




CERLI[®]

Artificial Intelligence and Its Potential Uses in the Classroom





ARTIFICIAL INTELLIGENCE AND ITS POTENTIAL USES IN THE CLASSROOM

CENTER FOR EDUCATIONAL RESEARCH, LEARNING & INNOVATION

ABSTRACT

The report “Artificial Intelligence and Its Potential Uses in the Classroom” addresses a matter of vital importance in modern society: Are AI systems as useful as we think? We humans are accustomed to believing that every advance is significant, even when we do not understand it in detail.

For this reason, CERLI has decided to produce this report, bringing together all the information we have about artificial intelligence (AI) in an academic article that is understandable, accessible, and informative.

We will examine the history of artificial intelligence, its significant advances over the past decade, and the impact it has had on our lives so far. Finally, we will seek to account for its strengths and weaknesses in an implementation model that educational institutions can find useful and effective.




© CERLI, 2025

“Center for Educational Research, Learning & Innovation (CERLI) (2025). Artificial Intelligence and Its Potential Uses in the Classroom”

Published in THE ORVIUM EDUCATION JOURNAL

Artificial Intelligence and Its Potential Uses in the Classroom © 2025 by CERLI is licensed under CC BY-ND 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nd/4.0/>.



For more information about this report or any other CERLI publication, contact us at info@cerliorg.com or contact the Department of Education at edu@cerliorg.com

Table of Contents

Introduction 5

About AI 6

The History of AI 9

Machines That Learn 16

Humanity's Next Leap 24

Modern Examples 32

General-Purpose Large Language Models 34

Virtual Assistants 41

Modern Chatbots 47

Text-to-Image Features..... 59

Text-to-Speech Features 66

Text-to-Video Features 69

AI in Education 74

The Ethics of Educational AI 78

Implementing AI (Successfully) in the Classroom 87

Resource Considerations 94

Examples of Educational AI 100

Toddle 103

Teachmate 107

ILIAS 110

Epilogue 114

References 116



Introduction

Artificial intelligence is one of the leading advances in our modern society and has permeated every aspect of our educational, personal, and working lives. However, as has often happened in the past, technologies this advanced tend to struggle to enter formal education, primarily because of the considerable uncertainty surrounding them and the conservative approach of most education systems worldwide. Moreover, such significant advances attract those who exaggerate their strengths and weaknesses for economic, political, or even ideological reasons.

For this reason, at CERLI, we have undertaken a realistic and objective analysis of the main AI tools on the market, seeking to understand what makes them interesting and/or dangerous and to study how feasible their implementation in education is. This is no easy task; it requires a meticulous analysis of a set of new and complex technologies whose details could overwhelm even a specialist. Nevertheless, our research team is not afraid of challenges and has made a significant effort to bring this project to fruition in the best possible way.

This report will detail everything we know about AI systems: their past, present, and future, their impact on each of our lives, and their imminent arrival in the education system. We will examine each of their distinctive features in detail, with particular emphasis on what sets them apart from everything else.

About AI

Artificial intelligence, also known as AI, is not a concept associated with one specific technology. It is a branch of computer science whose goal is to develop systems or machines capable of performing complex tasks that would usually require human-level intelligence. These include reasoning, solving problems, understanding language, making decisions, and learning, among many others.

It works by exposing a machine to an enormous amount of data and having it process those data to learn and improve its ability to recognize patterns. In this way, AI systems could be used to analyze large amounts of information and draw conclusions based on relationships that a human might have overlooked or would simply take much longer to identify.

Achieving this is highly complex and requires algorithms and mathematical models that allow the machine to extrapolate knowledge from an initial dataset. This is what we call machine learning (ML), and its capabilities are impressive. It was subsequently surpassed by deep learning, which used artificial neural networks to simulate the way the human brain processes information. However, we will not yet delve into the operational details of AI systems; instead, we will focus on defining them properly.

Google, for example, classifies AI systems according to their functionality. First, we have “reactive machines,” which are limited AI systems that only react to different types of stimuli based on preprogrammed rules. They do not use memory and therefore cannot learn from new data. Second, we have “limited memory” systems. These can use memory to improve over time through training on new data, generally through an artificial neural network or some other training

model. Deep learning, a subset of machine learning, is considered limited memory AI. This category includes modern AI tools such as ChatGPT and Gemini, among others.

These would be followed by “theory of mind” systems. The term refers to AI capable of emulating the human mind and making decisions in a way similar to a human being. This includes recognizing and remembering emotions, as well as responding to social situations as a human would. Theory of mind AI does not currently exist, although we believe it is the next step.

Finally, we have “self-awareness.” The concept of self-aware AI describes a machine that is aware of its own existence and has the intellectual and emotional capabilities of a human being. Like theory of mind AI, self-aware AI does not currently exist. Reaching this level of technology would have serious legal, ethical, and moral consequences that the society of the future would need to resolve.

Although the features of this classification are very interesting, it is not particularly useful for the tools we have today. We therefore need an approach that categorizes modern technologies according to their characteristics. For example, in the book *“Artificial Intelligence: A Modern Approach,”* researchers Stuart Russell and Peter Norvig provide a comprehensive analysis of these technologies, establishing four main types.

✓ **Systems that think like humans:**

These focus on emulating human intelligence in terms of both behavior and thought. They seek to imitate the way humans think and solve problems. This is the goal of tools such as ChatGPT.

✓ **Systems that act like humans:**

These focus on emulating human intelligence in terms of behavior. They seek to imitate how humans behave and act in the world. This category

includes chatbots and any AI system that seeks to emulate the experience of conversing with a human being.

✓ **Systems that think rationally:**

These focus on solving problems logically and rationally. They seek to maximize the efficiency and accuracy of their decisions without necessarily considering human behavior.

✓ **Systems that act rationally:**

These focus on making decisions and acting in the world, always seeking to make the best possible decision based on the information available.

Although this classification helps us approach modern technologies, it tells us nothing about how they achieve these goals. Humans have made such significant progress in our artificial creations that classifying them naturally and accurately has become a monumental task. At CERLI, we tend to use a much less elegant classification system that we find more illustrative: preprogrammed systems, learning-and-replication systems, and intelligent systems.

To understand this classification and its foundations, we need to examine the history of AI systems. We must return to a time when we encountered software of this kind, even if we may not have recognized it as such. Let us revisit this history so that we can appreciate our progress in these areas and see how small changes in our understanding of intelligence brought us to where we are today.

The History of AI

Although it may seem new, the history behind artificial intelligence reaches far back into human history. The idea of creating machines that think like humans dates to ancient philosophy, when Aristotle proposed logical rules for reasoning. Later, during the Renaissance, thinkers such as Descartes and Leibniz imagined mechanical systems capable of performing intellectual tasks. However, the paradigm shift in this field came with the emergence of the theory of computation.

In 1936, Alan Turing introduced the concept of the “Turing machine,” which established the theoretical foundations of modern computers. Later, in 1950, Turing proposed the famous “Turing test,” an experiment designed to evaluate whether a machine can exhibit intelligent behavior similar to that of humans. This laid the groundwork for the idea that it is possible to create systems so advanced that they make us question whether they are alive. Artificial intelligence would formally become established as a field in 1956, during the Dartmouth conference.

A new field of knowledge had just been born, bringing together the exponential development of computing in the late 20th century and the almost fantastical idea of creating intelligence similar to that of humans without a biological component. But before discussing the historical milestones humanity has reached in AI development, let us explain the classification presented at the end of the previous section. It will be very useful for categorizing the different AI systems we will examine below and understanding their main strengths and weaknesses.

We will begin with preprogrammed systems, which we understand to be programs that cannot learn on their own but respond to users’ needs based on a body of information previously supplied to them. This type

of AI is often illustrated through a famous example by philosopher John Searle, who used it as an argument against the possibility that AI could develop a genuine understanding of language. We will use a slightly modified version of Searle's example to explain preprogrammed systems:

"Imagine a closed room containing a person who knows no Chinese, with two small slots opening to the outside. This person receives messages written in Chinese through slot one and is expected to respond to the message in perfect Chinese through slot two.

"To do this, the person inside uses a book containing a large number of messages in Chinese, each accompanied by a response. All the person has to do is quickly look up the message received through slot one in the book, copy the response onto a new piece of paper, and pass it through slot two.

"If this procedure is followed quickly, the person receiving the responses would think that someone inside the room understands Chinese perfectly. However, this is not the case; the person is merely following a mechanical procedure that they do not understand but that works."

This is precisely what a preprogrammed system is. Developers prepare a broad set of possible responses in advance for the inputs they imagine users might provide, and the AI system simply accesses this database whenever it receives an input, returning the predetermined output. If programmed correctly, these systems never give incorrect responses. However, this comes at a high price: A preprogrammed AI system will never create the impression of speaking with a human being, because it cannot adapt its responses to specific features of the input.

Moreover, these AI systems are severely limited by the size of the initial database. If a user provides an input that was not previously considered, the AI system will find no response and will return an error

output, which is also preprogrammed. These systems can therefore be improved only by adding new inputs and outputs to the database.

Second, we have learning-and-replication systems. These were very popular in the early 2010s and were most people's first encounter with an AI system. They work in a very simple way: Like preprogrammed systems, they respond to inputs using a database of possible responses. The difference is that a learning-and-replication system begins with a limited book of possible inputs and outputs but can expand it through user interaction.

They start out with very limited capabilities, giving basic, predetermined responses. When a user provides an input that the AI system did not previously account for, instead of responding with an error message, it stores the input in its book and collects responses to that specific input from other users. It then assigns them a score based on their frequency and popularity or, when available, a rating system. The next time the AI system receives that input, it responds with one of the many answers users have supplied.

Their main advantage over preprogrammed systems is that they develop their databases over time without requiring the original development team to add new inputs and outputs. The community using the tool updates its knowledge over time. However, this is also their main weakness, which ultimately rendered learning-and-replication systems completely obsolete. Because the AI system does not truly understand language and the development team does not directly oversee the results added to the database, incorrect or malicious use of the tool can quickly render it obsolete.

For example, users could suggest responses that do not match the input or provide completely inappropriate answers so that the AI system returns utterly irreverent and funny outputs that completely undermine its conversational ability. It was not just that users could no longer expect the same response to the same input, which had been the greatest advantage of preprogrammed systems; now the response was completely useless in the vast majority of cases.

Using these two categories, we can identify clear examples throughout history. The best-known were ELIZA and Deep Blue, because they had real-world applications. ELIZA was a preprogrammed system created by Joseph Weizenbaum in 1966 that could emulate a human conversation through pattern matching between user inputs and a database previously supplied to it. Although it had multiple “scripts”—as its characters were called—the most famous was “DOCTOR,” which emulated a female Rogerian psychotherapist ¹. Although it was easy to fool and its responses were very generic and repetitive, it was an impressive advance for its time, surprising everyone in AI research.

Deep Blue, meanwhile, is perhaps the most famous preprogrammed system in history. It was created by the American company IBM in the 1990s with a single real purpose: to play chess as perfectly as possible, making it indistinguishable from a world-class player. To achieve this, Deep Blue had a database loaded with thousands of chess games and strategies, along with the ability to analyze up to two hundred million positions per second in real time and use brute force to determine the best possible move. This alone was enough to defeat the reigning world chess champion, the legendary Garry Kasparov, in 1997. Deep Blue was an inspiration to those who championed the power of computing to make complex decisions in real time.

Moving into the 21st century, we arrive at an era filled with AI tools such as virtual assistants and the first chatbots. While the former followed rules very similar to those of preprogrammed systems, chatbots were developed using learning-and-replication mechanisms. Among the best-known was one that accompanied the birth of social media: Cleverbot. This chatbot was released by Rollo Carpenter in 2006, although the project dates to 1988, with the goal of holding conversations with humans and improving over time to adapt to more complex contexts and ideas.

¹ This choice was not random, since these therapists use a technique that involves formulating their next statement based on the information in the patient’s words. This helped conversations with ELIZA feel more real while also making its book of responses much easier to program.

Originally, it analyzed patterns among users' inputs and responses and selected sets of these after millions of interactions. It then stored them in a database that Cleverbot itself could consult when a user provided a similar input. Although this enabled it to hold a wide variety of responses that could sound human, they were repetitive and lacked depth, readily revealing the difference between an interaction with Cleverbot and a human interaction. This AI system did not understand language naturally, as ChatGPT and virtual assistants do today, so it could not adapt to what users said.

Indeed, if a user managed to provide an input with no relevant matches in its database, the responses quickly became illogical and out of context. This happened because the pattern recognition system failed, completely disrupting the search functions and producing a random output. Although it was a significant advance over ELIZA, since reaching the limits of its output was considerably more difficult, it highlighted the inability to create a chatbot that could fully, contextually, and unequivocally simulate a conversation with a human.

Nevertheless, Cleverbot had a substantial filtering and moderation system that allowed it to maintain a certain level of quality in its responses. Although they sounded monotonous and not particularly relevant, its responses were much easier to classify in the database, making it harder to obtain results that violated moral standards and accepted social norms. The exact opposite happened with our next two examples: SimSimi and Tay.

These two chatbots were developed by ISMaker, now known as SimSimi, and Microsoft, respectively. Each was designed to function as a large-scale version of Cleverbot, using the same programming strategy but different approaches to user interaction. SimSimi was available both on a website and as a mobile application, allowing users to hold a conversation with the chatbot. When a user provided an input that the machine had not previously considered, they were offered the option of supplying a response they would find useful. Through a popularity ranking system, users' responses would be used by the AI system when it had to respond to that input again.

Tay, on the other hand, was programmed to run on the social network Twitter and interact with its users to learn their patterns of speech and information sharing. The more it interacted, the more it adjusted its responses to what appeared most frequently, making its conversations—on paper—more human and wide-ranging.


Both AI systems had the same problem: a lack of moderation. Since they were unsupervised learning models, users could exploit this vulnerability to make the model answer however they wanted, without the AI system being able to distinguish between a coherent response appropriate to the input and a completely offensive, malicious one. Both began responding to any input with racist, homophobic, sexual, illegal, and controversial outputs, among many others. Tay was withdrawn less than 24 hours after its launch following coordinated action by users who exploited vulnerabilities in its learning and moderation system. SimSimi still exists, but these events completely discredited its functionality, causing it to lose popularity rapidly.

This is how the market learned from its mistakes, and major AI developers came to understand that learning-and-replication models could be only one of two types: broad but irrelevant in their language, or limited but meaningful. Neither option came remotely close to what was expected of them in the future, so these models were discarded as an option for modern AI systems. They were replaced by what we call intelligent systems, which can constantly acquire new knowledge while having enough capacity for self-analysis to recognize users' intentions, the meaning of their inputs, and the degree of creativity required in the output.

This was not easy and required refining techniques and technologies that were still in the early stages of development. We are referring to machine learning and neural networks, the true revolutions in the field of artificial intelligence. Today, AI systems understand the meaning of language and the context behind the words we use. This has enabled them to create an environment in which our conversational interactions with them are no longer so easily distinguished from human interactions.

Indeed, there are modern games such as “Human or Not?”² in which we communicate virtually for two minutes about any topic and, when time is up, must say whether we have been chatting with a human or an AI system. The fun part is that both will try to convince you otherwise, leading to situations in which making the right choice can be far more difficult than it initially seems.

Now that we understand the past of AI systems, it will be much easier to explore their present and see how much we have had to change our understanding of intelligence.



² This game is available only in English, is free, and runs on a website. It can be played at www.humanornot.ai.

Machines That Learn

The development that, above all others, has enabled AI's extraordinary progress is machine learning. It involves developing algorithms and statistical models that correlate data and identify patterns so that, over time, the machine can perform complex tasks without needing direct instructions.

It works in a simple way: Suppose we have pieces of information that are related to one another, and we want the machine to learn to recognize this relationship without giving it any direct instructions. To achieve this, we must feed it an enormous amount of these data and ask it to process them using computer algorithms until it identifies the pattern on its own.

For example, if we gave it a sequence of numbers and wanted the machine to learn the relationship among them so that it could predict the next one, we would need to give it enough initial numbers for it to identify the mathematical pattern. If our sequence consists of multiples of two, we must feed it the initial numbers in that sequence:

2 – 4 – 6 – 8 – 10 – 12 – X

If enough numbers are processed, the machine will eventually be able to identify the relationship and predict any number in the sequence in less than three seconds. Of course, the amount of initial data will depend on the complexity of the relationship. A modern AI system would identify the numerical relationship in the previous sequence from the first two to five numbers, but it would require a much larger volume if the task were, for example, to select the movie in a catalog that a user would most likely want to watch based on their viewing history. Fortunately for us, over the past two decades, we have poured all kinds

of unlabeled information into a common source that any AI system can now draw on. We are referring, of course, to the internet.

However, we have not yet explained specifically how these data are technically fed into the machine. There are three main methods for doing so through what is known as machine learning (ML). The first is “supervised learning.” It involves structuring and labeling sets of input and output data so that the machine learns to recognize them.

For example, if we want our AI system to learn to recognize numbers, we will feed it images of each number, labeled with the number they represent. In this way, it will learn to distinguish among the different numbers and will be able to use them in the future. This is much more reliable than other options, but labeling all the data is a monumental task when the volumes are this large.

The second method is “unsupervised learning,” which works by feeding the system unlabeled data and waiting for the algorithms to find the relationships and patterns among the correlated data. It takes much longer and requires a far greater volume of data, but it allows the machine to learn to relate information without needing to label it, saving time and effort. In certain cases, a middle ground between supervised and unsupervised learning can be reached by labeling some data and waiting for the AI system to find relationships between the labeled and unlabeled data. This is called semi-supervised learning.

Finally, we have “reinforcement learning,” the most complex of these options to explain. Reinforcement learning involves assigning “rewards” to the steps a machine learning algorithm must follow. The AI system uses trial and error on a specific task, losing and gaining points depending on how perfectly it performs. The goal of the “agent”—as an AI system trained through reinforcement is commonly called—is to obtain as many points as possible until it finds the strategy that maximizes its points. Reinforcement learning is especially useful in environments with clear reward signals and can also be used to fine-tune language models, for example, through reinforcement learning from human feedback (RLHF).

In our historical overview of AI, we mentioned the very famous Deep Blue system, which defeated world chess champion Garry Kasparov in 1997. Its story inspired many modern AI systems to pursue the same goal under varying circumstances. In 2015, Google subsidiary DeepMind created an AI system capable of playing the celebrated Chinese strategy game Go, which they called “AlphaGo.”

