

Generative AI: Challenges to higher education

Sencer Yeralan^{1*}, Laura Ancona Lee²

¹ Philosophy, ADA University, Azerbaijan

² Academic Internationalization, Paulista University (UNIP), Brazil

*Corresponding author E-mail: sencer@yeralan.org

Received Jun. 12, 2023

Revised Sep. 2, 2023

Accepted Sep. 8, 2023

Abstract

Generative Artificial Intelligence has rapidly expanded its footprint of use in educational institutions. It has been embraced by students, faculty, and staff alike. The technology is capable of carrying out a sustained sequence of interactive dialogs and creating reasonably meaningful text. Not surprisingly it seems to be routinely used by faculty to generate questions and assignments, by students to submit assignments and aid in self-learning, and administration to create manuals, memoranda, and policy documents. With its potential to lead to significant social innovation, tethering on the verge of becoming a disruptive technology, it seems most unlikely that it will fade away without being fully enfolded into almost all aspects of academic and pedagogical activity. While it is early to predict the exact place of this technology in education, we present thoughts to aid deliberations and give a brief review of the opportunities and challenges.

2022.

Published by ARDA.

Keywords: Generative artificial intelligence, AI, Higher education, Thought experiments

1. Introduction

Admittedly quite pedantic to the more savvy reader, we start with an attempt to demystify key concepts related to Artificial Intelligence (AI). Artificial Neural Networks (ANN), by and large, form the basis of current artificial intelligence models. They were first conceptualized in the mid-20th century with the idea of mimicking the structure of the human brain. The brain consists of many neurons connected through synapses in a complex network. These neurons carry electrical charges, which led to the hypothesis that a similar network could produce electrical activity that may result in some form of intelligence, including perception, thought processes, consciousness, self-consciousness, and intelligence.

The early circuits consisted of multiple inputs and one output, similar to an analog computer where input and output values were determined by voltages at each terminal. An input signal representing its strength by its voltage is modified by an input function and then brought to a summer (a simple summer circuit may be constructed with an operational amplifier [1]). The summer then adds up the voltages to produce a single output. These functions could amplify or attenuate, in a linear or nonlinear fashion, the input voltages. Once a single neuron unit was built, these neurons could be connected to form a collection of inputs and outputs of more

complex networks. Typically, there would be a collection of input neurons, and a collection of output neurons, as given in Figure 1.




Figure 1. A simple artificial neural network

To use the neural network, the specific inputs of the application would have to be written as a vector of values. These values encode the features of the input. The output could also be a collection of values. Take for example, an application to identify songs. The signal from the microphone could be examined at 100 millisecond intervals. The frequency at each epoch could be taken as the inputs. For a 20 second song segment there would be 200 inputs. The output would be one of the songs that the application could identify.

In most applications, the input is a feature set rather than raw data. In the example given, instead of frequencies, we could first extract the chords, and then sample the input as a chord progression. Then, the sampling could be done using longer time intervals, say seconds. A 200-input neural network could then take as an input a song segment of 200 seconds. This application could then identify the song being played, since each song would be expected to have a unique cord progression over time.

In this example, input vector then gives the specific frequency (or chord) at each time interval. The output is one of many known songs. The circuit identifies which song is being played.

The prerequisite for the neural network to properly function is the determination of the input functions. The process of determining these functions is similar to curve fitting using the *least-squares* method [2,3]. Some parameters have to be estimated so that they best relate the output to the input. Interest in neural networks dwindled when it was discovered that a single neural circuit, called the *perceptron*, could not mimic the behavior of an XOR gate, notwithstanding that the implementation of AND gates and OR gates was rather straightforward. However, interest was rekindled when it was recognized that XOR behavior or other types of behavior that were not possible with simple networks could be implemented provided that additional hidden layers were inserted between the input layers and output layers. Neural networks that employ many hidden layers is referred to as *deep neural networks*.




Figure 2. An artificial neural network with a hidden layer

Determining the specific values and forms of the input activation functions is a daunting task, but it can be done through an iterative process where inputs and outputs are presented to the new network. A mathematical search process is then implemented, where the values of activation function parameters that best relate the given outputs to the inputs are sought. This process is called training the neural network, and as said, conceptually it is a curve-fitting exercise.

The logic of neural networks is straightforward. No neuron has any intelligence. Instead, it is assumed that intelligence is an emergent behavior that occurs when a large number of neurons interact. This approach acknowledges that each perceptron is not a component of intelligence but rather intelligence is emergent. There are good thought experiments such as the Game from the 1960s [3] and the Chinese Room from the 1980s [4].

