Use of *Deep Learning* in Oversight

Luís André Dutra e Silva is a public servant at the Federal Court of Accounts in Brazil, Bachelor in Computer Science from Uniceub, with a certification in Software Engineering from IEEE and in Projects Management from Stanford University.

ABSTRACT

This article describes how the deep learning technique may be applied to actions of external control and the fight against corruption. Historic facts, which show the evolution of this technique, are presented in this article, as well as the functioning of the artificial and biological neural networks and a set of application examples.

Keywords: machine learning, deep learning, neural networks, fight against corruption, control actions, algorithms.

1. INTRODUCTION

Among the Machine Learning techniques (subject of the body of knowledge on Artificial Intelligence based on algorithms that use a large number of examples for the training of computer models), *Deep Learning* has been highlighted in recent years. A set of techniques that utilizes deep artificial neural networks with many intermediate layers between the input and the output layers (LECUN et al, 2015).

The technological edge of this approach is the excellent results obtained in problem solving, results that exceed even the performance of the best specialists in certain areas of knowledge such as the recognition of locations and semantic features in images;

victory in strategy games more complex than chess; and human beings overcoming verbal comprehension psychometric tests.

The following is a brief history of this technique and its application in oversight and the fight against corruption, as well as a description of the functioning of neural networks.

2. HISTORY

The connectionist models of automatic learning were initially implemented by the industry during the 1950s, due to the emergence of large computer systems. The first attempts to implement these models, however, were not successful, due to limitations of the processing power of computers at that time, coupled with the lack of theoretical foundation that supports the execution of the technique. Among the failed initiatives, we can mention the simulation of "electronic brains" by Nathanial Rochester of IBM's research labs.

Another factor that contributed to reduce interest of the scientific community in this issue was the publication in 1969 of an article by American scientist Marvin Minsky. The article showed that the perceptron, the most primitive artificial neuron, would be inappropriate to reproduce all the basic logic operations like, for example, the "exclusive or" logic function essential for solving Boolean equations.

Despite the difficulties experienced, in the mid-1970s, with the improved processing power of computers, a new algorithm called *backpropagation* (WERBOS, 1974), which demonstrated the possibility of connectionist models to faithfully reproduce all logic operations performed by the human mind, depending only on the number of artificial neurons used and the number of layers designed for specific purposes.

Parallel to the efforts of scientists and engineers to mimic the biological mechanisms of intelligence, in the late 1970s, neuroscientist Vernon Benjamin Mountcastle reached a fundamental discovery in this matter. He demonstrated that the neocortex has a single learning algorithm, which is repeated in all regions of the brain in columnar structures (MOUNTCASTLE, 1978).

Despite the availability of neural network algorithms during the 1980s, only in the mid-1990s came the first research using this technique applied to the analysis of corporate fraud. Kurt Fanning (1995) showed that self-organizing neural networks might be used to predict frauds in the financial reporting by companies.

During the following decade, there were significant advances in low cost parallel computing devices, a fact that enabled the training of highly complex models in infinitesimally less time than previously available by sequential means. In 2004, Jeff Hawkins, a former engineer at Intel and founder of Palm Inc., argued that the structures discovered by Mountcastle function as small pattern recognizers, whichcan be interconnected in order to learn any concept and even to perform predictions and generalizations about something not experienced (HAWKINS, 2004). In addition, in this decade, new forms of boot parameters caused great enthusiasm among proponents of *Deep Learning*, a technique seen before only as limited to data storage and unable to generalize its predictions.

After 2010, numerous scientific papers were written demonstrating the applicability of neural models that assist large organizations in their corporate governance actions through fraud detection in financial transactions.

A few years later, in 2012, Kurzweil reinforced Hawkins' theory. He has shown that connectionist models, when combined with statistical models, detect temporal patterns and implement the central idea of the organization in hierarchical layers, necessary to pattern recognizers similar to those discovered by Mountcastle.

For the purpose of the prediction of irregular activities, the training of artificial models based on their biological counterparts most often occurs without supervision. Thus, the active neurons, which represent the learned concepts, are determined competitively and cannot be directly translated into something that is already known, but need the help of experts to determine the emerging meanings of such a learning process.

Therefore, its most direct applications are in the online form of anomaly detection and in the development of concepts that would traditionally be based on the Artificial Intelligence methods, available on a large scale since the 1960s, to combat corporate fraud, and not in Deep Learning techniques based on a large number of actual and continuous examples.

More recently, in 2014, there were random connection disposal methods, which inserted the noise required for the non-occurrence of simple memorization, allowing the construction of models with exponentially higher than the previous performance, in terms of generalization of behavior (HINTON et al, 2014). The motivation of this connection disposal method is the process that occurs in sexual reproduction, wherein the couple's genes are combined with a small random mutation in the genetic code transmitted by heredity, this being the more efficient means in the evolution of more advanced organisms.

