Chapter 52
Business Applications of Deep Learning

Armando Vieira
Redoctopus, UK

ABSTRACT

Deep Learning (DL) took Artificial Intelligence (AI) by storm and has infiltrated into business at an unprecedented rate. Access to vast amounts of data extensive computational power and a new wave of efficient learning algorithms, helped Artificial Neural Networks to achieve state-of-the-art results in almost all AI challenges. DL is the cornerstone technology behind products for image recognition and video annotation, voice recognition, personal assistants, automated translation and autonomous vehicles. DL works similarly to the brain by extracting high-level, complex abstractions from data in a hierarchical and discriminative or generative way. The implications of DL supported AI in business is tremendous, shaking to the foundations many industries. In this chapter, I present the most significant algorithms and applications, including Natural Language Processing (NLP), image and video processing and finance.

1. INTRODUCTION

Artificial Intelligence (AI) is a relatively new area of research – that started in the 50's – marked by some successes and many failures. The initial enthusiasm originated in the materialization of the first electronic computer, soon fade away with the realization that most problems that the brain do in a blink are in fact very hard to solve by machines. These problems include, locomotion in uncontrolled environments, language translation, or voice and image recognition. Despite many attempts, it also become clear that the traditional approach to solve complex mathematical equations was insufficient to solve the most basic situations that a 2 years old toddler had no difficulty – like understand basic language concepts.

This lead to the so-called long "AI winter", where many researchers simply gave up of creating machines with human level cognitive capabilities – despite some successes in between, like the IBM machine Deep Blue that become the best Chess player in the world, or the application of neural networks for hand writing digits recognition in late 80's.

DOI: 10.4018/978-1-7998-0414-7.ch052

But AI is today one of the most exciting research fields with plenty of practical applications, from autonomous vehicles to drug discovery, robotics, language translation and games. Challenges that seemed insurmountable just a decade ago have been solved – sometimes with super-human accuracy - and are now present in products and ubiquitous applications, like voice recognition, navigation systems, facial emotion detection and even in art creation, like music and painting. For the first time, AI is leaving the research labs materializing in products that may seems to have emerged from Science Fiction movies.

How this revolution became possible in such a short period of time? What changed in recent years that puts us closer to the General AI dream? The answer is more a gradual improvement of algorithms and hardware than a single breakthrough. But certainly on the top of the list is Deep Neural Networks, or commonly referred Deep Learning (DL) (I. Goodfellow 2006).

Artificial Neural Networks is hardly a new field. They were around for about 50 years and got some practical recognition after mid 80's with the introduction of a method (Backpropagation) that allowed training of multiple layers neural networks. However the true birth on Deep Learning may be traced to the year of 2006, when Geoffrey Hinton (R. S. G. E. Hinton 2006) presented an algorithm to efficiently train deep neural networks in an unsupervised way – data without labels. They called them Deep Belief Networks, and consisting of staked Restrictive Boltzmann Machines. They differ from previous networks since they were generative models capable to learn the statistical properties of data being presented without any supervision.

Inspired by the depth structure of the brain, deep learning architectures have revolutionized the approach to data analysis (R. S. G. E. Hinton 2006) (I. Goodfellow 2006) (Schmidhuber 2015, Schmidhuber 2015). Deep Learning networks have won a paramount number of hard machine learning contests, from voice recognition, image classification, Natural Language Processing (NLP) to time-series prediction – sometimes by a large margin. Traditionally AI relied on heavily handcrafted features, for instance, to have decent results in image classification, several pre-processing techniques have to be applied, like filters, edge detection, etc. The beauty of DL is that most, if not all, features can be learned automatically from the data – provide enough (sometimes million) training data examples are available.

Being essentially non-supervised machines, deep neural architectures can be exponentially more efficient than shallow ones. The number of computational elements is only limited by the number of training samples – which can be of the order of billions. Deep models can be trained with hundreds of millions of weights and therefore tend to outperform shallow models such as SVMs. Moreover, theoreti-

Figure 1. DL has a higher capacity to learn from large amounts of data