2. Natural language processing

Natural Language Processing (NLP) is implemented in a relatively large neural network with a sufficient number of hidden layers. Until recently, it was not known whether such a network could mimic a person speaking a natural language, or if one could convincingly interact with the system using natural language inputs and outputs [5].

Training the new network for NLP, especially the Large Language Models (LLM) that are now commonly used, requires a significant amount of computing power, quite a large neural network, and many hidden layers. However, the general rather simplistic logic remains the same. OpenAI's product, ChatGPT, was an ambitious project that required building a LLM and training it with the help of many human trainers. GPT stands for Generative Pre-trained Transformer. Essentially the neural network looks at the words and picks the best words to follow a given incomplete line. This is an oversimplification, as there are many different aspects to this process. For example, there may be alternative endings to a given initial sub-sentence, so a determination is needed as to how to generate the remainder of the sentence. Similarly, the system checks for each word in a sentence to see if it is consistent with the statement taken as a whole. Nonetheless, the system picks words and phrases and creates a text output. The neural network possesses the rules, just like other grammar checkers, and possesses the capability, just like other spell checkers, to make the output grammatically consistent with proper punctuation. When the system is run, it was noticed that the speed is sufficient for it to be perceived as an almost-human companion. One can tell interacting with ChatGPT that it is not a human, thereby not passing the Turing test fully. This, however, does not prevent it from proving to be a useful tool. As such, it behaves like a co-pilot or a companion that can take on some of the more tedious tasks in creating text.

Another interesting and useful feature that makes ChatGPT preferable to search engines is that it uses the so-called *Open Session* and incorporates past conversations into its generation of new text. Therefore, one should not use ChatGPT just to submit a query and look at the output, as one would normally do with a search engine. Rather, one could use ChatGPT in an iterative manner, asking for something, submitting a query, looking at the output, and refining the query. For example, one may ask ChatGPT to write a chord progression for a 12-bar blues song and later refine it by asking to make the chords jazzier or to transpose the progression to another key. Such interactive episodic search is perceived by humans as communicating with another human being, enhancing the appearance that ChatGPT is a co-worker or a co-pilot.

The fact that ChatGPT is trained using material on the web makes it somewhat of an alternative to a keyword-based search engine. ChatGPT should be considered as a versatile engine. There are many possible ways that it could be incorporated into an online experience. For example, one may have three open tabs while searching for a given topic. Then, one may ask a ChatGPT plugin of their browser to summarize the information on these three tabs and present it in text form. This is a task that could be done by a co-worker, so clearly, it improves the search experience and simplifies going back and forth between tabs to manually collect information. Currently, there is a flurry of activity in implementing useful plugins for browsers and other common applications.

3. Uses of AI in higher education

The use of AI in higher education is not a new phenomenon [6], [7], [8]. It has raised some concerns, predominantly as a fear that students may take advantage of AI-based tools to complete their academic assignments with little or no genuine effort, and with no understanding. Professors may also use such tools to generate multiple-choice quizzes or presentations for their classes.

Despite the European Union (EU) AI Act, the first AI regulation on artificial intelligence, it was rather surprising that Italy banned the use of Generative AI (GenAI), specifically ChatGPT by OpenAI. The EU Parliament would like AI to be safe, transparent, traceable, non-discriminatory and environmentally friendly, overseen by people, rather than by automation, to prevent harmful outcomes. The rationale for the Italian ban was on grounds of privacy. Much anecdotal evidence exists regarding how, by and large, most higher education institutions also had a negative reaction to the use of GenAI by students, faculty, or staff.

While AI-based tools can be useful for brainstorming and generating initial drafts, there is concern about the potential for academic dishonesty. Plagiarism checkers have become more sophisticated in detecting content that has been generated by AI systems. However, it is still possible to submit AI generated output as genuine work. As usual, technology to create and to detect advance concurrently, each feeding into the development of the other.

Despite these concerns, GenAI can be a valuable tool to enhance the efficiency of the writing process. One can ask GenAI to provide alternative responses and several versions of an initial draft. This may enhance and simplify the initial brainstorming phase and enrich the thinking process. Therefore, the idea that rather than banning GenAI outright, universities should consider ways to regulate their use has merit.

