output biases may go undetected, with individuals not recognizing their presence, leading to reinforcement and reproduction. Students benefit from learning how to critically engage with the outputs produced to verify their accuracy, and whether they take a singular stance and fail to acknowledge alternative viewpoints when present, or seemingly incorporate extreme views. As learners develop critical analysis skills, they may challenge any assumptions they hold of outputs generated by AI as being free of bias and learn how to use the technology in ways that benefit them, skills they can take with them post-graduation.

#### The Thorny Issue of Student Access

Another concern with LLMs is student access which can pose a barrier to the types of transformative learning that we described in this essay. The pandemic provided a reminder that we still live with a deep digital divide. Although access to the internet has grown extensively in the last two decades with smartphones, remote schooling conditions called attention to the discrepancy of connectivity and devices between the digital haves and have nots. Remote working conditions threw into sharp relief the number of students who still lack access to broadband let alone devices and peripherals and home environments that would enable them to work productively and with ergonomic integrity on screens (Auxier & Anderson, 2020; Correia, 2020; Francis & Weller, 2022). We will likely soon see a slew of educational applications and learning technologies powered by large language models. Just a few months after the release of ChatGPT, Khan Academy began a pilot to integrate GPT-4 into its learning platform (Kahn, 2023). Although Kahn Academy announced its program with the aim of providing "equal access" to AI learning tools, current generations of LLMs tend to take enormous amounts of computational resources to run, and it remains to be seen whether chatbot tutors, for example, threaten to exacerbate inequalities or help remedy them. There are several unknown variables when it comes to language models like ChatGPT that will heavily influence access to them. Some of those are socio-political (as discussed), but here we'd like to discuss important economic variables that include the computational resources that

the models need to operate and the systems of monetization that have yet to be worked out that will provide access to them.

Many of the most popular computing applications of the last decade ostensibly appeared to be free for the user. If someone secured a smartphone and data connection, they could access productivity applications such as email, cloud-based word processing software, and a suite of social media applications. However, these "free" programs have required vast digital infrastructures and server farms to run, and we paid for this infrastructure—and produced profits for the companies—through the data we created by using them. We surrendered our privacy. Social media applications were not the products we were using. Our data profiles were the products, and social media companies sold them to advertisers, a process Shoshanna Zuboff (2019) has documented at length in *The Age of Surveillance Capitalism*. Social media companies thus tried to maximize their applications for engagement, as the more eyeball time on their app, the more data they received about our preferences, and the more valuable the applications tended to be for advertisers.

The first LLMs released to consumers do not function on a similar economy, however. The companies that run them can collect data on users, including the inputs used to prompt the models, but as of yet, the limited hints we have from large technology companies suggest that data exhaust does not appear to be a particularly efficient way for companies to profit from them. For example, Sam Altan, CEO of OpenAI as of this writing, stated in a 2023 congressional hearing that OpenAI has not been designing its chatbots to maximize user engagement in the way that social media companies do, largely because OpenAI does not have the computational infrastructure to support heavy use. This could certainly change. Snapchat, for example, has integrated a chatbot for users to interact with, and with other companies looking to infuse chatbots into our most intimate spheres, it may eventually be the case that users divulge interests and dimensions of their personality that they would not have otherwise through social media usage. It is possible we subsidize our own chatbot use by surrendering more of our privacy.

For the time being, however, language models remain expensive to develop and run, which could limit use as well as the development of open-source models (Heaven, 2023). And the most popular emerging model does not run by selling consumer data but rather by subscription. As of this writing, users who register for an OpenAI account can use the free model of ChatGPT (based on a GPT-3.5 family of models), or they can pay a monthly fee for a more powerful version (based on the GPT-4 model). Google and Microsoft have begun releasing models that are free to use (through Bing chat and through the Google Cloud productivity suite), but they have deep pocketbooks to do so, and long-term free access to robust models is not guaranteed unless engineers continue to find ways to reduce the cost of the inferences the machines make. And so rather than using social media as an analogy for access, which is all about "free use" via the exploitation of user data, we might think about gaming as an analogy for how access to language models might play out. Currently, a smartphone and WiFi connection will allow one to play any number of free mobile games that run on advertisements. However, access to the most powerful games with the best graphics still requires subscription costs, dedicated hardware (gaming consoles or powerful computers with graphics cards), and a good deal of bandwidth. This has produced a sharper divide in access to gaming than in social media.

