

Complete Healthcare Compliance Manual

Artificial Intelligence and Compliance Programs

By Alan Brill,^[1] CISSP, CFE, CIPP/US, FAAFS

What Is Artificial Intelligence in Relation to Compliance Programs?

Artificial intelligence (AI) is simply the application of computer processing to simulate the actions of a person. One of the earliest AI systems was, in fact, a medical application called “MYCIN.” The program was designed to diagnose bacterial infections and recommend appropriate medications, with the dosage adjusted for the patient’s body weight. Viewed from current technology, MYCIN was quite primitive, using an inference engine with approximately 600 rules derived from interviews with expert human diagnosticians. MYCIN was originally written as part of a doctoral dissertation at Stanford University and was never used in actual medical practice for legal and ethical reasons (along with limitations related to the technology of the day.) But it formed the basis for continued experimentation and development.

There are aspects of AI that are continuously evolving, but there are some basic terms that are worth understanding.

Machine Learning

This is a subset of AI in which the computer’s algorithms (essentially the AI computer program) are able to modify the computer’s actions with the objective of improving through experience. In many settings (medicine, aviation, or automobiles for example), learning through actual experience could be counterproductive. For example, imagine stating that a number of airplane crashes happened because the airplane’s computer program hadn’t learned to deal with unexpected turbulence yet. So, typically machine-learning systems are given what is called “training data” in order to learn how to function. Provided with data and outcomes, the software should be able to modify processing to result in better—or more accurate—performance.

Rule-Based Machine Learning

This involves systems that evolve a set of rules by which a program makes decisions. MYCIN, for example, had hundreds of rules. In a rule-based system, the program uses its experience to identify which rules are more or less useful and to modify the rules or the weights given to them to improve processing outcomes.

Deep-Learning Systems

These systems are generally characterized as having multiple layers of processing, using layers that go from general to specific analysis, that are often being applied to large networks of unstructured data. An example might be a system that is designed to read human handwriting. Clearly, experience tells us that this isn’t easy, as there are as many variations in handwriting as there are people. But there are generalizations that can be used to do some preliminary analysis (for example, that a given character is uppercase) that can lead to deeper analysis to try to figure out which character is being represented.

Cognitive Computing

This is generally thought of as an alternative name for AI. There is no widely accepted definition, but you may run into the term as a synonym for AI.

Computer Vision

This is a subset of AI that focuses on how computers use digital images (still or video) in their processing. An assembly line for drug packaging can use computer vision technologies, for example, to inspect sterile vials of injectable medication to ensure that labels have been affixed and that the top is properly sealed. This can be done at the speed of the assembly line, with a mechanical “kicker” used to eject vials not meeting the specifications.

Natural Language Processing (NLP)

This is the part of AI that focuses on enabling interactions with humans by interpreting their language. It includes automated language understanding and interpretation, automated language generation, speech recognition, and responding with spoken responses. In the past few years, this has gone from the lab to millions of homes, with digital assistants like Siri and Alexa ready to listen and respond to requests. In many cases, the vendors of these systems seek user’s permission to use recordings of these interactions to improve the system’s performance. This has been recognized as a privacy issue. In at least one case, recordings of interactions with a digital assistant have been subpoenaed in connection with a murder trial.

Chatbots

These are very similar to natural language processors, although they were developed to replace human operators in online text-based chat systems. For example, a chat system could be fielded to answer routine questions and to forward difficult or complex ones to human operators, thus reducing the workload on the humans. In some cases, these can use text-to-speech processing to enable spoken responses.

Graphics Processing Units (GPUs)

These are specialized processors operating within computers designed to process image data. A GPU could be used to create the images displayed on a computer’s screen. However, these powerful units have been used for many other purposes. A current example is that GPUs are often used to process cryptocurrency transactions (a process known as mining, which can be very profitable). Specialized computers using massive numbers of GPUs have been developed as mining machines for cryptocurrency processing.

Internet of Things

This is a term that refers to the abundance of devices that can connect to a network that are not traditional computers (or smartphones or tablets). Ranging from smart lightbulbs to cameras to refrigerators, they enable remote control and monitoring of connected devices. There has been enormous growth in the number of medical devices that can connect to a network. Unfortunately, there are serious security concerns that have resulted in Food and Drug Administration (FDA) warnings relating to several devices, including network-connected infusion pumps.^[2]

Application Programming Interfaces (APIs)

This refers to the connections between devices and the rules by which these connections are made and interpreted. So, for example, if an AI-based analytic engine is to be given access to a particular database, an API defines the way the systems interact, how requests are made, and how they are responded to.

AI and the Compliance Function

When it comes to AI, compliance professionals are presented with what could be characterized as a double-edged sword. On one hand, AI represents an opportunity for compliance professionals to automate certain compliance activities. AI software can perform a compliance function within a given automated function. For example, an AI system could be instructed to issue a report (or email or text message) to a compliance officer if certain values are exceeded or fall below a specified threshold. If regular reports from multiple people are required, the system can monitor whether it has received the reports. It can be programmed to send a notice to those who have not made their report, and eventually to the compliance officer if reports are not received within a specified time period. The system can adjust processing based on an individual reporter's performance. So, for example, more leeway might be given to someone who always files their reports on time versus someone who is frequently late.

For compliance officers, using AI represents what might be called a force multiplier, in that it enables compliance tasks to be assigned to a machine rather than requiring a human to track and identify reports not received on a timely basis. Because typical budgets for compliance are never enough to do everything a compliance officer might like, automating some processes can make those resources go further, which can be a valuable part of the overall compliance process in an organization.

On the other hand, AI software cannot exist in a vacuum. It needs to be properly controlled and carefully examined by a compliance professional. This person should be involved in the development or adoption of the AI software, along with its customization and testing. Compliance professionals should not underestimate the importance of being involved in testing. Problems with data used to train the system can produce results that might seem completely appropriate to the AI technical team, but may be recognized by compliance specialists as reflecting, for example, inherent biases that may be implicit in the training data, which is often historic in nature and may have been obtained from periods where various issues (like racial or gender bias) may not have been recognized. The technical people involved in the AI development process may not be sensitive to these issues. Compliance professionals must be—and can serve as—a vital system of checks and balances to assure that old problems are not carried forward into the new AI-based system.

AI and deep-learning systems can impact the traditional compliance function. Compliance professionals can both protect the organization from AI-related problems and take advantage of AI's potential capability to enhance and serve the compliance function.

This document is only available to subscribers. Please log in or purchase access.

[Purchase Login](#)