The use of any tool may be deemed “good” or “bad”, “honest” or “dishonest” relative to the circumstances and individuals. Could it be that the first cutting objects seen in archaeological museums were used only for hunting and never for wars or murders? Does dishonesty imply ignorance? The ethical use of learning tools benefits those who want to learn. The discussion about the shortcomings of the system that allow deceit and dishonesty is not sufficient to ban GenAI in education.

A comprehensive posture regarding GenAI requires a thorough understanding of many components of this rather elaborate construct. There are not only technical issues, but pedagogical, ethical, commercial, sociological, and political concerns, all woven together into a rather entwined tapestry. It is safe to submit that more time and thought are needed to arrive at a generally accepted position regarding the place of AI in higher education. To this end, in the remainder of this study, we will present a few views and dwell on thought experiments to help clarify and better understand the scope of this issue.

4. Deliberations and thought experiments

How GenAI may be incorporated into higher education is not a simple proposition. It requires substantial debate. It also requires a systems view that re-evaluates the purpose and role of higher education. The task is hindered by the fact that most university administrators are ill-equipped to fully comprehend and appreciate the significance of these new technologies. This is especially worrisome at neo-liberal corporate universities [9].

We now present a few thoughts and point out historical parallels from a wide range of domains in hopes of promoting the much-needed debate on this topic.

4.1. The sand heap paradox

The Sand Heap Paradox raises concerns about the use of GenAI language models in academia and higher education. The paradox asks, “When does a sand heap become a non-heap as you remove grains of sand one by one?”. Clearly, when there is only one grain of sand left, it is no longer a heap. However, in the beginning, it is a heap. So, when exactly does that transition occur from a heap to something else?

Similarly, when does the use of GenAI in a word processor become academically unacceptable? Word processors already have spell-checking tools.

Spell-checking was once done manually requiring some mental effort. There is very little concern about the use of built-in spelling tools, dictionaries, thesauruses, etc. Nowadays, word processors also look at grammar and suggest improvements, and there does not seem to be much academic concern about using grammar checking word processors. On the contrary, most professors would encourage the use of such tools, as would industry. Word processors are also good at checking for punctuation mistakes and suggesting corrections or improvements. But when a word processor suggests rephrasing using a GenAI, it often becomes subject to banning.

However, it is important to note that the capabilities of language processors vary widely, from no help at all (like a typewriter) to a capable GenAI-driven word processor. Therefore, it becomes difficult, almost paradoxical, to draw a line beyond which the use of GenAI becomes academically unacceptable.

4.2. The arts and crafts movement

The Arts and Crafts movement of the mid to late 1800s emerged after the Industrial Revolution [10]. With the advent of industrialization, it became possible to manufacture goods at a much lower cost than before. However, these products lacked some of the artistic elements that were present in their handcrafted predecessors. As a result, artisans and crafts people, who were adversely affected by industrialization, initiated a movement that claimed that industrialization was harming the artistic and aesthetic aspects of products. This movement later gave way to Art Nouveau and Art Deco.

In the past, having an industrial product rather than a handcrafted one may have been frowned upon. However, today, this concern seems to have all but disappeared. People are happy to wear clothing with the insignia of the manufacturer, and all such clothing, whether it's footwear, jackets, or trousers, comes from an industrial process. This does not seem to bother consumers, rather, people with brand loyalty proudly display such products. Society has gone from considering machine-crafted products inferior to finding them not only acceptable but also preferable.

Could a parallel be drawn here to the possibility of people in the future considering handwritten memoranda or essays as inferior to those that are artificially created by GenAI? Just as we no longer know how to sew our clothes, might we reach a point where not knowing how to write our own essays will not bother us? Here, we are not questioning whether this would be a good thing or a bad thing. We simply pose the possibility of such a future à la the *Brave New World* [10]. Will there be a time when the use of GenAI is so widespread that human-written essays are no longer preferred? It's an interesting question that raises concerns about the role of technology in our lives.

4.3. Legitimate uses of GenAI

Given that there are legitimate uses for a companion, co-writer, or co-pilot, outright banning GenAI does not seem to be a good idea. Nor is it likely to be accepted by those who would like to use its beneficial contributions to improve their efficiency and effectiveness. There are several examples of how GenAI could be used in a positive and productive manner, such as quizzing oneself on a given subject to prepare for an exam or asking GenAI how best to prepare for a given assignment (see the Appendix).