3. OPERATION

In 1981, Mountcastle received the Nobel Prize for Medicine for his discovery (considered the Rosetta Stone of neuroscience) that the fundamental structure of the cerebral neocortex is a minicolumn with about a hundred neurons arranged in six distinct layers and that the human brain has about 100 million of these small structures throughout its

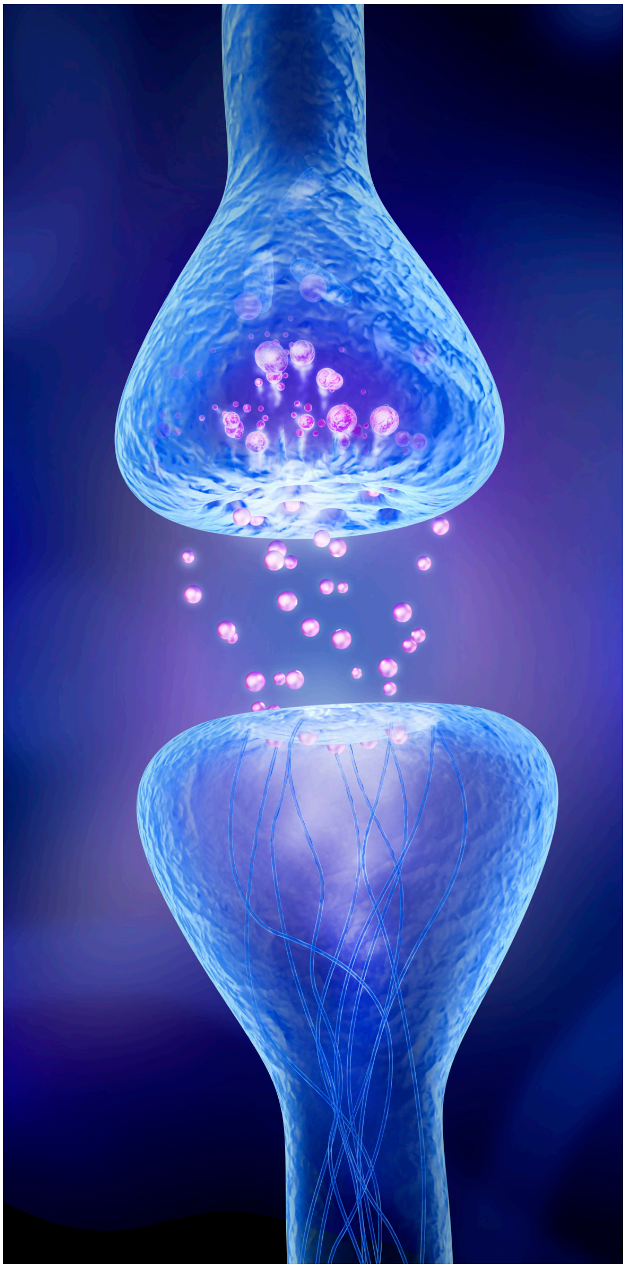

extent (MOUNTCASTLE, 1978). There are connections in both directions, which means, therefore, that the basic unit of our brain is a recurring and modularized network because each of the minicolumns connects to hundreds of other similar ones, forming a column of up to 70,000 neurons whose approximate size is comparable to that of a pinhead. The columns are grouped to form specialized tissues in a given cognitive function, and these tissues, in turn, can be connected to any other module, regardless of the distance that separates them, because there are axons that can go from one extreme to another in the brain.

According to Hawkins (2004), the lower layer of the neocortex has a much larger number of connections (synapses) with terminals originating from axons from other regions of the nervous system than connections that originate from this layer. In this manner, the level representation that the first layer offers is connected to temporal events from the senses. From the second layer, however, progressively persistent concepts and independent of temporal changes are formed. Finally, the neuron activation of the sixth layer occurs whenever patterns are present in the input connections and such activations persist for the duration of the exposure to these standards.

In accordance with a persistent cognitive function of representing concepts, the sixth layer has a much larger number of terminals coming out than the number of synapses connected to dendrites, which are connected to this layer. In addition, a very small number of sixth layer neurons is activated for a given concept and, empirically, it has been proven that always the same neurons are activated for any particular concept (ZADOR, 2000). This demonstration was carried out in a scientific experiment in which persons subject to monitoring of the brain by magnetic resonance imaging, upon displaying the picture of a specific person, always had the same activation of the sixth layer neurons.

Moreover, the simultaneous presence of different patterns in different regions of the neocortex in the first layer makes these patterns enhance the weights (synapses) that associate them. Thus, the neocortex functions as an associative memory in which a pattern or part of it, activates neurons which represent related standards. According to this associative architecture, the function of the sixth layer neurons connections to the others is to make predictions about which standards will succeed those which are present in the immediately inferior connections, reinforcing the sequence of predicted events that were learned by experience.

It is also important to consider the fact that the hierarchical structure of the layers of the visual cortex inspired the creation of convolutional neural networks (CNNs). This type of implementation has proven to be the most suitable for visual pattern recognition (LECUN, 1995). There is nevertheless a fundamental difference between CNNs and biological neural networks: CNNs are generally unidirectional (FFNs), while the natural neural networks are recurring (RNNs).

