


How prepared is your business to make the most of AI?

To make the most of AI, you need to know the basics: what data do you need, and do you have the right people in place?



January 1, 2019

By [Javier Zamora](#) and Pedro Herrera

Artificial intelligence (AI), once the preserve of dystopian novels and sci-fi movies, is now a hot topic among business leaders wondering how it might transform or impact their business models. But what does “artificial intelligence” mean?

The late John McCarthy, credited with coining the term in the 1950s, defined AI as “the science and engineering of making intelligent machines, especially intelligent computer programs.”

But what is intelligence? And can it be separated from human intelligence? McCarthy’s contemporary, Marvin Minsky, believed this was extremely difficult, partly because “words we use to describe our minds (like ‘consciousness,’ ‘learning’ or ‘memory’) are suitcase-like jumbles of different ideas formed long ago, before ‘computer science’ appeared” and which are inextricably linked to human beings.

For our purposes, we will use psychologist Howard Gardner’s definition of intelligence as “the ability to solve problems, or to create products, that are valued within one or more cultural settings.” This definition expresses how AI can help organizations create new value propositions, as we will explain in this article.

A brief history of AI

In 1950, the English mathematician Alan Turing first devised a method of inquiry — known as the Turing Test — for gauging whether or not a computer was capable of thinking like a human being. The official beginnings of AI as a discipline didn’t arrive until six years later, however, when a select group of scientists and academics, including McCarthy, Minsky and Nobel laureate Herbert Simon, gathered at Dartmouth College in Hanover, New Hampshire, over the course of a summer to advance a research agenda on artificial intelligence. Among their lofty ambitions was “to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.”

During AI’s formative years, researchers focused on developing “expert systems” — computer programs that took data and information and then made if-then inferences, similar to human reasoning processes, to do basic problem-solving. While expert systems form the basis of some applications used today — for prediction, diagnosis or monitoring — they still relied on trained programmers to encode the rules. And given the infinite number of real-life scenarios that could arise, it became virtually impossible to program a system that would work in every possible scenario. By the 1980s, results had fallen so short of expectations that the period became known as an “AI winter,” as funding and interest in the field dried up.

Subsequent evolutions and technological improvements, both in computing capacity and

learning algorithms, have brought us to the current situation. We are seeing an important resurgence of AI, based mainly on “statistical learning,” that is, learning from data obtained through observation. Vitally, this alleviates the impossibility of anticipating everything that could happen when modeling a scenario and the changes that could occur over time, and makes it easier to model solutions to a problem.

Most of the AI systems used by companies today are based on statistical learning — or what is called “machine learning.” Like any automatic learning process, machine learning tries to learn from experience — the data — to model, predict or control something using computational, mathematical based algorithms. These systems are useful for classifying input data, such as recognizing and categorizing images, or making predictions like how much money a customer will spend over time or the likelihood that he or she will stop being a customer. Still, because their use is limited to the problem for which they were designed and trained, these systems are known as “narrow” or “weak” AI.

More recently, machine learning algorithms have come closer to imitating the functioning of the human brain. The neuronal interconnectivity of the brain allows us to solve complex problems, while the brain’s plasticity — the endless process of training and adaptation it undergoes with each new action it carries out — facilitates continuous learning. Artificial neural networks try to mimic these two features of the human brain. Most interesting is their ability to engage in parallel computing and “backpropagation” (the backward propagation of errors). When differences are detected between the result obtained and the expected result, the system is able to adjust itself and modify the approach it takes next time.

Machine learning algorithms based on neural networks have evolved to the point of being able to perform very complex learning processes. This is called “deep learning,” and it allows for applications such as visual or speech recognition. With ever more advanced algorithms, like convolutional neural networks, deep learning allows you to resolve highly complex, multidimensional problems involving very large volumes of data.

Despite these advances, current machine learning and AI tools in general still have limitations when it comes to reasoning or abstract thought — what would be considered “strong” AI. Getting the machines to be able to solve complex problems and act with plasticity in a fully automated way is extremely complex and, without doubt, AI’s greatest challenge.

The consensus is that for AI to surpass human intelligence — a moment known as “singularity” — new ideas are needed to overcome the restraints holding back perception, learning, reasoning and abstraction, on the path toward true “strong” AI. Ray Kurzweil, director of engineering at Google, predicts this will occur in the next 30 years — but people have been saying “within 30 years” for the last 30 years.

The boost of digital density

AI is wholly dependent on the ability to store and process large volumes of data efficiently. Without the confluence of the following factors, the recent advancement of machine learning tools would have been unthinkable.