GenAI can be used in multiple ways. For example, if you are studying for an exam, you can ask GenAI to give you a quiz at your level. You can specify if it is an undergraduate or graduate course and ask ten multiple choice questions. You then input your answers and GenAI will then tell you which questions you answered correctly. From there, you have a couple of options. You can ask for the answers to the questions you answered incorrectly and study those subjects, or you may ask for suggestions on what subjects to study to make up for those incorrect answers. After a few iterations of this process, you will have a personalized and targeted study program, with GenAI as a smart tutor providing custom-tailored support.

This indirect second-order request demonstrates another leap in perceived intelligence for GenAI. Rather than asking it to do something directly, one can ask it to tell you how best to ask GenAI to do something. This removes the burden of understanding, articulating, and comprehending a strategy to use the technology. Ultimately, the human will make the decision as to how best to use the technology, but there is a difference between starting from scratch and starting with a set of options or ideas.

In this case, one could argue that the benefit of GenAI is to overcome the initial hurdle of organizing initial thoughts and quickly arrive at a set of options that could be used as a launching pad. While the option picked to go forward with may not be among those initial options, starting with options simplifies the burden and allows one to quickly focus on the task at hand.

4.4. A thought experiment

As the field of GenAI continues to advance, it raises questions about the role of human input in the text generation process. Just as we have progressed from programming computers in machine code to writing code in high-level languages, could we also move towards a similar evolution in text generation?

Consider the idea of a student preparing for an essay. Instead of spending time on the actual text, could the student simply provide the key ideas and concepts to a GenAI platform and let the software compose the text itself? This approach would be similar to writing code in a high-level language and allowing the compiler to generate the machine code. Just as we assume the receiving party can compile and execute our code, we could also assume that the GenAI can compose the text based on our input.

While this approach may seem unconventional, it could be a way to streamline the writing process and minimize the energy required to convey the message. Companies now store standard “ChatGPT prompts” rather than email messages. Here, ChatGPT acts as a macro to generate the desired text. The focus shifts from the actual text to the underlying ideas and concepts. Ultimately, the text format could be seen as an intermediate step, similar to machine code, that is not necessary for conveying the unique ideas. This approach could be a promising avenue for exploring the capabilities of GenAI in education and beyond.

We now have computer science students who have never been exposed to machine code. Is it possible that in the future, people without a drivers license ride alone in self-driving cars? If so, how then could one criticize the generation of text by those who only know how to seed or prompt GenAI but do not know how to write in the traditional manner?

5. The road forward in academia

The use of GenAI in academia has become a hotly debated topic. Should we ban GenAI or should we embrace it and ask that students do nothing but generate their essays using such tools? The answer is not a one-size-fits-all solution, and it is too early to make any definitive rules regarding GenAI in academia.

After all, after graduation, students will undoubtedly use GenAI throughout their careers, where GenAI may be seamlessly integrated into many products. It may also come to pass that society views GenAI as something that is not only acceptable but preferred, as was the case with the Arts and Crafts Movement. We need to be cautious in making any specific decisions and procedures regarding GenAI.

What remains is that academia should understand that it has a unique role in the road forward. We academicians are responsible for producing the next generation of citizenry and workforce. We need to ask people to familiarize themselves with AI and truly understand the technology. It may be a good starting point for us to compose a list of properties, concerns, benefits, and threats of GenAI in an inclusive and transparent manner. We need to have a comprehensive dialogue so that all different aspects of this complicated and far-reaching issue could be addressed in the fullness of time.

Once all such aspects are understood and evaluated, we should attempt to come up with fixed processes that determine how students and faculty may use GenAI. Any early attempt to ban the use or any quick solutions,

such as relying on plagiarism checkers to disallow GenAI, is probably premature at this point. On the other hand, we may consider teaching the best way to use GenAI as a copilot, a co-worker, a companion, or a digital twin. Such a practice may better prepare the student for the latter part of the 21st century, which promises to be much different from the professional practices of today.

Ultimately, the road forward in academia regarding GenAI should be guided by caution, understanding, and a comprehensive dialogue. We need to prepare ourselves for the potential benefits and threats of GenAI. We should not rush to any premature decisions or solutions. With proper preparation and understanding, we can ensure that the next generation of citizenry and workforce is well-equipped to navigate the complex world that will most likely include GenAI.

6. Concluding remarks

The fast pace of acceptance of GenAI seems to point to a phenomenon not likely to disappear. If GenAI is to be incorporated into facets of future work, academia must address the issue and take a leadership role in defining the scope and boundaries of its use. All ethical concerns must be weighed in light of the potential of GenAI as a productivity tool and an efficiency multiplier. Early attempts to limit or even ban its use would seem counterproductive. Similar to the computerization and digitalization that left their indelible marks on pedagogy, the authors expect GenAI to follow a similar fate, but one decidedly faster in its spread and acceptance.

