

December 17, 2014, 1:38 PM ET

The Confusing Landscape of ‘Cognitive Computing’

By **Thomas H. Davenport**

There is little doubt that computers are increasingly capable of doing things that humans could once do exclusively. Making decisions, abstracting insights from data, recognizing images, learning from text—these are clearly now machine activities as well as human ones, and in some cases machines are already faster and better than we are at them.

But what specific technologies are we dealing with here, and what the heck do we call them? “Artificial intelligence” has been bandied about for a while, and it’s become an umbrella term for a lot of different tools. Perhaps its only problem is that it’s old, and we have heard so many times of its rise (and fall) that many have grown weary and skeptical of the term.

The newest umbrella term is “cognitive computing,” suggesting that we are finally developing computers that can mimic the human brain. The only problem with this term comes when you have an extended conversation with a neuroscientist, and you realize that we still know very little about how the brain works. There are articles [attesting to our ignorance here](#) and [here](#). So to refer to these smart machines as examples of cognitive computing is not really very smart.

The truth is that there are a variety of technologies that comprise the current landscape of smart machines (the overview term that I find least problematic). Each is suited to a different set of purposes, and it’s pretty rare that more than one is integrated in a particular application. Unlike the human brain—which can perform a variety of cognitive tasks, a particular computer system can only do one type of task. Here’s an incomplete alphabetical list of the systems that can undertake cognitive tasks:

Analytics—Tools that do mathematical or statistical analysis on structured data—typically numbers. These have grown in power and sophistication over recent years, and can now be used to drive a variety of decision types. They are also increasingly embedded in other systems and business processes.

Complex event processing—This type of system takes as inputs a variety of real-time data sources and types about events, and then combines them to determine their significance or to take action. It takes in data, transforms it as necessary, analyzes it to detect trends and patterns, and takes necessary actions. CEP systems are widely used in algorithmic stock trading, in which a system might monitor a variety of economic and social indicators, and then determine that a stock trade would be economically beneficial. CEP is also used to detect credit card fraud—ideally before it is successfully committed.

Image recognition—Early systems for recognizing images were quite limited. But now that computers have become a lot more powerful, they can identify more complex images, including specific faces and types of animals. Google was able, for example, to build a complex image recognition [that could identify cats in videos](#).

Machine learning/neural networks—These are somewhat automated approaches to analytics. They create models to fit data and improve as they learn. There are various forms of machine learning approaches, including neural networks, Bayesian classifiers, decision trees, support vector machines, and so forth. The differences between these are generally perceptible only to specialists.

Natural language processing—These tools take text or speech as input, and increasingly can “read” or extract meaning from it. [IBM](#) Corp.’s original Watson falls into this category, as does Apple Inc.’s Siri. Since much of human experience is represented in language, this is a powerful category, and the one most likely to be described as “cognitive computing.”

Rules and business rules—Rules express logic in a structured language—typically an “if/then” structure. They were the primary programming approach for so-called “expert systems,” a branch of artificial intelligence. Business rules express operational approaches to business in this structure. A business rule might specify how customers are to be treated (e.g., a customer returning an item doesn’t have to go through a credit check), or when certain quantitative thresholds are reached (e.g., a mortgage loan can be given if the loan-to-value ratio for the house in question is less than 20%).

These and other categories of smart machines can now address almost any topic on which there is data or recorded expertise. They’re all useful, but they’re not universally useful. Since each tool is suited only to a particular purpose, managers increasingly need to be familiar with the tools and how they fit particular applications. If you’ve got a problem that can be structured in a set of rules, you don’t want to hire Watson for that job. If you have a bunch of data in rows and columns, a rules engine won’t help you much.

One thing that human brains—at least some of them—are good at is seeing the big picture. That can include looking over the variety of technology alternatives for cognitively-oriented situations, and selecting the right one. We could probably have a system—I am envisioning a set of rules—that asks the business user a set of questions about the desired application, and then recommends a particular technology. But we don’t have that yet. So in the short run we will have to rely not on smart machines, but on smart humans, for this purpose.

Thomas H. Davenport is a Distinguished Professor at Babson College, a Research Fellow at the Center for Digital Business, Director of Research at the International Institute for Analytics, and a Senior Advisor to Deloitte Analytics.