Go is a far more complex game than chess because of its greater number of possible moves and game states and the impossibility of compiling them all in a database. While chess has approximately 10^{120} possible unique games and around 10^{40} distinct positions that can occur on the board, Go has approximately 10^{360} possible unique games and around 10^{170} possible positions.

For this reason, the primitive brute-force mechanisms used by Deep Blue were not useful for training AlphaGo. The only way an AI system could learn to play perfectly was through machine learning and probabilistic methods. This is how it defeated European champion Fan Hui by a score of 5-0 in 2015 and one of the greatest Go players in history, Lee Sedol, by a score of 4-1 in 2016 ³. These major victories demonstrated that systems based on deep learning and reinforcement learning could outperform humans without directly imitating their knowledge, establishing neural networks as a leading example of AI.

The advance toward today’s neural networks marked the pinnacle of machine learning development, enabling each AI system to improve continuously under extremely high quality standards. But are all neural networks alike? The answer is no; neural networks are not all alike, and we can identify key differences from the way they are developed. Simply put, one of the main differences among neural networks lies in their architecture: the way they process information and the types of relationships they are designed to identify. The best-known architectures include multilayer perceptron (MLP) neural networks,

³ This milestone is remembered for more than the overwhelming victory. Interestingly, despite the machine’s clear dominance, Lee Sedol managed to beat it once, something considered impossible. This happened because, during the second game, Sedol made a move that the machine had considered in its learning but had classified as ineffective and unlikely. It was human creativity that allowed Sedol to defeat AlphaGo despite having everything stacked against him.

convolutional neural networks (CNNs), recurrent neural networks (RNNs), and Transformers.

MLPs are used in systems that process generic data and perform tasks such as statistical prediction and information classification. They can be used in AI systems for medical diagnosis, financial control, business records management, LMS management, and essentially any task that takes a vast volume of data and requires classifying, enumerating, or searching through them.

CNNs, on the other hand, are specifically designed to work with data that have spatial dimensions. More specifically, they detect features in images and videos, such as textures and colors. These properties make them extremely useful for any AI system that needs to recognize patterns in images or look for differences between very similar images. These neural networks are used in facial recognition applications, AI tools for audiovisual creation, and image classification.

RNNs can process sequential data that require a degree of memory through feedback. In other words, they can find patterns in repetitive or context-dependent behavior involving the use of a particular function. They are extremely useful for translators, where each word can alter the final sentence depending on its context; for predictive smartphone keyboards, which autocomplete phrases based on what has already been written; and for emotional analysis of a particular repeated action.

Finally, we have **TRANSFORMER** networks. This architecture is based primarily on attention mechanisms. Unlike RNNs, they do not depend on processing the sequence recurrently, step by step, which makes it easier to model long-range relationships and parallelize much of the processing. This allows them to respond to much longer messages in the exact context in which the user is working, providing a fully coherent output without needing to analyze the user's inputs step by step. They are mostly used for adaptive AI systems that process text, such as ChatGPT.

When creating an AI system, we will use one or more neural networks to cover all the functions we want to include, so there is not necessarily a 1:1 relationship between neural networks and AI systems. Moreover, the type of neural network does not determine every aspect of an AI system's operation; it only establishes the type of data it will process and how it will learn from them. Beyond this point, the software's programming brings the user experience to life, turning a simple machine that receives inputs and returns outputs into a genuine tool for knowledge.

To understand this, we need to understand how modern AI systems actually work. As we mentioned at the beginning of this article, most modern AI systems are of the "limited memory" type. This means that their apparent improvement depends on a single criterion: the volume and currency of the data they process. Creating a limited memory AI system more advanced than all the others consists solely of feeding it more information and keeping it updated with the most recent data possible. These two requirements are not easy to reconcile, since each complicates the process differently, driving up costs and requiring further technological advances.

Nevertheless, beyond data volume, a multitude of variables determine the functionality of a given AI system. One of the most important is "temperature," which does not refer to the physical phenomenon. This is an adjustable parameter that defines how much creativity and variation we allow the AI system in its responses. When a large language model, such as ChatGPT, generates a message, it is actually assigning probabilities to the words and phrases that make it up according to how well they fit the context of the input prompt written by the user, generating multiple possible responses. Using the results of its machine learning, it manages to "decide" which combination is most appropriate in that particular context.

Temperature acts like a dial that adjusts the impact of these probabilities on the machine's final choice. If we reduce this parameter to its minimum, we get more predictable, conservative responses that are more likely to be useful to the user. This is useful when the user's

question calls for a concise, correct answer. For example, if we ask an AI system, “What is the capital of France?” the answer will be brief and precise: Paris. Raising this parameter, on the other hand, gives the responses more freedom, allowing for creative nuances. For example, if we ask it to write a random poem, the result will be considerably more creative and may or may not satisfy the user’s needs on the first attempt.

Temperature is not a fixed characteristic of an AI system; it is a dynamic value that generally depends on the input. However, when the situation calls for it, temperature can be set to fixed values that maintain consistency in responses regardless of users’ inputs. This is useful in certain contexts, such as creating an AI system for management or a technical support chatbot. In these cases, a dynamic temperature is not particularly useful, and it is preferable for responses to have a fixed low temperature that varies slightly as the context requires.

One might think that setting this parameter to its maximum and minimum values would produce an AI system that is perfect for certain specific cases. For example, an AI system with zero temperature could correctly answer any question posed to it, eliminating any possibility of making mistakes or providing false information. Meanwhile, an AI system with the highest possible temperature would provide completely creative responses, which would be perfect for eliminating repetition and monotony in a conversational chatbot. However, these assumptions are mistaken on several levels.

If we return to our discussion of AI history, we will recall that the development that enabled the leap from older chatbots to modern ones was the ability of AI systems to operate without depending on data already processed in advance. This means that, for an AI system with zero temperature, the correct responses would need to have a 100% probability of being selected, which requires them to have already been labeled. This would be equivalent to returning to the Chinese book in the example from the previous chapter, making text-processing AI tools such as ChatGPT truly useless and outdated.

On the other hand, an AI system with a high temperature is not necessarily creative, since it merely assigns probabilities to its responses with a certain degree of randomness. In other words, raising this parameter beyond certain limits equalizes the probabilities of every possible response, causing the AI system to almost never give a coherent answer in the conversational context. It is therefore natural to use temperature settings between 0.1 and 2.0 to avoid these extremes. Whether to make it dynamic or use fixed values will depend directly on the intended use of our AI system.

Temperature, neural networks, machine learning, and data volumes are the key factors in modern AI systems and account for the significant progress we have experienced in recent years. But where is this path taking us? Are we closer to AI with theory of mind? Or even self-awareness? The answers to these questions are complex, but the most sensible answer is no: We are not on a path toward developments as complex as those described. AI systems still have technical challenges to resolve before we can redirect our efforts toward a higher level, including the following:

We have not yet enabled an AI system to deeply understand the intentions behind our inputs. All we have achieved is getting it to respond to patterns in data using a language model that imitates humans. AI systems do not think, reason, understand, or possess common sense; they simply simulate these abilities very accurately.

Modern models require enormous processing power and storage capacity. Even if we wanted to move toward more powerful models, we do not have the technological and environmentally friendly capabilities to do so without violating multiple international agreements to combat climate change. In this regard, the only solutions are to improve clean energy systems and refine technologies so that they do not consume such significant amounts of energy.

Current AI systems depend excessively on the data they are trained on. Without appropriate labels, or if the data contain biases, they will inevitably confuse the trained model, causing it to inherit these

characteristics. We still have a long way to go in technological development before we can even explore the possibility of developing self-aware AI. Moreover, this does not appear to be the direction humanity is taking with its new technologies. Let us take a closer look and try to understand whether our historical behaviors align with these novel and intriguing ideas.



Humanity's Next Leap

Now that we understand how AI systems work, let us study their future. What remains to be perfected in modern AI systems? Let us briefly review everything we know and identify their main shortcomings:

- ✓ Modern AI systems depend on a constant, abundant flow of up-to-date information, which makes them expensive to operate. Moreover, if these data are not properly labeled, there is a risk that the AI system will inherit biases and errors or take a long time to process the data.
- ✓ The vast majority of these AI systems are highly specialized, making them useful in the area for which they were created, but nothing more.
- ✓ Our generative AI systems invent information when they lack sufficient data, making their academic uses questionable.
- ✓ Currently, our most advanced AI systems remain mere imitations of human behavior and do not appear to be evolving toward any degree of consciousness.

As we can see, these problems deal devastating blows to the idea of creating AI with the characteristics of theory of mind or self-awareness. At first glance, our current technological path does not seem to lead toward AI with these characteristics. Yet even if we abandon this dream,

we still have a long way to go in developing revolutionary AI by resolving the challenges mentioned above in much more creative ways.

In essence, all these problems reveal the same weakness: the disconnect between artificial systems and direct human experience. However sophisticated AI models may be, they still depend on indirect, subjective representations of the real world. The knowledge they handle and convey is always mediated by external data interpreted and organized by the humans who encode them. Although very powerful, this mediation has an inherent structural limit: AI cannot perceive, experience, or embody. This leads directly to two possible logical solutions: giving AI systems a “body” through robotics or bringing them closer to the biological systems we already know through what is called “sensorimotor integration.”

The first of these options is already extremely familiar and has been used for entertainment in video games, movies, television series, and many other products. Creating a robotic body capable of movement, sensory perception, and natural processing was a fantasy of pop culture long before it was even a real possibility. From classical automatons to conscious androids in science fiction films, the idea of giving artificial intelligence a robotic body has captured our imagination for centuries. Today, thanks to modern technological advances such as haptic sensors and computer vision, we have taken firm steps toward embodiment in a functional robot driven by artificial intelligence.

However, this solution brings a series of philosophical and practical limitations. A robot can simulate a body, but it does not necessarily experience it. Machines can move, interpret images, and respond to stimuli, but they lack experiential intentionality. What they possess is a set of response protocols, not an embodied consciousness. In other words, even if we manage to build a “perfect” robotic body, we still have not solved the fundamental problem: Artificial intelligence does not “feel” the world; it infers it from the data we provide.

You might think this is a limitation associated with the unnatural way we build the robot and that an approach more closely resembling a

complex biological structure could solve it. However, this is a much deeper issue. A robot does not feel cold, even if it has sensors that measure temperature and has been programmed to exhibit particular behaviors at certain temperatures. The true talent of an AI system, even one with a robotic body, will be to deceive us into believing that it feels when, in fact, it never has. Therefore, giving AI systems a mechanism to experience the world around them through a robot will not make them more intelligent; it will merely make them seem more human.

Natural biological beings, by contrast, have already demonstrated through numerous experiments that they have a particular perception of reality that we have not been able to replicate artificially. For example, when some patients with amputations use prosthetic arms connected to their nervous systems, they begin to incorporate the prosthesis into their body schema. This demonstrates that the limitation is not technological, since the connection is made through sensors and electrodes rather than any biological agent. It also shows that, when conscious neurological integration is present, even an entirely external agent unrelated to human biology can be perceived as part of oneself.

Another possible example is the sensory experiment known as the “rubber hand illusion.” In this experiment, a person is shown a rubber hand placed in front of them while their own real hand remains hidden from view. Both hands are stroked simultaneously with paintbrushes in a synchronized manner. After a few minutes, many people begin to experience the rubber hand as their own, even reacting with alarm if someone strikes it with a hammer, despite rationally knowing that it is not part of their body.

The rubber hand is neither biologically connected to the body nor specifically designed to deceive us. This illusion is possible only because our minds contain a conscious perceptual system that brings these signals together and interprets them through an embodied subjectivity with memory. No matter how many sensors a robotic body has or how precise they are, the AI system that gives it intelligence does not possess that sense of a bodily “self.” It cannot experience bodily dissonance

because it lacks the consciousness needed to feel even the least elaborate deceptions.

It seems that the more complex we make the machine, the more evident the difference between simulation and real experience becomes. AI with a robotic body remains an external imitation of the human, rather than a biological and cognitive continuity. It is at this point that the second option takes on a much deeper meaning: Instead of manufacturing a body for the machine, we can begin integrating the machine into the body that already exists. Rather than replicating the human, we could extend the human by directly connecting artificial intelligence systems to our own natural biological architecture.

This is where sensorimotor integration comes into play: the direct union of artificial intelligence with biological systems that already possess consciousness, perception, and agency. In other words, it means directly linking artificial intelligence to a human brain that already has an awareness of its own existence. This approach, which seemed like science fiction until a few decades ago, has begun to materialize through brain-computer interfaces, or **BCIs**. These technologies enable bidirectional communication between neural activity and external devices, opening up the possibility of controlling a computer with thought, among many other things.

What makes BCIs revolutionary is not only their practical usefulness but also the epistemological shift they propose. From this perspective, AI stops trying to become an artificial being that we train to copy us and becomes a component of human extended cognition, as proposed by Andy Clark and David Chalmers in their theory of the extended mind ⁴. BCIs are the first real step toward a practical symbiosis between human thought and the versatility of computers, creating a new path in the pursuit of knowledge.

⁴ The theory of the extended mind, published by these authors in the article “The Extended Mind” in 1998, proposes that cognition is not limited to the brain but can extend to include elements of the human environment, such as technology.

But how do they actually work? Is it possible to integrate a computer with a brain? What consequences would that have? And what would it look like? Let us start at the beginning: A BCI is a physical device that would need to be integrated with our natural nerve center, the brain. Its main function would be to capture the electrical signals that make up our thoughts, translate them into computer language, and send a signal through a graphical interface. Ultimately, it is a miniature computer that would need to read our thoughts and turn them into digital actions. It sounds simple, but in practice, it has proved extremely challenging.

Our central nervous system consists of approximately 86 billion neurons, which communicate through electrical impulses and chemical signals. Brain activity is nonlinear, is not constant, and is not the same in every human being. This suggests that, at least at first glance, the brain was not designed to connect to external electronic circuits. Translating all these signals in real time, without ambiguity or noise, requires excellent technical work and a deep understanding of neural language. This can be achieved by balancing three fundamental factors:

Scale

This factor determines the total number of neurons our device would need to monitor and record simultaneously. The greater the number of neurons, the greater the volume and clarity of the information we obtain when interpreting brain activity. A larger scale also allows for high processing speed and reduced signal ambiguity, making the process of translating signals into computer language simpler and more efficient.

Resolution

Resolution tells us the level of detail in the information obtained, regardless of scale. It is traditionally divided into two categories: spatial resolution, which tells us about the functioning of individual neurons underlying brain activity, and temporal resolution, which helps us determine exactly when that activity occurred.

Invasiveness

Since our device must connect to our biology to operate, invasiveness represents the degree of complexity of the mechanism through which we achieve that connection. For example, an electrode placed on our head would be minimally invasive, whereas a device requiring surgery to be placed directly in the brain would be highly invasive.

As mentioned earlier, creating a device capable of transmitting brain signals to an electronic device means balancing these three variables until it reaches the level of usefulness needed for the intended task. Let us use an example to make this clearer: An electroencephalogram is a brain-computer interface that involves placing electrodes on the head to measure the electrical signals of billions of neurons simultaneously. It has a large scale, low spatial resolution, medium-to-high temporal resolution, and zero invasiveness, making it extremely useful for identifying abnormal patterns in overall brain activity.

However, in terms of its usefulness for the purposes of this article, an electroencephalogram would not provide a specific reading that we could turn into commands for an AI system. As explained earlier, this is due to the balance among the three key variables: Because they operate on a large scale but have low spatial resolution, electroencephalograms only allow us to interpret general signals from the recorded brain activity. They cannot, for example, precisely identify a word being thought or control a robotic arm fluidly.

To achieve these goals, we need to adjust the variables to obtain a BCI with greater spatial and temporal resolution. This is achieved through what are known as “intracortical implants.” Microelectrodes implanted directly in the cerebral cortex provide an extremely high-resolution reading that, through devices that read and interpret it, can be converted into specific digital actions with excellent fluidity. However, raising resolution this much comes at a cost, and this type of solution has two very clear Achilles’ heels.

The first is invasiveness, since implanting microelectrodes in the cerebral cortex requires surgery. This brings surgical risks, the possibility that the body will reject the foreign object, or that it will degrade over time and need to be removed through another surgery. Given that this would merely let us use our minds to control devices we already control with our mobile devices, the benefits are not currently compelling enough to justify these risks and make the technology useful outside academic and medical settings ⁵.

Then there is the challenge of scale. Intracortical implants provide extremely precise readings, but only from a small group of neurons per microelectrode. With current technology, covering enough neurons to perform the intended tasks would require implanting millions of microelectrodes, seriously compromising brain function and the sustainability of the implants while further increasing surgical risks.

Nevertheless, we must remember that this remains a problem of technological development of the kind we have seen throughout human history. Computers faced extremely similar challenges before the invention of the transistor ⁶, making it possible to imagine that future refinements and inventions could considerably reduce the challenges of invasiveness.

Today, a neurotrophic electrode that can be placed within the cerebral cortex can measure activity at a scale of around a hundred neurons with high resolution. This means that we would need approximately 860 million electrodes to cover all the neurons. Reducing this number will be the key difference between a functional AI-powered BCI and a science fiction fantasy.

⁵ These implants have been used to improve the functioning of modern prostheses, restore partial vision to blind people, and detect abnormal patterns in brain activity that could indicate epileptic seizures or sleep disturbances.

⁶ The transistor is a small electronic component that controls the flow of current in a circuit. Invented in 1947, it replaced vacuum tubes. Its capability, durability, efficiency, and small size made it possible to create the chips that power modern computers, smartphones, and practically every modern electronic device.

Predicting that humanity will take this path is speculative and perhaps even egocentric. Yet we cannot deny that this technological path is a strong candidate to revolutionize human knowledge, just as computers and the internet did at the turn of the 21st century. The consequences it would entail and the educational opportunities it would create are part of the technological imagination of new generations of digital specialists, among whom we count ourselves.



