

# A University Leader's Glossary for AI and Machine Learning

REVISED 2023

# Foreword

By Dr. Mildred Garcia

---

The past few months have revealed just how profound an impact artificial intelligence may have on everyday life. At colleges and universities across the country, the rapid rise of generative AI tools like ChatGPT has already reshaped the way that students, faculty, and administrators interact with technology. The emergence of powerful AI platforms comes at a time when institutions were already navigating profound transformation in the wake of the pandemic — and it has launched an entirely new conversation about the changing role of artificial intelligence in the higher education landscape

But while generative AI and large language learning models like ChatGPT have been the subject of an overwhelming volume of media coverage in the past few months, the truth is that AI is nothing new in higher education. In fact, American higher education has a long history with artificial intelligence. Most scholars trace the birth of AI research to a summer workshop on the campus of Dartmouth College that brought together leaders from institutions across the country to discuss the rapidly emerging field. More recently, public colleges and universities like SUNY Albany and Morgan State University have made significant investments in AI research and development. Most, if not all, of these institutions are already implementing AI-enabled tools on campus as well — both in and out of the

classroom. Think of software that calculates the ideal schedule for every student, or AI-powered coaches that use text message reminders and recommendations to help keep learners on track.

What's changed in the past few months is generative AI has made the profound impact of artificial intelligence easier to recognize with our own eyes. Within a matter of months, AI tools have become much more ubiquitous, easier to use, and out in the open. Any student and faculty member can tap into a nearly limitless resource just by typing an address into their web browser. Perhaps most importantly, AI is demonstrating the potential to help colleges and universities make significant progress in providing students with the sort of timely, responsive support that is proven to increase persistence and retention. Of course, it's also causing anxiety about what happens when technology this powerful is left unchecked, and igniting debate about issues ranging from academic integrity to algorithmic bias.

It's now incumbent upon institutional leaders to navigate this changing landscape and tap the potential of AI to help more students succeed — but too often, those leaders don't have the information they need to make informed decisions.

---

As with all new technologies, the more the general public starts talking about it, the more oversimplification and misinformation enters the discourse. That means it's more urgent than ever for administrators and decision-makers to understand what AI really means and how it actually works. It was hard enough to parse all the risks and opportunities of artificial intelligence before ChatGPT. Imagine how hard it could become in the future with AI-powered technologies influencing nearly every facet of daily life.

We hope that this glossary will help. Drawing on the experience and extensive research of the technologists and AI experts at Mainstay, it lays out key terms related to AI – and why they matter – in a format designed for higher education leaders. At a time when it's never been more important to understand the changing role of artificial intelligence, we hope it helps to create a knowledge base that enables institutional decision-makers to navigate this fast-changing landscape. We look forward to continuing the conversations that this document will spark, and we are excited to continue to learn about the ways that AI, when implemented thoughtfully and intentionally, can transform the college and university experience for the better.



Mildred Garcia is the former president of the American Association of State Colleges and Universities in Washington, D.C.. She is currently Chancellor of the California State University system.

## Introduction

Four years ago, we released a guide for institutional leaders seeking to better understand the growing role of artificial intelligence in [every corner of higher education](#). From improving the efficiency of sprinkler systems to supporting students with [virtual teaching assistants](#), AI was already a near-ubiquitous presence on campus before ChatGPT took the world by storm.

Since then, a growing body of research has shown that AI can play a critical role in helping institutions tackle pernicious challenges – from [“summer melt”](#) to [student retention](#), [academic success](#), and [well-being](#). That's more important than ever at a time when colleges and universities are grappling with shifting demographics, declining enrollment, and the need to not just respond to, but also anticipate, the needs of an increasingly diverse student population.

But the landscape of AI is also changing fast. ChatGPT and other generative AI tools are quickly [growing in popularity](#), reshaping conversations around instruction and academic integrity, and accelerating shifts in the world of work that awaits students after they graduate.

