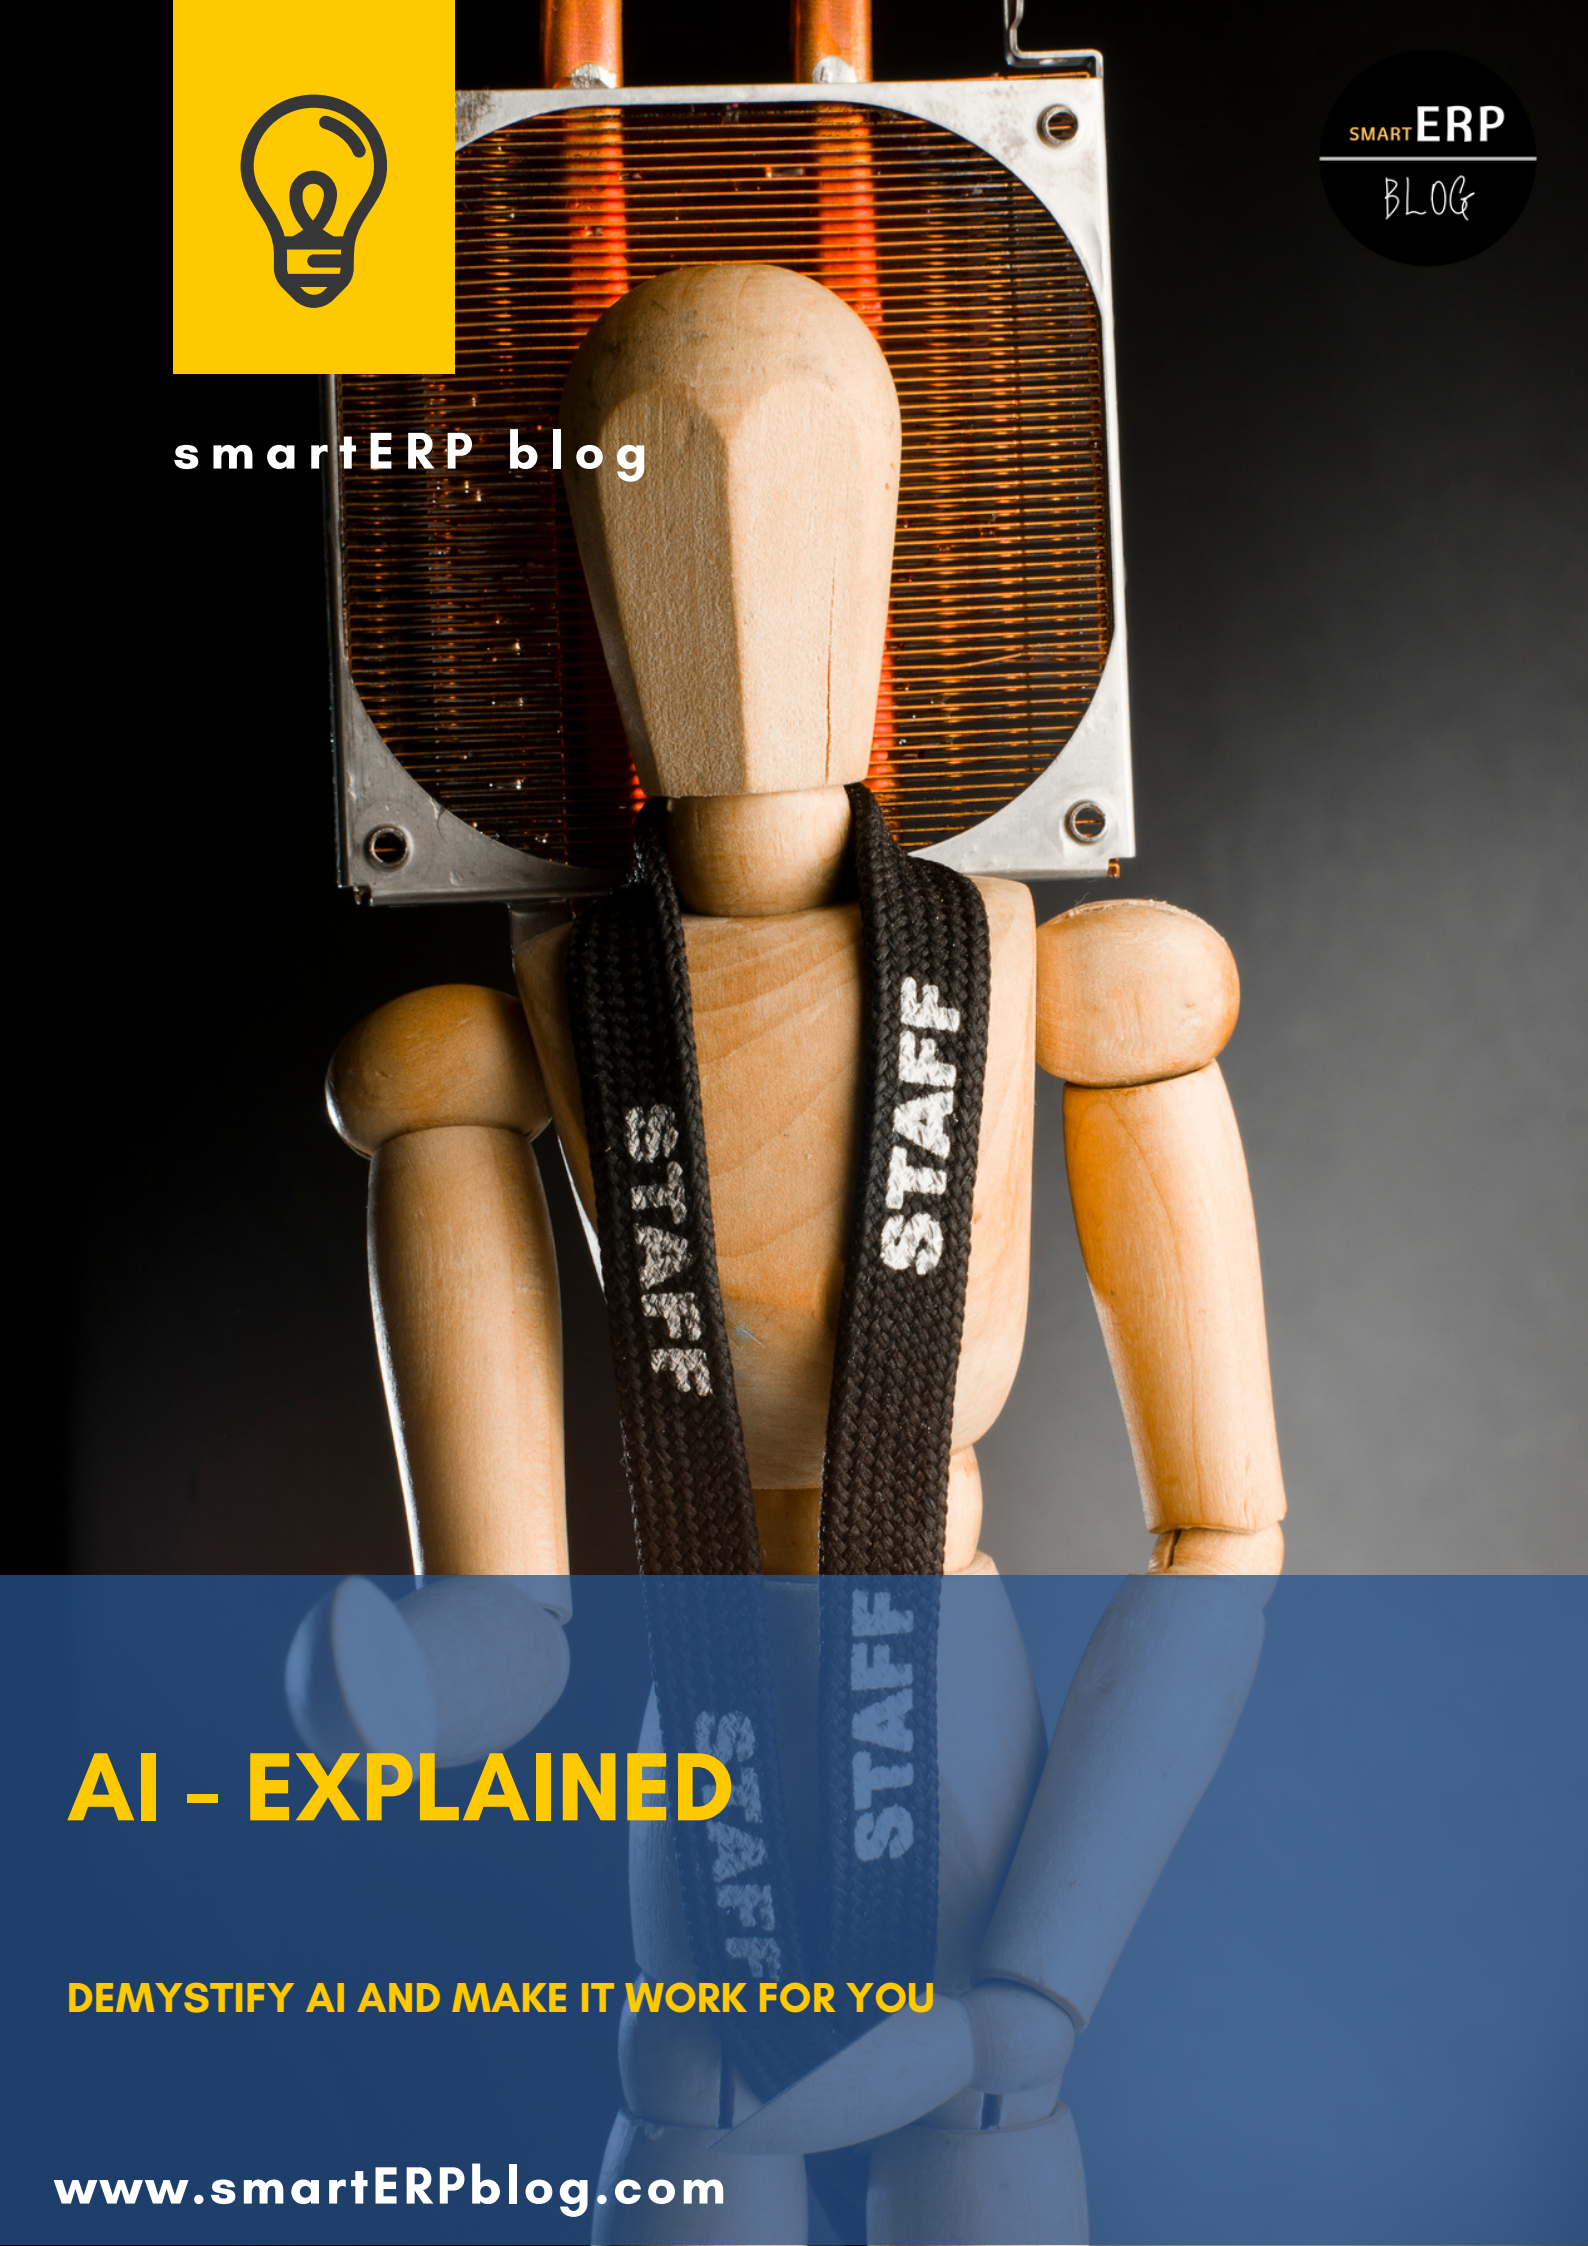

smartERP blog

AI - EXPLAINED

DEMYSTIFY AI AND MAKE IT WORK FOR YOU

www.smartERPblog.com

Introduction

Artificial Intelligence has become a major topic. Hype, fantasy, total nonsense, and boring reality is all part of the mix.