AI engineers are hard at work trying to improve the performance of models so that they use fewer computational resources. They have been fine-tuning LLMs and using reinforcement learning from human feedback to sharpen LLM functioning, but the amount of processing power needed to run most language models remains high compared to many consumer-facing applications. If engineers cannot produce powerful models that use fewer resources, we may see a situation where students who can afford to pay have access to the most powerful models and students who cannot pay have access to stripped-down slower models with degraded performance. When it comes to education, we need to account for the language models that will be fine-tuned for specific domains and specific educational purposes. We can use a generic language model like ChatGPT for free as an all-purpose tutor (as long as we are careful about its issues with hallucinations). But we will likely soon see specialized models (likely at a cost) trained to help students with specific disciplines and competencies that perform much better in that domain than a generic all-purpose model.

Access to LLMs is also an international concern. According to Martindale (2023), ChatGPT has been banned in several countries, including Russia, China, North Korea, Cuba, Iran, and Italy. China's ban, in particular, stems from concerns related to its impact on values, mainstream views, and national security issues (Fan, 2023). Italy initially banned ChatGPT due to concerns regarding the collection and storage of personal data and the potential exposure of misinformation to young people. However, the ban was later lifted after OpenAI released a new form that allows European Union (EU) users to request the removal of personal data and developed a new tool to verify users' ages during signup (Robertson, 2023). However, the story does not end there. Currently, Canada and some EU countries have opened or are considering opening investigations into ChatGPT's practices in their respective countries. Furthermore, educational institutions such as New York City Public Schools, the Los Angeles Unified School District, Sciences Po in Paris, France, and RV University in Bengaluru, India, have also banned Chat GPT. The rationale behind these bans revolves around the belief that such AI tools lack the ability to foster critical thinking and problem-solving skills, which are considered crucial for academic and lifelong success (Castillo, 2023).

#### Transformative Learning for Students from Low Income Backgrounds

Currently, it is incredibly difficult to enumerate the benefits and challenges of student access to generative AI tools. We are in the midst of an AI arms race with billions of dollars being invested and exorbitant claims being made about the potential for AI to revolutionize everything from the economy to higher education. The landscape of AI has many unresolved factors that could shape the trajectory and general usefulness of LLMs for students, including security and privacy, legality, the development of plugins and complementary technologies, and model accuracy (Laquintano et al., 2023).

There is already evidence that students find language models useful for completing their coursework. In our study, students who identified as Pell-eligible, meaning that they qualified for Pell grants due to their expected family contribution, described a variety of benefits:

Some of us come to college not knowing how to write academic papers, and chat GPT taught me how. Now I can write an academic paper on my own. It's a good guide to teach certain writing skills!

It provides a quick response at any time of the day. If I am studying at 6pm at night, my professor is not having office hours then, so if I am very confused, I will sometimes ask AI for help. It is usually a last resort if I am in a time crunch and cannot find the information in the textbook.

The AI gives more in depth information on how to answer the question. Some of the processes I haven't been taught in class how to use and so the AI acts as another source for my learning.

[L]imits barriers to learning/education with same availability of free resources

*AI* makes learning a lot easier. It breaks down complex text into easily understandable sentences. It also helps you plan for the future and activities.

Whether or not students' usage enhances or disrupts their learning is an open question. In the best-case scenario, we could potentially see the emergence of accessible and highly personalized AI systems that can assess, guide, and provide feedback on student learning with highly qualified human instructors remaining in the loop. In other scenarios, the tools might remain mediocre and short circuit student learning by acting as personalized cheating assistants that grease the path for those who can afford access to higher education.

Higher education will likely face significant challenges over the next few years as it helps students learn to use AI to augment their learning and not displace it. The generative AI gold rush will

likely soon produce a marketing explosion of tools created specifically for students in higher education, and institutions will be in the position of having to assess the quality of those tools and determine what kinds of learning can be offloaded to machines and what kinds of learning require human feedback. Institutions will have to think carefully and in discipline-specific ways about access to these tools and how and why to invest institutional resources in new educational applications. On the one hand, we need to be cautious and deliberate about the efficacy of the tools in which we invest. On the other hand, if we make no institutional investments, we will potentially create a situation where students with higher socioeconomic status (SES) can subscribe to more powerful tools while students with lower SES do what they can with less powerful "free" applications.

