accenture

ACCENTURE FEDERAL SERVICES

HOW MACHINES THINK
(& why government executives should care)

What makes Artificial Intelligence (AI) a unique technology is its ability to "learn" from experience. Instead of programming applications with predefined rules to perform specific functions, data is used to teach the application to complete a desired task. This underlying approach – machine learning – relies on large scale processing of structured or unstructured data to build intelligence, enabling the computer to make sense of data, recognize patterns and ultimately to support desired actions or conclusions.

Given the importance of machine learning within AI, knowing how it operates can be valuable to even the non-data-scientists in understanding how they can use AI and how the technology may evolve in the near future. Here, we discuss different approaches to machine learning, potential mathematical models and problem types, and how to train and test specific algorithms.

LEARNING MODES

Machine learning can be used to interpret data in one of three ways: Supervised learning, unsupervised learning and reinforcement learning.

Supervised

This method works well when there is a lot of well-understood data readily available. When we know what decisions have emerged from the data before, we can teach the machines to replicate that decision-making process. It's analogous to a rookie cop's ride along with a veteran officer: The AI effectively rides alongside the existing process to see how it's done.

When ample data is on hand, supervised learning is typically the preferred mode for implementing machine learning. It's easy to measure success when we know what the outcomes ought to look like. Even when the preferred outcomes are unclear, supervised learning can at least take an initial pass at the data by applying "pattern matching" processes. In this mode, machine learning can find similarities in the data that enable it to establish generalities and draw initial conclusions.

Unsupervised

Many agencies and other entities are awash in data. They're ingesting sensor information, user metrics and a vast diversity of other data points, but haven't yet begun to interpret all that information or put it toward actionable ends.

In this scenario – big swaths of untested data – the unsupervised learning model asks the computer to first seek out patterns. By identifying these clusters or groupings, the machines can reduce the data to more limited set. In a simple example, unsupervised learning might distinguish between cars and trucks, in which case operators would cut by half the number of vehicles in the data set that might be tested for cargo loading limits.

By identifying patterns and groupings, unsupervised learning can lay the groundwork for initial actions that in turn form the basis for future supervised learning iterations.

Reinforcement

Suppose you have no data but still want to train the machines to make good decisions. Reinforcement learning makes this possible. How do people learn not to touch a hot stove? By touching once and not doing it again. The same premise applies here. Suppose researchers wanted to simulate a driverless car. Each time the computer took a "wrong" action – braking suddenly, or turning too sharply – the researchers could give the machine learning negative feedback. Positive choices likewise get "rewarded" in the system. In this way it is possible to build up a knowledge base and a system of decision making, even without having actual production data on hand.

HOW DOES MACHINE LEARNING WORK?