Against that backdrop, it's more important than ever for university leaders to understand the principles that underpin artificial intelligence — so they can better design a student experience that both responds to and anticipates the needs of an ever-changing world.

---

In short, college and university leaders are facing growing pressure to get smart on artificial intelligence. And that starts with a basic understanding of what AI actually is, and how it works. Here's a glossary of key AI terms designed to help higher education decision-makers navigate the landscape.

### Why do we call it intelligence?

As early as the 1950s, when the term “artificial intelligence” was coined, engineers at Dartmouth realized that instead of telling the machine what to do in each circumstance, they could instead focus on a single prediction problem: “What would a human do?”

Prediction is not just one of the many things your brain does. It is the primary function of the neocortex, and the foundation of intelligence. The human cortex is, simply put, an organ of prediction.

It is important to note that humans (really, all animals) and machines are good at different aspects of prediction. A jet and a bird both fly, but they do so in fundamentally different ways. In “Prediction Machines,” researchers Ajay Agarwal, Joshua Gans, and Avi Goldfarb highlight similar differences in describing the potential – and limitations – of AI in cancer diagnosis. Human pathologists, as it turns out, typically make the right call when identifying cancer; machines are better at identifying the absence of cancer than its presence. In short, humans and AI make different types of mistakes. Recognizing these different abilities, and combining human and machine prediction,

---

enabled doctors to overcome these weaknesses — and, in doing so, dramatically reduce the error rate.

## **Workflow Implications**

As anyone who's tried to use ChatGPT can attest, the benefits of AI rarely accrue from simply dropping an AI tool into a workflow to increase the productivity of a given task. Deriving real benefit from implementing an AI tool often requires rethinking, or “re engineering,” the entire workflow. Similar to the advent of the personal computer, we are just at the beginning of seeing widespread productivity gains from AI.

Consider the arrival of the spreadsheet, which diminished the returns of being able to perform many calculations quickly on a calculator. At the same time, it increased the returns of knowing how to ask the right questions — to take advantage of spreadsheets' ability to efficiently analyze many different scenarios. The same is true for ChatGPT, which has given rise to an entirely new job category: “[prompt engineers](#)” who can train AI tools to produce more effective responses.

The most profound implication of AI in higher education is, likewise, rooted in AI's role as a complement to the work of counselors, admissions officers, and student affairs personnel. Institutional leaders should take care not to delegate AI strategy to their IT department, or “silo” AI within a single office. The most powerful AI tools go beyond enhancing the productivity of tasks performed.

---

Instead, the proliferation of predictions creates an imperative to rethink, and perhaps shift, an institution's strategy and approach.

## Glossary

**Machine Learning (ML):** The process that enables systems to automatically learn and improve from experience without being explicitly programmed. Machine learning generally refers to the development of computer programs that can access data and use it to learn and make predictions for themselves.

**Why It Matters:** Everything we think of as AI starts with machine learning. A program's ability to learn on its own is the foundation of artificial intelligence. That said, while all AI is machine learning, not all machine learning is AI.

**Artificial Intelligence (AI):** Machine learning that resembles what we consider human intelligence. There are two broad categories: "narrow" AI (or ANI) is programmed to perform a single task, such as predicting the weather or playing chess. "General" AI (or AGI) can successfully perform any intellectual task that a human being can.

**Why It Matters:** Currently, general AI only exists in the realm of science fiction (think Her or 2001: A Space Odyssey). That said, the landscape of narrow AI is becoming ever more sophisticated, with [Star Trek-like translation](#), personal assistants like Alexa and Siri, and increasingly precise [credit card fraud detection](#). Some evidence suggests that the latest iteration of ChatGPT, GPT-4, maybe approaching [general AI](#).