- **Increase of computational capacity.** Thanks to the exponential function of Moore's Law, it has become easier and cheaper to develop the highly sophisticated and complex computing needed to run machine learning systems. In addition, the use of graphics processing units (GPUs) instead of central processing units (CPUs) is better suited to the distributed computing of neural networks.
- **Digitization of the physical world.** The reduction in technology costs also led to a surge in the quantity of data being produced by organizations, people and objects through the internet of things. This is what is known as "digital density." Think of all the data generated just for an online purchase. We can know who buys what, at what time and how they paid for it; what they looked at before buying, how many times they checked back before buying, and numerous other variables that can feed prediction models of purchasing behavior and classify customers by type. Such big data are very useful in training machine learning systems.
- **Better algorithms.** New programming techniques that distribute processing work across different nodes or computing units offer constantly scalable performance and make it easier to train complex models with many variables and many layers of neurons (deep learning). The quantity of learning data needed for AI models can be significant both in terms of depth (the number of events under analysis) and dimensions (the number of variables that can a priori determine a result).
- **Open software.** This has significantly democratized access to AI applications, some of which are now driven by large user communities, such as OpenAI or Scikit-learn. Others are being developed by companies that want to create large user ecosystems, such as Google's TensorFlow or Microsoft's Cognitive Services.

Speaking the language of data

As a manager, you need a basic understanding of the concepts behind machine learning so that you can communicate effectively with the experts who implement AI projects in your organization. Here are the essentials:

- **Training data.** These are the values obtained from past cases or those that are happening in real time. They will differ depending on the variables selected. In fact, the selection of variables will determine how good a set of data is and is key to obtaining an efficient automatic learning system. Imagine you want to develop a prediction system for how much someone might like a particular movie based on the films he or she has

already enjoyed. The training data would be the previously viewed films, sorted by the degree of satisfaction they produced. Each film would contain a series of description variables (genre, length, year, language, distributor, etc.). The training data can also come from public access databases, such as ImageNet, which contains more than 14 million images, each with a description to train image recognition systems.

- **Supervised vs. unsupervised learning.** To explain the relationship between input and output variables, we use supervised models. If, on the other hand, we want to detect patterns of behavior using an available set of variables, we use unsupervised models. Most companies employ supervised learning processes that respond to a specific business problem. For example, in an email app, a supervised model can classify emails that are junk or not interesting, and separate them from the rest. The model is fed further information by the users themselves whenever they correct a wrongly classified email. In innovation processes, unsupervised models are usually used that do not respond to any concrete objective, such as detecting groups of similar customers to improve personalization in the design of new products for different segments.
- **Learning by reinforcement.** This technique is similar to supervised learning but with one major difference: the algorithm uses a reward or punishment system. It is often used when there are few training data or when the data are generated by interacting with the environment. A case in point would be to program a machine to learn a game to the point that it can obtain the maximum score without explicitly teaching it the rules.
- **Prediction vs. classification.** Supervised learning processes differ according to the type of variable you want to obtain. When the variable is a characteristic, it is called a classification model; when you want to predict a number or value, it is called a prediction or regression model. The latest advances in machine learning have significantly reduced the cost of prediction and classification, leading to new value propositions.

The limits of AI

Knowing the real possibilities of machine learning helps not only to spot promising investment opportunities, but also to carefully manage expectations about AI-related projects, thereby mitigating the disappointment the results can often produce. Here are some of the limits of AI:


- **It requires training data.** Algorithm-based learning techniques essentially involve training a model with data. Although there are learning systems that generate their own training data, such as AlfaGo Zero (an AI program that taught itself to play the ancient Chinese game of Go), they are limited in their scope and not suitable for most business situations. Managers often wonder how much information they need to train a model.

That usually depends on the particularities of the case and the complexity of the algorithm they want to use. Working with large volumes of data does not guarantee good results. Also important are the selection of variables and the quality and reliability of the training data. An altogether better question to ask is what problem you can solve with the data available. If you have little information, it's better to use algorithms that are "resistant" to learning with few data, which logically tend to be less reliable. The key is to use common sense. For example, to build a model that forecasts the future weekly demand of certain products, you need not only the historical weekly sales data for each product but also the time series of variables such as holidays, weather, macroeconomic indicators or the sales channel.

- **It's not very plastic.** Even an AI system trained correctly with past data may fail in the present if circumstances change. For example, a regular customer may stop buying your products if a competitor offers a better service at a lower price. AI systems still have limited capacity to adapt to circumstances like that, due to their lack of cerebral plasticity.
- **It can be overfitted.** All supervised learning processes allow you to measure the degree of reliability of a resulting model, by comparing the prediction with past reality. The more the two coincide, the more a model can be trusted. To gauge a model's reliability, data can be subdivided into several groups according to the stated purpose, such as training, validation or testing. But total confidence today does not guarantee total reliability tomorrow. Many models suffer from "overfitting," whereby your model becomes too well-adapted to the random "noise" of your specific training data, leading to misleading insights or faulty predictions that fail to take account of the variance in new, real-world circumstances.

