

What isn't AI?

Originally published at <https://www.aiaa.net/artificial-intelligence/articles/what-isnt-ai>



It's AI to the rescue! Or to the end of the world. Actually for quite some time to come, it's probably neither

Opinion Those old enough to have survived some of the more famous digital buzzwords - such as 'CGI', 'cloud' and even 'digital' - might well regard the current media and business obsession around AI with circumspection.

As a buzzword, AI is generating at least as much boardroom fog as 'cloud' was ten years ago - adoption that's very often without definition or understanding of applicability: the view from the outermost margin of the crowd that has gathered around the subject.

It is possible, just possible, that AI will neither affect your field, even indirectly, nor prove to be a reasonable or necessary IT proposition within your business. At a glance, this applies to the dentistry profession, which is likely to be far more affected by tandem tech evolutions such as Augmented Reality and new innovations in 3D printing.

It's also possible that business does not really understand what Artificial Intelligence is in practical terms at this stage of its commercial evolution, which would seem to be a practical shortcoming in applying it.

So in order to understand what products are actually available, or might be available sooner rather than later, let's assume a hierarchy of abstraction for AI, from the most pragmatic to the most academic and abstract.

1: Fixed branching

At the bottom is scripting and branching: a process, mechanical or analytical, chooses one of, say, six roads ahead based on predetermined conditions or user input. Anyone who has ever created or even seen a flow-chart understands the process, and that obtaining the correct result may give the superficial illusion of agency, when it really represents simple switching.

A great deal of research may have gone into the switching, but by the time we are using it, the system is usually 'baked'. We pass through such processes ourselves at simpler traffic light systems, self-service supermarket tillers and automated customer service systems.

We understood these principles very well even [in the 19th century](#), and their potential shortcomings have been mapped [in real life](#) and [fiction](#) alike.

These systems have elaborate but 'fixed' gears. You can buy this right now.

2: Reactive branching

Above branching and scripting is responsive or reactive branching, where systems react to live data about external conditions.

Here the background work for the process will have included testing how the parameters of external data affect the decisions within the process. Data such as weather (aviation and motorway traffic), machine conditions (tire and braking systems), stock market fluctuations (Automated Trading Systems) and Twitter content (sentiment analysis and security systems) can have their output bolted on to the process for decision guidance.

The background work to create a safe template for reactive branching needs to be more exacting, since unanticipated results [can be catastrophic](#).

These systems have fixed gears which work in concert with other systems which also have fixed gears. You can buy this right now, but *caveat emptor*.

3: Self-learning systems

Systems based on supervised or unsupervised machine learning vary radically in complexity and scope, but at best are expected to learn something for themselves about the processes they must undertake, and the ways in which they will analyze and interpret incoming data.

But in most cases self-learning systems operate within far more rigid parameters, with their human creators left to devise optimizations based on performance.

These systems have complex gear matrices; some of the gears may become less or more utilized as analysis begins to reveal the most efficient processes for the assigned task. You can invest in this right now, as a 'long bet'.

4: Self-building systems

At the most abstract and pure level of AI research, self-building systems - among which famously sits the theoretical [Von Neumann Machine](#) - are capable of changing how they work and, perhaps, of building other machines or system models. A self-building system left free to redefine its own ethical philosophy or purpose is typically the fodder of sensationalist news headlines about AI research, as well as of science-fiction.

The level of branching possible with such an autonomous machine presents a problem that's on a par with meteorological data analysis (one of the most unpredictable among the hard sciences). The possibilities become fractal, and perhaps dangerously recursive.

These systems can design or decommission their own gears. You can buy the (sci-fi) book right now.

Recommender systems

So what can this kind of stratification do to help us penetrate the 'pixie dust' factor of AI as a business buzzword?

For one thing it explains the difference between an algorithm and an active Machine Learning process. Machine Learning (live, slow, expensive, flexible) can analyze large and otherwise ungovernable data sources in order to demonstrate what processes of data analysis would be suitable for an algorithm ('baked', fast, cheap, relatively inflexible). But it's the algorithm that will do the work, and within very fixed parameters.

One practical example is recommender systems, one of the hottest topics in AI. By observing mean relationships between the viewing choices of an individual and the historical choices of others (where they intersect), Netflix has developed recommender algorithms for its customers.

However these algorithms are [susceptible](#) to eccentric customer choices in small datasets (for instance in a new country launch where there are little initial data available), and have only recently evolved from country-specific criteria, not least because the system must not recommend content which is available on the Netflix network but not available in that customer's catalogue.

And that's an algorithm — call it AI if you like, but it barely falls into the second category of reactive branching, in so far as it will adapt to 'live' data but never changes its fundamental principle: that 'birds of a feather flock together' rather than 'opposites attract'.

There's nothing in the framework that could allow such a thought to occur, or to make any change in how the algorithm works even if it did. It's Starbucks-style AI, if it's AI at all.

Certainly there is no active Machine Learning network operating on all the data at any one time, even in the smallest Netflix licensing region. AI simply strikes a template, like a parent sending a child out into the world.

Technological limitations for Neural Networks

Anyone who regularly reads AI research papers will be familiar with the constant undercurrent of anticipation among researchers for the day when Deep Learning systems (such as Convolutional Neural Networks) can be deployed at scale in an affordable way and with the kind of low latency necessary to be viable in live consumer-level or business-level systems.

As new dedicated AI hardware vies with GPU-based frameworks to push the field ahead, it's a real prospect, but still a distant one.

Chatbots - the empty suits of AI

It's only when a company applies level #2 AI technology to level #3/4 AI propositions that the developmental gulf between them becomes clear. One of the most infamous recent examples of this was Microsoft's experiment in creating a Twitterbot capable of genuinely learning behavior from others. In short order, it became [a rabid racist](#).

Because of the [strictures of the Turing Test](#), chatbots of this kind have risen in public perception as a kind of litmus paper for the state of the art in AI. However the father of AI did not mistake the illusion of intelligence for actual intelligence and agency, and one wonders what he would think of the popular inference derived from his thought experiment.

Chatbots belong, at best, in the higher levels of fixed branching or the lower levels of reactive branching; they perform relatively cheap party tricks, fueled by voluminous datasets and crowd-pleasing pattern recognition algorithms; but at the current level of facility, their lack of real intelligence or resourcefulness becomes clear early on, placing them only marginally ahead of automated customer systems.

The state of the AI market in 2017

The current business interest in AI resides firmly within reactive branching, with the focus of research dominated by image recognition systems and pattern recognition in general.

Within that scope, the impetus is still towards scripting and automation; the prospect of machine agency may sell the tickets, but it's the promise of salary reductions and increased dividends that's funding a great deal of the research.

Images: [Wikimedia Commons](#)