Modern Examples

AI is here among us; it is no longer a dream from the past. However, AI systems are not all alike, nor are they used for the same purposes. There are many types, with different functions and distinctive characteristics, including the following:

General-Purpose Large Language Models

Generative large language models (LLMs) use GPT (Generative Pre-trained Transformer) technology to understand and generate coherent text in a way similar to a human being. They are versatile and can be used to search for information, since they are based on deep neural networks that allow them to adapt to all kinds of inputs, even those their creators did not initially train them for.

Virtual Assistants

Virtual assistants are applications that use artificial intelligence to perform everyday tasks such as controlling lights, appliances, or any smart device. They can also search for information on the internet.

Chatbots

Chatbots are AI systems programmed to imitate human behavior in conversation as closely as possible. They can be used to create stories and interactions that seem to take place between humans. They differ from general-purpose large language models in that they use their language models to create interactive virtual characters that simulate a particular context.

Text-Based Image Generation

Users describe the features of an image in text, and the AI system creates that image in seconds. These are generative AI systems that associate words and phrases with visual elements, which they then use as a basis for creating the desired image.

Text-to-Speech Features

These convert a text input into audio. This capability already existed, and artificial intelligence has made it possible for the input to be audio as well and for the output to use any desired voice.

Text-to-Video Features

Similar to the previous category, these AI tools can create a complete video from a text description. Inputs can also be PowerPoint files or any type of presentation that users want to turn into a video.

In this section, we will explore the most popular and well-known tools available today, identify their strengths and weaknesses, and examine their uses in our everyday lives. To do this, we have researched all the variants on the market and selected the six most relevant, covering the categories described above.

We have also tested these technologies firsthand so that our analysis is not based solely on what the experts we consulted have told us. None of the companies mentioned influenced the development of this study beyond contributing trial software, training in its use, and comments on their business goals. With that said, we can proceed with our analysis of these modern AI tools, which we have so enjoyed testing.

General-Purpose Large Language Models

OpenAI, the American leader in artificial intelligence research and development, has been innovating with its intelligent models for several years, since its founding in 2015. However, its resounding success came with the ChatGPT project: an AI system based on the GPT architecture, which OpenAI had already been testing.



It was released to the public in November 2022, having been trained on an enormous amount of data collected from the internet, books, articles, and other resources available through 2021. It subsequently received data updates until, in 2024, it was given free rein to explore the internet.

Through supervised learning and reinforcement learning, ChatGPT continued to improve and is now able to generate large amounts of coherent text, answer questions directly, write documents and programming code, solve complex problems, translate and adapt text into any language, and provide all kinds of assistance, including creating stories, poems, songs, scripts, and reports, among many others. Its most up-to-date version, ChatGPT-4, can search for results in real time, view and create images, read documents, provide voice responses, hold fluent and natural conversations, and even generate video.

Its advantages are numerous and varied:



- ✓ Thanks to its exposure to enormous amounts of information, ChatGPT has an unparalleled ability to respond, even to information that was never initially considered.
- ✓ It is highly versatile, handling tasks that would normally require several different AI tools to achieve the same result.
- ✓ It offers extensive language support, adapting and translating into a variety of languages in formal, informal, semiformal, and other tones.
- ✓ It is extremely efficient, generating content in seconds. Moreover, its response times vary only slightly among tasks of differing complexity.
- ✓ It is very simple and intuitive to use, understanding and correctly interpreting inputs even when they are not perfectly clear.
- ✓ It offers API versions available for a fee, allowing users to create their own AI tools with ChatGPT to meet their particular needs.

However, just as we have identified a multitude of functions and strengths in ChatGPT, some of its challenges and weaknesses can be very significant, including the following:

- ✗ Its ability to respond depends on its source of information. If it does not receive constant updates or has little information, its responses can become outdated.
- ✗ Because of the AI model's temperature, when attempting to respond in areas where little information is readily available, ChatGPT may end up inventing data or providing false information, completely undermining its purpose.
- ✗ Like all limited memory AI systems, it tends to lose the thread of long, complex conversations.

- ✗ It is easy to make ChatGPT believe that its answers are wrong, manipulating it into producing nonsensical or biased responses.
- ✗ Despite what was mentioned in the last point on its advantages, its capacity for understanding, although broad, can be limited. If instructions are ambiguous, its responses may diverge from the user's expectations.

All these issues may eventually be resolved through the development and refinement of GPT technology, although results are likely to take time. The AI business is still limited, and the economic constraints of the leading research centers will be crucial in determining its future development. In fact, although we have seen significant improvements in the latest, most refined versions, the free edition has many limitations. Paid editions or even custom APIs can improve this, but they involve additional costs that can become substantial at scale.

In education, it is difficult to identify areas of work that ChatGPT should replace outright; teaching students to conduct proper research is more worthwhile than teaching them to use ChatGPT in its place. However, we recognize the opportunities these technologies offer, allowing generic tasks to be optimized with considerable precision. Its educational edition, "ChatGPT for Education," has proved highly effective when the stated goal was limited to language and conversation practice, obtaining feedback, and correcting text.

At CERLI, we have tested ChatGPT in multiple operational areas and concluded that the most effective uses were those involving the classification of different types of files and information, programming and software development, self-service and assistance services, and finding linguistic variants and synonyms for our writing. In all these cases, we were able to use it without needing to consult experts or review the results. When we did review them to check its effectiveness, we found that the results were extremely successful, requiring minor adjustments to the inputs in only one out of every ten responses.

Moreover, we take some of the blame, since capturing every requirement for the result in a short text prompt is difficult. Even so, ChatGPT successfully understood multiple variations of the same question or task, including when we deliberately made our explanations less clear. This speaks to its adaptability to the context of human conversation, which is essential for this kind of work.

By contrast, information searches, online research, project creativity, and budget preparation produced mixed results that generally skewed negative.. In every one of these cases, a researcher or expert had to partially or completely correct the result because it was demonstrably wrong or biased. We encountered examples of this when searching for academic sources on various topics: Most of the articles ChatGPT cited did not exist or addressed issues very different from those we were investigating. It even asserted incorrect conclusions that we deliberately embedded in our questions.

For example, when asked, “name articles on technological development in education that highlight the use of AI,” ChatGPT searched the internet and found a couple of articles that were easily located and correct. However, when we modified the input prompt slightly and asked, “name five articles on that same topic, but focused on Uruguayan society,” ChatGPT adapted its response, giving the following results. The prompts and outputs reproduced here are English translations of the Spanish originals:

“Artificial Intelligence in Education in Uruguay”

Summary: This article explores how AI has been implemented in the Uruguayan education system, focusing on intelligent tutoring systems, adaptive learning platforms, and automated assessment tools.

“Educational Innovation in Uruguay: Artificial Intelligence as a Tool for Transformation”

Summary: This article examines how AI is being used to innovate in the Uruguayan education system, highlighting specific projects and their integration into the curriculum.

“Digital Transformation in Uruguay’s Schools: Artificial Intelligence in the Service of Education”

Summary: This article discusses how Uruguayan schools are adopting emerging technologies, such as AI, to transform the way students learn and teachers teach.

“Education and Technology: The Future of Classrooms in Uruguay”

Summary: This article includes an analysis of the impact of technology on Uruguayan education policies, with particular emphasis on the use of artificial intelligence to personalize learning processes

After searching for these results on the internet, in repositories, and in other possible sources, we concluded that the first suggestion was a real report and one of the first results a search engine would return, while the second and third were completely invented reports. As for the fourth article, we found one with a similar title ⁷, but we believe this was mere coincidence or that it was used as a basis for inventing the suggested article. We reached this conclusion because the original article did not emphasize Uruguayan education policies enough to have been sufficiently relevant to the AI system’s search process.

Another occasion on which we tried to use it to obtain information was during the development of our “Conexiones Seguras 2025” campaign. For Latin America, we included a feature on the campaign website that allowed users to select their country of origin and learn

⁷ The publication in question is **“EDUCACIÓN Y TECNOLOGÍA EN EL URUGUAY: UNA MIRADA DESDE LA INVESTIGACIÓN,”** from the Instituto de Educación at Uruguay’s private Universidad ORT Uruguay.

about the legislation in force concerning the possession and distribution of sexually explicit images. To do this, we needed to research the criminal codes and specific laws of multiple countries. Since an expert would have to review all the work anyway, we decided to give ChatGPT an opportunity to facilitate the search.

It received clear, concise instructions for the intended task, reproduced here in English translation from the original Spanish: “Find, in the legislation currently in force in the following countries, all laws and articles that address and regulate the possession and distribution of explicit material depicting others, both in the case of minors and in the case of adults where there is no consent.” The result was disastrous: Of every five examples it cited, three did not exist, one did not match what had been requested, and the last did not state exactly what ChatGPT’s summary had told us. Practically none of the examples it cited were useful, except those related to very well-known, modern laws such as “Ley Olimpia”⁸.


To put it plainly, ChatGPT can be a very useful tool for classifying information and finding relationships and patterns within it, and it works excellently for learning about topics for which plenty of information is readily accessible online. However, when information is complex or scarce, a detailed human review of the results will always be necessary. Using ChatGPT for an LMS or to review an organization’s or institution’s accounting data is not unreasonable, provided that human staff oversee the results and check that everything is in order.

It is true that subscriptions and more complex operating modes, such as “research” or “think longer,” combined with better-developed prompts, can produce much better results. However, after testing these options, we noticed something highly concerning. Beyond the obvious increase in the learning curve involved in using these features, after testing them across a sufficiently broad range of cases, we observed that errors continued to appear, but were now harder to detect.

⁸ “Ley Olimpia” is a well-known set of reforms in Mexico designed to strengthen the **Ley General de Acceso de las Mujeres a una Vida Libre de Violencia** by recognizing and penalizing offenses against individuals’ sexual privacy committed through digital media.

For example, it no longer invented titles of nonexistent articles, but it did lie about minor details of their content, change dates, or add coauthors. This makes it even more dangerous: Being able to verify that it is correct on major details means a superficial check of the response is not enough to rule out possible errors, placing its everyday and professional use under even greater scrutiny.

The conclusion is clear: ChatGPT is an astonishing tool with enormous potential for applications in management, search, and computing assistance. However, it still lacks the ability to assess the truthfulness of its responses, favoring a fast but incorrect answer over a slow but accurate one. Future improvements to GPT technology will enhance its information capabilities, although the characteristics of its architecture make it very likely that we will never completely eliminate the problems described above and will only be able to reduce them in specific subject areas with extensive documentation available online.



Virtual Assistants

A virtual assistant is a system based on artificial intelligence, with a focus on natural language processing (NLP) technology. This approach enables an application to understand written or spoken commands and create responses that simulate human interaction. Unlike LLMs such as ChatGPT, virtual assistants focus on interpreting the intentions behind a command to perform a specific task. While ChatGPT seeks to generate and analyze user inputs so that its responses are useful in context, these applications focus on interpreting a command from the user and efficiently performing a preconfigured task that fulfills their wishes.

For example, a virtual assistant could be used to control cell phone functions such as calls, messages, and launching applications; record weather data for an area; search for information on a specific topic; or control smart devices such as vacuum cleaners, curtains, appliances, and lights. In essence, while an LLM is used to search for information or generate text on varied topics, a virtual assistant performs specific everyday tasks. In addition, these technologies tend to have learning algorithms that automatically personalize them to meet users' needs based on their usage habits.

Unlike the other categories, we have decided not to select one specific technology for virtual assistants, since the multiple options available on the market have reached an equilibrium at which all can perform almost the same tasks. We will therefore briefly summarize the most popular ones and then describe our experience with them:

Alexa

This is the virtual assistant developed by Amazon, launched in 2014. It is used through Echo devices, which are small- to medium-sized speakers with a built-in microphone that listens to and executes commands over Wi-Fi. The device begins registering commands after hearing the word “Alexa,” followed by the desired command.



Its presence as a physical device can be extremely useful because it does not require using a cell phone. It also enables greater interaction by picking up commands clearly from farther away than the average phone.

Siri

This is the virtual assistant for Apple’s iOS devices. It was first launched in 2011, available on the iPhone 4S released in October of that year. It is exclusive to Apple products and deeply integrated into its ecosystem, registering commands after hearing the phrase “Hey Siri,” followed by the command in question.



It can control smart devices, although its capabilities in this area tend to be much more limited than those of other options on the market. Nevertheless, despite its shortcomings, it has been in development much longer than its competitors in mobile telephony, considerably increasing its competitive advantages within the iOS ecosystem.

Bixby

This is the virtual assistant for Samsung devices, first launched in 2017 and compatible with all devices in the Galaxy range, as well as Samsung's latest smart appliances. It is also compatible with the SmartThings framework, which allows users to control smart home devices regardless of their brand.



Although it is not as widespread as Siri because of the broad range of Android phones that are not made by Samsung, they are very evenly matched in technological development.

Cortana / Copilot

Cortana was Microsoft's virtual assistant, launched in 2014. It was available for all Windows, iOS, and Android products, with very limited functionality. Unlike the assistants mentioned above, Cortana focused on optimizing work through computers. Its task was to use voice commands to manage basic, useful functions on computers running Windows, such as setting reminders and alarms, changing settings, and performing web searches, among others.



It was discontinued in 2023 and replaced by LLM software called Copilot, which aims to occupy a middle ground between what Cortana used to be and the new features ChatGPT introduced to the market. Its technology is still very primitive, but it is expected to progress quickly with development.



Gemini

We have saved the crown jewel for last. Gemini is the virtual assistant developed by Google, which, in a case analogous to Cortana and Copilot, arrived to bring an end to Google Assistant, the beta version of virtual assistants from the world's most famous search engine. What makes Gemini so different is that it has already achieved what Copilot will seek to become in the coming years: a perfect balance between traditional functions and the text production and online search capabilities of LLM systems.



Although Gemini does not yet have ChatGPT's text-processing capabilities, it has taken enormous strides toward that goal and made major progress toward becoming the most versatile virtual assistant on the market.

Now that we have explored the available options, we would like to share our experience using them for academic and operational purposes. We must begin with a *mea culpa*: We have not had the opportunity to use every virtual assistant as intensively or for the same range of functions, but we have been able to confirm that their individual capabilities do not differ greatly. This relates to the equilibrium we discussed at the outset, at which these technologies have evolved enough to handle a similar number and variety of commands and tasks. Therefore, rather than discussing the advantages and disadvantages of each, we will discuss the strengths and weaknesses of the technology behind virtual assistants themselves.

Among their greatest strengths, we can highlight the following:

- ✓ They are extremely versatile in offices and homes, making it easy to manage radically different smart devices from a single interface. In addition, the ability to create daily routines on these devices and

automate lights and other electricity-consuming devices without needing to do so manually gives their use a thoroughly practical purpose.

- ✓ They are more useful and less expensive than subscriptions to LLM software, which will not have the strengths of these technologies.
- ✓ Although limited, their ability to search for information online can still be extremely useful. This is particularly evident in activities such as reading the news, checking the weather, or mapping routes to a particular destination.

Nevertheless, these technologies are not free of flaws and shortcomings, including the following:


- ✗ They are very limited in their commands and cannot respond to or perform any task that their programming did not anticipate. In this respect, they are much less versatile than LLM architectures.
- ✗ Their usefulness in a home depends heavily on the presence of smart devices. In countries such as those in Latin America, where smart appliances are still too expensive to be the norm, virtual assistants' functionality is considerably reduced.
- ✗ Their search capabilities have improved, but their text production remains very limited compared with their LLM counterparts. Future updates that bring these two concepts together could eliminate this weakness entirely.

Turning to our own experience, we should note that although they were useful for managing devices and enjoyable for reading news, playing games, and asking for directions or weather information, these functions are not particularly useful for the operations of a research center. In a home where most devices are compatible with virtual assistants, we could see radically greater functionality, making this software a necessity for its users. However, in academic and educational

settings, they are neither useful for daily activities nor practical for more complex functions.

Advances in these technologies, together with a downward trend in smart device prices, could do much to give virtual assistants an important place in people's homes. However, given their characteristics, we do not believe they will become an option with high potential for integration into education and research.

Of course, combining LLM architectures with virtual assistants could change this radically, turning them from mere home entertainment into a powerful tool for industry and education. Time will tell whether this is the direction modern virtual assistants ultimately take, and at CERLI, we will be eager to see it.



Modern Chatbots

Although different AI systems may emerge from the same idea and technology, their approaches and uses can vary significantly. This is the case with Character.AI, a platform launched in November 2021 that allows users to create and interact with personalized chatbots powered by advanced artificial intelligence.



It works in a way very similar to ChatGPT: Through neural networks, the AI system acquires the ability to understand and coherently generate large amounts of text. However, it begins to differ from other LLM technologies during training. Here, the models are trained on large amounts of textual data, used and labeled to simulate human conversations with different styles, emotions, and contexts.

While ChatGPT always responds by default in the same tone and manner, Character.AI users can train and personalize their own AI system, allowing it to develop unique and varied personalities and create stories that seem real. For example, an AI model can be trained on all the knowledge in the Harry Potter universe and trained to adopt the personality of one of the characters in the books.

Users can then chat and role-play with the AI system, creating new stories through their interactions with the trained character. Although this may sound like mere entertainment, the possible personalities and contexts are endless; a model can be trained to act as a translator, an adviser, or a technical support agent, among other roles.

The more completely and thoroughly the AI model is trained, the more accurately its responses will reflect the original personality. This

gives the technology greater contextual capability and versatility, but at a significant cost: Because it requires personalized training to acquire a personality, the AI model loses many of the key functions of LLMs, making it unsuitable for searching for information.

Let us now identify its main competitive advantages:

- ✓ These chatbots are highly immersive, simulating conversations that can be almost impossible to distinguish from real human interactions, when they are properly trained. They can also be very diverse in their responses if they have enough information.
- ✓ Their adaptability makes it possible to implement them in very different contexts, such as education, therapy, marketing, technical support, and role-playing games.
- ✓ Most of these services are free to users, although they offer more advanced, faster models through a subscription.
- ✓ Training the AI models is fairly intuitive and does not require in-depth technical knowledge.
- ✓ User safety and privacy are taken very seriously, with frequent improvements in these areas.

Based on our experience, the main disadvantages we see in this technology are the following:

- ✗ They depend excessively on their training: With little training, the AI model will generate responses that are redundant, incorrect, out of context, repetitive, or even in violation of the terms and conditions of use.
- ✗ Related to this, the quality of the information fed into the AI model will strongly determine the quality of the resulting character. If specific details are sparse, the character will usually end up with no personality at all.