Is AI right for me?

Use this example, shown here for a customer service process, to make a similar evaluation for your own business problem.



The importance of ends

Merely employing AI technologies, such as machine learning, does not make a company an AI organization, just as having a website does not make you an online business. What matters is not so much the technology per se as its impact on the business model and how it can enable the organization to achieve a particular goal. Clearly, machine learning — like the internet of things, big data, mobility and the cloud — should be treated as another manifestation of our increasingly digitally dense world, one that allows organizations to redefine their value propositions and, by extension, their business models.

New value propositions can use technology for four main types of interactions, which can be used in combination:

- automation, to eliminate manual processes;
- anticipation, to predict events;
- coordination, to participate in heterogeneous platforms or ecosystems;
- customization, to better meet customer needs.

Before adopting machine learning in your company, you must determine whether it plays an

important role in any of these four interactions. This analysis should be carried out during an initial learning phase.

Next comes the implementation phase, in which the new value proposition — to replace or improve existing processes, reduce costs or generate new revenue — is put into action.

Then comes the discovery phase, in which other value propositions are explored through a process of unsupervised learning.

During the learning phase, you must define the new value proposition and decide which processes to make more efficient with the application of machine learning. To do this, you must answer these questions:

- What processes can be improved or replaced by new ones?
- What training data sets do you need?
- How can you measure whether these techniques are being correctly applied?

The next step is to build a minimum viable product to learn from metrics and validate the original hypothesis, iterating throughout the process. Normally, at this stage, you should use supervised learning methods to obtain prediction and classification models.

What talent do I need?

Just as you evaluate the suitability of AI, you must also evaluate your organization's readiness and ability to undertake AI projects. In recent years, the role of data scientist has gained importance but faces several challenges.

- **Lack of knowledge specific to your business.** A data scientist hired by a large company may have a broad knowledge of the sector, but it can take years to grasp the subtleties of every business area and the different opportunities that AI can provide to each.
- **Communication barriers.** It is important to know how to explain the possibilities and limitations of AI in simple terms and at the same time interpret the real needs of business associates.
- **Access to data.** Most of the training data will consist of the information generated by the company's different applications. Accessing these data requires detailed technical knowledge of how and where to obtain them.

To improve your chances of success, we recommend developing the following roles and profiles within your organization:

- **Business translator.** A key figure, the business translator must interpret business challenges, opportunities and areas of improvement, and translate them into actionable proposals that can be addressed through AI learning. The business translator is also

responsible for communicating the results to each relevant business area and serving as an interpreter for lay users, using visuals to show how AI is being used and suggesting possible new uses.

- **Data scientist.** This post requires training in applied mathematics as well as sound knowledge of programming languages and database management. The ideal profile should have strong practical and creative skills, especially when it comes to selecting the data sets.
- **Data engineer.** This person provides the data scientist with access to the necessary training data. This work is usually easier when the company has a centralized repository of information. Even so, the task of combing through the myriad business concepts, features and data points can be daunting. The data engineer also needs to have a clear understanding of the regulatory requirements for personal data processing, for which he or she will need to work closely with the company's data protection officer.
- **Model integrator.** Supervised learning processes generate models with rules that allow you to tweak the input variables in order to obtain the desired output. If we adopt standard formats for these models, we can include them in transactional applications, such as enterprise resource planning (ERP), and thus apply machine learning to any process in the company. This responsibility falls to the model integrator, who must have an in-depth knowledge of both the model and the company's applications. He or she must oversee this process of integration, as well as upgrade the model through a process of continuous learning.

Data, your most precious resource

AI algorithms are already available to most companies. Soon they will become an abundant commodity, as has already happened with other information systems. This is why data are the most precious resource of all. This has several implications.

First, you need to prepare your technological infrastructures to systematically and appropriately capture and store data in a central repository.

Second, you must rethink your strategy and make sure data are being used to inform your value proposition.

Finally, you must work on data governance. This will require new profiles and knowledge to understand the implications and scope of data management. In this way, you will be able to generate a virtuous circle and get the most out of AI.

This article is published in [IESE Business School Insight No. 151 \(Jan.-April 2019\)](#).

This content is exclusively for personal use. If you wish to use any of this material for academic or teaching purposes, please go to [IESE Publishing](#) where you can purchase a special PDF version of “[How prepared is your business to make the most of AI?](#)” (ART-3222-E) as well as the [full magazine](#) in which it appears, in English or in Spanish.



www.iese.edu/insight