---

**Neural Network:** AI is based on the idea of neural networks -- essentially artificial brains built from silicon. These networks provide a structural, layered approach to processing data, based on the way the human brain works: each layer of processing (made up of artificial "neurons") provides the input for the next layer. Through clustering and classification, these layers recognize patterns that can, for example, differentiate a photo of a dog from a cat. It's worth noting that even the most cutting-edge ML experts do not fully understand what is encoded in an individual neuron of a neural net.

**Why It Matters:** Neural networks are what enable AI to "learn" from user behavior. AI technologies typically start with very little information, but the neural network enables them to build expertise over time based on the data they take in. For instance, the more an AI receives questions from students about FAFSA completion details, the more quickly it will recognize how to respond to those questions.

**Deep Learning:** Every single recent breakthrough in ML and AI is a result of deep learning with neural networks. Deep learning is a subset of Machine Learning which employs networks capable of learning from data that is unstructured or unlabeled. This type of learning uses neural networks to extract increasingly subtle and complex patterns, allowing for more sophisticated tools like accurate speech and facial recognition.

**Why It Matters:** Deep learning is particularly

---

important for AI tools like chatbots, which interact directly with humans using natural language. Deep Learning allows AI to learn how to [interpret natural language](#) like abbreviations and emojis, enabling it to more effectively communicate with students. Other news worthy applications include [board games](#) and [video games](#).

**Chatbot:** The user interface for most non-AI computer programs is a mouse, keyboard, or touch screen. Chatbots provide a different kind of user interface for AI systems, one that uses speech (either spoken or typed). These AI programs range in sophistication from relatively simple and rule-based (e.g., providing a canned response to a specific question) to more complex and AI-enabled (able to parse human language and learn from previous conversations to improve accuracy constantly).

**Why It Matters:** It is easy to build a simple chatbot, but complex to build a genuine AI chatbot — which is why the arrival of ChatGPT has taken the world by storm. Because they can respond to a nearly limitless number of users at once, chatbots have the potential to provide real-time support at unprecedented scale — which, in the context of higher education, is helping institutions boost enrollment and student success while enabling advisors to focus on students who need more hands-on, personalized guidance.

**Large Language Models:** A language model is a statistical tool that determines the probability

---

of a given sequence of words. Provide a language model with a bunch of text, and the model determines a “reasonable continuation” (in the [words of Stephen Wolfram](#)) based on a corpus of sentences in its database. Of course, the more existing text that a language model has to review, the better it will perform — hence, the advent of large language models (LLMs), which are based on millions or even billions of sentences.

**Why It Matters:** ChatGPT has thousands of lifetimes' worth of text in its mechanical “brain”, so it's no wonder that it's such an effective tool. But at its core, all it's doing is remembering and predicting. It remembers a giant corpus of text — a huge portion of the internet — that it's been trained on. And it predicts responses by adding one word at a time to fit whatever prompt you provide. That means it can provide incredibly impressive responses to questions. It also occasionally hallucinates and down right lies. It doesn't know what year it is and blatantly fabricates facts with utter confidence. That's what makes ChatGPT and other LLMs so powerful, and so potentially harmful.

**Emergent Abilities:** In the context of artificial intelligence, an emergent ability is one that is not present in smaller models but “emerges” in larger models. Importantly, what makes emergent abilities special is that they can't necessarily be predicted by evaluating the performance of smaller models. That's why ChatGPT's ability to compose a sonnet or produce a delicious recipe is so surprising — until it started to happen, researchers

---

weren't sure it would be possible at all.

**Why It Matters:** The more sophisticated AI becomes, and the larger the models we use to train it, the more emergent abilities we're likely to discover. Consider the way that LLMs can translate from one language to another: they're trained on a vast amount of text data from many languages, allowing them to learn those languages' patterns and structures. The models use that knowledge of grammar, vocabulary, and context to comprehend the meaning of a sentence in one language and then generate an equivalent sentence in another language.