On another research front, related to an architecture closer to reality, the recurring artificial neural networks, as well as the neocortex, are stimulated primarily by temporal patterns (GRAVES, 2012). Thus, the sequence of successive stimuli allows them to be classified automatically so they can be associated to other sequences previously learned. It is a proven fact, in this case, that the need that the human mind has to learn sequential patterns. For example, the sequence of notes of a melody can be easily remembered. However, our mind can hardly recall the reverse sequence of musical notes.

This limitation does not occur on artificial neural networks, since recurring two-way models (BRNN) can learn any temporal or spatial sequence in both directions. Such models are proven to be more efficient in recognition sequences such as speech, writing and successive events that may be related to any irregularities or unlawful activity.

4. APPLICATIONS IN OVERSIGHT

As an example of an application in textual bases, a simple traditional neural network (*Multi Layer Perceptron*) with few intermediate layers is able to classify high precision types of decisions contained in judgments of the Court. This classification is fundamental for defining a context for the subsequent extraction of attributes (named entities) of the various deliberations that need to be continuously monitored by specialists.

Considering the other extreme, an application on photographic bases, in order to cover a large number of audited public works, convolutional neural networks can be used to monitor, using images obtained by remote sensing, the progress of project implementation. This type of application can perform the comparison of images related to the various stages of each project and indicate possible delays or technical non-conformities with the specifications.

When it comes to detecting anomalies applied to fraud discovery in agreements, for example, an artificial neural network, in unsupervised mode, may recognize abnormal situations in its execution after receiving as input thousands of normal situations that do not represent irregularities. Therefore, if there is a discrepancy between the realized resource flows and the expected execution of the projects, this method can indicate, with high probability, the occurrence of illicit activities.

In the case of treatment sequences, a recurrent network can be trained with the temporal series of price offers and attributes of objects of bids, learning to identify sequences that represent irregularities in the bidding process. This is possible thanks to the large number of examples of previous sequences, which were classified as irregular or not by experts. Therefore, this type of solution would contain the consolidated knowledge of many professionals over decades of experience.

Neural networks can also be used for routing and the classification of irregularities in the Special Rendering of Accounts processes through the recognition of textual and logical patterns in documents from different sources of unstructured data. This process consists in performing unsupervised training, carried out in a large quantity of documents with the purpose of finding semantic clusters that can later be associated with groups of irregularities in these processes. After the association of such groups to the meanings identified by experts, it is possible to create neural models supervised able to perform the same type of classification and routing of new documents, which were not part of the initial training.

5. CONCLUSION

Therefore, the use of Deep Learning solutions to aid in the fight against corruption can cause a gain of sufficient scale to cover a much greater number of cases of irregularities than what is possible to achieve today by a simple sampling related to the materiality of the resources involved. However, of course, the role of the experts of each area involved cannot be discarded, since the situations found automatically do not represent deterministic indications, but indications with a probability associated with the greater or lesser degree of certainty that represent important findings to be overseen by TCU and other oversight bodies.

REFERENCES

MINSKY, Marvin; PAPERT, Seymour. Perceptrons: An Introduction to Computational Geometry. Cambridge: The MIT Press, 1969.

WERBOS, Paul J. Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences. Cambridge MA: Harvard University Press, 1974.

MOUNTCASTLE, Vernon B. An Organizing Principle for Cerebral Function: The Unit Model and the Distributed System. Cambridge: MIT Press, 1978.

LECUN, Y; BENGIO, Y. Convolutional Networks for Images, Speech, and Time-Series. The Handbook of Brain Theory and Neural Networks. Cambridge: MIT Press, 1995.

FANNING, K.; COGGER, K.O.; SRIVASTAVA, R. Detection of Management Fraud: A Neural Network Approach. Nova York: IEEE, 1995

ZADOR, Anthony. Neural Representation and The Cortical Code. San Francisco: Annual Reviews, 2000.

HAWKINS, Jeff; BLAKESLEE, Sandra. On Intelligence. Nova York: Times Books, 2004.

LIVNAT, A; PAPADIMITRIOU, C.; PIPPENGER, N., FELDMAN, M.W. Sex, mixability, and modularity. Proceedings of the National Academy of Sciences. Redwood City: Highwire Press, 2010. Volume 107. p. 1452–1457.

KURZWEIL, Ray. How to Create a Mind: The Secret of Human Thought Revealed. Nova York: Viking Books, 2012.

GRAVES, Alex. Supervised Sequence Labelling with Recurrent Neural Networks. Nova York: Springer 2012.

HINTON, Geoffrey; SRIVASTAVA, Nitish; KRIZHEVSKY, Alex; SUTSKEVER, Ilya; SALAKHUTDINOV, Ruslan. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research. Brookline: Microtome Publishing, 2014. Volume 15. p. 1929-1958.

LECUN, Y; BENGIO, Y; HINTON, G. Deep learning. Nature. Nova York: Macmillan Publishers Limited, 2015. Volume 521. p. 436-444.