AI has been around for many decades, and it already surrounds us in countless forms in just about any aspect of our lives. Globally there are millions of programmers creating AI solutions daily. Most AI solutions are small, unassuming and people just take them for granted, not even realising they are dealing with Artificial Intelligence.

AI brings many qualities to the table in an ever more complex and rapidly changing world. Because technology is becoming a determining factor of success it is a huge advantage if we build a basic understanding of what AI actually is, and what we can realistically expect from it. This doc aims to give a quick and simple overview from someone who has been building and working with all kinds of AI for the last nearly forty years.

What is Artificial Intelligence

Actually, any kind of computing can be referred to as Artificial Intelligence that performs tasks we previously associated with human cognitive skills. Apps can play chess, GPS can find routes to physical destinations, computers can pilot an aeroplane, monitor our heart problems, help us transfer money or suggest movies that we will probably like. It is a form of intelligence that a dog, a child, or even most adults could not show.

Using 1s and 0s, and applying simple logic to them enable computers to perform tasks that fall right into the cognitive skills category. Storing and recalling information, performing calculations and reaching conclusions even with outside stimuli from the environment. They can interact with the physical world too through electric switches, motors and sensors. These are traits we normally associate with intelligence.

There is more than one way to create Artificial Intelligence. From the many approaches two kinds of AI proved especially useful.

The first kind we could call “0/1 logic computers”, referring to the fact that to program them we use two binary numbers and three basic logical operators, AND, OR plus NOT. (Fortunately modern programming languages make this quite simple, and we can produce working code using a strict set of words like IF, THEN, WHILE and so on.)

The vast majority of AI created fall into this first kind. They follow human instructions to the letter. They are predictable and testable to a large extent. In certain tasks they are superior to human minds simply because of immaculate calculations, and the ability to store and process astronomical amounts of data.