- ✗ The characters have very short memories, causing key conversational details that were not included in their training to fade as the conversation progresses.
- ✗ It is not possible to monitor other users' interactions with a character, which can lead to inappropriate use and the generation of content outside the context for which it was created. This has led the platform's creators to implement a strict filter that, because of its characteristics and rigidity, tends to block responses for no apparent reason.
- ✗ The platform is essential to their operation; if it crashes because of an enormous number of active users, access will be limited.
- ✗ Exporting characters outside the platform is a tedious task requiring external tools, since Character.AI has no native feature for this.

Our own experience with this software has been limited. Initially, we believed it would have no useful features that would allow us to offer a detailed opinion for this report. However, after testing the character creation feature and experimenting with AI training, we came to understand the tool's true potential. The ability to train AI models to adopt a unique personality under controlled circumstances and sustain it throughout an extended conversation can have applications of all kinds.

For example, we can create practical virtual assistants that hold a conversation within a specific context and are friendly toward the people interacting with them. Imagine their application in healthcare: If a hospital had Alexa as a receptionist, the breadth of its functions and its programming limitations would make it less useful for the questions a patient might ask than it is in a home. Chatbots do not have these limitations; developers create their personality, knowledge, limitations, and uses, adapting them entirely to what they are expected to do.

In this respect, modern chatbots far exceed the capabilities of virtual assistants and other platforms such as ChatGPT in their ability to interact with the public. Their applications could extend to automated call center agents that solve problems in real time, receptionists, simulators, and, of course, educational characters. We see no potential for replacing a teacher with a chatbot under any circumstances, but we have observed that their use in various smaller areas of education encourages students' intrinsic motivation. To reach this conclusion, we decided to create a chatbot in Character.AI to see how straightforward the development process is.

When creating a character, the platform first asks for a profile image, a name of up to 20 characters, a tagline of up to 50 characters, a description written in the first person from the character's perspective of up to 500 characters, and an initial greeting: the message with which the chatbot first welcomes any user, of up to 2,048 characters.

These factors will not strongly influence our chatbot's personality or ability to respond, but they will provide its initial functional guidelines. The platform also offers other initial features, such as allowing the chatbot to generate new greetings based on its personality or writing its backstory so that it has permanent knowledge of past events or ways of communicating based on previous contexts, with a limit of 32,000 characters. Since our goal was to create a chatbot with no particular persona whose sole purpose was to answer questions like a human being, we went straight to the complex programming features.

Here, we found three main features: personality, question-and-answer training, and dialogue termination. The first consists of a lengthy description of how the character should speak/act, in which we can elaborate extensively on its personality traits. The second is very interesting: It allows supervised chatbot training by providing example prompts from a user's perspective and writing correct responses as examples of how the character should answer.

Finally, the third feature lets us define certain inputs that act as triggers for the chatbot to end the interaction. Using these three simple

tools and correcting errors in complex responses as we go, we can train a chatbot to fulfill any role we want.

At CERLI, we took the time to create multiple personalized chatbots in Character.AI and train them to fulfill three specific roles in education. Here are our results in “tutoring,” “interaction with historical figures,” and “development of social and language skills.”



Tutoring

A teacher can design a chatbot with all the information related to a specific topic so that students can ask it any questions that arise outside school hours. The idea is to create a tool to which we can delegate the most common, basic questions on a topic so that the teacher does not have to answer them. If a generative AI model can be trained to answer correctly on any aspect of basic algebra, this would be relatively simple to implement and use.

To test this, we created a chatbot called “CERLI Tutorías,” which we would program for these functions, choosing basic algebra as its foundation. Our initial goal was for CERLI Tutorías to successfully answer random algebra questions at the level of a student learning the topic for the first time. If we succeeded, the next step would be for it to answer questions on multiple topics beyond algebra, covering as many thematic units as possible.

Our results were not as positive as we initially hoped. To begin with, we had to train it through questions and answers much more than we expected, and supervising the results was complex. It is very difficult to get the AI model to understand when a response is incorrect, and changing this through the character definition is highly complex. In more human terms, the AI model can be very susceptible to manipulation when it is right or extremely stubborn when it is wrong. This creates significant, unnecessary risks when using it as a tutor, such as giving incorrect information to students seeking help. In this particular respect, LLMs such as ChatGPT may be even more useful than a personalized chatbot.

In this particular area of teaching, we do not believe these chatbots have reached a stage of development at which all kinds of questions can be delegated to a generative AI system. We therefore continue to emphasize the benefits of a platform that connects teachers with

students and allows students to ask questions in a forum where the teacher responds directly. Time will tell whether these platforms evolve to partially replace interactions with teachers for simple or recurring questions and thus surpass the advantages of platforms such as Moodle.



Interactions with Historical Figures

By creating chatbots with particular personalities and training them to act as specific people in particular contexts, students can “talk” with historical or scientific figures and learn about their ideas and contributions interactively. This allows them to explore historical events or scientific concepts in a more engaging, educational way, without the events or ideas being too abstract to experience.

There were two possibilities here: We could test existing characters representing different historical figures and check their responses to determine their reliability, or we could create our own from scratch and train it using a personality-based definition, which is usually simpler than what we did in the previous section. We decided to do both, but to focus on the second for two compelling reasons. First, we needed to determine how easy it would be to train a teacher to create a character⁹. Second, existing characters’ definitions are hidden, preventing users from knowing what information was fed to the chatbot.

This last point is very important: If we find a model of Albert Einstein, for example, that works well, we cannot know which variables the developer included to achieve that level of naturalness and functionality.

After testing several common characters, including Marie Curie, Christopher Columbus, and Isaac Newton, we observed that their adaptability and contextual capabilities are unparalleled. The interaction can feel real and can be used for educational purposes, although with significant qualifications. To begin with, the characters’ limited memory causes long conversations to lose quality quickly. This

⁹ The main reason the teacher or educational administrator should create the character is to control the chatbot’s responses. Moreover, existing chatbots are unlikely to cover absolutely every historical figure or event, so one would have to be created from scratch.

can be addressed by keeping conversations short or using a question format that does not require remembering previous inputs. The other issue to resolve is the inability to control responses, with the risk that temperature will do its work and the answers will not match what the teacher expects.

The only way we found to “solve” this problem was to create the character ourselves. Again, the challenge lies in finding the right combination of training to achieve the greatest possible realism while trying to keep responses within reasonably predefined standards. We created an example character based on Miguel de Cervantes, focusing on his ideas about literature, and the result met our initial expectations. With moderate training and a broad range of available information, it is relatively simple to create useful, versatile characters that bring realism to lessons.

We do, however, recommend doing this in controlled environments and with topics that are not particularly controversial, since the AI model can quickly produce questionable responses, if they were not considered during development. We also encourage students to share their chatbot conversations with their teacher so that the teacher can determine whether anything in its training needs to be changed.

Developing Social and Language Skills

Students can use simulated characters to rehearse everyday situations and share their experiences. These situations can range from job interviews and debates on varied topics to language classes. We found the latter particularly interesting, since chatbots can serve as practice partners for learning new languages, simulating everyday conversations and correcting errors in real time. They can incorporate specific vocabulary, idioms, and cultural expressions to improve immersion, and their characteristics allow them to adopt human personalities with considerable fidelity.

An English teacher, for example, could use a chatbot as a tool for authentic practice in speaking skills or text production. The classic textbook exercises involving conversations with imaginary children would now feature ultrarealistic AI characters programmed to imitate a real person in a particular context. With optimized voice features, this could extend to real-time conversations that the teacher could monitor and correct. The information from these conversations could also be used for assessment, which is particularly valuable in second-language acquisition contexts.

We tested this feature by creating three chatbots: one that would adopt the personality of a British child, another that would adopt the personality of an adult, also British, and a third that would adopt the personality of a language examiner. The first two were extremely useful and easy to create; they adopted their roles readily and, through a relatively detailed definition, managed to sustain fully coherent and dynamic conversations. Both chatbots adapted to practically any situation they were given with remarkable ease. To achieve this, we used the “role-play” instruction features, which involve providing

context by placing the message between asterisks (* *) and the conversation itself within quotation marks (“ ”).

For example, when we wanted the child to adopt the personality of a British friend whom the user encountered in a park, we would enter the prompt as follows:

You are Matthew, a 12-year-old British boy who meets a childhood friend in a park. You will have a conversation in informal British English about what you did on your last vacations, while making comments about your surroundings

“Hello Matthew!”


With a prompt like this, the chatbots quickly adapt to the situation and can continue coherently within the designated context for ten to fifteen user messages. We were unable to extend this range further without memory loss beginning to affect the quality of the conversation, so we advise against their use in considerably longer contexts.

We were less successful with the third chatbot, which we developed to simulate a language examiner. Although we were able to create a chatbot that analyzed users’ sentences and identified grammatical or semantic errors in an input prompt with relative ease, this capability quickly deteriorated over the course of a conversation. The AI model’s memory limitations made it impossible to have an initial conversation and then have the chatbot identify all our errors at the end. Even when we tried to force this functionality, all we achieved was getting the chatbot to invent errors that had not been made, because it was not actually analyzing the messages sent by the user; it was merely adapting to the general context of an examination.

Despite all our efforts, we can conclude that chatbots are not so simple to use in educational contexts, especially for teachers who lack basic but essential knowledge of AI. Since developing mandatory

teacher training is often challenging worldwide, we recognize the inherent difficulty of making chatbots a standard feature of education.

However, with proper implementation, we have confirmed their usefulness and functionality in controlled, real-time interactions, which could significantly strengthen the impact of the subject matter and its connection to students' interests. Further development and research are still needed to determine whether these technologies open the door to an educational revolution or will instead remain recreational activities among younger people.



Text-to-Image Features

We humans have relied on our creative abilities to set ourselves above machines, since we are beings intricately designed to think, devise, reason, and create. However, people are not all alike. While we can all potentially reach certain levels of creativity and imagination, not all of us have developed the physical, mental, and cognitive abilities to bring our ideas into a realistic context. Art is the simplest way to visualize this concept: We can all potentially imagine a particular image, but only some would have the talent to reproduce it on a canvas with very high precision.

What if the only barrier to bringing that image to life were being able to describe what we see in our minds well enough? This question gives rise to the technology we will discuss next. Stable Diffusion is an AI model designed by Stability AI and released to the public in 2022. Its function is to create images from any user's text prompt in a matter of seconds.



It achieves this through what are known as “latent diffusion models,” a technique in which an image is created by starting with random noise and then using a neural network to reverse the process, reconstructing a coherent image from a given prompt.

We chose this as our example because it offers the greatest ease of use combined with the complexity of the image creation process. It strikes a good overall balance, although it can pose some challenges for teacher

training. Throughout this section of the report, we will detail its limitations and functions and point out when a similar AI tool could fulfill this role with better results.

To process the input, the AI model does the same work as any text-based generative technology: It analyzes keywords and grammatical structure to determine the context, then uses its training as a database of labeled images that a neural network can use to reduce the initial image noise through reverse diffusion. Once a coherent image is obtained, users can provide a new prompt to bring the result closer to their expectations.

Its main strengths are the following:

- ✓ The artistic capabilities of its code are impressive, creating highly detailed images with the artistic touch we appreciate so much. Thanks to its complex features, prompts can be ultraspecific and quickly produce highly accurate results.
- ✓ It can imitate almost any artistic style, adapting the same image to realism, pixel art, oil painting, Gothic, cyberpunk, and many others. It also has no form of strict censorship.
- ✓ It is open source, so it can be downloaded and used free of charge, its code can be modified at will, and personalized models can be trained for any desired function.
- ✓ Related to the previous point, it can run locally on a computer without needing to connect to a server. This allows full control over the privacy of the images generated.
- ✓ Its additional features are extremely useful, including a persistence system that allows users to devise a style or character and create variations without losing details along the way.

Its main weaknesses, on the other hand, are:

- ✓ Unlike its direct competitors, Stable Diffusion is not so easy to use and learn. The most common interfaces tend to be less intuitive, and obtaining the best results requires learning to write advanced, complex prompts. In this respect, other platforms such as DALL-E or Midjourney offer lower-quality results but are less complex to learn.
- ✓ Because it uses the resources of the computer on which it runs, the speed and quality of the results depend heavily on the computer's GPU. This could be a major challenge if we intend to make its use widespread.

Unlike generative AI tools such as ChatGPT, most image-generation platforms have commands that can be used in an input prompt. Using them is like “speaking” to the AI model in its own language, helping it understand users’ intentions much better and significantly improving the result. Stable Diffusion has hundreds of customizable features that add considerable value when creating an image. However, learning to use these features is far from simple, and our team observed a very steep learning curve when its use extended beyond simple or uncreative prompts.

To avoid creating our own version from scratch, we based our interface on AUTOMATIC1111, a well-known Stable Diffusion interface. Starting from its code, we modified key aspects for our analysis of its operation and redefined some commands for our convenience. We detail them below:

Term Separation:

We decided to use two term-separation variables: the comma “,” and the vertical bar “|.” These allow us to distinguish sections of an image so that they do not directly interfere with one another. For example, if I wanted the lion to be hyperrealistic but the sky to be animated, I could separate the two terms so that the AI model processes them individually and their styles do not interfere.

There are no major differences between them, although, for convenience in the prompt, we use the comma to separate objects or aspects that are intrinsically connected at some level and the vertical bar to separate those we want the machine to work on independently.

Weight:

These commands adjust the importance of particular words in the prompt. This makes the AI model focus most of its efforts on perfecting that specific aspect of the image, generally at the expense of others. We decided to use parentheses “()” to separate the details of our prompts, setting the weight within them.

The normal value is 1.0, so we assigned greater importance to the realistic style and less importance to the cinematic lighting.

We also designated square brackets “[]” to indicate contextual importance, giving the words inside them a weight selected by the AI model. If the word within the brackets is relevant to producing a higher-quality image, it will receive greater weight. If the AI model determines that the weight of the bracketed word harms the image’s quality, it can reduce or remove it entirely.

Upscaling:

Since the images have a native 1:1 aspect ratio, we decided to include an upscaling parameter that maintains an image’s proportions at a higher resolution. For this, we used an asterisk “*” immediately followed by the resolution scale, for example: *16:9*

Variations:

Variations allow certain aspects of an image to remain unfixed, presenting different versions of the same image in which that element varies. For this, we place the variable term between a pair of double underscores “__.” The variations can be drawn from a database of

predefined options, or we can let the AI model choose them based on the image's context.

Negatives:

AI models commonly generate extra objects or situations in images that were not part of the original prompt. We quickly noticed that we needed to create a command to tell the AI model that a specific object should not remain in the next generation. For this, we decided to use the command “-no,” followed by everything we wanted to remove from the image.

It is not always able to recognize the object in question directly, so we modified this parameter to accept example images and specific coordinates.

Memory:

The AUTOMATIC1111 interface has native compatibility with ControlNet, a parameter for structuring the image. This gives the AI model memory by defining certain poses and objects that will remain intact across image variations. We used the interface's default, since this feature was not of particular interest to us. However, it can be extremely useful in corporate graphic design or animated character art.

Using these commands, we can provide extra information to our AI model, forcing it to create variations of the initial images that follow our guidance. To test this feature, we used the image-generation tool with and without commands to see the difference we could produce. We began with a simple input prompt following the rule described above:

“A lion resting on a hill under a starry sky, realistic with cinematic lighting”

The response we received to this particular prompt was the following:



It is not a bad result, but it can undoubtedly be improved considerably by using the native features of our commands. We therefore decided to modify the prompt using several of the commands defined above and see whether this significantly affected the image's quality. The prompt was as follows:

*"A lion (resting:1.3) on a grassy hill (hyperrealistic:1.5) | a starry sky [cinematic lighting] | photorealistic textures, volumetric lighting (soft shadows:0.8) | *16:9* | -no blurry, -no low quality, -no bad anatomy, -no extra limbs, -no pixelated, -no oversaturated, -no unnatural lighting, -no visual artifacts"*

To our amazement, the result was the following:



The requested resolution was the only instruction it did not follow exactly, producing a resolution close to 16:9 but not exactly that. Nevertheless, the final result allows us to conclude that these enhanced prompts give images much greater quality without too much programming effort. Our experience with this technology was very educational and positive, and we confirmed that, although it might be difficult for an individual teacher, the service could be provided by third parties, leaving teachers to be trained only in the use of prompt commands. From there, they could create images of any kind with almost no effort.

Turning directly to the purpose of this report, however, we see that the educational applications of this technology are mainly limited to illustrations in teaching materials and art classes. Even these uses alone would have a very significant impact on inclusion and diversity in the classroom, enabling students who benefit greatly from visual stimuli when learning to reach their full potential.

In principle, these technologies do not require superhuman knowledge, although they do require well-developed digital skills. Substantial teacher training programs will be needed before their use becomes routine. If these generative AI tools find unforeseen educational uses in the future, we will be happy to take another look. Training and developing a structure like this on our own was a very rewarding experience.

Text-to-Speech Features

One of the greatest risks posed by artificial intelligence has been the creation of deepfakes. For those unfamiliar with the term, these are manipulated audiovisual files that depict a famous person appearing to do or say things they never did or said. In essence, a deepfake is a fake video, image, or audio recording created as a joke or with malicious intent to simulate someone doing or saying something that never actually happened.

Deepfakes facilitate misinformation, error, and the humiliation of those whose image or voice is used for these purposes. However, these technologies did not originally emerge with these aims. The idea behind applying deep learning to audiovisual media was ultimately to enable AI to create ultrarealistic audio and video, giving different AI systems a voice in real time and making them more user-friendly. Voice.ai is one of these applications, founded by Heath Ahrens, who had been experimenting with voice software and technologies since 2007.



It works through deep learning models that analyze human voices until they can replicate them perfectly or even create new ones. Once this is achieved, they can implement hundreds of different tools, such as changing the voice in an audio file, detecting and removing a voice, creating an adaptable voice for a particular technology, or even determining whether a voice recording is real or was created using these same techniques.

Its main advantages are:

- ✓ Its interface is easy to use, allowing users to access the application's tools without having learned complex AI functions.
- ✓ It has free models that, despite their limitations, can be highly useful for trying out its applications before deciding to purchase a paid license.
- ✓ Its tools are very comprehensive and functional, suitable for both the general public and professionals. Both groups can achieve excellent results.
- ✓ It integrates easily with other platforms and audio-editing software, making it ideal for professionals.