One interesting observation relates to the concerns that GenAI would take jobs away from the workforce. It can be argued that tasks such as punctuation and pagination, expanding into a narrative key bullet point, etc. which can be delegated to an assistant, or a speech writer may easily be referred to GenAI without reservation. In fact, such practices would serve as testimony that burdening creative minds with these rather meaningless tasks was a waste, and that it is high time that such mistakes be left behind.

With regard to the use of GenAI that generates narratives from bullet points, the authors feel that GenAI provides an effective tool in teaching students to be precise, succinct, and focused, rather than pursuing glib ramblings. This is especially an important skill to be nurtured in engineering and sciences.

Declaration of competing interest

The authors declare that they have no known financial or non-financial competing interests in any material discussed in this paper.

Funding information

No funding was received from any financial organization to conduct this research.

Declaration of use of AI in the writing process

No AI tools were used in writing this manuscript.

References

- [1] S. Yeralan, Operational Amplifiers through Experimentation and Reflection: An Engineer's Journey. IUS Press, 2023.
- [2] J. H. Holland, Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. Bradford Books, 1992, isbn: 978-0262581110.
- [3] A. Dneprov, "Игра [The Game]", Знание—Сила [Knowledge—Power], с. 39—41, 1961. url: <http://q-bits.org/images/Dneprov.pdf>.
- [4] J. R. Searle, "Minds, brains, and programs", Behavioral and brain sciences, т. 3, с. 415—457, 1980.

- [5] I. Solaiman, M. Brundage, J. Clark и др., Release Strategies and the Social Impacts of Language Models, 2019. arXiv: 1908.09203 [cs.CL].
- [6] S. Amershi, M. Cakmak, W. B. Knox и Т. Kulesza, “Power to the People: The Role of Humans in Interactive Machine Learning”, AI Magazine, т. 35, № 4, с. 105—120, дек. 2014. doi: 10.1609/aimag.v35i4.2513. url: <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2513>.
- [7] J. Belda-Medina и J. R. Calvo-Ferrer, “Using Chatbots as AI Conversational Partners in Language Learning”, Applied Sciences, т. 12, № 17, 2022, issn: 2076-3417. doi: 10.3390/app12178427. url: <https://www.mdpi.com/2076-3417/12/17/8427>.
- [8] M. A. Kuhail, N. Alturki, S. Alramlawi и K. Alhejori, “Interacting with Educational Chatbots: A systematic Review”, Applied Sciences, т. 28, № 17, 2023, issn: 1573- 7608. doi: 10.1007/s10639-022-11177-3. url: <https://doi.org/10.1007/s10639-022-11177-3>.
- [9] A. Taylor, “Perspectives on the University as a Business: The Corporate Management Structure, Neoliberalism and Higher Education”, Journal for Critical Education Policy Studies, т. 15, № 1, с. 108—135, 2017.
- [10] R. P. Blakesley, The Arts and Crafts Movement. Phaidon Press, 2006, isbn: 978-0-7148-3849-6.

Appendix

This appendix illustrates a session with ChatGPT as a tutor.

Input query:

How can you help me to prepare for my college course final exam?

ChatGPT:

ChatGPT can assist you in preparing for your college course final exam in various ways.

Here are a few examples:

1. Quiz Preparation: You can ask ChatGPT to generate practice quizzes for your exam. You can specify the level of difficulty, the number of questions, and the topic of the questions. You can take the quiz and get instant feedback on your performance.
2. Study Recommendations: Based on your performance in the practice quizzes, you can ask ChatGPT to recommend study materials and resources to help you improve in the areas where you struggled. ChatGPT can suggest books, articles, videos, and other resources that are relevant to your exam.
3. Concept Clarification: If there are specific concepts that you are struggling with, you can ask ChatGPT to explain them to you in a clear and concise manner. ChatGPT can provide you with definitions, examples, and analogies to help you understand the concepts better.
4. Essay Writing: If your exam requires you to write an essay, you can ask ChatGPT to help you brainstorm ideas, structure your essay, and provide feedback on your writing. ChatGPT can suggest relevant sources, help you develop a thesis statement, and provide tips on how to write a compelling essay.

Overall, ChatGPT can be a useful tool in your exam preparation process by providing personalized and targeted assistance to help you achieve your academic goals.

This page intentionally left blank