A similar emergent ability can occur in biology, for LLMs that are trained on biological data to comprehend the language of mRNA, molecular biology, and the nervous system. By analyzing complex biological texts, LLMs can uncover hidden relationships and decipher the meaning behind genetic sequences, molecular interactions, and neurological processes. This deep understanding can aid in more accurate diagnoses, personalized medicine, and drug discovery.

Of course, this rapid evolution of AI comes with both opportunities and risks — and will require AI developers and policymakers alike to carefully consider the implications of the new models they're building.

**Generative AI:** A broad category that refers to AI that can create – or “generate” – data,

---

including text, images, video, or audio. ChatGPT is a text-based form of generative AI; other tools, such as DALL-E for visual art and Runway for video, are rapidly growing in popularity and prominence.

**Why It Matters:** The rise of generative AI represents a huge step in the evolution of artificial intelligence. Rather than only being able to make narrow, specific decisions based on input from humans, generative AI can create entirely new content. As with all developments in the world of AI, this opens up profound possibilities, but also new challenges and uncertainties — as we've already begun to see with the rise of ChatGPT.

Wharton Professor Ethan Mollick's [recent blog post](#) on generative AI offers a helpful framing for how to think of the technology: it's helpful to think of ChatGPT and other similar tools as more like interns than like robots — “weird, somewhat alien interns that work infinitely fast and sometimes lie to make you happy, but interns nonetheless.” That approach, which he details extensively in the post, can help users (whether administrators, faculty, students, or anyone else) engage and communicate with generative AI tools in ways that are both productive and less likely to be frustrating — or, at least, no more frustrating than working with an intern can be.



Task Category	What it Means
<b>Just Me Tasks</b>	Tasks for which AI is not useful and only gets in the way (at least for now).
<b>Delegated Tasks</b>	Tasks that are assigned to AI, but also have human review and oversight. Particularly effective for tasks that are low importance, but time consuming.
<b>Centaur Tasks</b>	Tasks in which AI is deeply integrated into workflow and work (e.g., technical writing that benefits from AI's ability to restate and streamline).
<b>Automated Tasks</b>	Tasks that are left completely to AI, with no human review or oversight required.

## Generative Pretrained Transformer

**(GPT):** A Generative Pre-Trained Transformer (GPT) is a type of language model specifically designed to generate text responses to a given text-based input. "Generative" refers to the model's ability to create original text. "Pre-trained" refers to the fact that the model has reviewed many thousands (or in the case of ChatGPT, billions) of sentences. "Transformer" refers to the specific type of deep learning that a GPT does, which involves a sophisticated approach to reviewing all of the data in its model.

**Why It Matters:** The phenomenon of ChatGPT has its roots in 2017, when a research paper called "[Attention Is All You Need](#)" detailed the potential of the transformer, and inspired a wave of innovation in AI, particularly generative AI and large language models.

Fast-forward a few years as the technology became more sophisticated, and in a matter of months after its introduction, the GPT model (which of course underpins ChatGPT) has become the best-known AI language model. In the coming months, it's likely to become the basis for all sorts of new uses of AI — which means that understanding how it works will be increasingly important for anyone navigating a crowded landscape of AI-enabled education technologies.

## Natural Language Processing (NLP):

The ability of a computer to parse accurate meaning from human conversation (or "natural language"). By using contextual clues, NLP can help machines make sense of what humans are trying to say -- for example, parsing the difference between trying to reach the accounting department and finding out the requirements for the accounting major.

**Why It Matters:** Natural language processing isn't new to higher education. In fact, admissions offices, beset by a torrent of questions when email became popular in the late nineties, were among the first to use NLP to streamline operations — and ensure rapid response to Gen X students who eschewed once-popular phones and call centers. Today, NLP can be the difference between a good user experience and a bad one. Just ask anyone who has been stuck on the phone with an automated system that can't seem to send them to the right department. Perhaps most importantly, NLP is getting better at learning the particular

---

linguistic stylings of today's students. Albeit nascent, advances in NLP hold the potential to parse almost any form of communication, meaning that virtual assistants or chatbots can understand a question even if it's asked in a string of emoji.

