AI GLOSSARY FOR EXECUTIVES

As Artificial Intelligence continues to proliferate in the business world, you'll need to be able to discuss it with both technical and non-technical audiences.

That means cutting through the hype and occasional confusion around concepts and technologies. Here's a primer on key AI terminology for IT and business leaders.

Artificial Intelligence (AI):

AI is a broad umbrella that encompasses multiple disciplines or technologies, including some of the terms listed below. Here's a straightforward definition shared with us recently by Bill Brock, VP of engineering at Very: "AI, simply stated, is the concept of machines being able to perform tasks that seemingly require human intelligence." Since AI is the term, let's add one other definition, this one from our recent Harvard Business Review Analytic Services report, "Executive's guide to AI": "Artificial intelligence is the science and engineering of making intelligent machines. This includes intelligence that is programmed and rules-based, as well as more advanced techniques such as machine learning, deep learning, and neural networks."

Machine Learning (ML):

One of the most common examples of AI in practice today comes from machine learning, which the HBR Analytic Services report defines as: "A way of achieving AI without complex programming, rules, and decision trees. Instead, with machine learning, data—often in very large amounts—is fed into an algorithm so the algorithm can train itself and learn." At the moment, this typically means...

Supervised ML:

When people refer to ML today, they're usually referring to supervised ML. This entails supplying the ML algorithms with data—usually lots of it. As Brock defined it for us, supervised ML "teaches the machine by giving information on the parameters of the desired categories and letting the algorithms decide how to classify them." The "supervised" basically means that this type of ML requires plenty of human input and monitoring, especially in terms of the datasets that power its algorithms.

Unsupervised ML:

This contrasts with unsupervised ML, which does not require training data to learn. This makes it more complex and currently less common, though that appears likely to change over time. "Unsupervised machine learning is already being used—or is in the process of being developed—for applications such as image recognition, cancer detection, music composition, robot navigation, autonomous driving, and many other innovations," Brock told us.

Deep Learning:

Deep learning is commonly referred to as a subset—aka "deeper"—form of machine learning. The HBR Analytic Services report offers this definition, adapted in part from Skymind.ai: "A form of machine learning but with more layers. Deep artificial neural networks are a set of algorithms that have set new records in accuracy for many important problems, such as image recognition, sound recognition, recommender systems, etc." As Skymind explains here: "Neural networks are a set of algorithms, modeled loosely after the human brain, that are designed to recognize patterns." To describe in a slightly different way, deep learning can be thought of as a way to use a set of more complex algorithms to enable a machine to mimic the human brain as much as possible.

Robotic Process Automation (RPA):

In plain terms, RPA is the use of software bots to automate repetitive, rules-based processes. Aaron Bultman, director of product at Nintex, recently explained it to us like this: "RPA is a form of business process automation that allows anyone to define a set of instructions for a robot or 'bot' to perform. RPA bots are capable of mimicking most human-computer interactions to carry out a ton of error-free tasks, at high volume and speed." People don't always agree about whether RPA should be considered AI, but this depends at least in part on how you define the latter umbrella category. For example, the HBR Analytic Services report notes that RPA should qualify, according to its definition of AI.

Explainable AI:

Explainable AI is an AI system where the outcomes or results can be understood in proper context by humans. This differs from the "black box" model of AI, where people cannot trace or understand how AI arrived at a particular result. Stephen Blum, CTO and cofounder of PubNub, recently shared this helpful comparison to help define the term: "Explainable AI can be equated to 'showing your work' in a math problem. All AI decision-making processes and machine learning doesn't take place in a black box—it's a transparent service, built with the ability to be dissected and understood by human practitioners."

Natural Language Processing (NLP):

NLP (like pretty much every other term in this glossary) is a hot topic because it pertains directly to the growing number of human-to-machine voice interactions in both professional and personal contexts. The folks at SAS have a good, clear definition: "Natural language processing (NLP) is a branch of artificial intelligence that helps computers understand, interpret and manipulate human language. NLP draws from many disciplines, including computer science and computational linguistics, in its pursuit to fill the gap between human communication and computer understanding."

Computer Vision:

If NLP is largely about how machines hear (and respond to) humans, computer vision could be thought of as how machines "see"—not just humans, but potentially any image or even video stream. Wikipedia has a succinct definition: "Computer vision is an interdisciplinary field that deals with how computers can be made to gain high-level understanding from digital images or videos. From the perspective of engineering, it seeks to automate tasks that the human visual system can do." Computer vision is likely to continue drawing attention not just for its productive applications, but also for potential risks such as AI bias and other concerns. (To be sure, AI bias is a potential issue across many types of AI systems).

AIOps:

Here's one AI segment that refers specifically to IT operations. According to Gartner's definition, "AIOps combines big data and machine learning to automate IT operations processes, including event correlation, anomaly detection, and causality determination." Put another way, it's about using AI and data to automatically identify or handle issues that would have once depended on a person to take care of manually. You can consider AIOps a more specific AI-driven category of the broader automation trend in IT, but remember that not all forms of automation would necessarily be considered AIOps.

If you want a deeper dive on several key AI topics and trends, consider these related reads:

AI in the enterprise: 8 myths, debunked »

AI vs. machine learning: What's the difference? »

What is explainable AI? »

How to explain Robotic Process Automation (RPA) in plain English »

How to explain machine learning in plain English »

How to explain deep learning in plain English »

AI bias: 9 questions IT leaders should ask »

An Executive's Guide to AI »