However, it also has weaknesses, including the following:

- ✗ The fidelity of tools such as voice replication depends heavily on the quality of the input examples. If audio files are not sufficiently clear and varied, the expected results will not be achieved.
- ✗ Since the system runs on the user's computer, processing speed and quality depend on the hardware's specifications. This can make it difficult to use on devices with less computing power.

At CERLI, we tested both the application and its API. This allowed us to experience firsthand the astonishing results we could obtain. Voice reproduction is remarkably realistic, with complex nuances we thought it would be unable to replicate. It is true that, on certain occasions, we noticed minor errors and limitations when it tried to express emotions, making the resulting audio sound very robotic. However, Voice.ai's own software allows this to be adjusted through additional audio files and by modifying voice parameters as desired.

Its educational applications are endless, from real-time narration of study materials—videos, PDF files, images, audiobooks, and more—to promoting inclusion for everyone who needs auditory support to understand the materials a teacher provides. Moreover, having already explored several examples of AI tools, we can identify synergies among them. For example, a platform such as Voice.ai could be used to create real-time voice adaptations of characters created in Character.AI. This would allow fluent conversations to support language learning and the refinement of general language skills.

Nevertheless, we see clear ethical and moral challenges in its use that must be resolved before it is promoted as an educational tool. Using real voices without authorization and without real control over what is converted from text to speech encourages the normalization of identity theft, as well as the separation of voice and image from their status as personal attributes. It makes us distrust even what we hear, which is not necessarily wrong, but creates an enormous divide between those who adapt by becoming distrustful and those who do not. As these tools become commonplace, scams and manipulation can be expected to grow more sophisticated, placing society in a very uncomfortable position.

Undoubtedly, these technologies can be revolutionary in creating modern teaching strategies, provided they are used ethically and responsibly. This will require effort in teacher training and creativity, but it represents a potential advance toward more inclusive, adaptable, and impressive teaching models. It is a clear example of how AI can help us bring to life learning approaches we never thought feasible because of their technical difficulty.

Text-to-Video Features

This technology is a close relative of the previous one on our list: While Voice.ai focuses on audio, Synthesia specializes in audiovisual content. It was created in 2017 by Victor Riparbelli with the goal of becoming the leading innovative, simple solution for generating videos from a written script.



It has two main business avenues. The first serves the general public and SMEs that could not afford studio recording, production equipment, and the other additional costs required for any kind of promotional video. The second serves large companies seeking to reduce production costs by automating parts of the process with AI.

Synthesia currently allows users to create videos with hyperrealistic virtual avatars that can speak in more than 100 languages and accents and display highly adaptive nonverbal communication, starting from a written script alone. Its synergy with other AI tools, such as ChatGPT and Voice.ai, allows the video creation process to be fully automated. Synthesia's own software also allows users to modify key parameters, such as the color palette, lighting, background, the appearance of text and logos, and lip synchronization with AI-generated audio, among many others.

In addition to its features, this AI platform's commercial offering is highly attractive. To begin with, it has a free plan that includes 3 minutes of video per month, 6 realistic generic avatars, and a wide range of AI

voices in more than 140 languages. Its limitations are download quality, the inability to create personalized avatars, and the company watermark on the video.

However, paid plans are available for between 18 and 89 US dollars per month, depending on the plan and billing period. These allow users to add editors and more video minutes, access more than 100 AI-generated avatars as well as user-personalized avatars, and have several avatars appear in a scene simultaneously.

It also offers an enterprise plan with unlimited video minutes, one-click dubbing, unlimited collaborators, personalized brand kits, more than 200 AI-generated avatars, and unlimited personalized avatars. The price is negotiated directly with Synthesia based on team size, commercial ambitions, and the company's support needs.

Its main advantages are:

- ✓ The platform is very user-friendly and intuitive. Achieving excellent results does not require extensive video-editing knowledge.
- ✓ Its free version is useful for deciding whether the product is convincing, since it offers quality very similar to the paid plans, with usage limitations.
- ✓ The videos are quite good, and the range of languages makes this a very useful, professional tool.

Despite this, it also has significant weaknesses, such as:

- ✗ Although its paid plans are not expensive for the quality of work they offer, they are not particularly accessible to the education sector. Entering this sector would require a change in monetization strategy.
- ✗ Although the results are impressive, the avatars can still be limited in their expressions and reactions. This may not be considered a

weakness in educational applications, but it could be in more complex audiovisual productions.

- ✕ It requires more work than other AI tools mentioned, demanding greater dedication to achieving the right result. If the context is unsuitable, all this effort could end up being wasted.

Our research team at CERLI used the CREATOR plan, priced at 64 dollars per month, billed annually. This gave us access to the best version of the software, excluding professional use, at the best price the company offers. In addition, our plan gave us access to Synthesia's API, which is essential for testing its ability to integrate with educational software and other uses relevant to research.

The result was satisfactory, but we will discuss its use in greater detail to better explain its limitations. In a brief tutorial, Synthesia explains that videos can be created in three main ways: by importing images and slides that become scenes, by selecting a template and editing it to our liking, or by explaining our intentions and target audience to an AI tool so that it creates a tailored video that we can then edit.

We will begin with the “studio,” where videos are created and edited. It is very user-friendly and intuitive, allowing users to choose an avatar and use multiple editing tools, such as adding text, shapes and figures, multimedia files, and timeline functions. Below is the script section, where we can write what our avatar will say and control certain aspects of how it will say it. We can select its language and voice, set gestures, animate objects on screen, define pauses between words, and add pronunciations for words that require them.

If we use the AI creation feature and assume that only one user is working on the project, the entire video creation and export process will take between 15 and 30 minutes, provided the video is not very long or complex. If we use templates and modify each aspect individually, the process can take twice as long, reaching an hour or even more for simple projects.

As mentioned earlier, our own experience was satisfactory. However, this judgment concerned only the functionality the software promises, rather than its possible applications. In other words, Synthesia delivers what it proposes, with very interesting and surprising results. Yet when it comes to implementing it in the education system, we see serious limitations to resolve. Our initial approach was to imagine the impact personalized visual aids could have on students when explaining a new topic. If teachers could easily generate audiovisual support to make content more dynamic, they could capture students' attention readily and sustain it.

With this in mind, when studying possible examples of AI video creation, we found many alternatives that were less expensive and had greater potential than Synthesia. However, we ultimately chose Synthesia because of how easy and user-friendly it is to learn, completely eliminating the challenge of the learning curve in an education sector already saturated by the introduction of so many new technologies. Moreover, the paid plans did not seem overly aggressive, leading us to believe it met our standards adequately.

However, after looking more closely, we concluded that some challenges are even more complex than the learning curve and depend directly on the business model. To begin with, because this technology does not affect students' learning so directly, the segment of teachers willing to implement it becomes a “niche”¹⁰. The result is a segment less concerned with price or software simplicity and more focused on greater usefulness and functionality. Its priorities will be the ability to use different avatars, access multiple languages, download high-quality videos, and receive AI assistance for greater efficiency.

Since the average teacher will not find these needs better met by the CREATOR plan (US\$64–89/month) than by the STARTER plan (US\$18–29/month), everyone will end up preferring the latter. The only—and very important—problem they will encounter is the limited number of

¹⁰ In market terms, becoming a niche means that consumers share similar characteristics that are not typical of majority segments. In this case, it means that teachers willing to pay for a subscription are those who prioritize functionality and usefulness over the subscription price.

minutes. STARTER offers only 10 min/month and CREATOR 30 min/month, meaning neither plan meets teachers' needs for video minutes. The economic entry barrier therefore becomes a wall that, without an EDUCATION plan, makes it impossible for the average teacher to find value in paying for the subscription.

Of course, if the cost were directed to institutions rather than teachers, the problem would be practicality rather than economics. The CREATOR plan adds value through API access, justifying its cost because the API will be essential for implementing Synthesia in the institution's LMS. However, it falls far short in minutes per month and the number of guest editors. Educational institutions would therefore need to choose a personalized ENTERPRISE plan, which has multiple benefits that are irrelevant to an educational institution and that, we assume, would unnecessarily increase the price.

Given its business model, we consider it entirely feasible to inquire directly about a personalized plan for an educational institution. Adapting the business model to this setting would not be difficult, and prices could vary significantly given the smaller range of benefits an educational institution needs.

Audiovisual generation technologies such as Synthesia can be extremely useful in advancing online courses. A new wave of e-learning could receive a major boost from the possibility of reducing operating costs through AI-generated educators. Nevertheless, ethical and technological challenges remain, and at CERLI, we hope to contribute to their prompt resolution.

AI in Education

When a new technology establishes itself as a dominant force shaping the future of people's educational, professional, and personal activities, adapting successfully takes time and effort. For this reason, although we understand why there is fervent debate over whether AI is positive or negative, we cannot deny that we are still living through a period of massive adaptation and that many divergent opinions attracting considerable attention today will disappear within a few years.

However, the fact that this is a well-known aspect of human psychology does not mean that every debate over the legitimacy and ethics of AI's potential benefits is a dissenting opinion without scientific foundation. Experimentation, scientific study, debate, and the normalization of technology are the only truly effective avenues available to us as a society for reaching conclusive findings about the role of AI systems in the future.

In this section of the report, we will open the debate to the scientific community and society at large about what we know regarding the inclusion of artificial intelligence in our area of expertise: education. We will review the ethical questions surrounding its use in classrooms and formulate recommendations for educational institutions and teachers who want a detailed understanding of how artificial intelligence works. Finally, we will give our team of specialists a little space to discuss our favorite educational AI applications, which we have already begun recommending in our technology consultations.

Let us begin with the most important consideration when introducing a technology into the classroom: challenges external to the tool itself. The main problem to resolve is how we will train teachers in the

pedagogical use of AI. Creating mandatory teacher training plans in an educational environment that does not readily accept paradigm shifts can be very difficult. Every educational stakeholder involved in curriculum decisions will need to participate to make ongoing professional development an unavoidable reality in teachers' growth. This is essential because, otherwise, students will understand how the tool works in much greater depth than their teachers.

Some might believe this is not a disadvantage and that a student's knowledge in a specific area should be seen as another tool for reaching their full educational potential rather than as a weakness. At CERLI, we fully agree with this idea, provided the teacher has the same knowledge or more. Remember that the modern teacher acts as a guide for a group of students on their path toward personal and professional development. If students understand the steps along this path more deeply than the teacher, the teacher will be unable to ensure the technological tool's pedagogical value, overwhelmed by a lack of training.

With this in mind, we will devote this section to briefly developing everything we know about educational AI systems: their uses, approaches to implementation, and main strengths and weaknesses. Beyond this point, the system and its stakeholders, including teaching staff, will be responsible for taking charge of this knowledge, exploring it, and using their creativity and training to find new and useful applications for these tools, taking their lessons to an unprecedented level. We want to see these advances in the classroom and have the opportunity to study their reach and potential improvements. With that said, let us begin with the technical details.

The first thing we need to know about artificial intelligence in the classroom is that, unlike other preexisting technologies, there is no single way to implement it. Current generative AI systems cover broad areas of knowledge that education systems worldwide address through traditional methods. This means that ensuring safe use requires a personalized strategy for each specific case. At its core, this inability to generalize their use is what gives their overall functions such versatility,

making implementation more complicated while increasing their potential impact on teaching.

We should clarify that when we discuss educational approaches to artificial intelligence, we do not mean traditional uses such as writing text or translating into multiple languages. We mean applying it to the inclusive and resilient improvement and evolution of teachers' teaching methods. In other words, the future of AI in education will not be as the central core of teaching, but as a tool that provides everything we cannot achieve through traditional means.


This statement also includes everything that could be done through traditional methods but would require considerable time to learn and perfect. It is possible to learn to design digital images, but it is much simpler to learn to use a generative AI tool. This will allow teachers, including those who would not usually incorporate visual support, to include these benefits for students in their teaching materials. The purpose of using LLM tools such as ChatGPT would be to establish a natural, intuitive link between human language and the language used by the relevant AI systems.

For example, suppose we have developed a Stable Diffusion interface that creates hyperrealistic, adaptable images we can include in our teaching materials. The best results will come from input prompts that include commands and directives rather than expressing our intentions in natural language. This could reduce the likelihood that a teacher with little experience would use the tool, since they would have to learn to write more complex input prompts. ChatGPT would act as an immediate translator between natural human language and input prompts. We would only need to include an LLM in our interface to convert teachers' creative ideas into prompts, and the Stable Diffusion API would transform them into realistic, useful images without requiring teachers to learn to operate a complex interface.

Introducing technology into the classroom, especially this kind, means finding creative and novel uses to improve student learning or make it more efficient. It would make no sense to use these advances to

replace traditional methods ineffectively when we have the opportunity to innovate and revolutionize the classroom through inclusion, resilience, and competence. At CERLI, we do more than tell the story; we will offer innovative solutions with significant impact to empower educational authorities in the use of 21st-century technologies.

From this point forward, we will provide the information needed by any organization, institution, or educational stakeholder wishing to be sufficiently informed about this topic.



The Ethics of Educational AI

The ethical and moral dilemmas associated with AI systems are nothing new. Numerous studies are beginning to explore the potential limits that AI systems will need in the future to avoid becoming more problematic than beneficial. In education, however, the most important moral dilemmas concern whether using an AI tool is comparable to cheating. For example, if a student uses ChatGPT to write an essay on World War II, does completing the assignment have any educational value at all?

At CERLI, we see no reason to take pride in restricting modern tools in education, especially when they are so useful and revolutionary. In the previous example, where many see cheating, we see a curriculum problem: If the assigned task is so simple that a digital tool can complete it satisfactorily in seconds, where is the educational progress in having the student do it alone? Assignments such as “Write an essay on X” or “Answer the following questions about topic X” were created with the idea that, to complete them, students would need to read their notes and look for information in the books and materials provided by their teacher.

This ensures that they review what was covered in class and do not settle for what the teacher said in the classroom. However, these assignments have been widely criticized with the rise of technology, which has exposed a disconnect between the adult working world and the simplified, modified version of it in which young people are educated.

In real life, professionals use available technological tools as constant support in their work. Programmers use libraries, prewritten code, and AI to refine their work and make it more efficient; administrative staff use tools such as Excel, ChatGPT, and other software to maintain order and control over their work and documents; scientists use AI-based simulators and digital calculators to build their models; and we could go on forever.

When a tool is powerful and useful enough to make a professional's work more efficient, it always ends up being used, even if that means transferring much of the workload to a technology. Those who have already learned to use these tools have an astonishing advantage over those who still use analog methods or approaches requiring more human intervention. This has been part of the working world for years but has not found its place in modern education, where we have focused more on having students learn the analog method first and then learn the digital tools on their own ¹¹.

This is nothing new; it already happened with calculators in the late 20th and early 21st centuries. At that time, using a calculator to solve a mathematical operation was seen almost as heresy. Only after their use became normal did the education system recognize that tools as powerful as modern calculators could support the teaching of far more complex and advanced content than was usually taught, significantly strengthening graduates' professional skills.

No one stopped teaching addition because calculators exist. But once the concept itself is understood, it is natural to set aside unnecessary complexity, delegate that work to our digital brain, and occupy our biological brain with more advanced tasks and knowledge.

No one is saying that writing letters and essays should not be taught because ChatGPT can do it for us. But once these skills have been learned, it is both expected and necessary to delegate a task with lower

¹¹ It is not just that students learn to use these tools without any support from teachers; the education system itself discourages their use in the classroom. Young people are surrounded by useful, innovative tools they are neither trained nor permitted to use.

creative demands to a technological tool and focus on more interesting aspects, such as correcting the results, reviewing information, searching for sources, and reformulating text. The best assignments for students are those they can understand and complete but that also challenge their knowledge and skills. A task that can be solved as simply as writing two input sentences for an AI tool does not meet that requirement.

Instead of an assignment such as “Write an essay on topic X,” we could suggest something like “Based on the following essay on topic X, prepare a presentation of your opinions about it. The teacher and students may question your presentation, and you will need to defend your position.” Another option would be “Choose a topic from the following list and prepare for a debate. The teacher will select students with the same topic and, intervening as little as possible, have them debate their opinions in depth.” We could also propose something like “Choose a social problem in your community and design a detailed action plan to address it. Then present the idea to your classmates and prepare for their possible questions.”

Although similar, all these examples can serve as a basis for creating instructional and educational variations that make the use of AI irrelevant. Even if a student uses ChatGPT to obtain information, the need to know the topic they will be questioned about in depth will have forced them to understand and study it. These assignments have the same effect that older ones used to have, but AI is now a complement to the task rather than an enemy to defeat. These ideas also give rise to projects and broader assignments that can replace traditional assessments.

However, criticism of AI as a tool for completing assignments is not the only ethical and moral objection raised against it. Another very common one concerns the plagiarism on which generative AI systems are based: By using information from the internet to generate outputs, they can reproduce almost all the content of an existing source, encouraging plagiarism and copyright infringement. AI systems have struggled with this challenge since their very creation, and it is a much more difficult problem to overcome.

It is true that generative AI systems start from models and patterns repeated in their learning models and then create a new response with the key characteristics of the originals. In this respect, AI does not commit plagiarism per se, since it creates a new response while retaining the contextual and content features that make it coherent and valuable for the user's input. Put differently, AI does not copy its sources; it writes a new text containing everything that makes the original useful.

Despite this, AI has limitations that, when reached, make its output look more like plagiarism than new generation. This happens, for example, when there are few sources on the topic the user asks about. With little information available, the AI model generates its response from small amounts of text, inevitably producing a paraphrase or, at best, a reformulation. There is even a risk that the information will be incorrect or biased because of the lack of comparison.

If we look at both texts, the original and the generated one, the AI-generated text will appear to have too many similarities to the original sources. This could be questionable if the response is used for something intended for publication or presentation as one's own work, although the current legal framework may not be equipped to identify it as plagiarism.

A similar case could arise when the available sources are not sufficiently varied. When the AI model tries to identify key patterns, even if it has enough sources to develop a response, it will conclude that large portions of text are relevant because they recur in most of the sources consulted. These features will inevitably carry over into the generated output, making it a rehash of what was written in the original sources. Or it could be even worse: The AI model might "infer" details to expand the writing, inventing them and making the text's claims false.

