

Why isn't AI delivering?

*Business communities have high hopes for artificial intelligence (AI), but will the current stage of evolution of information technology (IT) meet their expectations? Most likely not. **Ivan Belik** and **Derrick Neufeld** write that modern AI development is, to a large extent, based on machine learning methods that effectively meet current business needs but have serious limitations in addressing upcoming IT business challenges. History, they say, teaches us that optimism and heightened interest in AI technologies are sure to be followed by a period of frustration and decline in investments in AI – in other words, an AI winter.*

“Artificial intelligence is going to change business forever.” This has become a common refrain in business and academic communities, accompanied by endless speculation and discussion on how AI affects business and the world economy. Thanks to massive media coverage of both the positive and negative impacts of artificial intelligence on business and society, AI is now a household term. However, due to the flurry of speculations about AI as inevitable, many objective attempts to justify the business world's expectations of AI are frequently overshadowed.

Since the 1950s, AI has experienced several winters, periods of reduced business interest in AI due to unfulfilled expectations. Aside from a few minor episodes, the two “coldest” (i.e., most severe) winters were in the late 1970s and early 1990s. The former was caused primarily by unfulfilled hopes pinned on speech-understanding technologies, as early attempts to train computers to recognise and translate spoken language were unsuccessful, and the latter by the collapse of the market for AI-based expert systems, as most companies specialised in the commercial use of computer systems that mimic human decision-making failed.

The current strong interest in AI by business is likely to transition seamlessly into a new AI winter thanks to the limitations of machine learning (ML), the nucleus of modern AI, frequently referred to as weak or narrow AI. In simple terms, weak AI can be interpreted as a machine learning-driven IT product focused on large-scale data processing to solve a particular business problem. Over the past decade, a variety of cutting-edge ML algorithms have made weak AI successful in addressing many business challenges related to data-clustering and classification. Nevertheless, modern ML, as we know it today, will most probably not constitute the core of the much-anticipated AI, frequently referred to as strong AI.

Unlike weak AI, which focuses on accomplishing narrow analytical tasks, strong AI aims to mimic human cognitive abilities. Strong AI is often considered a concept that synthesises at least three basic abilities of future IT solutions, such as generalisation, inventiveness, and consciousness. The capability of modern ML, when viewed as the foundation of strong AI, to realise each of these abilities in future AI developments is highly controversial. Why this is so is easier to understand through concrete, illustrative examples.

Generalisation

Generalisation is the ability to adjust to new data, i.e., to react appropriately to data never seen before. A variety of modern ML approaches, such as deep learning, aim to mimic human ability to generalise. Results are frequently impressive but still a far cry from even the most basic human capabilities. A good example is a classic video game called Breakout (or Brick Breaker), where players attempt to break down a wall of bricks by hitting a bouncing ball with a moving paddle. Suppose there are two players, a five-year-old child and a machine (i.e., computer), who have never played Breakout. Obviously, we can teach the child to play the game, and we can train the machine to play it as well, based on the game's data.

Now, consider a modified version of Breakout, called Pong, with two vertically moving paddles (a basic simulation of table tennis). A five-year-old able to play Breakout does not need to be taught how to play Pong because humans can generalise, that is, they can intuitively understand the new game based on previous knowledge about the first game. The machine, however, has to learn how to play Pong from scratch using new data, even though Pong's gaming process is analogous to that of Breakout. The machine is not able to generalise. It cannot utilise its Breakout "skills" to play Pong. Even with cutting-edge ML algorithms, the machine will require a new round of data-driven training processes, even though the task does not significantly deviate from a task that the machine has solved successfully in the past. Technically, retraining the machine to solve a new task is not a problem, but at the current stage of AI evolution, the very ML approach to data-driven generalisation to solving similar or related tasks cannot compete with the human ability to generalise.

Inventiveness

Inventiveness can be characterised as the ability to be creative. It is even more challenging for ML algorithms to "learn" how to innovate than to generalise. The potential ML-driven capability of AI to innovate is based on simulations of human abilities to explore non-existent needs, formalise them into problems, and search for potential solutions. We can illustrate this with an example using Lego. Consider a five-year-old child who has never seen Lego before. First, we show the child how to join the Lego bricks vertically. Then we ask the child to construct something original, whatever the child can imagine. Most probably, the child will demonstrate creativity by quickly joining bricks horizontally and in other directions, based on the human ability to generalise.

However, for machines, the problem of constructing something original is challenging. ML algorithms initially require extra training to "understand" how to join Lego bricks horizontally. Then, even if we teach the machine how to join bricks in all directions based on data from real-world objects, we will most likely end up with one of the following results: The machine will either try to replicate existing objects from the real world or it will generate obscure objects through randomisation. In either case, the results can hardly be called original or innovative. The root of the problem is the limited capability of ML-driven machines to accurately reproduce human creativity. Modern machines function based on probabilistic and optimisation ML techniques that have a hard time replicating even a simplified version of the human ability to innovate.

Consciousness

Consciousness can be seen as self-awareness. Humans have certain innate patterns of behaviour and ways of assessing the environment. They are able to understand how their thoughts, actions, and emotions coincide or not with personal behavioural standards. Self-awareness is inherent to human nature, so simulating it is highly challenging, but it can be illustrated by a video game in which players can choose their character. Of course, children and machines, as players, can both simply choose at random, but we are most interested in how non-random choices are made. Due to inherent consciousness, a child can choose a character that coincides with their personal thoughts, feelings, and behaviours. In this case, the selection process is driven by the human ability to associate personal gaming preferences with the traits of in-game characters.

However, a machine chooses its character differently, as the machine has no inherent personality traits that could form the basis for its gaming preferences. Certainly, the machine could be trained, using the child's data, to pick a character in a way very similar to the child, but the machine's choice still would be based not on self-awareness but its ability to process data and imitate the characteristics and behaviour of the child. Machines can be taught to analyse large-scale data much faster and more accurately than humans, but processing speed and quality have nothing to do with the machine's self-awareness.

To conclude, the ability of modern ML to cope with the challenges of strong AI is becoming increasingly questionable. In this regard, the main driver of the coming AI winter is growing disappointment over the expectations placed on modern ML compared to its actual ability to solve the emerging problems of strong AI. This means we can expect to weather another AI winter before we can develop a new generation of ML methods (or perhaps even a completely new computing paradigm) and emerge into an "AI spring."



Notes:

- *This blog post represents the views of the author(s), not the position of LSE Business Review or the London School of Economics and Political Science.*
- *Featured [image](#) by [geralt](#), under a [Pixabay](#) licence*
- *When you leave a comment, you're agreeing to our [Comment Policy](#).*