Student access is a primary consideration as we create institutional policies related to generative AI. We've already seen some schools move to prohibit access to LLMs by blocking them on computer networks. Several later reversed their policies. Such policies can prove untenable over time, but they do have the immediate effect of granting access to LLMs only to students who can secure access in other ways with their own devices and data connections. In some respects, we've already been through a version of this problem. With the explosion of student textbook costs, we have already seen in the last decades inequitable access to learning resources, as low-income students struggle through classes by renting textbooks or by borrowing them from the library in ways that are constrained by space and time. High income students can buy textbooks, create marginalia, and have complete access to them all semester. Institutions must be certain that whatever policies put in place related to generative AI do not replicate the same mistakes we have made when it comes to other student learning resources.

#### **Concluding Remarks**

As an educational developer who oversees a center for teaching and learning and studies inclusive pedagogy; a director of an academic learning center for multi-language learners; and an English professor and scholar who studies generative AI and directs a college writing program, we are invested in supporting a diversity of learners as generative AI evolves. The widespread integration of AI presents many opportunities for transformative learning. As we argue in this essay, these tools hold much promise in fostering equitable learning, but they also have their potential pitfalls. Higher education will learn much in the coming years as instructors and institutions at large continue to critically examine their teaching practices to support a diversity of learners.

#### Acknowledgments

We shared several of the ideas within this article during a summer professional development series on generative AI in higher education. We are deeply appreciative of the Lehigh Valley Association of Independent Colleges for sponsoring such conversations.

#### References

- Auxier, B. and Anderson, M. (2020) As Schools Close Due to the Coronavirus, Some U.S. Students Face a Digital 'Homework Gap.' Pew Research Center. https://www.pewresearch.org/facttank/2020/03/16/as-schools-close-due-to-the-coronavisus-some-u-s-students-face-a-digitalhomework-gap/
- Baker-Bell, A. (2020). *Linguistic justice: Black language, literacy, identity, and pedagogy*. Routledge. https://doi.org/10.4324/9781315147383
- Barr, R. B. and Tagg, J. (1995). From Teaching to Learning: A New Paradigm for Undergraduate Education. *Higher Education*, 60. <u>https://digitalcommons.unomaha.edu/slcehighered/60</u>

- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623). https://doi.org/10.1145/3442188.3445922
- Broussard, M. (2023). *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*. MIT Press. https://doi.org/10.7551/mitpress/14234.001.0001
- Byrd, A. (2023). Truth-Telling: Critical inquiries and the corpus texts that train them. *Composition studies*, *51.1*.
- Canagarajah, A. S. (2006). The place of world Englishes in composition: Pluralization continued. *College* composition and communication, 57.4, 586-619.
- Castillo, E. (2023). These schools and colleges have banned Chat GPT and similar AI tools. *Best Colleges*. <u>https://www.bestcolleges.com/news/schools-colleges-banned-chat-gpt-similar-aitools/#:~:text=The%20most%20notable%20K%2D12,on%20content%20safety%20and%20accur acy.</u>
- Correia, A. P. (2020). Healing the digital divide during the COVID-19 pandemic. *Quarterly Review of Distance Education*, 21(1).
- Crawford, K. (2021). *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence.* Yale University Press.
- Deignan, A., & Potter, L. (2004). A corpus study of metaphors and metonyms in English and Italian. *Journal of Pragmatics, 36*, 1231-1252. <u>https://doi.org/10.1016/j.pragma.2003.10.010</u>
- Egan, K. B. (1999). Speaking: A critical skill and a challenge. Calico Journal, 16, 277-293.
- Fan, S (2023). OpenAI回应 | ChatGPT为何不向所有中国用户开放注册? [OpenAI Response |Why is ChatGPT not available for registration to all Chinese users?] *Sina*. https://finance.sina.cn/chanjing/gsxw/2023-02-17/detail-imyfywyt6841773.d.html
- Francis, D. V., & Weller, C. E. (2022). Economic inequality, the digital divide, and remote learning during COVID-19. *The Review of Black Political Economy*, 49(1), 41-60.
- Gee, J. P. (1989). Literacy, discourse, and linguistics: Introduction. *Journal of Education*, 171, 5-17. https://doi.org/10.1177/002205748917100101
- Heaven, Will Douglas. (2023). The open-source AI boom is built on Big Tech's handouts. How long will it last? *MIT Technology Review*. <u>https://www.technologyreview.com/2023/05/12/1072950 /open-source-ai-google-openai-eleuther-meta/</u>
- Inoue, A. B. (2015). *Antiracist writing assessment ecologies: Teaching and assessing writing for a socially just future*. Parlor Press LLC. <u>https://doi.org/10.37514/per-b.2015.0698</u>
- Kahn, Sal. (2023, March 14). "Harnessing GPT-4 so that all students benefit. A nonprofit
- approach for equal access." <u>https://blog.khanacademy.org/harnessing-ai-so-that-all-students-benefit-a-nonprofit-approach-for-equal-access/</u>
- Laquintano, T., Schnitzler, C. & Vee, A. (2023). Introduction to teaching with text generation technologies. In A. Vee, T. Laquintano, & C. Schnitzler (Eds.), *TextGenEd: Teaching with Text Generation Technologies*. The WAC Clearinghouse. <u>https://doi.org/10.37514/TWR-J.2023.1.1.02</u>
- Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against nonnative English writers. arXiv:2304.02819. <u>https://arxiv.org/pdf/2304.02819.pdf</u> <u>https://doi.org/10.1016/j.patter.2023.100779</u>
- Liu, S., & Yu, G. (2022). L2 learners' engagement with automated feedback: An eye-tracking study. Language Learning & Technology, 26, 78-105.
- Making English Fun (2023, March 14). How to use ChatGPT in an ESL English classroom?<u>https://makingenglishfun.com/2023/03/14/how-to-use-chatgpt-in-an-esl-english-classroom/</u>
- Martindale, J. (2023, April 12). These are the countries where ChatGPT is currently banned. *Digitaltrends*. <u>https://www.digitaltrends.com/computing/these-countries-chatgpt-banned/</u>