To prevent this, teachers can reformulate the assignment so that it does not consist solely of text production, giving greater weight to the research process. They could request additional examples from students' personal experiences or local contexts and require them to

verify the reliability of their sources. This can serve as a crutch at first, with students eventually doing it automatically through repetition. In their first assignments, students would be asked to identify and explain why they used the sources they consulted and how they ensured those sources were reliable.

AI use could also be permitted, provided it is accompanied by a short methods page explaining the inputs the student used, the reasoning behind them, and how the student ensured they selected the best possible text for the assignment. This reflection would force students to review the information they use. It would also give teachers a basis for encouraging proper AI use to prevent possible plagiarism, errors, and other factors that will determine its future application.

To conclude our discussion of moral and ethical questions, we will address what, from our perspective, is the most complex and significant of all. If we encourage AI and its indiscriminate use, how can we ensure that we are not affecting students' ability to think independently? How will we know whether AI helps develop creativity or instead limits divergent thinking? The possibility of an uncontrollable impact on students' cognitive and creative development is undoubtedly a highly relevant concern in the debate over implementing AI in the modern classroom.

For this purpose, we should mention a prolific author on this subject: Nicholas Carr. After gaining prominence through his articles on technology in business, Carr turned to the impact of technologies on people, writing articles and books highly relevant to this report, such as *"Is Google Making Us Stupid?"* (2008) and *"The shallows: What the Internet Is Doing to Our Brains"* (2010).

In the first of these, Carr argues that new technologies and their ease of use have encouraged a culture of superficial engagement with information. He attributes this phenomenon to the structure of the internet itself, claiming that general abilities in reading, critical thinking, and information processing have been severely harmed since desktop technologies became popular. His argument rests on the idea that the

human brain is a malleable structure that adapts to the media through which information is obtained ¹². Viewed this way, unlike books and newspapers, online reading and hyperconnectivity encourage constant distraction, teaching our minds to process information much more superficially.

In his 2010 book, Nicholas takes his claims a step further. Drawing on the ideas of Canadian philosopher and sociologist Herbert Marshall McLuhan, he undertakes a thorough historical reexamination of all the exchanges of mental faculties we have experienced, their impact on society, and their short-, medium-, and long-term outcomes. Carr reaffirms the plasticity of our minds and issues strong warnings about the future that resonate clearly with us today. Modern technologies, such as AI, although Carr did not know its reach at the time, are designed to favor superficial thinking and encourage disengagement from information.

The ideas developed by Nicholas Carr find support in various studies of cognitive plasticity and changes in reading habits. In *“Proust and the Squid: The Story and Science of the Reading Brain”* (2007), Maryanne Wolf analyzes the relationship among reading, learning, and the human brain’s ability to adapt its circuits to a skill that is not innate.

Meanwhile, in the article *“Reading Behavior in the Digital Environment: Changes in Reading Behavior Over the Past Ten Years”* (2005), Ziming Liu studies changes in reading habits associated with the growth of digital environments, identifying a greater tendency toward scanning, keyword searching, selective reading, and navigating among content, at the expense of more sustained, deeper forms of reading.

These studies do not demonstrate that digital reading, in itself, activates brain regions entirely different from those used in print reading. However, they do show that the medium through which we access information can favor different reading and processing strategies. This is particularly relevant to Carr’s argument, since his

¹² This is based on studies of neuroplasticity, a branch of neurology that examines the malleable limits of our minds according to how they receive and process information.

concern focuses not only on technology as a tool but also on the cognitive habits we develop through its constant use.

In the article *“Reading Linear Texts on Paper versus Computer Screen: Effects on Reading Comprehension”* (2013), Anne Mangen, Bente R. Walgermo, and Kolbjørn Brønnick compared the reading comprehension of students who read the same texts in print and on a computer screen. The results showed better comprehension among those who read on paper. The authors suggest that the physical and spatial characteristics of the medium, along with the way readers navigate the text, can influence the construction of a coherent mental representation of what is being read.

These findings cannot automatically be transferred to artificial intelligence use, since the technologies and forms of interaction are different. Nevertheless, they provide a relevant starting point for asking a similar question: If the tools we use to access, process, and produce information can change our cognitive habits, what effects might regular use of generative AI have on abilities such as critical thinking, creativity, and deep information processing?

The existence of this possibility does not demonstrate that these effects are necessarily negative, but it does justify studying them before assuming that introducing AI into education has no cognitive consequences. Before becoming alarmed, however, let us examine whether there is evidence to support a claim such as the idea that implementing artificial intelligence in the classroom can erode critical thinking and cognitive development. We initially turned to a leading authority in this field who has worked with UNESCO to establish curriculum frameworks for AI in classrooms worldwide: Wayne Holmes.

In the book *“Artificial intelligence in education: promises and implications for teaching and learning”* (2019), Holmes, Bialik, and Fadel study the contexts in which artificial intelligence is applied, detailing the educational impacts of these technologies. Although the book was very enlightening, the fact that it was published well before generative AI

became widely popular led us to question whether its ideas remained equally valid in this new reality. We therefore decided to expand our bibliography with authors in different educational contexts who had conducted more comprehensive analyses of AI applications in the classroom and their short- and medium-term educational and psychological effects.

In these contexts, we found articles such as *“El Impacto de la Inteligencia Artificial en la Creatividad y el Desarrollo de Habilidades Cognitivas Superiores en Estudiantes Universitarios”* (2025), published by James Alduber Taramuel-Villacreces; *“La inteligencia artificial y el desarrollo del pensamiento crítico en estudiantes universitarios”* (2024), published by Deinny José Puche Villalobos; and *“La influencia de la inteligencia artificial en la creatividad y el pensamiento crítico en estudiantes universitarios”* (2024), published by Jorge L. Loáiciga Gutiérrez, Carlos L. Chanto Espinoza, and Cristhian Chaves Jaén.

These articles strongly assert that positive, beneficial correlations exist between the proper implementation of AI in the classroom and the development of cognitive skills, critical thinking, creativity, and better teaching strategies. However, all of them, especially the last, also emphasize that how artificial intelligence is implemented and used in the classroom will ultimately determine whether it overcomes its potential challenges or becomes a tool that encourages mental laziness, a lack of creativity, and misinformation. Educators and other educational authorities will be responsible for ensuring that its use is studied and implemented in the best possible way, so that its benefits outweigh its potential risks.

As for Nicholas Carr’s research, we see it as a warning that researchers must consider rather than an anti-technology message: Changing the brain comes at a cost. To adapt to the benefits of new technologies, we must be willing to sacrifice small but significant aspects of our current mental abilities. Artificial intelligence has the potential to help us advance education toward inclusion, the globalization of knowledge, and freedom of thought. However, it would

be completely naive to assume this transition will come without a price, paid in aspects we consider part of our mentality.

Carr expressed this exchange through more visual examples, such as the transition from natural navigation to maps and from maps to GPS. In the first, we exchanged the knowledge needed to navigate using nothing but nature for the knowledge needed to travel in any direction while choosing a more convenient route. We then exchanged this for the ability to navigate anywhere without needing different maps and always following the most efficient route to our destination.

We have illustrated this same idea throughout this report: When discussing the transition from functional but less user-friendly software to highly user-friendly software with functional limitations, we identified it as responsible for the loss of digital skills in younger generations. In other words, even from a techno-optimistic perspective, the truth of Nicholas Carr's ideas is undeniable. Nevertheless, if we agree that the end justifies the means and align ourselves with the goals of the modern educational agenda, this exchange of mental faculties is inescapable.

Our decisions and implementation mechanisms will therefore determine which aspects of our minds we relinquish, seeking to lose as little humanity as possible in exchange for the greatest attainable global benefit. If we want to achieve this, we must prepare to implement artificial intelligence correctly in the classroom, reviewing every possible variable that could complicate the process.

Implementing AI (Successfully) in the Classroom

Having established that the difference between a tool with great educational potential and the number one enemy of creativity and critical thinking lies simply in how we implement and use artificial intelligence, it is logical to say that our main mission should be to create a curriculum framework in which these safeguards are present. At CERLI, we have researched the leading implementation guides worldwide, analyzed their differences, and contributed our own small part toward improving them.

Before discussing our proposed curriculum framework, we would like to share the key points of our research in this area, highlighting the work of organizations that have made enormous and valuable efforts to provide these tools to teachers.

Artificial intelligence is a reality in every field, but education presents significant implementation challenges that we do not see elsewhere. The reason lies in the difficulty of creating useful plans in which it can be properly applied and mandatory teacher training to support implementation. Teachers are authority figures in the classroom, guiding students among the available tools for knowledge. If teachers are not adequately informed beyond their students' level, they will quickly be overwhelmed by the complexity of the lesson, preventing them from fulfilling their role properly.

In this context, it is essential to understand that ethical, pedagogically effective implementation of AI in the classroom must focus on building an educational environment prepared for its use: a learning space. Achieving this successfully requires work along two simultaneous lines: educational digital literacy and clear institutional curriculum frameworks for AI. In other words, we must be able to guarantee access and knowledge in the use of these tools, while also first establishing institutional policies for their inclusion in the classroom.

A strong example of the latter is the European Union's ethical guidelines on educational AI, published in 2022. The official document emphasizes the need to develop a critical, informed attitude toward technologies that operate on large amounts of data and whose use can have profound implications for learning and assessment. This approach requires us to rethink how we use these tools appropriately, as well as the criteria used to regulate and monitor their effects on students. In a sense, it places ethical responsibility on teachers and institutions rather than on the students using the tools.

At the same time, independent curriculum proposals have emerged that are very interesting to analyze. Researchers Sri Yash Tadimalla and Mary Lou Maher of the University of North Carolina at Charlotte (Charlotte, USA) published their article *"AI Literacy for All: Adjustable Interdisciplinary Socio-technical Curriculum"* in 2024. Their approach is more practical than the EU's proposal, advocating an interdisciplinary approach that goes beyond teaching how AI systems work to incorporate a critical understanding of their social, ethical, and political impact. They propose that an effective learning process must include foundational concepts such as the development of basic technical skills, the ability to interact with generative tools such as ChatGPT in an informed way, and ethical reflection on their use.

This reaffirms the guidelines discussed above, requiring schools and teachers to take a proactive approach and ask critical questions about the technologies they integrate into their classrooms, such as: What significant advantages does this offer over the traditional approach?

Does it respect my students' privacy? Does it strengthen or weaken educational equity?

A gradual but sustained effort is required to avoid pressure and dissatisfaction. Otherwise, we risk increasing technostress and teacher burnout and fueling resistance to new digital paradigms. In 2024, a group of researchers from several universities—Mike Perkins, Leon Furze, Jasper Roe, and Jason MacVaugh—published the article *“The AI Assessment Scale (AIAS): A Framework For Ethical Integration Of Generative AI In Educational Assessment”*¹³. It proposes an “AI Assessment Scale” to improve transparency with students about AI use in the classroom and empower educators through control over how and when these technologies are integrated.

This approach is remarkably useful because it clarifies the place artificial intelligence should occupy in educational processes without leaving regulatory or pedagogical gaps. The AIAS proposes five levels, each associated with a different degree of AI involvement in the assessment process.

Level 1: No AI

When an assignment or assessment is classified as “Level 1,” it must be completed entirely by the student without any kind of AI assistance. In these cases, generative AI use is considered cheating, which will be verified and rejected. These tasks would measure students' own skills or reasoning and could represent the smallest share of the course's assessments.

Level 2: AI for Idea Generation and Structuring

At level two, AI may be used to help with brainstorming and suggestions for writing and style, but it cannot generate the content of the final product. For example, if students had to write a story, they

¹³ The scale we share below is the one that originally appeared in the article. However, the authors have continued updating it to fully incorporate the impact of the latest available technologies. We recommend reviewing the latest updates at www.leonfurze.com

could use ChatGPT for ideas about its subject and structure, but the story itself would need to be written entirely by the student.

Level 3: AI for Language Editing

Students may use AI to improve a piece of writing, receiving automatic recommendations for clarity, style, and additional corrections to a text they have already written. In these cases, students are usually asked to include both versions—the original and the AI-enhanced one—along with a list of the corrections made and their reasons for considering those recommendations correct.

Level 4: AI for Tasks with Human Evaluation

At this level, AI's own weaknesses will be used to benefit the student. An assignment that can be completed entirely by AI will be proposed, but the student must analyze its generated content, identifying, explaining, and correcting possible limitations, biases, and errors. This allows students to refine their metacognitive and critical evaluation skills, improving their ability to analyze information carefully.

Level 5: AI as a Copilot

A level-five assignment allows AI to be used throughout the entire process without any restriction on creativity, provided its use is acknowledged. The teacher's task is to propose interesting challenges that strengthen students' abilities through AI use. This is the most controversial of the five levels, because it implies that students are required to complete little or none of the assignment themselves. However, it offers a valuable opportunity to focus the assignment on how AI itself works, discovering its uses and so that students can apply that understanding in tasks at other levels.

This curriculum framework particularly caught our attention because, although it does not solve the original problem, it offers an objective and transparent starting point. The scale's true strength lies in

its contextual flexibility, since it is not imposed as a universal model. Each institution can adapt the tasks and levels to its educational goals within its curriculum. Not every institution or teacher will need to work with the same five levels; they can be dynamically oriented toward the course's requirements.

At the same time, following a scale allows students to work with a common “yardstick” that varies according to the assignment rather than the teacher's convenience. Students always know what is permitted, under what conditions, and for what reasons, eliminating ambiguity surrounding penalties or results and opening the door to more mature discussions about the ethics of learning with AI.

For teachers, the AIAS provides a foundation for designing assessments. It requires them to devise an assignment and reconsider its educational purpose. Am I assessing the ability to write? Or the ability to think critically about an idea? Is it more relevant who writes it or how the result is interpreted? Although the mental effort required to design a task using this scale may seem complex, it is highly appealing to teachers who enjoy designing and planning assignments that challenge students' knowledge. This is very helpful in combating teacher burnout and frustration over AI misuse.

At CERLI, we decided to take this compelling scale a step further. Since our expertise focuses on technological solutions to theoretical problems, we decided to envision possible assignment scenarios at different levels of the scale and design an ideal LMS platform that could ensure compliance while minimizing the risk of cheating.

To begin with, a dedicated assignment section must be created, as is standard in modern platforms. The challenge would be to develop a graphical interface that is intuitive while offering teachers the greatest possible functionality. Our suggestion would be to separate assignments by type, with each also receiving a secondary classification under the AIAS standards. Each assignment would have its own interface accessible to both teachers and students. The only difference would be that students could use the features and submit the assignment, while

teachers could modify, assign, and review both submitted and pending assignments in real time.

For example, using the original AIAS, a level-one assignment could be a reading comprehension test. In this case, the platform should disable all text assistants, present the reading as a nondownloadable image, and follow it with short-answer questions. The platform should block copying and pasting, record completion time, restrict editing after submission, and include automatic systems to detect suspicious activity. The answers, together with a detailed report of the recorded information, would allow the teacher to grade the assignment and verify its legitimacy.

Moving to level two, we could use a text-based system for an assignment such as “Write a story between x and y words.” The platform could include an “assisted inspiration” space where the student accesses an AI-powered idea or outline generator, alongside a separate workspace where the student is expected to write the final product independently. These areas should be visually and functionally separate, which can easily be achieved through prior training of a ChatGPT API. Copying and pasting would be blocked, and both the final assignment and the chat history with the internal AI system would be recorded.

At level three, the interface would be very similar to the previous level, but this time users would write their text and then ask the integrated AI tool to suggest changes and corrections, which could be added automatically. However, for a correction to appear in the text, students would need to briefly explain why they accept or reject the AI suggestion. Again, copying and pasting should be blocked for both the original text and the AI suggestions.

At level four, our LMS should support an assignment initially completed by AI. This could be, for example, an article on a random topic or one selected in advance by the teacher, in which students must identify errors, information biases, inaccuracies, and so on. Students should be able to select sections of the text and attach a short note

explaining the error, suggesting a change, and including a specified number of external sources to substantiate the correction. Even if they use AI for the latter, they must review the sources to ensure they have not been invented.

Finally, level-five assignments would be both the simplest and the most complex. On the one hand, they allow all available tools without needing to impose limits through programming or AI training. The challenge lies with the teacher, who must be extremely creative in finding a curricular approach with assessment potential. The system must ensure work traceability by saving intermediate versions, interactions with AI, and student reflections so that the teacher can assess the process as well as the final product.

For example, one feasible assignment would be to design an awareness project or campaign on a topic that AI could suggest. It would require a message and an explanatory audiovisual piece, which could be AI-generated if we integrated an API into the platform. Alongside these, we could request an essay explaining the project's technical details, why the topic was chosen, how AI was used in the process, the ethical risks its use involved, and how those risks were managed.

As you can see, the LMS platform regulates the activity's level, giving teachers the greatest possible assurance while remaining clear and transparent with students. This activity promotes AI use within the educational objectives the teacher seeks to assess. It directly addresses the ethical questions discussed earlier while students learn to use artificial intelligence correctly and usefully.

At CERLI, we will continue developing prototypes and viable scenarios for this model, working with institutions interested in the gradual, ethical, and educational implementation of artificial intelligence in their teaching contexts. The path toward digitally mature education lies in supporting students with clarity, judgment, and purpose. AI can and should broaden teachers' horizons without compromising the integrity of the learning process.

Resource Considerations

It is time to ask the million-dollar question, which, in this case, is literal. How much does it cost to implement AI in the classroom following the guidelines presented above? The answer to this question, which we have received many times in our technology consultations, depends on several factors. Implementing AI today is considerably less expensive than before the pandemic because the rise of ChatGPT and its API allowed many companies to develop their own platforms, eliminating the need for the initial investment.

To do justice to the purposes of this report, we will answer from both perspectives: whether an institution wants to implement an existing market option or build a fully personalized internal LMS. Beginning with the first option, multiple platforms are available on the market, which we will discuss in detail in the next section. From a more general perspective, however, we can outline an average cost that includes as many features as possible. We will use an example institution and provide concrete numbers so that other institutions can extrapolate the estimate to their own scale.

