

A UNIVERSITY LEADER'S

Glossary for AI & Machine Learning





Foreword

A little over a decade ago, graduation rates at Georgia State University hovered around 30%. Seven out of ten students who came to us left with no degree. Numbers were even lower for underrepresented minorities and low-income students. Through the systematic use of analytics and a commitment to institutional change, Georgia State has been able to raise our graduation rates by more than 70% and eliminate disparities in completion rates based on race, ethnicity and income level.

While significant gains were being made in graduation rates, another problem was growing: “summer melt.” Summer melt is the percent of confirmed incoming freshmen who never show up for fall classes. This is a pernicious challenge nationally, especially for students from underserved, urban backgrounds. As many as 25 percent of confirmed college freshmen from some urban school districts never make it to a single college class. They simply melt out of the college-going population in the summer between finishing high school and starting college, most often casualties of the battery of bureaucratic tasks—financial forms, signatures, immunizations records, loan applications—that they are asked to complete.

What if students had a personal guide to help them navigate this complex bureaucracy? In 2016, Georgia State partnered with AdmitHub to develop a novel answer: Pounce. Named after Georgia State’s mascot, Pounce is an AI-enhanced chatbot available to students 24/7 on their smart devices. With the support of AI, Pounce provides students immediate and conversational answers to almost 3,000 commonly-asked questions and nudges students to complete important tasks. In its first three months of operation, Pounce exchanged more than 200,000 texts with Georgia State’s incoming freshman class of 2016, helping them to pose and to resolve problems in real time. With the help of Pounce, summer melt at Georgia State is down by 37%, meaning that 350 additional students—mostly low income—are in their seats ready to go the first day of fall classes who were sitting out the college experience just three years earlier.

New technologies have the power to transform. For the first time, they may allow us to deliver real-time, individualized support to all students—not just the privileged few who enroll at elite institutions with low student-to-faculty ratios.

To realize the promise of technology in helping students navigate and complete their college education, it’s critical that we establish a common lexicon to bridge the divide between technology and practice. This glossary describes and defines terms that institutional leaders should know. It puts new technologies in a context that will be familiar to institutional leaders. Learning the language of emerging technologies is a preliminary but essential step in exploring how these tools can fulfill our shared mission of supporting the success of students from all backgrounds.

Tim Renick

Senior Vice President for Student Success, Georgia State University



Introduction

Artificial intelligence (AI), it seems, is infiltrating every corner of higher education. From improving the efficiency of sprinkler systems to supporting students with virtual teaching assistants, AI has quickly become a near-ubiquitous presence on campus.

Perhaps with good reason. Institutions are being asked to do more with less, as they grapple with shifting demographics and not just responding to, but anticipating, the needs of today's students as they navigate the complexity of financial aid and admissions, campus life and course schedule.

And early returns suggest that AI can play a role in helping institutions tackle pernicious challenges, from “summer melt” to student engagement in class.

In response, a growing number of products are touting AI and machine learning as part of their sales pitch. But what's the difference between AI that can actually help students succeed and simple decision trees or marketing ploys? How can institutional decision-makers cut through the jargon to separate the sort of AI that can actually learn over time, and play a role in improved outcomes?

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.

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, engineers recognized they could instead focus on a single prediction problem: “What would a human do?”

Prediction is not just one of the things your brain does. It is the primary function of the neocortex, and the foundation of intelligence. The cortex is 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, overcame these weaknesses -- and, in doing so, dramatically reduced the error rate.

WORKFLOW IMPLICATIONS

Of course, the benefits of AI rarely accrue from simply dropping an AI tool into a workflow to increasing the productivity of that task. Deriving a real benefit from implementing an AI tool can require rethinking, or “reengineering,” the entire workflow. Similar to the advent of the personal computer, it will take time to see widespread productivity gains from AI.

The arrival of the spreadsheet, for example, diminished the returns to being able to perform many calculations quickly on a calculator. At the same time, it increased the returns to being good at asking the right questions in order to fully take advantage of the technology’s ability to efficiently analyze many different scenarios.

The most profound implication of AI in higher education will, likewise, stem from the role of AI in complementing the work of counselors, admissions officers, and student affairs personnel. Institutional leaders should take care not to delegate AI strategy to their IT department. 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 organization’s strategy and approach.



Glossary

Here's a glossary of key AI terms designed to help higher education decision-makers navigate the landscape.

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 for themselves and make predictions.

***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, AGI 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. We don't have self-driving cars yet, but we do have *Star Trek*-like translation, personal assistants like Alexa and Siri, and increasingly precise credit card fraud detection. Narrow AI is invisibly making our lives significantly better, which is why we shouldn't underestimate its potential in higher education.*



Neural Network

AI is based on the idea of neural networks — essentially an artificial brain 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 network.

***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 convolutional neural networks. Deep learning is a subset of machine learning which employs networks capable of learning from data that is unstructured or unlabeled. It uses neural networks to extract increasingly subtle and complex patterns, allowing for things 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 newsworthy 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, like Siri and Alexa, provide a different kind of user interface for AI systems 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. ([There's a whole website about it.](#)) A chatbot's ability to understand run-on questions, handle misspellings, and deal with emojis are all key indicators of a chatbot powered by AI.*

Because they can respond to a nearly limitless number of students at once, chatbots have the potential to provide unprecedented support in real-time — which can streamline processes like admissions and enrollment and enable advisors and support staff to focus on students who need more hands-on, personalized support.

Generative Adversarial Networks (GANs)

A recent breakthrough that enables machine learning programs to learn with little or no data. GANs are a combination of two neural networks that compete against each other at extremely fast speeds, enabling them to learn more quickly than a human ever could. The highest profile example of this is Alpha Go, which recently beat the reigning champion of the board game Go four games to one. Another famous example, of course, is [WarGames](#).

Why it Matters: *AI that uses generative adversarial networks can improve much faster, and with less data, than through traditional machine learning. This means it can more quickly adapt to your particular campus or students.*



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 emojis.*

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 this technology, the breakthroughs often go to those 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 other non-literal semantic usage, 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 occurs when the return to machines handling all functions is greater than the returns to including humans in the process. A job is augmented when machines take over some, but not all, tasks.

***Why it Matters:** A speedy reply to a support question is much less frustrating than spending minutes or hours on hold. By automating the easy questions, AI frees up humans' time to focus on the unusual or difficult questions and cases.*