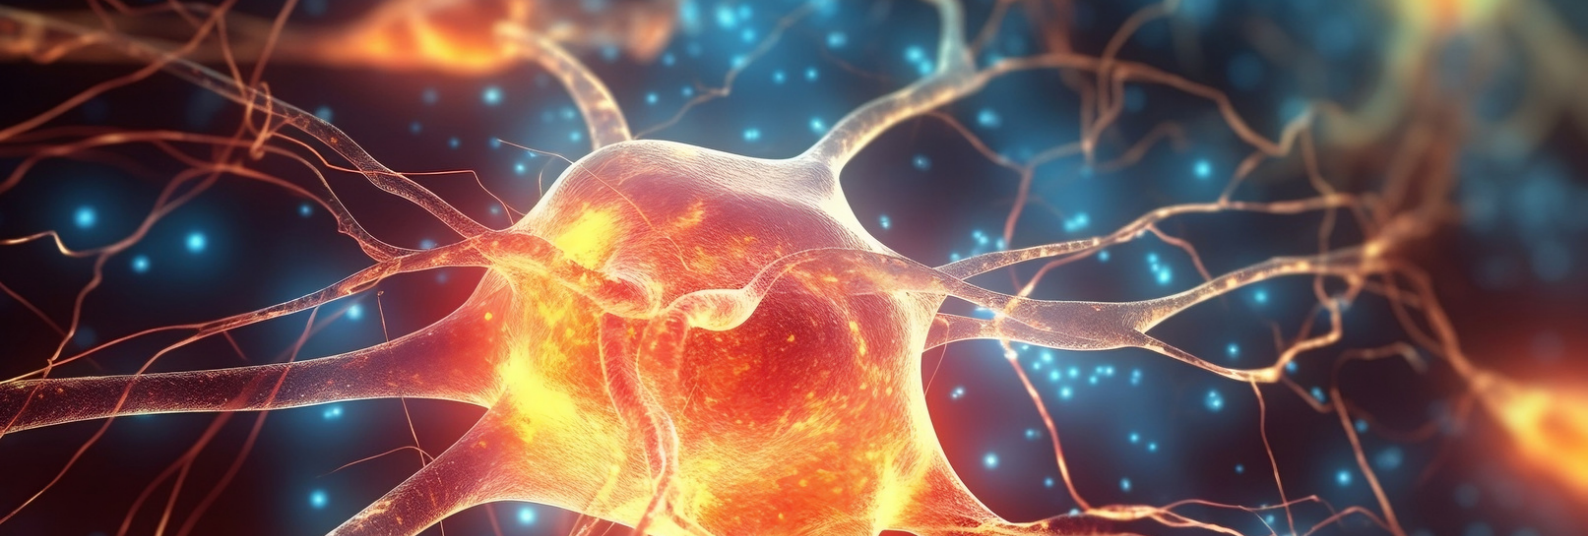

Mimicking the brain in computers

However, “0/1 logic computers’” are limited by their flawless logic and precision. This makes them clumsy in many areas of the human realm. This is where neural networks come to the rescue as the second kind of popular AI approach.

When people use the word AI, often they refer to neural networks because of certain structural resemblance to the brain. Neural networks are not new at all, but they require an insane amount of computational power to perform even the simplest of useful tasks. Until recently our computers were not powerful enough to run them at this scale.

Neural networks simulate the work of individual “neurons” joined together in a network. The human brain itself has around 100 billion neurons where each neuron acts like a micro computer. One neuron on its own doesn’t do much, but connect them in a complex interconnected network (see image at the top of the page) and things start to happen. An individual neuron can be connected to as many as 100 000 other neurons, and with 100 billion neurons doing their job simultaneously, it becomes very complex and powerful.

The fundamental principles of neural networks are quite simple and unassuming. When a neuron gets excited, other neurons connected to it get a signal. Every neuron then processes all of their incoming signals through their specific algorithms and settings. Based on this the individual neuron determines whether to get excited or not, which in turn will be processed by further “watching” neurons.

At any given moment waves of billions of signals splash around simultaneously in our brain. To “learn” the individual neurons change their settings based on the “feedback” they receive from their respected outside “world”. The end result is what we experience as humans - we can store information, recall memories, assess situations, reach decisions and execute them.

Neural networks

“0/1 logic computers” are stubbornly detail focused and this extreme precision makes them terrible at recognising general patterns. It is extremely laborious to prepare them for every variation while even very small and simple networks of neurons are excellent in recognising and categorising similar “data sets”.

For example, our brain only has to see a horse a few times to be able to recognise other horses forever - in just about any environment, from just about any angle. Neural networks inherently act the same way and they can’t help but categorise and group things through similarities. Because of this they are very good at generalisation and tolerating small variations in data.

To create individual “neurons” in a neural network, and to “connect” them to each other we use “0/1 logic code”. However, once we have created the neural network, we mainly interact with it as a whole rather than the individual neuron.

Neural networks cannot be programmed through instructions. They have to be trained, just like the brain. We send it stimuli and train it through feedback. Using a previous example, we can give it feedback like - “no darling, that is not a horse, that is a donkey.”

The number of “neurons” and connections in a neural network determine how many different things it can distinguish between and how many functions it can perform. Also, the structure of the connections and the algorithms of the individual “neuron” determine what kind of work it can support. In our own brain different kinds of neurons and connection structures support vision, languages, motion, decision making and so on.

While the fundamental principles of neural networks are simple, they require big, powerful computers to make all those billions of calculations every second. It also took us a lot of experimentation to find useful neural algorithms and connection structures. Once up and running we could start the long and tedious training that enables a neural network to do useful things like processing natural human language.