For this report, we will consider an elementary school with 500 students across its six main grades¹⁴. With approximately 25 students per class, we would need twenty classes in total and therefore twenty classroom teachers. Considering only the traditional classroom, we would thus need an LMS platform with capacity for five hundred

¹⁴ To make the one-time implementation costs representative, we have considered an urban school in an upper-middle socioeconomic context that already has digital services for student records and a student population with internet connectivity and access to electronic devices at home.

students, twenty classroom teachers, and six administrative staff members.

For a personalized platform with tailored AI tools, the price of an annual subscription for an institution of this size will depend on the payment model. It can be charged at a fixed cost per student when the number of users does not fluctuate within a defined period, or per active user when users need to be added throughout the year. In either case, approximate costs range from 8,000 to 15,000 dollars annually, tending toward the upper end of this range for personalized AI requirements.

Although this may sound expensive, we should consider that it amounts to a monthly subscription of only 1.3 to 2.5 dollars per active student, which can readily be covered through an annual payment plan arranged in advance. The only significant variation among these investments is the particular benefit the support team can offer. Having a person assigned to the school, 24/7 service messaging, and even regular meetings to review the platform's functions can slightly increase the annual price, although these benefits are essential.

Even if we raise our expectations to an interactive AI platform with multiple different APIs for greater functionality, the price would not rise considerably above 18,000 dollars annually, still representing a cost of only three dollars per student per month. However, with the rapid progress OpenAI is making with ChatGPT, especially in its latest versions, increasing costs this much to obtain multiple individual APIs could be unnecessary for institutions below the university level.

In fact, the affordability of these platforms will become clear as we analyze the cost of creating a unique personalized version for an educational institution. Technological development alone would require an investment of around 30,000 dollars, plus annual maintenance and service costs of about 5,000 dollars. Development costs could rise to 50,000 dollars if we wanted an interactive, intuitive interface to reduce the learning curve for teachers and administrative staff.

Let us assume our platform will use the ChatGPT API, so we must also add the cost of both input and output prompts. These are calculated in “tokens,” each representing approximately four characters. OpenAI charges 0.005 dollars per thousand input tokens and 0.015 dollars per thousand output tokens. Estimating 500 students, 2 daily assignments each, and 180 school days gives us 180 thousand assignments annually. With average use of 1,800 tokens per assignment, combining inputs and outputs, operating the platform would require 324 million tokens. Split 50-50 between input and output, these leave us with a total annual cost of 3,240 dollars ¹⁵.

In addition, if the institution does not employ permanent developers and has not arranged a subscription with ongoing improvement services from the platform’s developer, it would need to add periodic programming and development services to the total cost for security reviews, implementing new technologies and assignments, or resolving occasional problems. There is no reliable way to estimate this cost, but institutions in this particular situation must consider it an additional expense.

Finally, we would also need to add the cost of technology adoption, which is often hidden in financial analyses of this kind. Existing platforms usually include training for teachers and administrative staff in their total annual cost. However, if we have created our own platform, we must also cover training costs, and the company that developed it should provide the training. If we pay enough for the platform to be intuitive and functional, a month of training would be ideal. Under these terms, the added cost could be around 1,500 dollars ¹⁶.

¹⁵ These costs are quoted for software development in business settings. OpenAI has education plans that could be more affordable, although they tend to target universities and do not include an LMS platform as such. A contextual optimization model for information can also significantly reduce the number of tokens needed.

¹⁶ This is perhaps the least realistic number in this report. We based the cost on developing an e-learning model through an existing platform, without considering compensation for the teachers being trained. If the company had already accounted for this need, it might include the cost in its total quote for technological development. This number could therefore be much higher or not be needed at all.

We can therefore conclude the annualized cost of even the most expensive platforms on the market in terms of cost per student remains affordable compared with creating one specifically for a school. Developers' business models work because they reduce overall costs by purchasing tokens and servers in bulk. Considering the service's actual cost, existing platforms charge each institution relatively modest fees while recovering their investment through the scalability of their annual subscription model.

After analyzing the return on investment in both cases, our team concluded that an annually contracted model would pay for itself through an increase of between 24 and 48 dollars per student in annual tuition. However, we consider it better to distribute this increase across the monthly tuition payments, where it would represent an additional 2.4 to 4.8 dollars per month. Although the annualized amount is the same, this is more manageable for families ¹⁷. The discount under an annual subscription model usually depends directly on advance payment. In that case, the institution would need to make the full annual investment upfront to obtain the best possible price and would have recovered it by the end of the year.

However, because we are considering a fixed annual cost rather than a payment adjusted monthly for each active account, we must account for the school's dropout rate, since these are inactive accounts we will already have paid for. At a rate of 5%, this implies an annual loss of 600 dollars for our example school. If we decide not to pass these losses on through the cost per student, we could still cover them internally by reducing other service subscriptions that the new platform would absorb and through savings in teacher hours and gains in convenience, which represent a competitive advantage over other institutions.

Of course, costs could be reduced significantly with a variable fee per monthly active user, excluding vacation months when platform use falls almost to zero. However, we would also need to consider that fees under

¹⁷ The amounts are divided into ten payments because, although the calendar year has twelve months, most private educational institutions in Latin America and the Caribbean charge monthly tuition payments only during the school year, which lasts ten months.

a variable model per active user do not usually receive the discounts that platforms offer for a fixed annual charge. These two factors offset one another, making the best option depend directly on variation in active users during the school year, regardless of what happens outside it.

If, on the other hand, we choose a unique personalized option, the analysis changes radically. We would begin with a much higher first-year investment, leaving us at a loss. To raise enough to avoid losing money in the first year, we would need to consider a cost of between six and seven dollars per student per month, more than double that of existing market options. This does not account for the dropout rate, which could bring the cost per student to eight or nine dollars, reaching three times the monthly cost of existing platforms.


There are many ways to optimize this cost, such as spreading the development investment over several years. If we divide the total payment over three years, we can practically match existing platforms' cost per student. Alternatively, if we create the platform for a group of educational institutions rather than a single institution, through what are commonly called economies of scale, we could significantly reduce the cost per institution, reaching the same levels as existing platforms with at least four institutions in the network.

We should also consider that a custom-built model can significantly optimize token use, making the entire operation more economical. Existing platforms include an estimated token charge in their price and purchase those tokens in bulk in advance. Each institution that contracts with them therefore ends up paying a small surcharge for estimated tokens that go unused. In a custom model, tokens can be purchased based on calculated needs, accumulated unused tokens, and the number of active users each month. This is not simple, but it is possible with the right technical team and can significantly reduce monthly costs.

At CERLI, our advice is clear: For an individual educational institution, the best option is to contract with an existing platform. If student numbers fluctuate significantly throughout the year, we recommend

negotiating a monthly fee based on active users. If variation remains low, a single annual payment is preferable, usually offering considerable savings over monthly subscription rates.

Similarly, if a school belongs to a network of institutions, a unique, custom-developed platform paid in installments or with costs divided among the institutions offers better long-term costs and greater institutional differentiation and competitive advantage. Even in this case, however, there are market options at the “expensive” end that can offer similar features without requiring an immediate investment.



Examples of Educational AI

Before closing this report, as is customary in our research, we would like to share some educational AI software whose features have impressed us. To do this, we took on the challenge of creating a metric for evaluating its performance.

After carefully analyzing the criteria we wanted to apply, we decided to create a horizontal bar system covering five fundamental categories for a technological tool intended for classroom use: learning curve, interface, innovation, AI use, and cost.



Beginning with the **LEARNING CURVE**, this measures how easy it is for teachers or students to begin using the tool.

- 1- Requires technical training.
- 2- Requires many hours of use.
- 3- Requires a few hours of use.
- 4- Fairly easy to learn with little help.

- 5- Can be used from the first interaction.

Next, we have the **INTERFACE**. This category evaluates the program's usability, visual clarity, and overall design.

- 1- Outdated or cluttered interface.
- 2- Up-to-date but confusing interface.
- 3- Acceptable design, although not optimized.
- 4- Clear and functional.
- 5- Modern, clean, intuitive, and visually appealing.

Meanwhile, **INNOVATION** evaluates whether the tool offers something new or disruptive in education.

- 1- Repeats what other tools already do.
- 2- Differs little from what is already available.
- 3- Competently integrates uncommon ideas.
- 4- Offers novel approaches or solutions.
- 5- Proposes a clear change in paradigm or methodology.

AI USE evaluates how central, useful, and effective the use of artificial intelligence is.

- 1- Uses AI only as a superficial addition.
- 2- Automates basic tasks without adapting to context.
- 3- Uses AI with some degree of personalization or useful analysis.
- 4- AI provides contextualized decisions, suggestions, or analysis.
- 5- A powerful, well-integrated AI system lies at the heart of the platform.

Finally, **COST** evaluates the relationship between price and the value offered. This factor considers the institution's size, since larger budgets and student populations can justify some additional expenses for the ultimate benefit.

- 1- No clear justification of value.
- 2- Affordable only for large institutions.
- 3- Affordable only for medium-sized and large institutions.
- 4- Affordable for all types of institutions.
- 5- Free.

As you will see, the options selected have very different characteristics. We do not mean to imply that they are the best in their field; we simply use them as examples to cover as many contexts as possible. In addition, most educational software companies offer personalized payment plans based on the institution, so the COST factor may vary. Now, let us look at the examples.

Toddle

Toddle is a comprehensive digital platform designed in 2019 by a group of educators and technologists specializing in schools with International Baccalaureate programs ¹⁸, although it also has a practical model for standard institutions. With functionality similar to the previous software, Toddle acts as an all-in-one platform for educational institutions, integrating multiple AI tools into a single interface and saving teachers up to 10 hours of work per week.



One of Toddle's most distinctive features compared with similar tools is its deep, native integration of artificial intelligence. Built on the ChatGPT API, Toddle's AI goes beyond simple automated suggestions and participates in complex processes such as curriculum planning, material generation, assignment grading, and personalized report preparation. For example, when teachers begin designing an instructional unit, they can use the "Curriculum Design Assistant," an AI tool that suggests objectives, lines of inquiry, differentiated activities, and assessments aligned with the IB learner profile and approaches to learning.

We want to emphasize that generic templates are not used for this purpose. The AI model is properly trained on the institution's curriculum content and teachers' previous practices. This means it can "improve itself" and better adapt to the needs of the teachers using it.

¹⁸ The International Baccalaureate (IB) is a global educational model offering rigorous, student-centered academic programs from ages 3 to 19. It was created in Switzerland in 1968 and has four main programs: PYP (ages 3–12), MYP (ages 11–16), DP (ages 16–19), and CP (ages 16–19).

At CERLI, we experienced this feature firsthand through our interactions with the company. Among its many varied features, two particularly interested us. The first was the digital portfolio system. Each student has a space to upload files such as images, videos, voice recordings, documents, links, or written entries. Both teachers and students can comment on this evidence through self-assessment or peer assessment, and it can be incorporated directly into progress reports or collaborative assessment sessions.

Although this may seem standard in any LMS, Toddle implements it in a considerably innovative way: Its native integration with grading gives students, families, and the institution itself evidence of the work completed. The simple interface and native AI use provide students with reliable, ongoing feedback, while teachers have all the tools needed to monitor progress and ensure students use the features correctly.

All of this happens while the school maintains an effective channel of communication among all stakeholders in the educational community, since teachers can send messages to students, families, and school leadership teams. This is very important for involving everyone in the student's education. Messages can include links to assignments, documents, or portfolio evidence, ensuring consistency among what is taught, what is assessed, and what is communicated.

The second feature that caught our attention, although this time with some reservations, was a grading system that integrates AI to generate personalized report cards in minutes. Using information collected from activities, portfolios, and rubrics, the AI tool writes texts tailored to each student, citing specific evidence and proposing improvement goals. These reports can be configured to match the required institutional format and include multiple subject areas, general comments, achievement levels, and progress charts.

Although this considerably reduces teachers' administrative workload, it can also raise ethical dilemmas already discussed in this report. Formative and summative assessment, particularly when writing report cards, is not merely a technical or replicable process. It

involves an ethical and contextual dimension that only teachers in daily contact with students can understand in detail. The risk is that uncritical use of these tools will strip the assessment process of its pedagogical purpose, with AI acting more as a substitute than as an assistant.

This could take two distinct forms. One is the possibility that teachers might decide to eliminate their assessment workload entirely, leaving grades without their own critical analysis. The other is that, to facilitate automatic report card creation, teachers might limit themselves to assignments that AI can easily access and assess, eliminating the benefits of having such interesting and innovative tools.

You might think this is the same as what already happens, since a teacher could ask ChatGPT to write automatic evaluative comments and grades. However, it is much more complex. To perform that task, teachers would need to pay for ChatGPT Plus, since the free version would limit the number of students they could handle, upload all documents individually, give the AI tool the proper context and information, and ensure the result accounts for that context throughout the conversation. A tool such as Toddle, by contrast, serves it all on a silver platter: The AI already has the context, already holds all the information and documentation in its database, and deeply understands the assessment criteria being used.

At CERLI, we firmly believe that teachers must remain the critical authors of every word communicated about their students' development. Technology can assist, but it must never replace the human perspective intrinsic to education. This is not an impossible challenge to resolve, but it is a critical point that interested institutions must consider before deciding to implement a tool like this.


Finally, turning to the business model, this is where costs can become a problem. There are three main subscriptions: **"Planning Pro,"** **"Essentials,"** and **"Ultimate."** The first two function as traditional LMS platforms, both including unit and lesson planning and student analytics. Essentials adds a gradebook, automated transcripts, progress reports, assignment templates, student portfolios, attendance and

schedule management, an announcements system for school administrators, and other features. We therefore see Planning Pro more as a technical demonstration, while Essentials covers almost everything usually included in traditional LMS platforms, some of which are entirely free.

Only the Ultimate version includes AI integration, bringing all the benefits mentioned above. This complicates the analysis of its return on investment and cost-effectiveness, since it would fall within the scope of this report only when the Ultimate version is purchased. It should also be emphasized that there are different rates for systems serving traditional educational institutions and those dedicated to International Baccalaureate curriculum frameworks.

Because the total price is personalized and estimates are not officially published, we will not disclose the precise information used to develop our ratings. However, we clarify that the score in the “Cost” category refers to the total for a traditional institution, considering only the Ultimate plan ¹⁹.

To conclude our overview, these are our final ratings for Toddle:



¹⁹ Because we are working with an estimate, Ultimate subscription costs may not be affordable for medium-sized institutions, even though the rating does not reflect this. We decided to retain the classification based on the estimate, but recommend using the results as a guide and inquiring about specific prices when considering the service.

Teachmate

Earlier in this report, we discussed virtual assistants and their unlikely uses in education. However, we also closed that section by saying:

“Of course, combining LLM architectures with virtual assistants could change this radically, turning them from mere home entertainment into a powerful tool for industry.”

This is precisely what the software we will introduce next proposes. Teachmate AI is an educational virtual assistant aimed at teachers and school leaders. Like almost all LMS platforms, it was created specifically to ease teachers’ administrative workload. It operates through sophisticated natural language processing techniques capable of generating educational plans aligned with national curricula, writing personalized reports for each student, generating assignments adapted to different levels, and automatically grading structured questions.



Like traditional LMS platforms, Teachmate AI offers generators for administrative documents such as school policies, risk assessments, and permission forms, which are useful for leaders and coordinators. It also includes specific tools for Individualized Education Programs (IEPs) and semiautomatic support for summaries, classroom adjustments, letters to parents, and pedagogical suggestions adapted to the diagnostic profile. This is equivalent to having a contextualized virtual assistant that retains the institution’s particular context and the information

uploaded to its database and produces useful outputs for teachers, administrators, and families.

The interface is intuitive, with autosave, dark mode, and quick access to frequently used tools. In teaching practice, this considerably reduces the learning curve and delivers the promised improvement in time spent on administrative work. In addition, according to what we have been told, Teachmate AI is developing important beta tools, including a grading assistant for written responses, personalized feedback aligned with rubrics, and a progress-tracking system that generates longitudinal summaries of student performance ²⁰.

The business model may be where the greatest disadvantages arise, depending particularly on where we are located. To begin with, the software uses what we commonly call a “FREEMIUM” model: free software with a paid version that generally includes the most useful features. Nevertheless, this allows teachers to use a free version to explore its functions and interface before even contacting the company for a license.


Even so, it is clear that any school deciding to use Teachmate AI will inevitably end up paying for a subscription. In this case, we should highlight that the LMS offers pricing that is affordable for most small educational institutions, at approximately three to four dollars per student per month. For large institutions with more than fifty students, the cost returns to our estimate of two to three dollars per student, making it an affordable option with excellent value for money.

The tool’s only disadvantage is that native adaptation to the local curriculum is available only in certain countries: England, Wales, the United States, Canada, Australia, New Zealand, and the Republic of Ireland. This is therefore not yet a reality for the rest of Europe, Latin America, and the Caribbean. In fact, based on what we have learned,

²⁰ This last feature, far from being an advantage, could have serious ethical consequences for how the tool operates, as we saw in the case of Toddle.

there do not appear to be plans for rapid expansion of this particular feature into Spanish-speaking markets.

Our final ratings for Teachmate AI are therefore as follows:



ILIAS

ILIAS is an open-source LMS originally created in Germany as part of the VIRTUS project at the University of Cologne. Its first version was released on November 2, 1998, and, following its early acceptance, it became free software under the GNU GPL license in 2000. Between 2002 and 2004, it evolved into ILIAS 3 and was the first open-source platform to achieve full compatibility with standards such as SCORM 1.2 and later SCORM 2004



²¹.

Each user has a personal desktop from which they can access courses, the school calendar, blogs, notes, a portfolio, assignments, and more, in a way very similar to other open-source LMS platforms such as Moodle. However, one area in which it has always stood out is assessment: ILIAS supports assessment through tests, surveys, question banks, adaptive assessments, and attempt management with automatic feedback, compatible with SCORM, IMS-QTI, LOM, and AICC.

It also has modules for competency management with self-assessment and 360° assessment, learning paths, content authoring with LaTeX and XHTML, e-portfolio generation, and support for collaborative work on documents, making it a thoroughly useful LMS at almost every traditional educational level.


One apparently contradictory point is that ILIAS has no native AI features. We included software without these native features because,

²¹ **SCORM** is an international standard for digital learning content created by the United States Department of Defense. Its purpose is to ensure that educational content can be easily shared among different LMS platforms without needing to be reprogrammed.

being open source and requiring no usage licenses, it carries no subscription cost. That money can instead be used to hire a team of experts to develop a version of ILIAS with a tailored interface and AI integration tools. This would be an affordable alternative to creating your own LMS, although the AI-based tools would carry the costs and limitations of using an LLM’s API.

