The smart guide to AI

Introduction

Everyone is talking about artificial intelligence (AI). More than 80% of organisations believe AI is a "strategic priority" for their business and many are scrambling to adopt it as they attempt to carve out a competitive advantage over their rivals.

But in this haste, it's easy to lose sight of what AI can actually do for a business – and it can be hard to identify the best way to get started. That's where this guide comes in. It offers clear and practical perspectives on potential AI applications and provides advice on the steps you should take to get started, to ensure you make the best possible decisions for your organisation.

Get to grips with your options

At its heart, AI is machines acting more intelligently. That can manifest itself as anything from a program that detects people's faces in an image, to a solution that automatically processes millions of medical documents to suggest the most up-to-date techniques to treat patients. It could even be a self-driving car.

Before you can begin to leverage these innumerable possibilities, there are fundamental decisions to make about how you intend to run and manage your AI. There are two main options – and the route you select will depend on your business's needs, industry and strategy.

Cloud-based solutions

- Made up of many different components and based upon multiple technologies
- Can process billions of bytes of structured and unstructured data, enabling systems to learn, understand, reason and provide insights based on that information
- Ideal for businesses with minimal data security concerns and that want to move information to and from the cloud
- No need to invest in and manage data centres

On-premise infrastructure

- Made up of hardware and supported software, enabling you to create and run custom cognitive applications from scratch, free of concerns about power and performance
- Businesses can rest assured that their data is safe and compliant, by remaining in full control of their datacentre
- Provides businesses with the most popular machine learning frameworks to enable the rapid deployment of high-performance AI

Understand why the time is now

AI is compute-intensive. Computers must be "trained" to understand different datasets, reason with the information and create an output, be that answering a query, spotting a pattern or completing an automated task. Such is the power required, that until recently it just wasn't a practical approach – particularly for organisations looking to create their own AI infrastructure on premise. But today, AI is exploding thanks to:

The rise of the graphics processing unit (GPU)

In order for computational processes to become as "intelligent" as humans, they need large datasets to learn from – larger than has ever been possible. And help has come from an unexpected quarter. The technology that powers high-end graphics cards, traditionally favoured by gamers and CGI professionals, is specifically designed to perform millions of calculations in parallel: exactly what AI needs too. GPUs have helped usher in the latest generation of AI by providing the power to process the compute-heavy tasks of "training" artificial neural networks to analyse, understand and reason with data. This enables them to learn faster – and perform tasks more accurately.

This has helped drive the discipline of "deep learning". It builds on "machine learning", which is an algorithmic application of AI that enables computers to analyse and learn from large datasets. Deep learning takes that even further, using machine learning to spot nonlinear patterns in vast amounts of data, in a way akin to the human brain. Without the power of GPUs to drive it further, AI might still be confined to the realms of sci-fi.

Greater scalability

GPUs have aided the rise of AI and deep learning, but hurdles remain. It can sometimes take weeks to train AI models to a sufficient level of confidence, because many of the most popular deep learning frameworks struggle to run across multiple servers, even when using four or even eight GPUs. However, the best server and AI infrastructure enables businesses to cluster servers with ease and leverage the power of hundreds of GPUs, making it more than 50 times faster to teach AI models.

Ease of use

If it's to really take off and evolve, AI can't be reserved for data scientists. The very best AI solutions and tools will use the technology to automate processes such as code writing and server cluster management, to take the heavy lifting away from developers. This will enable businesses with even the most rudimentary knowledge of AI to dive in, create applications, explore the potential of AI – and innovate.
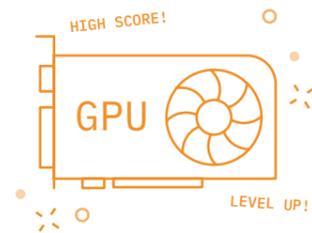
Consider the impact

Before you start your AI journey, it's important to ensure you are prepared to cope with the new requirements of running AI solutions in your business. Here are some key things to consider before adopting AI:

1. What problems do you want to solve?

2. What business goals do you want to achieve?

3. Is your data organised and ready to be utilised in AI?

4. Do you need to bring in outside talent or upskill existing staff?

5. Does your data comply with the relevant data regulations?
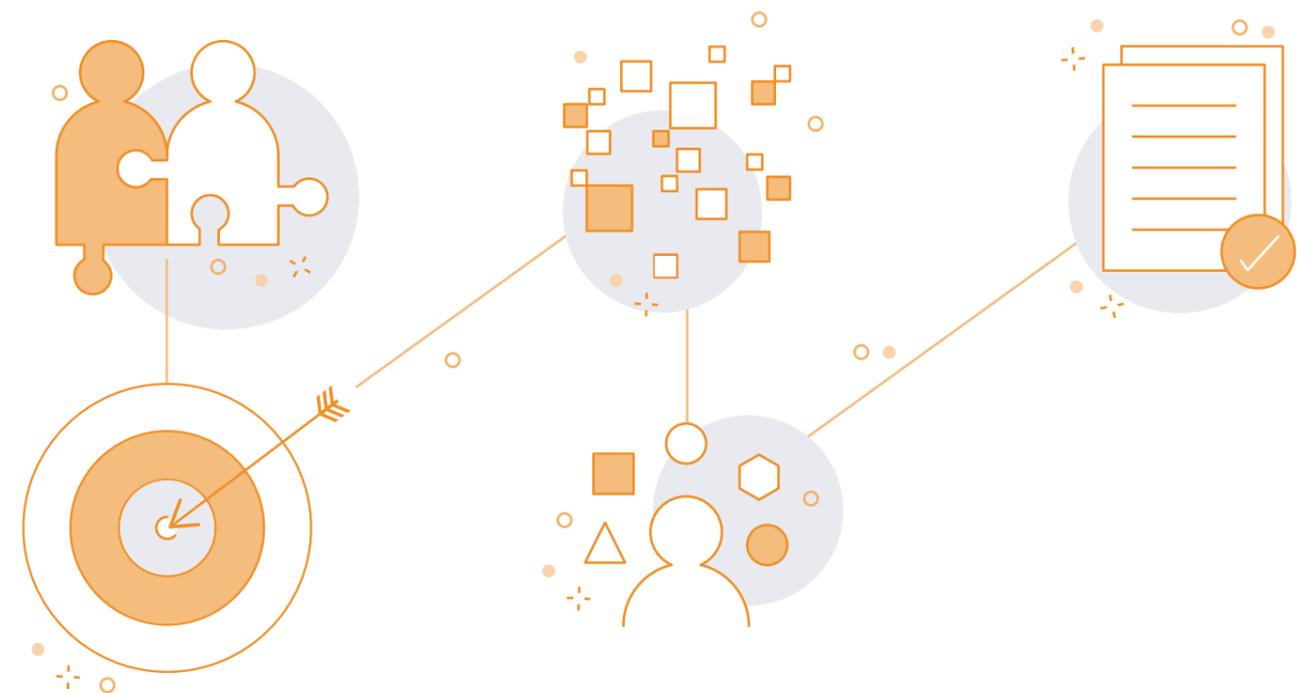
With the volume of data produced every day only set to grow, it's crucial to start to manage it, understand it and make it work harder for your business. Implementing AI throughout your organisation can enable you to do just that, while freeing up time to dedicate to business-critical tasks and ultimately provide better experiences for your customers.

IBM can help you start your AI journey with a holistic hardware and software solution that enables you to access the most popular machine- and deep-learning frameworks and build applications that fit your specific requirements. Call us on 0333 210 0131 today to learn more.