Intelligence is in the eye of the beholder

It is very easy to be deceived by appearances. Indeed, our very ability to categorise and recognise things based on their similarities tricks us when it comes to judging AI. We spot a trait and just assume the rest is there. We see AI produce a sentence and we assume it has the same understanding and sensations that a human would.

This is not the case at all. The human brain has many specialised functionalities and most will have to perform in harmonious collaboration to create even the simplest of a normally functioning human. At the moment we are unaware of most of these essential functionalities of the brain, let alone understand them.

With AI maybe we have reached a brave 1% of all of this, but reaching more is not straight forward at all. Our very own mental limitations may prove to be an impossible barrier to go much further. It is not simply a question of volume.

We have to be careful when judging AI's intelligence or creativity. We assume more than there actually is.

AI can give you a definition of "apple", it can draw a picture of it in countless different styles, it can give you the equivalent translation in different languages, it can even generate a poem on the subject.

But all through this "apple" means nothing to it. It was trained to find a definition on the Internet or in a database. It was trained to find an image and then generate similar images from lines, colours and pixels. It was trained to generate strings of words similar to what people respond well to. Still there is no equivalent to what "apple" means to a human. It was programmed to beat you at chess, but it doesn't know what chess is. It doesn't have to understand the poem or text it generates, as long the humans who read it respond positively to them.

In essence it is not the thoughts and responses of AI that we actually judge but the interpretation and responses our very own brain creates. Hold that thought - like beauty, intelligence is in the eye of the beholder.

Dangers of AI

Popular films, books and articles attest to our deeply engraved fear of aliens and Artificial Intelligence, but it may have more to do with our own psyche than actual reality. At the same time unfounded assumptions, simple misunderstandings pose real dangers.

It is easy to overestimate AI's abilities - the same way as it is easy to overestimate fellow humans. They will perform seamlessly in one area just to let us down somewhere else because of their inadequate lack of understanding and skills. It is dangerous to rely on something which is simply not there.

Also, we often call in AI to help us with complex problems. But this complexity inherently introduces the element of the unforeseen. You change one seemingly benign thing and it might result in the failure of the system somewhere else. This is acceptable if we are aware of it. Complex technology is like any power tool. You must handle it with discipline and care.

Finally, it is tempting to use AI to do work instead of us only to miss the actual outcome we seek. Say, the student writes an essay for a school assignment using AI. But the desired outcome of the exercise is not words of wisdom on a piece of paper, but the necessary changes the student's brain should go through to build higher level skills and understanding. In this example the AI only hindered us from becoming better and more able.

Before we trust and delegate any number of tasks and jobs to AI we should be clear about the actual outcome we want to achieve, and whether we have a good enough understanding of what is going to happen.

Technology’s purpose and promise

Technology’s only acceptable purpose is to help us humans have better lives. For hundreds of thousands of years we’ve been using technology to perform tasks, do work that saved us time or extended our abilities.

Without technology our current population and its way of life would be impossible. Today more work is done by “technology” than humans. With the further rise of AI this will be taken to dramatically new levels.

This also means that no business or individual can compete successfully in life without the support of technology. We should fully embrace AI in innumerable fields and let it help us with the boring, the dangerous, the unbearable and the impossible.

It is a good way to think of AI as your symbiotic extension, your assistant or faithful servant. It could be watching your heart and alerting doctors before it is too late. It could be processing invoices for you, it could keep track of what you need to learn next to be ready for your exams. The list is endless, in some way technology can support almost every human activity.

However, there is a fine line between being augmented by essential tools, applications and programs, and being overwhelmed and crippled by surrounding technology.

There is no simple guidance on how technology should be applied. It has become so vast, powerful and intimate, that no one can afford to keep a distance and rely only on the advice of experts. In this modern age we all have to develop a basic savvy and understanding of what technology is and how it can be applied beneficially in our business and life.