McMurtrie, B. (May 26, 2023). How ChatGPT can help or hurt students with disabilities. *The Chronicle* of Higher Education. <u>https://www.chronicle.com/article/how-chatgpt-could-help-or-hurt-</u> <u>students-with-</u>

disabilities?utm\_source=Iterable&utm\_medium=email&utm\_campaign=campaign\_6923362\_nl\_ Academe-Today\_date\_20230530&cid=at&source=&sourceid=

Mezirow, J. (2008). An overview on transformative learning. Lifelong Learn. 40–54.

- Mezirow, J. (1994). Understanding Transformation Theory, Adult Education Quarterly, 44(4). 222-232. https://doi.org/10.1177/074171369404400403
- Oversight of A.I.: Rules for Artificial Intelligence. Hearing Before the Subcommittee on Privacy, Technology, and the Law of the Senate Committee on the Judiciary. 118<sup>th</sup> Congress (2023, May 16). Testimony of Sam Altman. https://www.judiciary.senate.gov/committeeactivity/hearings/oversight-of-ai-rules-for-artificial-intelligence
- Owusu-Ansah, A. L. (2023). Defining moments, definitive programs, and the continued erasure of missing people. *Composition studies*. *51.1*.
- Robertson, A. (2023, April 23). ChatGPT returns to Italy after ban. *The Verge*. <u>https://www.theverge.com/2023/4/28/23702883/chatgpt-italy-ban-lifted-gpdp-data-protection-age-verification</u>
- Roussel, S. (2011). A computer assisted method to track listening strategies in second language learning. *ReCALL*, 23, 98-116. <u>https://doi.org/10.1017/s0958344011000036</u>
- Singer, J. (n.d.). Reflections on Neurodiversity: Neurodiversity: Definition and Discussion. https://neurodiversity2.blogspot.com/p/what.html
- Smitherman, G. (1977). *Talkin and testifyin: The language of Black America*. Wayne State University Press.
- Vohs, K. D., & Baumeister, R. F. (2004). Understanding self-regulation. In R. F. Baumeister & K. D. Vohs (Eds.), *Handbook of self-regulation (1-9)*. New York, NY, The Guilford Press.
- Young, V. A. (2010). Should writers use they own English? *Iowa Journal of Cultural Studies*, *12*(1). <u>https://doi.org/10.17077/2168-569x.1095</u>
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. New York: Public Affairs.