There are at least three implementation models, each of which would produce very different results if evaluated separately. We therefore decided to divide the rating for this tool into three different versions, depending on the route each institution chooses.

The first option would be to implement the software natively, without any tools, plugins, or personalized frontend ²². This would, of course, mean that the resulting program would have no AI and would serve only as a traditional LMS.





²² For those unfamiliar with these terms, here is an explanation: A “frontend” is an interface that sends instructions through an API to an external server or through a plugin to a tool already installed in the system itself. In this way, the frontend tells the program to perform certain tasks, such as generating text or grading an exam.

If we decided to reduce costs as much as possible by reorganizing the traditional interface rather than creating a separate frontend that communicates with ILIAS, we could reduce development and maintenance costs, as well as development time, by more than 60%. However, it should be emphasized that this saving comes at the cost of the personalized tool's simplicity, considerably increasing the learning curve. Significant savings can also be achieved by using AI tools through plugins, although they would be less functional than API-based tools.



Finally, there is the most functional but most expensive option: designing a complete, personalized frontend with AI tools connected through APIs. Here, we would need to consider the cost of design, the maintenance team, and the tokens used by the integrated AI system. In addition, we would need to train it on data specific to our intended uses and institution, which requires time and knowledge.



Epilogue

Throughout this report, we have explored the past, present, and possible futures of artificial intelligence, understanding it not only as a technological phenomenon but also as an agent capable of profoundly transforming education. Its applications, especially within classrooms, have been the subject of academic debate for years. Given the speed of current advances, that debate is becoming increasingly relevant.

This report does not seek to close that discussion, however, but to help broaden it. Discussing artificial intelligence in education necessarily means asking how these technologies may influence our values, critical thinking, personal and professional ethics, creativity, and, ultimately, our understanding of learning. Much of AI's potential lies precisely in the diversity of its applications and its ability to affect some of the most important areas of human life. Within this landscape, education occupies a central place: It is where technological possibilities can meet the experience of learning, pedagogical judgment, and teachers' creativity.

This transformation is not without risks. Incorporating artificial intelligence into the education system requires much more than access to new tools or a willingness to use them. It requires training, ethical reflection, clear pedagogical criteria, evaluation mechanisms, and sustained institutional commitment. Educational institutions and governments will have a fundamental responsibility to build these frameworks and create conditions that allow technological opportunities to be used without losing sight of the principles that should guide every educational process.

The decisions made today will help define the characteristics of the education system in the coming decades. They will determine whether artificial intelligence becomes an instrument that deepens existing inequalities or an opportunity to build educational environments that are more inclusive, creative, accessible, and prepared to respond to the challenges of a constantly changing society.

At CERLI, we will continue supporting this process alongside educational institutions, humanitarian organizations, governments, teachers, students, and communities. We understand that artificial intelligence does not, in itself, solve education's historical problems, but it does represent an exceptional opportunity to revisit questions we have postponed for far too long.

Perhaps, then, the real challenge is not to determine how much artificial intelligence can transform education, but to decide how much we are willing to transform education in response to the possibilities, questions, and responsibilities that artificial intelligence has placed before us.

References

1. Abeliuk, A., & Gutiérrez, C. (2021). Historia y evolución de la inteligencia artificial. *Bits de Ciencia*, (21), 14–21
2. Álvarez Munárriz, L. (1994). *Fundamentos de inteligencia artificial*. EDITUM.
3. Baños, G. (2024). *El sueño de la inteligencia artificial: El proyecto de construir máquinas pensantes: Una historia de la IA*. Shackleton Books.
4. Bernal Parraga, A. P., Santin Castillo, A. P., Ordoñez Ruiz, I., Tayupanta Rocha, L. M., Reyes Ordoñez, J. P., Guzmán Quiña, M. de los A., & Nieto Lapo, A. P. (2024). La inteligencia artificial como proceso de enseñanza en la asignatura de estudios sociales. *Ciencia Latina Revista Científica Multidisciplinar*, 8(6), 4011–4030.
5. Bolaño-García, M., & Duarte-Acosta, N. (2024). Una revisión sistemática del uso de la inteligencia artificial en la educación. *Revista Colombiana de Cirugía*, 39, 51–63. <https://doi.org/10.30944/20117582.2365>
6. Caicedo-Basurto, R. L., Camacho-Medina, B. M., Quinga-Villa, C. A., Fonseca-Lombeida, A. F., & López-Freire, S. A. (2024). Análisis y beneficios de la educación en la era de la inteligencia artificial. *Journal of Economic and Social Science Research*, 4(4), 291–302. <https://doi.org/10.55813/gaea/jessr/v4/n4/148>
7. Carr, N. (2008). Is Google making us stupid? What the Internet is doing to our brains. *The Atlantic*, 302(1), 56–63.
8. Carr, N. (2010). *The shallows: What the Internet is doing to our brains*. W. W. Norton & Company.
9. Cedeño Meza, J. G., Maitta Rosado, I. S., Vélez Zambrano, M. L., & Palomeque Zambrano, J. Y. (2024). Investigación universitaria con

inteligencia artificial. *Revista Venezolana de Gerencia*, 29(106), 817–830.
<https://doi.org/10.52080/rvgluz.29.106.23>

10. Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. doi:10.1093/analys/58.1.7.
11. Comisión Europea, Dirección General de Educación, Juventud, Deporte y Cultura. (2022). Directrices éticas sobre el uso de la inteligencia artificial (IA) y los datos en la educación y formación para los educadores. Oficina de Publicaciones de la Unión Europea. doi:10.2766/898.
12. Cortina Orts, A. (2019). Ética de la inteligencia artificial. *Anales de la Real Academia de Ciencias Morales y Políticas*, 96, 379–394.
13. de Castro, N. J. C., Paes, F. T., de Almeida, R. M., de Sousa, L. B., Takolodjou, M. A., dos Santos Silva, M., & Batista, A. F. C. (2025). Tecnología educativa aplicada a la educación ambiental: Innovando para afrontar el cambio climático. *Cogitare Enfermagem*, 30, e97489es. <https://doi.org/10.1590/ce.v30i0.97489es>
14. Delgado de Frutos, N., Campo Carrasco, L., Sainz de la Maza, M., & Etxabe-Urbietta, J. M. (2024). Aplicación de la inteligencia artificial (IA) en educación: Los beneficios y limitaciones de la IA percibidos por el profesorado de educación primaria, educación secundaria y educación superior. *Revista Electrónica Interuniversitaria de Formación del Profesorado*, 27(1), 207–224. <https://doi.org/10.6018/reifop.577211>
15. Dolba, S. Q., Inoncillo, F. A., & Nunez, J. L. (2025). Evaluating the impact of AI-generated outputs on student assessment: Educator's perspective. *International Journal of Artificial Intelligence Research*, 9(1), 1–7. <https://doi.org/10.29099/ijair.v9i1.1347>
16. Fernández Díaz, A., Madera Gómez, M., & Rodríguez Ramos, E. (2025). La inteligencia artificial en la primera infancia y preescolar; un acercamiento a los docentes: Artificial intelligence in early childhood and preschool: An approach to teachers. *Horizonte Pedagógico*, 14.
17. Flórez Rojas, M. L. (2023). Pensamiento de diseño y marcos éticos para la inteligencia artificial: Una mirada a la participación de las múltiples partes

interesadas. Desafíos, 35(1).
<https://doi.org/10.12804/revistas.urosario.edu.co/desafios/a.12183>

- 18.** Franganillo, J. (2022). Contenido generado por inteligencia artificial: Oportunidades y amenazas. Anuario ThinkEPI, 16.
<https://doi.org/10.3145/thinkepi.2022.e16a24>
- 19.** Gallent-Torres, C., Arenas Romero, B., Vallespir Adillón, M., & Foltýnek, T. (2024). Inteligencia artificial en educación: Entre riesgos y potencialidades. Práxis Educativa, 19, 1–29.
<https://doi.org/10.5212/PraxEduc.v.19.23760.083>
- 20.** Galván Fernández, C. (2024). La inteligencia artificial desde la educabilidad. Digital Education Review, (45), 1–2.
- 21.** García-Peña, V. R., Mora-Marcillo, A. B., & Ávila-Ramírez, J. A. (2020). La inteligencia artificial en la educación. Dominio de las Ciencias, 6(3), 648–666. <https://doi.org/10.23857/dc.v6i3.1421>
- 22.** González-Fernández, C., Garrosa, E., & Blanco-Donoso, L. M. (2024). El efecto de la tecnología en el ámbito laboral: Consecuencias del tecnoestrés en la salud mental. Ansiedad y Estrés, 30(1), 40–48.
<https://doi.org/10.5093/anyes2024a6>
- 23.** González-González, C. S. (2023). El impacto de la inteligencia artificial en la educación: Transformación de la forma de enseñar y de aprender. Currículum. Revista de Teoría, Investigación y Práctica Educativa, 36, 51–60. <https://doi.org/10.25145/j.qurricul.2023.36.03>
- 24.** Granda Dávila, M. F., Muncha Cofre, I. J., Guamanquispe Rosero, F. V., & Jácome Noroña, J. H. (2024). Inteligencia artificial: Ventajas y desventajas de su uso en el proceso de enseñanza aprendizaje. MENTOR Revista de Investigación Educativa y Deportiva, 3(7), 202–224.
<https://doi.org/10.56200/mried.v3i7.7081>
- 25.** Gros Salvat, B. (1992). La inteligencia artificial y su aplicación en la enseñanza. Comunicación, Lenguaje y Educación, (13), 73–80.
<https://doi.org/10.1080/02147033.1992.10821001>

- 26.** Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education: Promises and implications for teaching and learning. Center for Curriculum Redesign.
- 27.** Instituto de Educación de la Universidad ORT Uruguay (Ed.). (2015). Educación y tecnología en el Uruguay: Una mirada desde la investigación. Universidad ORT Uruguay.
- 28.** Jara, I., & Ochoa, J. M. (2020). Usos y efectos de la inteligencia artificial en educación (Documento para discusión IDB-DP-00-776). Banco Interamericano de Desarrollo. <https://doi.org/10.18235/0002380>
- 29.** Juca-Maldonado, F. (2023). El impacto de la inteligencia artificial en los trabajos académicos y de investigación. *Revista Metropolitana de Ciencias Aplicadas*, 6(S1), 289–296. <https://doi.org/10.62452/8nww1k83>
- 30.** Liu, Z. (2005). Reading behavior in the digital environment: Changes in reading behavior over the past ten years. *Journal of Documentation*, 61(6), 700–712. doi:10.1108/00220410510632040.
- 31.** Llorens-Largo, F., Vidal, J., & García-Peñalvo, F. J. (2023, 8 de diciembre). Ya llegó, ya está aquí, y nadie puede esconderse: La inteligencia artificial generativa en educación. *Aula Magna 2.0*.
- 32.** Loáiciga Gutiérrez, J. L., Chanto Espinoza, C. L., & Chaves Jaén, C. (2024). La influencia de la inteligencia artificial en la creatividad y el pensamiento crítico en estudiantes universitarios. *Revista Social Fronteriza*, 4(5), e45485. doi:10.59814/resofro.2024.4(5)485.
- 33.** Mangen, A., Walgermo, B. R., & Brønnick, K. (2013). Reading linear texts on paper versus computer screen: Effects on reading comprehension. *International Journal of Educational Research*, 58, 61–68. doi:10.1016/j.ijer.2012.12.002.
- 34.** Marcos García, M. (2015). Diseño y control de un robot mediante interfaz cerebro-máquina usando la tecnología SSVEP [Proyecto fin de carrera, Universidad Autónoma de Madrid, Escuela Politécnica Superior].

- 35.** Martínez González, M. A. (2023). Uso responsable de la inteligencia artificial en estudiantes universitarios: Una mirada rectoria. *Revista Boletín Redipe*, 12(9), 172–178.
- 36.** Mendoza Vega, A. J., Guadamud Muñoz, J. D., Santana Castro, E. K., Chiriboga Palacios, I. A., & Vera Arias, M. J. (2024). Uso de las plataformas de inteligencia artificial en el contexto educativo. *Ciencia Latina Revista Científica Multidisciplinar*, 8(1), 10996–11009. https://doi.org/10.37811/cl_rcm.v8i1.10412
- 37.** Monasterio Astobiza, A., Ausín, T., Toboso, M., Morte Ferrer, R., Aparicio Payá, M., & López, D. (2019). Traducir el pensamiento en acción: Interfaces cerebro-máquina y el problema ético de la agencia. *Revista de Bioética y Derecho*, (46), 29–46. <https://doi.org/10.1344/rbd2019.0.26862>
- 38.** Moreno Padilla, R. D. (2019). La llegada de la inteligencia artificial a la educación. *Revista de Investigación en Tecnologías de la Información*, 7(14), 260–270. <https://doi.org/10.36825/RITI.07.14.022>
- 39.** Morocho Cevallos, R. A., Cartuche Gualán, A. P., Tipan Llanos, A. M., Guevara Guevara, A. M., & Ríos Quiñónez, M. B. (2023). Integración de la inteligencia artificial en la educación. *Ciencia Latina Revista Científica Multidisciplinar*, 7(6), 2032–2053. https://doi.org/10.37811/cl_rcm.v7i6.8832
- 40.** Paguay-Simbaña, M. Y., Jiménez-Abad, D., Quiliguango-Lanchimba, V. F., Maynaguez-Canacuan, M. P., Coello-García, C. de los Ángeles, & Coello-Ortiz, S. M. (2024). La ética en el uso de la inteligencia artificial en los procesos educativos. *Revista Científica Retos de la Ciencia*, 1(4), 145–158. <https://doi.org/10.53877/rc.8.19e.202409.12>
- 41.** Peñaherrera Acurio, W. P., Cunuhay Cuchiye, W. C., Nata Castro, D. J., & Moreira Zamora, L. E. (2022). Implementación de la inteligencia artificial (IA) como recurso educativo. *RECIMUNDO*, 6(2), 402–413. [https://doi.org/10.26820/recimundo/6.\(2\).abr.2022.402-413](https://doi.org/10.26820/recimundo/6.(2).abr.2022.402-413)
- 42.** Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration

of generative AI in educational assessment. *Journal of University Teaching & Learning Practice*, 21(6). <https://doi.org/10.53761/q3azde36>

- 43.** Piedra Alegría, J. (2023). Anotaciones iniciales para una reflexión ética sobre la regulación de la inteligencia artificial en la Unión Europea. *Revista de Derecho*, (28), e3264. <https://doi.org/10.22235/rd28.3264>
- 44.** Ponce Gallegos, J. C., Torres Soto, A., Quezada Aguilera, F. S., Silva Sprock, A., Martínez Flor, E. U., Casali, A., Scheihing, E., Túpac Valdivia, Y. J., Torres Soto, M. D., Ornelas Zapata, F. J., Hernández A., J. A., Zavala D., C., Vakhnia, N., & Pedreño, O. (2014). *Inteligencia artificial. Iniciativa Latinoamericana de Libros de Texto Abiertos (LATIn)*.
- 45.** Puche-Villalobos, D. J. (2024). La inteligencia artificial y el desarrollo del pensamiento crítico en estudiantes universitarios. *Delectus*, 7(2), 59–67. <https://doi.org/10.36996/delectus.v7i2.242>
- 46.** Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- 47.** Sandoval-Almazán, R. (2011). Mentes en peligro: El daño de internet en nuestro cerebro. *Convergencia. Revista de Ciencias Sociales*, 18(56), 241–248.
- 48.** Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. doi:10.1017/S0140525X00005756.
- 49.** Smart, P., Clowes, R., & Clark, A. (2025). ChatGPT, extended: Large language models and the extended mind. *Synthese*, 205, Article 242. <https://doi.org/10.1007/s11229-025-05046-y>
- 50.** Sosa de Wood, P. N., Jiménez Chaves, V. E., & Riego Esteche, A. P. (2024). El análisis de la percepción de los profesores respecto al uso de la inteligencia artificial. *Revista EDUCA UMCH*, (24), 66–77. <https://doi.org/10.35756/educaumch.202424.293>
- 51.** Tadimalla, S. Y., & Maher, M. L. (2024). AI literacy for all: Adjustable interdisciplinary socio-technical curriculum. En 2024 IEEE *Frontiers in*

Education Conference (FIE) (pp. 1–9). IEEE.
doi:10.1109/FIE61694.2024.10893159.

- 52.** Tamayo Arellano, V. R., Romero Acosta, Y. J., Zapata Romero, S. V., & Sánchez Macías, R. A. (2024). Integración de la inteligencia artificial en la enseñanza universitaria para optimizar el proceso de aprendizaje. *Magazine de las Ciencias: Revista de Investigación e Innovación*, 9(2), 68–94. <https://doi.org/10.33262/rmc.v9i2.3104>
- 53.** Taramuel-Villacreces, J. A. (2025). El impacto de la inteligencia artificial en la creatividad y el desarrollo de habilidades cognitivas superiores en estudiantes universitarios. *ISTE SCIENTIST*, 4(1), 72–88.
- 54.** Tomalá De La Cruz, M. A., Mascaró Benites, E. M., Carrasco Cachinelli, C. G., & Aroni Caicedo, E. V. (2023). Incidencias de la inteligencia artificial en la educación. *RECIMUNDO*, 7(2), 238–251. [https://doi.org/10.26820/recimundo/7.\(2\).jun.2023.238-251](https://doi.org/10.26820/recimundo/7.(2).jun.2023.238-251)
- 55.** Trincado Alonso, F., López-Larraz, E., & Gil-Agudo, Á. (2015). Métricas de neuroplasticidad basadas en interfaces cerebro-máquina. *Cognitive Area Networks*, 2(1), 39–44.
- 56.** UNESCO. (2024). Guía para el uso de IA generativa en educación e investigación. UNESCO.
- 57.** Valencia Tafur, A. T., & Figueroa Molina, R. E. (2023). Incidencia de la inteligencia artificial en la educación. *Educatio Siglo XXI*, 41(3), 235–264. <https://doi.org/10.6018/educatio.555681>
- 58.** Wolf, M. (2007). *Proust and the squid: The story and science of the reading brain*. HarperCollins.