In machine learning you start with data

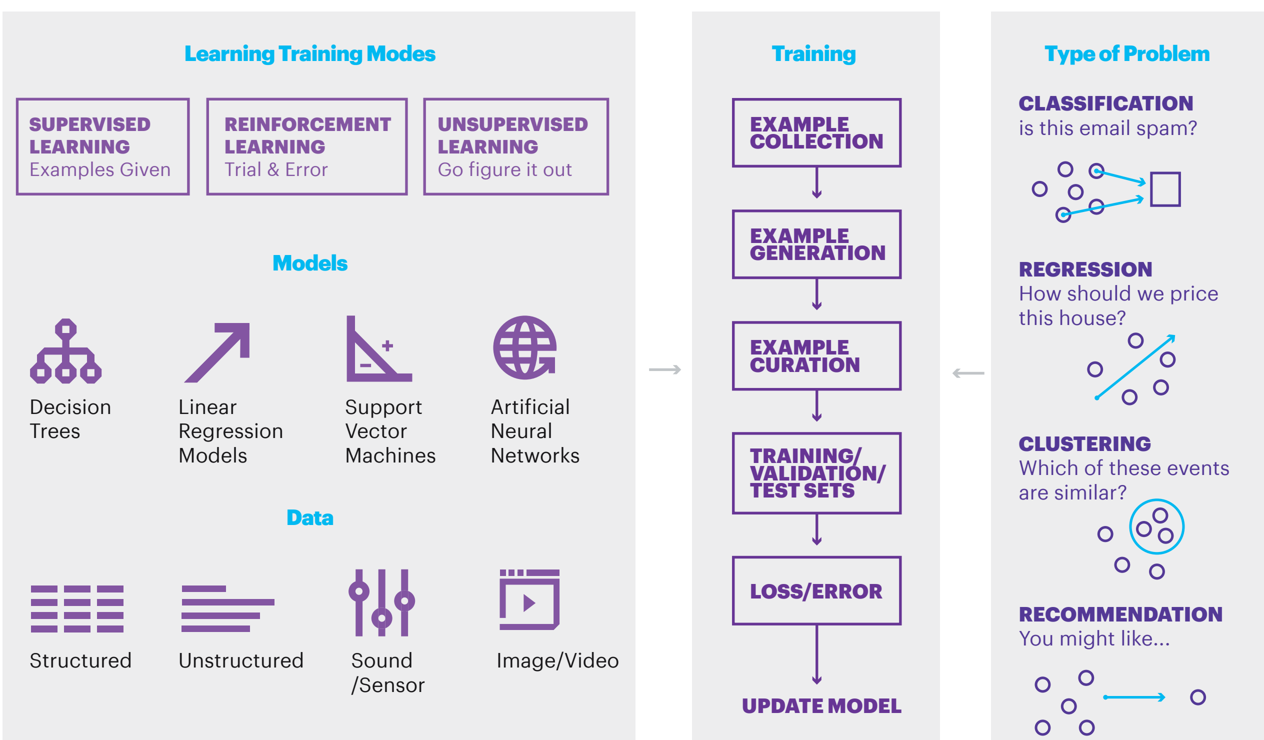

MODELS & DATA

There are multiple mathematical models available to help computers build connections between data and outcomes: Decision trees, linear regression models, support vector machines and artificial neural networks.

The important thing to know is that these various modes represent escalating levels of complexity. They get progressively harder to explain to the layperson. For an agency that may be called upon to justify the outcomes of its AI implementation, a simpler model may offer a more accessible narrative.

It's also worth noting that designers of machine learning have the ability to draw from diverse forms of data. In addition to structured data, where each formatted in a way that is easy for the computer to understand, machine learning also can cull from sensor feeds, audio data and video images. The definition of "data," in the context of machine learning, runs broad and deep.

TRAINING

Training processes lie at the heart of the machine learning enterprise. This may begin with "collection" – culling data from existing feeds and repositories – or it may start with "generation," creating new data specifically in order to drive AI.

Government may be especially well-situated to benefit from the use of large-scale data sets. The extensive nature of some government data repositories creates a unique opportunity to establish especially valuable training datasets.

A curation phase will identify within the existing data those specific examples that are most representative of the desired process and its outcomes. By curating the best, more representative data samples, developers are able to maximize the machines' learning capabilities.

Along with training comes validation testing. In this phase developers go beyond the curated data, giving the machines the chance to interpret unfamiliar data sets. By examining previously untried data, they are able validate the machine's interpretive abilities.

To round out the training developers will look at data loss and errors in the outcomes, then feed that information back into the system, effectively enabling the computers to learn from their own mistakes. As the cycle is repeated, the outcomes are refined and enhanced.

A challenge in using training data is that people are prone to making mistakes in labeling the data. This means that AI performance may be limited by the accuracy of the underlying dataset. Given the subjective nature of these assessments, having multiple sources contribute to the training datasets can help to improve quality. Inter-annotator agreement (IAA) is a method used to assess the quality of their labeling by measuring level of consensus between sources. For example, the IAA score can be used to identify the source of error– was it a recognition error or a classification issue? — to improve performance.

> **The White House has declared that "[h]igh quality datasets are critically important for training many types of AI systems. The Federal government facilitates AI innovation by investing in shared public datasets. The American AI Initiative calls for agencies to make Federal data, models, and computing resources more available to America's AI R&D experts, researchers, and industries."**
>
> **For federal agencies, understanding how AI learns from data can help them identify government datasets with the greatest potential to foster widespread innovation. This insight can also be used to understand the type or level of prep required to make data broadly usable.**

TYPE OF PROBLEM

In considering how best to implement machine learning, it's helpful to think about the type of problem at hand.

Machine learning can be used as a classification tool: Is this email spam? It can be used in a regression model, which offers a predictive capability. It can apply simple clustering, grouping together similar data points. Or it can be leveraged to offer recommendations: A citizen visiting a government web site might be steered toward a particular service or a needed form based on answers to filtering questions.

The type of business problem at hand typically will determine the way in which machine learning is applied. In fact, agencies will likely encounter a series of decision points on the road to machine learning. What types of data are available? What processes are we seeking to improve? What expertise do we have on hand to interpret the AI outputs? By seeking thoughtful answers to these and other related questions, government agencies will be best able to position themselves for a successful machine learning implementation.