**Supervised Learning:** A type of machine learning in which human programmers map training data (for example, historical student records) to a single correct output (for example, whether a student graduated on time). Then, the algorithm finds the patterns and mathematical relationships that connect the inputs and outputs -- so it can predict what the output will most likely be for inputs that it's never seen before.

**Why It Matters:** Supervised learning is critical for higher education institutions that have access to anonymous student information but aren't always sure how to use it effectively. By using supervised learning, institutional leaders can provide AI with information and let the program find insights that are, in many cases, counterintuitive. Because data plays such an enormous role in the supervised learning process, the breakthroughs often come from departments with the greatest access to data. If you're evaluating a vendor's machine learning credentials, start by asking about their unique data set.

**Turing Test:** A concept proposed in the 1950s to test the sophistication of artificial intelligence by determining whether answers

---

provided by a computer are distinguishable from those provided by a human. In 2014, a chatbot "beat" the Turing Test for the first time, convincing users it was a 13-year old boy.

**Why It Matters:** Passing the Turing Test is an indicator that a machine can understand human language and the flow of human logic – including slang, metaphors, idioms, or euphemisms – and respond appropriately. This is often an aspirational standard, but coming close means a better user experience for humans interacting with AI.

**Augmentation vs Automation:** Automation describes the phenomenon in which the return-on-investment of machines handling all functions of a particular process is greater than the return of including humans in the process. A job is augmented when machines take over some, but not all, tasks.

**Why It Matters:** Successfully implementing AI in a higher education context often depends on the right balance of automation and augmentation. By automating the easy questions, AI frees up humans' time to focus on the unusual or difficult questions and cases. A speedy reply to a support question is much less frustrating than spending minutes or hours on hold. But it's just as critical to create channels for machines to elevate more complex or personal questions to human advisors – who can engage with students in ways that even the most sophisticated AI isn't able to do. Simply put, the most effective uses of AI stem

---

from not just thinking about what it can do for your institution — but also considering what your institution and AI can do together.

In the words of Accenture's Paul Daugherty, "Those that use machines merely to replace humans will eventually stall, whereas those that think of innovative ways for machines to augment humans will become the leaders of their industries."

**Prompt Engineering:** The act of developing the best "prompt" to get a Large Language Model or other AI database to accomplish a desired task. As more and more ChatGPT users are coming to understand, making the most of generative AI takes much more than telling it to "do X" and expecting that to be enough. A recent [New York Times article](#) recommends a few so-called "golden prompts" designed to get the most out of ChatGPT, including telling the chatbot to "act as if" it is an expert in a given field (e.g., personal fitness; tutoring), and encouraging it to ask for more information.

**Why It Matters:** Effective collaboration between humans and AI requires sophisticated, specific prompts that recognize both what AI can do well — and where it still has a lot to learn. Prompt engineering has quickly become a [hot job](#), and savvy educators are finding ways to incorporate it into the classroom. Think of an assignment like "Write an essay about Great Gatsby with the help of ChatGPT. You will be graded not only on the essay you produce, but

your full dialogue with ChatGPT that helped you get there."

## About AASCU

*The American Association of State Colleges and Universities (AASCU) is a Washington, D.C.-based higher education association of 350 public colleges, universities, and systems whose members share a learning- and teaching-centered culture, a historic commitment to underserved student populations, and a dedication to research and creativity that advances their regions' economic progress and cultural development. These are institutions Delivering America's Promise.*

## About Mainstay

*At Mainstay, we believe one conversation can spark a brighter future. Our engagement platform makes it easy for institutions to start and measure conversations that drive positive learner outcomes at scale. Everything we do – from our rigorous research methods to the architecture of our conversational AI – is designed with our mission in mind: to spark progress with conversation.*