

Harmon on BPM: AI and Business Processes, Paul Harmon. July 3, 2018

If you are reading business publications, and especially if you are reading IT publications, you are encountering one story on Artificial Intelligence (AI) after another.

I encountered this same phenomena in the early 80s, when an earlier round of commercial AI bloomed, was talked about everywhere, and then faded. Lots of new technologies look exciting, and some even result in some wonderful case studies, but fade when its determined that they can't scale up or they prove deficient in some way. In the 80s AI companies were going to build expert systems, systems that would do the kind of analysis and make the kinds of recommendations one expected from human experts. In fact some systems were built that were as good or better than human experts. Some AI systems built in the Eighties are still in use. Most, however, proved too expensive to maintain and the interest in expert systems faded by the early Nineties. I mention this last round of commercial AI simply to warn everyone to maintain a sense of perspective.

That said, I am much more hopeful about the future of today's AI systems. One of the limits imposed on AI in the Eighties was a lack of sufficient computer power. Today's computers are much more efficient and much faster. Another limiting factor was the need to base decision making on rules. Humans had to create the rules and then, when things changed, the rules needed to be changed. If one was dealing with a large expert system with thousands of rules this was a very expensive and time consuming process. Most of today's expert systems rely on deep neural networks, which in turn use probabilities, and can learn for themselves by studying examples.

In other Columns, I have reported on the impressive success of Google's AlphaGO, an application that defeated the world champion GO player in 2017. When AlphaGO first played the European champion, in 2016, it was impressive, but everyone expected, based on AlphaGO's play that Lee Sedol, the world champion, would have little difficulty winning the match. What they didn't understand was that AlphaGO could play itself. Each game took minutes, and then AlphaGo was able to play millions of games between its two championship matches and was many times better when it defeated Lee Sedol in 2017.

When commercial AI became popular in the Eighties, it was almost entirely focused on expert systems – systems that performed analysis and decision making tasks. Today, AI is being used to listen and speak in various natural languages; it is being used to read text and extract information from photos and documents, and is able to develop arguments. AI is being used in a wide variety of robotic devices. And it is being used to analyze and make decisions.

Consider one application that is being developed – the self-driving car. In spite of some bumps along the way, a large number of companies are pouring millions into the race to have the first commercially viable automated car. A self-driving car would use vision systems to scan its environment in real time as it moved a vehicle along roads. It would use GPS to understand its location and maps to determine how to move from place to place. It would speak with riders to determine where they wanted to go and to provide them with estimates of time, and alternatives, like where to get groceries or to eat. And it would use its decision capabilities to make decisions about how to navigate and what to do in various kinds of emergencies.

At the moment, the suite of computers and software applications required for a self-driving vehicle are more expensive than the car itself. It's likely that specific elements of the package will be sold, initially, as aids to human drivers, and that the first automatic cars will be used for intensive driving tasks, as taxis, or for business purposes. At the moment, however, most experts agree that it will only be a decade or so until we see self-driving cars on our roads.

I wrote another Column in which I talked about how IBM helped a Japanese insurance company automate a claims processing job. The application was able to listen to voice messages, read text and scan graphs to gather information on the medical claims of a patient. Then applying the rules of the contract, they system could determine the claim. This application replaced some 40 claims adjustors and significantly sped up the processing of claims making it a very cost-effective application. Lots of applications of this kind are being explored by large companies throughout the world.

Last week, in San Francisco, IBM demonstrated Project Debater, an application running on a supercomputer that debated with human opponents. In effect the application listened as a topic was announced (eg. Resolved that the government should fund space travel), gathered information, developed an argument, and then debated. The application and the human each took four minutes to present a position (pro or con), four minutes to respond to their opponent's arguments, and two minutes to sum up. Project Debater debated twice in San Francisco, and the audience judged that it won once and lost once. We predict that Project Debater will debate again, and that it will rapidly improve.

To see a video of a bit of the debate, check:

<https://www.youtube.com/watch?v=naQujxmg9gg>

Project Debater did not have any special knowledge of the debate topic and didn't know what topic would be discussed until it was announced. What it did have was lots of information on how to search and analyze data, how to construct arguments, and how to speak for four minutes in a convincing manner. Undoubtedly IBM will explore the debate topic for a bit longer to learn more, but clearly their goal is not a debater, as such, but a software application that can help human's consider problems and make rational decisions. One can imagine a loan advisor, a cancer treatment advisor, an ad campaign advisor, and many more in the near future.

This round of commercial AI is much more complex than the round we experienced in the 80s, but then, too, everyone is more prepared. Most IT people have now taken courses on AI technologies in college, and the new applications are being offered on a variety of friendly platforms. In the last decade or two, most large

companies have explored AI applications in the form of Analytic Engines that have helped them analyze massive data bases and are, in a sense, more sophisticated as a result. In most cases today's applications won't seek to replace experts, but to serve as assistants. Or, like self-driving cars, the software will come packaged in a hardware device that we already understand.

As we suggested earlier, today's AI may encounter limits and we may find in a decade or so that today's discussions were over-hyped. Frankly, however, having lived through this once, this round seems much better grounded. I suspect this is the real thing, and that companies that get on board early and figure out how to use AI to improve their business models are going to leave the competition far behind.

AUTHOR



Paul Harmon is the editor of www.bptrends.com and the author of *Business Process Change*. In the Eighties, he coauthored the best-selling AI book, *Expert Systems: AI for Business*, and edited the newsletter, *Expert System Strategies* for over a decade.