

AI: Why a Systems Engineering Approach Is Essential

A holistic view of artificial intelligence systems can help organizations to unlock the full potential of the emerging technology.

By David R. Martinez

Artificial intelligence (AI) is already turning industries on their head, and the technology is poised to make an even greater impact on the world in the years to come. Doctors are using AI tools to help with diagnostics, carmakers are working to make autonomous vehicles a widespread reality, and nearly all of us each day view online or mobile advertisements that were selected specifically for us by an algorithm.

Too often, though, business and IT leaders take a limited view of AI. They often focus almost exclusively on machine learning (ML) – sometimes even using “ML” as a synonym for “AI.” But AI technologies are, in fact, key enablers to complex systems. They require not only ML technologies, but also trustworthy data sensors and sources, appropriate data conditioning processes, and a balance between human and machine interactions. Bringing all of these disparate sub-components together requires a system engineering approach – an approach that is, unfortunately, lacking in many organizations’ views and implementations of AI.

To get the most out of their AI initiatives, business and IT leaders must consider the entire ecosystem surrounding their AI systems, and make plans to recruit and retain talented multi-disciplinary teams that can help them at every stage of development and deployment.

AI System Architecture

A comprehensive view of AI should include the following sub-components:

Sensors & Sources – There’s good news and bad news about data collection, and they’re both the same: We’re now collecting more data than ever before, with some observers estimating that 90 percent of all of the world’s data has been gathered in the past two years. And 80 percent of that information is in the form of unstructured data (for instance, photographs, videos, speech, texts, etc.), which would not scale if stored in standard

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

relational databases. All of this new data represents an enormous opportunity in the realm of AI, but it also represents a significant challenge.

Data Conditioning – With so much data at their fingertips, data scientists must devise new ways to pre-process data and eliminate “noise” so that ML algorithms can make sense of it all. Within this sub-component of an end-to-end AI system, data scientists use techniques to transform raw data into information. The resulting data information is then input into an ML sub-component to extract knowledge.

Machine Learning – Once structured and unstructured data have been passed onto ML tools, specialists can begin trying out different types of techniques to glean knowledge from the conditioned data. There are many classes of ML techniques, such as, unsupervised learning where no labeled data is needed, supervised learning requiring labeled data, and reinforcement learning well matched to cases where one can identify goals, actions, and rewards without necessarily requiring labeled data.

Human-Machine Teaming – Again, ML tools are often mistaken for the entirety of AI. But by taking a step back, it is easy to see that ML is just one sub-component along a chain. Even after machines have made their conclusions by converting information into knowledge, the process of actually acting on the derived knowledge from ML algorithms must be followed by an interaction between humans and machines as teams. For instance, an ML algorithm might take data from scans such as MRIs or x-rays to detect potentially cancerous cells. But after that, it will be up to human doctors to assess the accuracy of the ML tool’s assessment and conduct additional tests for verification. The human-machine teaming enables taking knowledge and deriving insight.

Users (Mission) – This is the part where users actually consumed the derived insights to decide what action must be taken. It’s where doctors adjust their treatment plans, or national defense services formulate or alter their courses of action, or driverless cars steer clear of obstacles. It’s the outcome – the whole purpose of an end-to-end AI system.

Modern Computing – Undergirding all of these processes are a number of modern computing technologies, including central processing units (CPUs), graphics processing units (GPUs), tensor processing units (TPUs), neuromorphic computing, quantum computing tools to name a few from

more matured to less matured computing technologies. The choice of computing technology depends on the amount of data, type of algorithms, and the computing environment. For example, an AI system can be deployed to a cloud computing environment, with more available size, weight and power budget, to edge computing, with more constrained size, weight, and power budget, or a hybrid approach.

Robust AI – This last sub-component of a holistic approach to AI is really a collection of smaller sub-systems that organizations must continually apply to their AI tools to ensure that they are as accurate and safe as possible. Most basically, organizations must ensure that their ML algorithms are “explainable” – that is, that users can at least loosely understand the derived knowledge before taking actions. Data and algorithms must also be closely monitored for bias; some facial recognition programs, for instance, have struggled to accurately identify people of color. Security tools and practices must be applied to prevent data and/or ML techniques from being taken over by adversaries. And policies and training must be put in place to ensure safe and ethical use of AI systems.

Looking Ahead

I often get asked, “What are the practical applications of AI for a given industry?” Typically, AI makes a lot of sense for use cases where people are performing routine tasks, where large amounts of data are involved (volume), where users are interested in drawing timely (velocity) predictions/classification or characterizations from the data, and where there’s a need for scalability from different types of data (variety).

Currently, most applications of AI are limited to producing what I call “content-based” insights. These are important and powerful use cases that make users more efficient and improve decision-making. But over time, AI will evolve to include more collaboration-based applications, with multiple human-machine teams working together. And eventually, AI systems will increasingly achieve context-based insights that more closely mirror human decision-making by being aided with more sophisticated intelligent machines.

To participate in this evolution, and to accrue the associated benefits, organizations must break with their old ways of thinking and embrace an

approach to AI that emphasizes multiple sub-components working together in an integrated system engineering way.

[David R. Martinez](#), co-instructor of the MIT Professional Education Course [Engineering Leadership in the Age of AI](#), is Lincoln Laboratory Fellow at MIT Lincoln Laboratory. His focus is on research and applications of artificial intelligence, high-performance computing, digital transformation, and teaching